Information and Communication in Mechanism Design

A thesis submitted for the degree of
"Doctor of Philosophy"

by

Liad Blumrosen

Submitted to the Senate of the Hebrew University

September 2006

This work was carried out under the supervision of

Professor Noam Nisan

To my parents, Rivka and Avraham Blumrosen

Abstract

The emergence of the Internet was the main trigger for the exploration of problems in the intersection of computer science and economics. Existing tools from the two disciplines were found to be insufficient for the design and analysis of these new environments, and a joint analysis of these fields appeared to be necessary. Examples for such new environments include large electronic commerce arenas, computerized stock markets, web search engines, Peer-to-Peer systems and web-based social networks. Mechanism Design is a mathematical theory that puts the foundations for constructing protocols for environments with selfish players. A mechanism is a protocol that determines the output and the monetary payments according to the players' "messages" or "actions". The mechanism aims to optimize some system-wise goal, although the preferences of the players may be in conflict with this goal.

A considerable part of the mechanism-design literature relies on the "revelation principle", saying that we can restrict the attention to direct-revelation mechanisms where players simply report their true private secrets. However, such direct-revelation mechanisms are rare in practice, due to various behavioral, technical or regulatory reasons. This dissertation considers environments where the revelation principle cannot be applied. Instead, the participants are restricted to using some common, feasible or natural communication patterns. We study the power and limitations of each communication pattern, and study the loss incurred relative to environments with unconstrained information transmission. The information revealed in some of these patterns also supports the implementation of the optimal outcomes in an equilibrium, and the information is proved to be insufficient for this task for other patterns. The presented results therefore offer guidelines for designing complex decision-making systems that involve strategic agents, and classify the proposed models according to their potential power.

The first part of the dissertation studies environments where the expressiveness allowed for the players is severely limited – a very small number of actions is available for each player. The discussion starts with a comprehensive analysis of single-item auctions, and shows that nearly optimal results can be achieved even with a very small number of players' actions. We then give a full description of both the socially-efficient auctions and the revenue-maximizing auctions, and we present a tight analysis of the loss incurred by the informational restrictions. Similar questions are also explored for a more general framework, for which a characterization of the optimal solutions is presented for a wide family of mechanism-design settings and the applicability of these results is demonstrated in various models. One of the main results in this part of the dissertation characterizes a general setting where the best solution for the information-theoretic restriction can be obtained in an equilibrium without additional informational cost.

The second part of this dissertation considers the most prominent problem on the interface of economics and computer science – *combinatorial auctions*. In such auctions, multiple heterogeneous

items are for sale. Unlike the analysis of single-item auctions in the first part of this dissertation, a huge amount of information is required even for guaranteeing a reasonable approximation of the optimal outcome in combinatorial auctions. Most of the suggested combinatorial auctions try to circumvent this informational problem by designing iterative auctions, where the information revelation is only partial. In this dissertation, we embark on a systematic analysis of the power and limitations of iterative combinatorial auctions. Most existing iterative combinatorial auctions are based on repeatedly suggesting prices for bundles of items, and querying the bidders for their "demand" under these prices. We prove a large number of results showing the boundaries of what can be achieved by auctions of this kind. We first focus on the power of different kinds of the most popular way to implement combinatorial auctions – auctions where the prices are ascending over time: we show that ascending combinatorial auctions that do not use *both* bundle (non-linear) prices and personalized (non-anonymous) prices can not achieve social efficiency for general bidder valuations, or even a reasonable fraction of the optimal result. We also show settings when the information elicited by ascending auction can reveal the socially-efficient allocation, but it cannot reveal the prices of the "VCG" scheme that enables implementing the outcome in an equilibrium. We then prove several results regarding iterative auctions in which the prices are not necessarily ascending, but only use a polynomial number of such "demand queries": (1) that such auctions can simulate several other natural types of queries; (2) that they can approximate the optimal allocation as well as generally possible using polynomial communication or computation, while weaker types of queries can not do so; (3) that such auctions can solve the linear-programming relaxation of the winner-determination problem in combinatorial auctions.

Contents

1 Introduction **1**
1.1 Computational Economics . . . 1
1.2 Overview of the Results . . . 3
1.2.1 Part I: Mechanism Design with Restricted Action Spaces . . . 3
1.2.2 Part II: The Power of Iterative Combinatorial Auctions . . . 6
1.3 Presentation and Prerequisites . . . 8
1.3.1 Prerequisites . . . 8
1.3.2 Structure of Thesis . . . 9
1.3.3 Bibliographic Notes . . . 9

2 Preliminaries: Mechanism Design **10**
2.1 Classic Mechanism Design . . . 10
2.1.1 Mechanisms, Strategies and Equilibria . . . 11
2.1.2 Implementation . . . 12
2.1.3 The Revelation Principle and Truthfulness . . . 13
2.1.4 VCG Mechanisms . . . 13
2.2 Single-Item Auctions . . . 14
2.2.1 Socially-Efficient Auctions . . . 15
2.2.2 Revenue-Maximizing Auctions . . . 16
2.3 Combinatorial Auctions . . . 17
2.3.1 Computational Hardness and Communication Complexity . . . 18
2.3.2 Ascending Combinatorial Auctions . . . 21

3 Auctions with Severely Bounded Communication **27**
3.1 Introduction . . . 27
3.2 The Model . . . 31
3.2.1 The Bidders and the Mechanism . . . 31
3.2.2 Optimality Criteria . . . 32
3.3 Optimal Mechanisms for Two Bidders . . . 33
3.3.1 The Efficiency of Priority Games . . . 35
3.3.2 Efficient 2-bidder Mechanisms with k Possible Bids . . . 38
3.3.3 Profit-Optimal 2-bidder Mechanisms with k Possible Bids . . . 39
3.4 Optimal Mechanisms for n Bidders with Two Possible Bids . . . 41
3.4.1 The Characterization of the Optimal Mechanisms . . . 41
3.5 Asymptotic Analysis of the Welfare and Profit Losses . . . 44

3.5.1 Asymptotic Bounds on the Welfare Loss . . . 44
3.5.2 Asymptotic Bounds on the Profit Loss . . . 46
3.5.3 Asymptotic Bounds for a Growing Number of Bidders . . . 47
3.6 Sequential Auctions . . . 48
3.7 Future Work . . . 51

4 Implementation with a Restricted Action Space **52**
4.1 Introduction . . . 52
4.1.1 Our Framework . . . 53
4.1.2 Our Contribution . . . 54
4.2 Model and Preliminaries . . . 57
4.2.1 Dominant-Strategy Implementation . . . 57
4.2.2 Restrictions on the Action Space . . . 59
4.3 Implementation with a Limited Number of Actions . . . 60
4.4 Asymptotic Analysis of the Social-Value Loss . . . 63
4.5 Optimal Mechanisms for Two Players and Two Alternatives . . . 65
4.5.1 Diagonal Allocations . . . 66
4.5.2 Mutually Maximizer Threshold Strategies . . . 67
4.5.3 The Optimal 2-Action 2-Player Mechanisms . . . 68
4.6 Applications . . . 69
4.6.1 Application 1: Public Goods . . . 69
4.6.2 Application 2: Signaling . . . 71
4.6.3 Application 3: Message Delivery in Networks . . . 72

5 Informational Limitations of Ascending Combinatorial Auctions **74**
5.1 Introduction . . . 74
5.2 The Model . . . 77
5.3 Limitations of Item-Price Ascending Auctions . . . 78
5.4 Limitations of Anonymous Ascending Auctions . . . 82
5.5 Conclusions . . . 83

6 More on the Power of Ascending Combinatorial Auctions **84**
6.1 Introduction . . . 84
6.2 Ascending Item-Price Auctions for Substitutes Valuations . . . 85
6.2.1 Learning Substitute Valuations . . . 86
6.2.2 Computing VCG Prices for Substitutes Valuations . . . 88
6.3 Separation Results for Different Types of Ascending Auctions . . . 89
6.4 Some Positive Results on the Power of Ascending Auctions . . . 90

7 On the Power of Iterative Auctions: Demand Queries **92**
7.1 Introduction . . . 92
7.1.1 Extant Work . . . 93
7.2 A Survey of Our Results . . . 94
7.2.1 Comparison of Query Types . . . 94
7.2.2 Welfare Approximation . . . 95
7.2.3 Demand Queries and Linear Programs . . . 96

7.2.4 Representing Bundle Prices . . . 97
7.3 The Model . . . 98
7.3.1 Iterative Auctions . . . 98
7.3.2 Demand Queries . . . 99
7.4 The Power of Different Types of Queries . . . 99
7.5 Approximating the Social Welfare with Value and Demand Queries . . . 100
7.6 Demand Queries and Linear Programming . . . 103
7.7 The Representation of Bundle Prices . . . 104

8 Conclusions **106**

A Preliminaries: Mechanism Design **108**
A.1 Combinatorial Auctions - Missing Proofs . . . 108

B Auctions with Severely Bounded Communication **109**
B.1 Efficient Mechanisms: Missing Proofs . . . 109
B.2 Optimal Symmetric 1-bit Mechanisms . . . 111
B.3 Asymptotic Analysis - Missing Proofs . . . 112
B.4 More Asymptotic Results . . . 113

C Implementation with a Restricted Action Space **116**
C.1 Missing Proofs from Section 4.3 . . . 116
C.2 Missing Proofs from Section 4.5 . . . 118

D Informational Limitations of Ascending Combinatorial Auctions **121**
D.1 Critical Price Levels . . . 121
D.2 Limitations of Item-Price Ascending Auctions . . . 122
D.3 Limitations of Anonymous Ascending Auctions . . . 126
D.4 Constructing Mutually-Intersecting Partitions . . . 126

E More on the Power of Ascending Combinatorial Auctions **128**
E.1 Separation Results for Different Types of Ascending Auctions . . . 128
E.2 More Positive Results on Ascending Auctions . . . 132

F On the Power of Iterative Auctions: Demand Queries **134**
F.1 Simulating Queries by Demand Queries . . . 134
F.2 Missing Proofs . . . 135

Bibliography **136**

Chapter 1

Introduction

1.1 Computational Economics

In recent years, the field of computer science has expanded towards other areas, creating new interdisciplinary fields like computational biology and quantum computing. One of the fascinating emerging disciplines lies in the intersection of computer science and social sciences, and especially in the interface with economic theory – *computational economics*.

The research that was traditionally carried out in computer science and the social sciences differs in many fundamental aspects: the questions asked, the methodologies that are used, the objectives of the research and in the terminology. Here we list few notable differences:

- **(Normative vs. informative approach.)** The main motivation behind the classic economic theory is to explain phenomena that occur in existing real-world scenarios by defining the right models or by an empirical study. Computer-science theory has an engineering vision at its foundation: how can we design computers or protocols, and to what extent these protocols can be improved.

- **(Self-interested vs. controlled parties.)** The participants in most economic settings have preferences that may be in conflict with the objective function of the designer. Therefore, any analysis of such environments must pay attention to the incentives of these players that will affect their behavior. Classic computer science, on the other hand, studied the behavior of a single computer or a private network where the protocol designer has full control over the behavior of the different components of the system.

- **(What is a reasonable solution?)** The main effort in economics is in overcoming the incentives issues and try predicting and explaining the behavior of the players. If a "good" solution is proved to exist, then the common assumption is that the "market forces" will converge to this solution. However, computing or converging to such solutions may be computationally intractable, making them less relevant in practice. Rather, computer scientists will attempt to obtain approximate solutions for the problems, and these types of solutions are rare in the economic research.

- **(Asymmetric information vs. common knowledge.)** The complexity of many economic or game-theoretic environments stems from the fact that different players have different levels of knowledge about the state of the world. In particular, the players may hold private information that is hidden from the other participants. The designed protocols in such systems cannot access the relevant input, and efforts should be exerted in order to extract the relevant information from the players. The computer-science literature mostly assumes perfect knowledge of the system's description, and unclarities are usually attributed to other reasons (e.g., time uncertainties in online algorithms). Additionally, computational environments raise motivations for new objectives and measures that usually do not have a natural economic meaning. One example is minimizing the "makespan" (the load on the busiest machine) in scheduling models.

- **(Worst-case vs. average case analysis.)** Economic theory usually tries to model the average and the common behavior and therefore assumes a Bayesian model with statistical priors over the possible states of the world. Computational issues were mostly studied from the opposite approach and studied a worst-case viewpoint for these questions, possibly trying to avoid unwanted extremal behavior of the protocols.

Despite these differences, some remarkable research has been conducted in recent years in the border of the two disciplines. This research was triggered by the rapid advances in information technology, that allowed the emergence of new computational environments that are composed of self-interested parties. These environments became central for the interactions between individual and organizations, and the most notable platform for such interactions is the World Wide Web. An enormous rapidly-growing amount of trade is now carried out electronically, e.g., in electronic commerce transactions, in computerized stock exchanges and for selling online advertisements. The Internet also provides the infrastructure to other prominent interactions like file-sharing systems and social communities. It seems that the existing tools and methodologies are insufficient for designing and analyzing these new environments. One cannot separately solve the economic and the computational problems since there are tradeoffs between the economic and the computational properties. Researchers must have deep mathematical understanding of both aspects in order to be able to build systems that will perform well in practice. This dissertation belongs to this line of research that tries to contribute to a theory that will lay the foundation for designing and understanding such systems in the future.

The young computational-economics literature has already defined several novel and interesting contributions. We now briefly mention few of the prominent concepts. This dissertation belongs the field of *Algorithmic Mechanism Design* [116] that tries to interleave computational considerations with th classic "Mechanism Design" theory that considers the design of economic mechanisms like auctions and voting systems. A paradigmatic problem in this field concerns complex resource-allocation problems is the *combinatorial auction* problem, and it models markets for selling multiple non-identical items (see Chapter 2.3 for a survey on combinatorial auctions). Another important line of research studied the inefficiency of equilibrium solutions (also known as the *"price of anarchy"*) in games (see the survey [132]), that is, the fraction of the optimal solution that may be captured when the players act selfishly. Much attention was given to such questions in networks, trying to analyze the results of *selfish routing* of packets in large-scale networks. Another major question has been extensively studied in recent years concerns the complexity of computing equilibria, or converging to equilibria, in various settings (e.g., [55, 43] and the references within), also in general-equilibrium models (e.g., [46]). Several recent papers studied methods for selling *digital goods* ([64]), goods with a zero marginal production costs like MP3 songs or digital images, and "online" settings where the algorithms should be "competitive" relative to the optimal decisions, although they do not have any information in advance about the order in which bidders arrive (e.g., [90, 71, 8]). Finally, a recent exciting line of research considers selling advertisements in search engines (e.g., [15, 100, 119]); in such environments there is an intriguing interaction between individuals that search for information, companies that ask to advertise their products or services, and the search engines themselves (like Google, Yahoo and MSN) that try to maximize revenue and compete with the other engines.

1.2 Overview of the Results

Every economic interaction is based on some sort of communication between the agents. The communication protocol is especially important when some collective decisions should be made that depends on the players' private information. Many patterns of communication can be encountered in such settings, like raising the hand in an auction, publish prices in a menu, sending packets of information over the Internet, announce a take-it-or-leave-it offer, or even burning a pile of cash in front of the other players.

From a basic theoretic point of view, it seems that there is no need for sophisticated means of communication in the design of economic mechanisms. The celebrated *revelation principle* argues that every result that is achievable in equilibrium among rational agent can essentially be obtained by *truthful* mechanisms in which the players sincerely report their private data to the mechanism, and the mechanism performs any manipulation of that data on behalf of the players. In practice, however, such direct revelation mechanism are uncommon due to several reasons. First, it is sometimes infeasible to exactly reveal the exact data of the players; the information may be too large to transmit or to be figured out by the players, or the interface that the players are using may have capacity limitations. Second, mechanisms with indirect revelation may have some desired properties, like simplicity and intuitiveness, that may encourage participation and rational behavior. In addition, the expressiveness of the different parties may be asymmetric; for example, a seller my publish a price (a "take-it-or-leave-it offer") and the only two actions available to the buyer are to buy the good or not. Finally, players are usually reluctant to reveal their true private data and would prefer to participate in mechanisms where they can reveal only the relevant data.

This dissertation centers on several natural settings where indirect revelation of the private data is reasonable. We try to describe the tradeoff between the expressiveness that is allowed for each player and the results that can be obtained. This is done by presenting hardness results and positive results for each setting. One central goal of this dissertation is to analyze the informational requirements for obtaining the desired results in an equilibrium, and whether they differ from the information needed for realizing the optimal solution in the non-strategic setting. In other words, does handling the players' incentives requires additional amount of information over the information needed for merely determining the desired outcome?

1.2.1 Part I: Mechanism Design with Restricted Action Spaces

In classic mechanism-design settings, each player has some privately-known data – his *type* – and, according to this data, the player chooses an *action* that reveals some information to the mechanism. The mechanism designer determines the output and the monetary payments given the players' actions, and aims to implement some global goal that depends on the players' types. For example, in a simple auction model, the type of each bidder is the amount of money he is willing to pay for the item, and the mechanism designer might attempt to allocate the item to the player with the highest value.

Most of the mechanism-design literature hides an implicit assumption that the direct revelation of the private data is possible, that is, that the action space is isomorphic to the type space. As mentioned, this assumption does not hold in many settings due to various reasons. One example is in the well known "signaling" model [146], where workers send signals about their productivity levels by acquiring education. In this example, one may expect to find a small number of discrete education levels, although the type space for each player may have infinite size. Similarly, consumers usually have a multitude of different types but firms only advertise and sell a small number of possible "packages" of products and services (like in selling insurance policies and cellular phones deals). The restrictions on the action space can also come from technical reasons. For instance, implementing auctions over large networks must use very limited bandwidth.

The first part of dissertation studies such restrictions on the action space. We first present a thorough analysis of these questions in single-item auctions, and then we extend the discussion to a wider family of mechanism-design models where the players have one-dimensional types. Informally, that means that the private value of each player is a single scalar the possible values of which can be linearly ordered in terms of preference.[1] We study a Bayesian model where the players' types are independently distributed according to commonly known priors, and we measure the quality of the mechanisms on average, and compare them to the results that could be obtained without any restrictions on the action space.

Single-Item Auctions

Consider an auctioneer that has a single item for sale among a set of bidders. Each bidder i gains some secret value v_i from receiving the item, and he aims to maximize his (quasi-linear) utility: v_i-p_i, where p_i is his payment for the item (the utility is zero when the player loses). Unlike standard models, we assume that each bidder only has k possible actions. A strategy for each player would specify how he selects an action according to his value. For example, if there are two possible actions "0" and "1" (i.e., $k = 2$), one possible strategy is "*bid 0 when the value is smaller than 1/3, and 1 otherwise*".

Two of the most reasonable and natural objectives for the social planner are *welfare* (or social efficiency) maximization, when the auctioneer aims to allocate the item to the bidder with the highest value, and *revenue* maximization, where the goal is to maximize the expected monetary payment that the seller receives. We would like to implement these objectives when the players use dominant strategies – this is a strong equilibrium concept requiring that the strategy of each player always chooses the best action for this player, regardless of the actions of the other players. Without restrictions on the action space these problems are essentially solved. Vickrey [147] characterized the *second-price auction* that maximizes the social welfare in dominant strategies. Myerson [108] characterized revenue-maximizing auctions, and showed that this optimal result can be achieved in dominant strategies under a technical regularity condition on the distribution functions.

In this dissertation we try to optimize the above two objectives, but with an action space that may be severely restricted. The bottom line of our results is that a very small number of actions can be sufficient for having nearly optimal results, and we describe the mechanisms that achieve the optimal results. Another surprising result is that the optimal mechanisms must discriminate

[1] We provide a general definition of environments with one-dimensional types in Chapter 4.

between the players, even though the players are ex-ante symmetric. Following are our main findings:

- *Welfare maximizing auctions:* we present a full characterization of the socially-efficient games both for the case of 2 bidders and k actions and for n bidders and 2 actions. For arbitrary numbers of bidders and actions (i.e., any n bidders and k actions), we construct mechanisms that achieve *asymptotically* optimal results - mechanisms with a welfare loss of $O(\frac{1}{k^2})$ that is proved to be a tight upper bound. The characterization of *the* optimal mechanisms for arbitrary number of bidders and actions remains an open question.

- *Revenue-maximizing auctions:* we fully characterize the revenue-maximizing 2-bidder mechanisms and 2-action mechanisms, and give a tight asymptotic analysis of the optimal revenue loss and describe mechanisms that obtain this optimal loss. This is done via a reduction of the revenue-maximizing problem to the welfare-maximization problem.

- *Sequential auctions:* we show that if the players are allowed to send parts of their messages sequentially, then better results can be obtained. However, we show that the improvement is not overwhelming – the transmitted information may be decreased only by a linear factor.

General Mechanism Design Models

We also consider a more general model where the social planner has to choose an abstract "alternative". Each alternative gains the social planner with a "social value", and he attempts to choose an alternative that maximizes this social value. The social value, however, may be a function of the players' types, and therefore the social planner should also motivate the players to reveal information on their types. In the single-item auction case, the possible alternatives are "Bidder 1 wins" with a corresponding social value v_1 and the second alternative is "Bidder 2 wins" with the social value v_2. Indeed, a welfare-maximizing auctioneer will choose the alternative that maximizes the social value. Another example is "public-good" project, where players have a benefit v_i from using a public good (e.g., a bridge) and the social planner would like to construct this public good only if the sum of benefits exceeds the construction cost c. The first alternative here is "build the bridge" with a social value of $v_1 + v_2 - c$, and the second alternative is "do not build" with a social value of 0.

Most of our results are proved for the wide family of models where the social-value functions are *multilinear* in the players types, that is, they are polynomials where the degree of each variable in each monomial is at most 1. Both the single-item auction example and the public-good example are multilinear, but note that the objective functions in this model are not restricted to welfare maximization (i.e., the goal is not necessarily to maximize the sum of the players values). We illustrate a non-welfare-maximization application in a model for routing messages in networks. We study a general mechanism-design model that captures and extends the properties of most of the existing models with one-dimensional types. The preferences of the players should hold a "single-crossing" property and the social-value functions should be compatible with these preferences for having a dominant-strategy equilibrium. Although similar models have been described in the literature, either explicitly or implicitly, we believe that our description of this model has value of its own, especially for handling different preferences over discrete alternatives.

We consider this general model when the action space of the players is limited. We first notice that the problem of characterizing the optimal mechanism is actually composed of two questions.

The first question, is how to overcome the information-theoretic problems that arise due to the limited expressiveness of the bidders. That is, what is the best way to use the allowed action space in order to achieve good expected results. The second question involves game-theoretic considerations, as we want to implement the desired results in an equilibrium. An interesting question is whether the additional requirement for an equilibrium degrades the performance of the system, and to what extent. Our first result states that for a wide family of environments – with multilinear social-value functions – implementing the information-theoretic optimum in dominant strategies requires no additional communication at all!

We then asymptotically analyze the loss in social value as a function of the number of possible actions, and show that the rate of $O(\frac{1}{k^2})$ holds for general multi-linear settings. Still, this is the best possible general upper bound.

We also give a full description of the optimal 2-player mechanisms, for any number of possible actions, and show that some of the properties of the optimal auctions do not extend to the general case. For example, the optimal mechanisms are not always symmetric, and, surprisingly, in some settings they do not exploit all the actions that are available to to the players.

Finally, we illustrate applications of our general result to signaling games, public-good models and routing in faulty networks.

1.2.2 Part II: The Power of Iterative Combinatorial Auctions

In combinatorial auctions, a set of heterogeneous indivisible goods is for sale. Each bidder may have a different value for every combination of these items. The goal is to partition the items among the bidders such that the social welfare (the total value of the bidders for the bundles they receive) is maximized. Combinatorial auctions abstract many important resource allocation problems. Examples include allocating truckload transportation, airport slots, radio spectrum, bus routes and various industrial procurement environments. The reader is referred to Chapter 2 Section 2.3 for a survey on combinatorial auctions, and to the recent books on this topic [103, 37]. Combinatorial auction design involves severe strategic and computational problems. One significant problem is the communication problem: Given that the private data that each bidder holds may be composed of exponential number of values, how can the auctioneer elicit information from the bidders in order to find the optimal, or nearly optimal, allocation?

Many *iterative* auctions have been suggested in the literature to overcome this problem (see the survey in [122]): the bidders do not disclose their exact preferences, but only partially report the relevant data to the auctioneer. Most of the designed mechanism in theory and in practice use prices for this task in the following way: at each stage of the auction, bidders are presented with a set of prices, and choose their demand - their most desired bundle under the published prices. Such scenario is called in the literature a "demand query". This part of this dissertation centers on the family of auctions that use demand queries. We have special interest in the popular family of auctions *ascending auctions*. In such auctions, the published prices can only be raised at each stage. The most prominent application that uses ascending auctions are radio-spectrum auctions for wireless communication conducted in the US [57, 125] and all over the world (see survey in [40]). Ascending auctions are also used for selling railroad tracks [32], airport slots [39] and, of course, in e-commerce web-sites like Ebay and Amazon [52, 1]. There are several reasons for the popularity of ascending auctions (see [38]). One reason is their intuitiveness, that may encourage them to participate in the auctions, behave in a rational way and have more trust in the procedure. Also, ascending auctions have the desired property that bidders are only required to reveal partial information about their preferences. In addition, ascending auctions may also help reducing the "winner's curse" [104], and are believed to raise more revenue for the seller ("if they are willing to pay so much for this, it may worth it"). See Cramton [38] for a survey on the advantages and disadvantages of ascending auctions.

The main difference between the various models of ascending combinatorial auctions is in their pricing methods. The pricing method is subject to an interesting debate among the planners of the FCC spectrum auctions in the US (see, e.g., [57]). Some models present only *item prices (linear prices)*, where the price of each bundle is the sum of the prices of its items. Others choose to use *bundle prices (non-linear prices)* in which each bundle is allowed to have a price of its own. In some auctions, the same *anonymous* prices are presented to all bidders, while in auctions that use *non-anonymous* (personalized) prices, each bidder is presented with a personalized set of prices.

The Informational Power of Ascending Auctions

In this dissertation we present a systematic analysis of ascending combinatorial auctions. We provide several strong negative results that characterize the boundaries of what can be achieved by different types of ascending auctions. We study which types of ascending auctions can determine the socially optimal allocation of the items, and which ascending auctions can also calculate equilibrium prices.

We study a general model, where the auctioneer may use all the information gathered during the auction. Unlike most of the existing work, we are not focusing on reaching particular types of equilibria at the final stage of the auctions. Our hardness results are robust in other senses as well: they do not depend on any computational assumptions on the bidders or the auctioneer, and they also do not depend on the strategic behavior of the bidders or on how "rational" they are. The results only analyze the information that may be elicited during the auction.

Following are our two main results on the power of ascending auctions. The first results says that no *item-price* ascending auction can determine the optimal allocation for any profile of bidder valuations. The second result shows that no *anonymous*-price ascending auction can perform this task. These results cast doubts on the applicability of several auction designs that either use item prices or anonymous prices (e.g., [86],[125] and [149]) and justifies the added complexity added in the auctions of Parkes and Ungar [124] and Ausubel and Milgrom [4]. These results solve two open problems from [20] and [21].

The above two negative are proved using simple constructions of combinatorial auctions with only two bidders. We then present extensions to these results, which require using more sophisticated combinatorial constructions.

We first show that item-price ascending auctions and anonymous ascending auctions cannot even approximate the optimal social welfare to a better fraction than $O(\frac{1}{\sqrt{m}})$, where m is the number of items in the auction. We even show that an exponential number of item-price ascending auctions (i.e., separate ascending trajectories of prices) may be needed in order to find the optimal allocation.

We also consider a well-studied sub-class of valuations - *substitutes* (or gross-substitutes) valuations. For any profile of bidders with substitutes valuations, the optimal allocation can be determined by a simple item-price ascending auction ([82, 45, 70]). We show that slight deviations from substitutes preferences imply that this property no longer exists: for any profile of substitutes bidders we can add a single bidder such that no ascending item-price auction can compute the socially-efficient allocation. While the efficient allocation can be computed by an ascending auction for substitutes valuations, [70] showed that Vickrey-Clarke-Groves (VCG) prices cannot be calculated by anonymous ascending auctions. We strengthen this result by showing that VCG prices cannot be found even by non-anonymous ascending auction that uses n different ascending-price trajectories (one per each bidder), solving an open question from [44]. We prove another interesting property of substitutes valuations, and show that although item-price ascending auctions can determine the optimal allocation for any number of bidders, they cannot exactly learn the valuation function of a single bidder. We actually prove the above properties for a subclass of substitutes valuations - valuations that can be defined as an aggregation of bidders that each one interested in a single item (denoted as OXS valuations in [93]), and prove that such valuations do have a succinct representation.

The Power if Iterative Auctions: demand queries

Several auctions recently proposed in the literature use demand queries whose prices are not ascending over time (e.g., [12, 47, 58, 91]). The complexity of answering such a demand query depends on the bidder's internal representation of his valuation. For some internal representations this may be computationally intractable, while for others it may be computationally trivial. It does seem though that in many realistic situations the bidders will not really have an explicit internal representation, but rather 'know" their valuation only in the sense of being able to answer such queries.

We give a thorough analysis of what can be achieved using a polynomial number of demand queries. We first present the first algorithm to achieve the (best-possible unless $P = NP$) $O(\sqrt{m})$ approximation. This is a deterministic algorithm that uses demand queries, but, unfortunately, is not incentive compatible. (Recent papers present *randomized* incentive-compatible approximation mechanisms with the same approximation ratio [91, 48].) We then compare the power of demand queries to several natural or well-studied types of queries, and show how we can simulate such queries using demand queries. Finally, we formally prove a fact that was pointed out in [117] that a polynomial number of demand queries enables solving the linear-programming relaxation of the winner determination problem in combinatorial auctions. This is a surprising fact, since such linear programs have an exponential number of variables and the "oracle" that is required by the algorithm turns out to be exactly a demand query. This result has been extensively used recently for the design of algorithms for combinatorial auctions (e.g.,[47, 58, 59, 91, 50]).

1.3 Presentation and Prerequisites

1.3.1 Prerequisites

This dissertation is self contained and should be accessible for computer scientists, economists and researchers in related fields. In particular, the presentation does not require previous knowledge in game theory or mechanism design, for which a detailed survey is given in Chapter 2. Previous knowledge of game theory may shed more light on the context of the results in this thesis, and the reader is referred to [118, 61] for background on game theory. The reader should be familiar, at least at the concept level, with elementary notions in theory of computing, like basic asymptotic analysis of functions (see, e.g., [35]), NP-completeness (e.g., [120]) and Linear programming (e.g., [80]). I made an effort to make this dissertation comfortable to all relevant audiences, and thus I tried to avoid the usage of terms with double meaning, especially in computer science and economics (like “efficiency”, “competitive”,”non determinism” etc.). I apologize for all the places in the text that I missed, and where such terms could be ambiguously interpreted.

1.3.2 Structure of Thesis

Chapter 2 surveys existing results in mechanisms design; it begins in Section 2.1 with a description of the foundations of mechanism design in a general setting, and then lists main results in auction theory, starting from single-item auctions (Section 2.2) and then discussing multi-item (“combinatorial”) auctions (Section 2.3) with an emphasis on iterative combinatorial auctions. This survey on mechanism design does not attempt to cover all the vast literature in this field, but rather tries to draw the background for the work presented in this thesis. For more details on mechanism design and auction theory see, e.g., [98, 84, 118, 79, 115].

The body of this dissertation divides into two parts. The first part concerns mechanism-design settings where the number of actions (or the “amount of communication”) is restricted for each player. We first give a comprehensive study of such restrictions on single-item auctions in Chapter 3, and a more general setting is discussed in Chapter 4. In the second part of the dissertation we systematically study the power of different natural families of iterative combinatorial auctions. Chapter 5 measures the quality of the information that can be elicited by ascending-price auctions, and Chapter 6 further studies ascending auctions, taking into account additional factors like incentives issues. Chapter 7 considers iterative price-based auctions whose prices are not necessarily ascending. Finally, some conclusions are given in Chapter 8.

All the chapters are self contained, and can be accessed separately (except for Chapter 6 that uses definitions from Chapter 5).

1.3.3 Bibliographic Notes

Chapter 3 is based on joint work with Noam Nisan and Ilya Segal [26, 30, 31]. Chapter 4 contains joint work with Michal Feldman [25]. Chapter 5 is joint work with Noam Nisan that appeared in [28, 29], and Chapter 6 is also joint work with Noam Nisan that partially appeared in [28] and is partially unpublished. Finally, Chapter 7 is joint work with Noam Nisan that appeared in [27]. Other work that I have done and published during my PhD studies includes papers with Moshe Babaioff [9, 10] and with Shahar Dobzinski [23].

Chapter 2

Preliminaries: Mechanism Design

Mechanism Design is a sub-field of game theory that studies how to design decision rules that aggregate the information held by multiple parties in environments where the players' private information and private actions are not publicly observable. In the absence of the input to the decision rule, and since the parties are self-interested, one has to design mechanisms that elicit the information from the participants by motivating them to follow the "rules" of the mechanism. The designed mechanism forms a *game*, and the players in this game are expected to act *rationally* in order to maximize their own benefit. If the mechanism designer can somehow predict the behavior of the players in the game, then it may be able to choose the mechanism that optimizes the collective objective function. Auctions are probably the most notable application of mechanism-design theory, and parts of the classic *auction theory* and some more recent results, both in economics and computer science, will be surveyed in Sections 2.2 and 2.3. Among the other applications of mechanism design we mention the design of voting systems, allocating public goods and contract theory.

2.1 Classic Mechanism Design

The mechanism-design literature is divided into two main sub-disciplines: with or without *money*. Mechanism without money refers to environments where the players preferences can take unrestricted forms. These preferences are usually represented by some *orderings* on the set of possible outcomes, rather than a quantitative measure of the players' payoffs. Such models are used, for instance, for modeling different voting environments - when protocols are needed for aggregating the preferences of the participants. Unfortunately, strong negative results are known for such models, showing that the concept of aggregation is not well defined (e.g., the Condorcet Paradox), that only trivial aggregation methods possess basic requirements (Arrow's impossibility theorem [2]), and that all mechanisms, except trivial ones, can be manipulated by the players (the Gibbard-Satterthwaite Theorem [63, 140]).

This thesis concentrates on "quasi-linear" environments where we have some special type of commodity – "money" – that has a constant marginal utility for all players, and that performs as medium for exchanging utilities between players. Such environments actually assume that the players have unlimited budgets, and ignore any "income effect" of trades (i.e., changes in the preferences of players as a consequence of changes in the total amount of money they can spend).

2.1.1 Mechanisms, Strategies and Equilibria

A mechanism is a game with incomplete information among a set of players N, where each player holds a privately known type θ_i drawn from the set Θ_i. The social planner that designs the mechanism needs to choose an alternative from a set of *alternatives* denoted by $\mathcal{A}$. Each player i has some benefit from each chosen alternative $A \in \mathcal{A}$, that may depend on his type: $v_i(\theta_i, A)$.[1] The social planner cannot observe the types of the players, and he receives information on these types from the players' actions.

Definition 2.1. *A mechanism is composed of:*

- A_i – *a set of possible actions for each player.*
- $a : \times_{i=1}^{n} A_i \to \mathcal{A}$ – *a decision function that outputs an alternative according to the players' actions.*
- $p : \times_{i=1}^{n} A_i \to \mathbb{R}^n$ – *a payment scheme that defines the payment for the players according to their actions ($p_i(\cdot)$ denotes the payment for player i).*

The rules of the mechanisms are common knowledge, and each player observes these rules and decide about a strategy. The *strategy* of each player decides how the player selects his action for every possible type, i.e., it is a function $s_i : \Theta_i \to A_i$. The *utility* of a player from the mechanism, when the players use the strategies $s_1, ..., s_n$, is quasi-linear:

$$u_i(\theta_i, s_1, ..., s_n) = v_i\left(\theta_i, a(s_1(\theta_1), ..., s_n(\theta_n))\right) - p_i(s_1(\theta_1), ..., s_n(\theta_n))$$

In Bayesian models, where the types are drawn from a commonly known distribution, the utility $u_i(\theta_1, s_1, ..., s_n)$ for player i will denote player i's *expected* utility over the types of the other players.

A strategy is a *dominant strategy*, if it is the best plan for determining the actions in the mechanism, regardless of the behavior of the other players. We use a standard notation to denote a vector when one item is excluded. For example, we denote the strategies of all the players except player i $(s_1, ..., s_{i-1}, s_{i+1}, ..., s_n)$ by s_{-i}.

Definition 2.2. *A strategy s_i is* dominant[2] *in a mechanism if for every profile of strategies of the other players s_{-i} and every other strategy s'_i we have, $u_i(\theta_i, s_i, s_{-i}) \geq u_i(\theta_i, s'_i, s_{-i})$.*

A key solution concept in the mechanism-design literature is the *dominant-strategy equilibrium.*

Definition 2.3. *The profile of strategies $s_1, ..., s_n$ is a* dominant-strategy equilibrium*, if for every player i, the strategy s_i is a dominant strategy.*

When a dominant-strategy equilibrium exists, it is indeed reasonable to expect bidders to play according to it. This is a strong equilibrium concept that does not assume any statistical prior distributions, and requires no coordination among the players nor assumptions of one player on the information held by the others. The prominent caveats of this equilibrium is that it does not exist

[1] In more general models the values of the players may also depend on the types of the other players ("inter-dependent types"). We will not treat such models in this dissertation (except an implicit treatment in Chapter 4).

[2] This notion is often called in the literature a *weakly* dominant strategy, since it is defined by weak inequalities.

in many settings, and that it may lead to socially-inefficient outcomes (the "Prisoners Dilemma" is one well-known example).

The following important solution concept appears in environments where the types of the players are distributed by a joint distribution function f which is known both to the auctioneer and to the players. In a *Bayesian-Nash* equilibrium, no player has a profitable deviation from his strategy, assuming that the other players are not deviating as well.

Definition 2.4. *A strategy s_i is a* best response *for player i to the profile of strategies s_{-i}, if no other strategy can strictly increase the (expected) payoff of player i, i.e., for any other strategy s_i' we have $u_i(\theta_i, s_i, s_{-i}) \geq u_i(\theta_i, s_i', s_{-i})$.*

A profile of strategies $(s_1, ..., s_n)$ is called a Bayesian-Nash Equilibrium *if for every player i, the strategy s_i is a best response to s_{-i}.*

One can think of Bayesian-Nash equilibria as Nash equilibria ([110]) where the possible actions for the players are actually functions – their strategies. Definition 2.4 defines a "pure" Nash equilibrium, in the since that the players are not allowed to randomize over strategies. Note that even if we allowed "mixed" strategies, Bayesian-Nash equilibria do not necessarily exist in general, since the action space and the payoff space must satisfy some convexity assumptions. More details can be found, e.g., in [98].

2.1.2 Implementation

The mechanism aims to satisfy a "collective" decision that depends on the types of the players – a *social-choice function*

Definition 2.5. *A social-choice function is a function $c : \Theta_1 \times ... \times \Theta_n \to \mathcal{A}$.*

If a mechanism always outputs, in equilibrium, the same alternative as specified in some social-choice function c, we say that this mechanism "implements" c. The "quality" of the implementation depends on the type of equilibrium by which c is implemented. This dissertation will focus on two types of implementation: dominant-strategy implementation and Bayesian-Nash implementation.

Definition 2.6. *Let $M = (\{A_i\}_{i \in N}, a(\cdot), p(\cdot))$ be a mechanism and let c be a social-choice function. We say the M implements c with dominant strategies (resp. Bayesian-Nash equilibrium) if there exists a dominant-strategy equilibrium (resp. Bayesian-Nash equilibrium) $(s_1, ..., s_n)$ such that the decisions of the mechanism are always compatible with c, that is, for every $\theta \in \Theta_1 \times ... \times \Theta_n$,*

$$a(s_1(\theta_1), ..., s_n(\theta_n)) = c(\theta_1, ..., \theta_n)$$

One important well-studied social-choice function is the one that maximizes the *social welfare*. The welfare-maximizing social-choice function chooses, for every profile of types, the alternative A that maximizes the total value gained by the players, i.e., $c(\theta) \in argmax_{B \in \mathcal{A}} \sum_{i=1}^{n} v_i(B, \theta)$. Intuitively, the social welfare measures how "content" is the whole society from choosing the alternative A. This function is not sensitive, for example, to whether the division of the welfare is "fair" among the players, or to the number of players the receive positive values. The alternative that maximizes the social welfare is often called the *socially-efficient* outcome, or the *Pareto-efficient* outcome or simply the *efficient* outcome.

One desired property from any equilibrium is that no player is forced to participate in the game. This property is known as *individual rationality.* A profile of strategies in a mechanism is said to hold *interim* individual rationality if the strategy of each player, after observing his type, gains him an higher expected utility than his utility from not participating in the game. A profile of strategies is *ex-post* individually-rational, if for every player, and for every type of this player, the player always gains a higher utility than his non-participation utility, regardless of the behavior of the other players.

2.1.3 The Revelation Principle and Truthfulness

The "holy grail" of mechanism-design theory is to design mechanisms that achieve good results in an equilibrium – *incentive-compatible* mechanisms. One key observation is that the attention can be restricted to *truthful* mechanisms, where the action space of the players is their type space ("direct-revelation" mechanisms) and reporting the true values is an equilibrium strategy. Truthfulness can be defined for Bayesian-Nash equilibria, for the stronger concept of dominant strategies or for other equilibrium concepts.

Definition 2.7. *A direct-revelation mechanism is called* **truthful** *in dominant strategies (resp. Bayesian-Nash equilibrium), if reporting the true types is a dominant strategy (resp. Bayesian-Nash equilibrium) for every player.*

Proposition 2.1. (The Revelation Principle) *Consider a mechanism where all the players act simultaneously, and that implements some social-choice function c in dominant strategies (resp. Bayesian-Nash equilibrium), then there exists a* truthful *mechanism that implements c in dominant strategies (resp. Bayesian-Nash equilibrium).*

The revelation principle is based on an easy observation: if the players can apply transformations on their types and send the result to the mechanism, then we should be able to design a mechanism where the players report their actual types and the same transformations are performed inside the mechanism. Given that a deviation from the equilibrium was not profitable in the original mechanism, it would clearly not be profitable in the direct-revelation mechanism since both mechanisms implement the same social-choice function.

This principle is fundamental in Mechanism-design theory, and appeared in various sorts in [108, 66, 42, 63, 109]. One of the main themes of this dissertation is that in many settings we cannot apply the revelation principle. In such settings, there are some exogenous restrictions on the mechanism that preclude the usage of the type space as an action space. Most of this dissertation tries to overcome situations where this principle does not hold, and tries to characterize indirect mechanisms that achieve good results.

2.1.4 VCG Mechanisms

The main positive result in Mechanism Design claims that welfare-maximizing social choice functions can *always* be implemented in dominant strategies. This can be done using the celebrated family of VCG mechanisms due to Vickrey, Clarke and Groves ([147, 34, 67]). Such mechanism use a clever payment scheme for which the objective function of each player coincides with the social-welfare function, and therefore, any deviation from truthfulness may decrease both the social welfare and the player's utility.

Definition 2.8. *A VCG mechanism is a direct-revelation mechanism that, for every profile of reported types $\overline{\theta_1}, ..., \overline{\theta_n}$, outputs:*

1. *A socially-efficient alternative* A^**, i.e.,* $A^* \in argmax_{B \in \mathcal{A}} \sum_{i=1}^{n} v_i(\overline{\theta_i}, B)$*.*

2. *The following payment for each player* i*:*

$$p_i(\overline{\theta}) = -\sum_{i \neq j} v_i(\overline{\theta_i}, A^*) + h_i(\overline{\theta}_{-i}) \tag{2.1}$$

where $h(\cdot)$ *may be any function that is independent of* $\overline{\theta_i}$*.*

Definition 2.8 leaves one degree of freedom that concerns the functions h_i. One well-known payment rule, known as the *Clarke tax* or the *Clarke pivot rule*, suggests using the following h_i's:

$$h_i(\theta_{-i}) = \max_{A \in \mathcal{A}} \sum_{j \neq i} v_i(\theta_i, A)$$

That is, $h_i(\theta_{-i})$ is the maximal welfare achievable when player i does not participate in the game. This payment rule, together with Equation 2.1, results in an intuitive meaning to the VCG payment scheme: each player pays the loss to the "society" that is incurred by his participation in the mechanism. Using the Clarke pivot rule, the VCG mechanism also satisfies two desirable properties: *individual rationality* (see above) and "no positive transfers" – that the payments are always non-negative (no player is paid for participating).

This strong result raises the question whether we can implement in dominant strategies social-choice functions that are not welfare maximizing. This question if of great interest in computer-science settings, where the optimal result is often computationally hard, and therefore only an approximate solution can be reached in practice, and some other objective functions exist that are not "economically-oriented", e.g., minimizing the "makespan" in scheduling systems. Unfortunately, the VCG scheme no longer implies incentive compatibility for social-choice functions that are not welfare maximizing ([116]), and no other general scheme is known that ensures truthful implementation. In fact, a result by Roberts [127] proves that in settings where the type space of the bidders is unrestricted, only welfare maximization or a weighted variant of it ("affine maximizers") can be implemented in dominant strategies. Realistic environments usually have a restricted type space, and several recent results characterized conditions for dominant-strategy implementability in settings with restricted preferences (e.g., [89, 68, 17, 134]). If the type space is further restricted such that the type of each player is one-dimensional (e.g., a number drawn from some real interval), then there are various examples for other social-choice functions that can be implemented and the characterization of such functions is well understood (implementation in one-dimensional domains will be treated in details in Section 2.2 and in Chapters 4 and 3). In the range between one-dimensional types and unrestricted types, the set of implementable social-choice functions is still unclear. Examples for non-trivial incentive-compatible mechanisms that do not use the VCG scheme are very rare (see, e.g., [12]).

2.2 Single-Item Auctions

An auction for a single item is probably the most studied model in mechanism design. It is a well-defined, "clean" model that enabled the researchers to develop the most fascinating results in this area. In Section 4, we will generalize the framework and discuss general mechanism-design environments with one-dimensional private values.

A seller wishes to sell a single indivisible item among a set N of n bidders. The type of each player is its monetary benefit from receiving the item, and will be called the bidder's *value* and it is denoted by the real number v_i.

Most of the literature on auctions centers on the following two social criteria for the auctioneer: (1) Social efficiency (or social welfare) maximization - allocate the item to the bidder with the highest value. (2) Revenue maximization - maximize the (expected) payments to the auctioneer. The problems of characterizing the socially-efficient and the revenue-maximizing auctions are essentially solved, and the elegant solutions to the two problems are surveyed below.

2.2.1 Socially-Efficient Auctions

We say that an auction is *socially efficient* if there is a Bayesian-Nash equilibrium in this auction that always allocates the item to the player with the *true* highest value. It turns out that social efficiency can even be accomplished with dominant-strategy equilibrium, as we saw in the previous section that welfare maximization can always be truthfully implemented by using the VCG scheme. Applying the VCG framework to single-item auctions results in the well known *Vickrey auction*, or the *second-price* auction:

Definition 2.9. *In a* second-price auction *all the bidders simultaneously report their real-number bids, and the bidder with the highest bid wins, and pay the second-highest bid. All the other bidders pay zero.*

Theorem 2.1. *([147]) The second-price auction is socially efficient, truthful in dominant strategies and ex-post individually rational.*

Another well-studied auction is the *first-price* auction, where the bidder with the highest bid wins and pays exactly his bid. It is easy to see that the first-price auction is not truthful: a winning bidder has the incentive to lower his bid in order to decrease his payment. However, it turns out that the first-price auction is socially-efficient when we consider a Bayesian model.

Assume that each value v_i is independently distributed over real interval $[\underline{a}, \overline{b}]$. according to the (cumulative) distribution function F_i, and assume that a corresponding density function f_i exists for every bidder i. In this case, first-price auctions admit an efficient Bayesian-Nash equilibrium:

Theorem 2.2. *First price auctions are socially-efficient, Bayesian-Nash incentive compatible and ex-post individually rational. When all the values are distributed by the same distribution function F, the equilibrium strategy of each bidder is*

$$s(v) = v - \frac{\int_{\underline{a}}^{\overline{b}} F(x)^{n-1} dx}{F(v)^{n-1}}$$

For example, when the bidders' values are distributed uniformly on $[0, 1]$, then bidding according to the strategy $s(v) = \frac{n-1}{n} v$ is a Bayesian-Nash equilibrium.

Most people are more familiar with the iterative versions of first- and second-price auctions. The *English auction* starts with a zero price which is raised until only a single bidder demands the item, and the item is allocated to this bidder at the final price. In the private-value model, this auction is equivalent to a second price auction. The *Dutch auction* starts with a very high price that decreases until some bidder demands the item. This bidder wins the item and pays the final price. Dutch auctions are equivalent to first-price auctions.[3]

[3]Notice that in more complicated information models, where the valuation of a bidder may be influenced by

2.2.2 Revenue-Maximizing Auctions

Second price auctions and first price auctions do not necessarily maximize the seller's revenue in single-item auctions. We first define revenue-maximizing auctions (known in the literature as *optimal auctions*. It is clear that the optimal revenue should be achieved in an equilibrium, and we will consider all the results that can be obtained in Bayesian-Nash equilibria. This definition must also take into account individual-rationality constraints, otherwise, for example, we can charge each bidder any amount of money.

Definition 2.10. *An auction achieves a revenue x if there is an interim individually-rational Bayesian-Nash equilibrium in this auction that obtains an expected revenue of x. An auction achieves maximal revenue, if it achieves the highest revenue over all auctions.*

A seminal paper by Myerson [108] contains a neat characterization of revenue-maximizing auctions: for maximizing the revenue, the item should always be handed to the bidder with the highest *virtual valuation*:

Definition 2.11. *([108]) When the bidder's value is distributed according to a probability density function f and a cumulative function F, the following function is called the bidder's* virtual valuation*:*

$$\widetilde{v}(v) = v - \frac{1 - F(v)}{f(v)}$$

In this model, it is convenient to consider the seller as one of the player whose virtual valuation is constant and equals to his reservation value for the item. The *virtual surplus* is therefore defined as the virtual value of the player that receives the item (which may also be the seller, if he does not allocate the item). For example, when the bidders valuations are distributed uniformly on $[0, 1]$, a bidder with a valuation v has a virtual valuation of $\widetilde{v}(v) = 2v - 1$. Therefore, if we had 2 bidders with values of 1/4 and 3/4, then the "virtual" world would have 3 players with values of -0.5, 0.5 and 0 (assuming that the item is worth 0 to the seller). Note that if all the players had negative virtual valuations (that is, a value smaller than 0.5), then the seller would keep the item, actually implying a reservation price of 0.5 in the auction.

Myerson observed the surprising fact that, in equilibrium, the expected virtual surplus equals the expected revenue.

Theorem 2.3. *([108]) Consider a normalized model where losing bidders pay their lowest valuation ($\underline{a}$). Let h be a direct-revelation mechanism, which is Bayesian-Nash incentive-compatible and interim individually rational. Then, the expected revenue in h equals the expected virtual surplus.*

The characterization of revenue maximizing auctions is immediately implied by the theorem above – the item should be allocated to the player (the bidder or the seller) with the highest virtual valuation. It follows that revenue maximization can be reduced to the well-understood welfare-maximization problem by replacing the bidders' values with their virtual values.

One important issue that should be addresses, is whether allocating the item to the player with the highest virtual valuation will result in an equilibrium. It turns out that if the virtual functions

---

the values of the other bidders, the participants in English auctions gain information that is not available to them in second-price sealed bid auctions and therefore these auctions are no more equivalent. The equivalence between first-price auctions and Dutch auctions is maintained even in such models since no new information is revealed until the winner is announced.

are strictly increasing functions of the bidders' values, then this allocation rule can be implemented in dominant strategies. This is because the "regularity" condition implies that the auction has a monotone allocation: when a bidder increases her bid, she cannot turn from a winner into a loser. This monotonicity property is central in the analysis of dominant-strategy implementation for one-dimensional domains, and will be further discussed in Section 4.

Definition 2.12. *([108]) A probability density function f and its cumulative function F are called* regular*, if the virtual valuation $\widetilde{v}(v) = v - \frac{1-F(v)}{f(v)}$ is monotone, strictly increasing function of v.*

Theorem 2.4. *When the values of the players are distributed by the same regular function, then maximizing the virtual surplus (and thus, maximizing the expected revenue) can be implemented in dominant strategies.*

One corollary is that when the values are symmetrically distributed with a regular distribution, the second-price action also maximizes the expected revenue.

Finally, Theorem 2.3 also implies one of the most remarkable results of mechanism design - the *revenue equivalence theorem.* Since the virtual valuation depends only on the allocation decisions in the auction, and not on the auction's payment scheme, it follows that the revenue in auctions is also solely determined by the allocation decisions in the mechanism. Therefore, two auctions that always allocate the item to the same bidder will necessarily have, *in equilibrium*, the same expected seller's revenue! For example, second-price auctions and first-price auctions will have the same expected revenue: the first achieves the revenue in a dominant-strategy truthful equilibrium, and the latter is achieved in a Bayesian-Nash equilibrium where bidders bid a fraction of their value, e.g., $\frac{n}{n-1}v_i$ for the uniform distribution. The intuition here is that people will adjust their behavior according to the rules of the mechanism: they will report lower bids when the auctioneer charges them with "higher" payments.

Definition 2.13. (The Revenue Equivalence Theorem) *Consider an auction model where the bidders values are independently drawn from always positive density functions f_i. If two auctions have Bayesian-Nash equilibria that have the same allocation decision for each profile of values, and where losing players pay their lowest possible value, then the two auctions have the same expected revenue.*

2.3 Combinatorial Auctions

The combinatorial auction setting is formalized as follows: there is a set of m indivisible items that are concurrently auctioned among n bidders. For the rest of this chapter we will use n and m in this way. The combinatorial character of the auction comes from the fact that bidders have preferences regarding subsets – bundles – of items. Formally, every bidder i has a valuation function v_i that describes his preferences in monetary terms:

Definition 2.14. *A valuation v is a real-valued function that for each subset S of items, $v(S)$ gives the value that bidder i obtains if he receives this bundle of items. A valuation must have "free disposal", i.e., be monotone: for $S \subseteq T$ we have that $v(S) \leq v(T)$, and it should be "normalized": $v(\emptyset) = 0$.*

The whole point of defining a valuation function is that the value of a bundle of items need not be equal to the sum of the values of the items in it. Specifically for sets S and T, $S \cap T = \emptyset$, we say that S and T are *complements* to each other (in v) if $v(S \cup T) > v(S) + v(T)$, and we say that S and T are *substitutes* if $v(S \cup T) < v(S) + v(T)$. Note that implicit in this definition is the assumption that there are "no externalities", i.e. a bidder only cares about the item that he receives and not about how the other items are allocated among the other bidders.

Definition 2.15. *An* allocation *of the items among the bidders is $S_1 ... S_n$ where $S_i \cap S_j = \emptyset$ for every $i \neq j$. The* social welfare *obtained by an allocation is $\sum_i v_i(S_i)$. A socially-efficient allocation is an allocation with maximum social welfare among all allocations.*

In our usual setting the valuation function v_i of bidder i is private information – unknown to the auctioneer or to the other bidders. Our usual goal will be to design a mechanism that will find the socially optimal allocation. What we really desire is a mechanism where this is found in equilibrium, but we will also consider the partial goal of just finding the optimal allocation regardless of strategic behavior of the bidders. One may certainly also attempt designing combinatorial auctions that maximize the auctioneer's revenue, but much less is known about this goal.

Combinatorial auctions are an abstraction of the allocation problem that occurs in many real-life applications. "Spectrum auctions", held world wide and, in particular, in the united states, received most attention (see, e.g., [56, 40, 41]). Generally speaking, in such auctions, bidders desire licenses covering the geographic area that they wish to operate in, with sufficient bandwidth. Most of the spectrum auctions held so far escaped the full complexity of combinatorial nature of the auction, by essentially holding a separate auction for each item (but usually in a clever simultaneous way). In such a format, bidders could not fully express their preferences, thus leading, presumably, to sub-optimal allocation of the licenses. In the case of FCC auctions, it has thus been decided to move to a format that will allow "combinatorial bidding", but the details are still under debate. Another common application area is in transportation. In this setting the auction is often "reversed" – a procurement auction – where the auctioneer needs to *buy* the set of items from many bidding *suppliers*. A common scenario is a company that needs to buy transportation services for a large number of "routes" from various transportation providers (e.g. trucking or shipping companies). Nice surveys on industrial applications of combinatorial auctions can be found in [135, 16]. Several commercial companies are operating complex combinatorial auctions for transportation services, and commonly report savings of many millions of dollars. As a final example, we wish to mention auctions for paths between two specified nodes in networks. The items sold are the edges of the network, and the players have the different connection requests between nodes. This example demonstrates how various types of problems can be viewed as special cases of combinatorial auctions.

This section covers the topics in combinatorial auctions that are relvenant for the content of this dissertation and centers on theoretical results. The reader is referred to the recent surveys [103, 37, 114] that present a comprehensive treatment of various aspects of combinatorial auctions.

2.3.1 Computational Hardness and Communication Complexity

There are multiple difficulties that we need to address:

- Computational complexity – the allocation problem is computationally hard (NP-complete) even for simple special cases. How do we handle this?
- Representation and Communication – Specifying a valuation in a combinatorial auction of m items, requires providing a value for each of the possible $2^m - 1$ non-empty subsets. How

can we even represent such valuations? A naive representation would thus require $2^m - 1$ real numbers to represent each possible bid. It is clear that this would be completely impractical for more than about two or three dozen items. The computational complexity can be effectively handled for much larger auctions, and thus the representation problem seems to be the bottleneck in practice.

- Strategies – How can we analyze the strategic behavior of the bidders? Can we design for such strategic behavior?

The combination of these difficulties, and the subtle interplay between them is what gives this problem its generic flavor, in some sense encompassing many of the issues found in algorithmic mechanism design in general.

The algorithmic problem of computing the optimal welfare is NP-complete even for bidders with very simple valuation functions that only desire a single bundle of items ("single-minded" bidders). Even achieving a better approximation than $m^{\frac{1}{2}-\epsilon}$ is NP-hard ,for every $\epsilon > 0$ [136]. The hardness of approximation is based on the hardness of approximation of clique size of [73], with the strong version as stated appearing in [151]. This hardness results is tight fo single-minded bidders, where a simple greedy algorithm achieves an $O(\sqrt{m})$ approximation; The $\sqrt{m}$ lower bound is tight also for general valuations, as we show in this dissertation (Chapter 7) an $\sqrt{m}$-approximation algorithm. This algorithm is deterministic but not incentive-compatible; but later work described *randomized* algorithms that matches this lower bound for general valuation with incentive-compatible mechanisms [91, 48].

For general valuation, one may attempt to overcome the communication problem by designing iterative mechanisms that allow partial revelation of the bidders preferences. A strong result by Nisan and Segal [117] shows that obtaining the optimal solution for the winner determination problem, or even any approximation better than $\min\{n, m^{\frac{1}{2}-\epsilon}$ requires the transmission of exponential number of bits of information. This holds even when all the bidders' values are either 0 or 1.

One possible solution for the computational and the commincation problems is focusing on restricted classes of valuations. Indeed, for some sub classes of valuations the optimal allocation can be computed in polynomial time. One prominent example is the case of "susbtitute valuations" (see definition later in this section), where the integrality gap of the linear programming relaxation of this problem is proved to be zero [18] and therefore the optimal solution can be found using the Ellipsoid method (see [80]). Approximation algorithm are also known for the specific classes of submodular and subadditive valuations (e.g., [47, 59]). A multitude of algorithms, auctions and heuristics that either work under reasonbale assumption or work well in practice appear in the recent surveys on the winner determination problem [94, 107, 137].

Bidding Languages

This subsection concerns the issue of the *representation* of bids in combinatorial auctions. Namely, we are looking for direct representations of valuations that will allow bidders to simply encode their valuation and send it to the auctioneer. The reader is referred to [113] for a comprehensive survey on bidding languages for combinatorial auctions.

When attempting to choose or design a bidding language we are faced with the same types of trade-offs common to all language design tasks: *expressiveness vs. simplicity.* On one hand we would like our language to express important valuations well, and on the other hand we would like it to be as simple as possible. One would expect the goals of expressiveness and simplicity to be relatively conflicting, as the more expressive a language is, the harder it becomes to handle it. A well chosen bidding language should aim to strike a good balance between these two goals.

The common bidding languages construct their bids from *combinations* of simple *atomic bids*. The usual atoms in such schemes are the "single-minded bids": (S, p) meaning an offer of p monetary units for the bundle S of items. Formally, the valuation represented by (S, p) is one where $v(T) = p$ for every $T \supseteq S$, and $v(T) = 0$ for all other T. Intuitively, bids can be combined by simply offering them together using two possible semantics. One considers the bids as totally independent, allowing and subset of them to be fulfilled ("OR bids"), and the other considers them to be mutually exclusive and allows only one of them to be filled ("XOR bids").

An *OR bid* is composed from an arbitrary number of atomic bids, i.e., a collection of pairs (S_i, p_i), where each S_i is a subset of the items, and p_i is the maximum price that he is willing to pay for that subset. Implicit here is that he is willing to obtain any number of disjoint atomic bids for the sum of their respective prices. Thus, an OR bid is equivalent to a set of separate atomic bids from different bidders. More formally, for a valuation $v = (S_1, p_1)\, OR ... OR\, (S_k, p_k)$, the value of $v(S)$ is defined to be the maximum over all possible valid collections W, of the value of $\sum_{i \in W} p_i$, where W is valid if for all $i \neq j \in W$, $S_i \cap S_j = \emptyset$.

A *XOR bid* ([136]) contains an arbitrary number of pairs (S_i, p_i), where S_i is a subset of the items, and p_i is the maximum price that he is willing to pay for that subset. Implicit here is that he is willing to obtain at most one of these bids. More formally, for a valuation $v = (S_1, p_1)\, XOR ... XOR\, (S_k, p_k)$, the value of $v(S)$ is defined to be $max_{i | S_i \subseteq S}\ p_i$ (we may also use the sign $\oplus$ instead of "XOR").

The *size of a bid* denotes the number of atomic bids in it. The following propositions (due to Nisan [111]) provide some intuition about the representational power of OR bids and of XOR bids:

Proposition 2.2 ([111]). *XOR bids can represent all valuations. OR bids can represent all bids that don't have any substitutabilities, i.e., those where for all $S \cap T = \emptyset$, $v(S \cup T) \geq v(S) + v(T)$, and only them.*

Proposition 2.3 ([111]). *A valuation is called "additive" if $v(S) = \sum_{j \in S}\ v(\{j\})$ for all S. A valuation is called "unit demand" if $v(S) = max_{j \in S}\ v(\{j\})$ for all S.*

- *Any additive valuation can be represented by OR bids of size m but requires XOR bids of size $2^m - 1$ (as long as all items j get positive value).*
- *Any unit-demand valuation can be represented by a XOR bid of size m, but no unit-demand valuation can be represented at all by an OR bid (except trivial ones that only give positive value to a single item).*

The Query Model

The last subsection presented ways of encoding valuations in bidding languages as to enable the bidders to directly send their valuation to the bidder. In this subsection, and in the next subsection, we consider indirect ways of sending information about the valuation: *iterative auctions*. In these, the auction protocol repeatedly interacts with the different bidders, aiming to adaptively elicit enough information about the bidders' preferences as to be able to find a good (optimal or close to optimal) allocation. The idea is that the adaptivity of the interaction with the bidders may allow pinpointing the information that is relevant to the current auction and not requiring full disclosure of bidders' valuations. This may not only reduce the amount of information transferred and all associated complexities, but also to preserve some privacy about the valuations, only disclosing the information that is really required.

Such iterative auctions can be modeled by considering the bidders as "black-boxes", represented by oracles, where the auctioneer repeatedly queries these oracles. In such models, we should specify the types of queries that are allowed by the auctioneer. The auctioneer would be required to be computationally efficient in two senses: the number of queries made to the bidders and the internal computations. Efficiency would mean polynomial running time in m (the number of items) even though each valuation is represented by 2^m numbers. The running time should also be polynomial in n (the number of bidders), and in the number of bits of precision of the real numbers involved in the valuations.

Our first step is to define the types of queries which we allow our auctioneer to make to the bidders. Probably the most straightforward query one could imagine is where a bidder reports his value for a specific bundle:

Value query: *The auctioneer presents a bundle S, the bidder reports his value $v(S)$ for this bundle.*

It turns out that value queries are pretty weak and are not expressive enough in many settings. Another natural and widely-used type of queries is the *demand query*, in which a set of prices is presented to the bidder, and the bidder responds with his most valuable bundle under the published prices.

Demand query *(with item prices[4]): The auctioneer presents a vector of item prices $p_1 ... p_m$; the bidder reports a bundle in his demand set under these prices, i.e., some set S that maximizes $v(S) - \sum_{i \in S} p_i$.*

Both value queries and demand queries were extensively studied in the recent literature. A systematic analsyis of such queries, and comparing them to other natural types of queries is given in Section 7 if this thesis (and was published in [27]). Among other results, Section 7 shows that value queries are signifcantly weaker than demand queries, and one cannot approximate the optimal social welfare to a better ratio than $O(\frac{m}{\sqrt{\log m}})$ using a polynomial number of value queries. Yet, several positive results use value queries. For instance, a 2-approximation to the social welfare can be achieved for submodular valuations [93], and they can help to efficiently eliciting the preferences in restricted, but natural, settings [139, 150].

Demand queries are widely used in practice, mainly in the framework of ascending-price auctions (see below). As shown in Section 7, they allow to efficiently solve the linear-programming relaxation of the winner determination problem; several recent algorithms use this fact [91, 47, 50, 59]. More work on the power and limitations of demand queries appears in [12, 21, 117], and relations to machine-learning theory can be found in [21, 87] and in the references within.

2.3.2 Ascending Combinatorial Auctions

This section concerns a large class of combinatorial auction designs which contains the vast majority of implemented or suggested ones: ascending auctions. These are a subclass of iterative auctions with demand queries in which the prices can only increase. I.e., in this class of auctions, the auctioneer publishes prices, initially set to zero (or some other minimum prices), and the bidders

[4] In this section we consider demand queries where the auctioneer presents item prices, and the price of a bundle is the sum of the prices of the items it contains. In Section 2.3.2, we also consider demand queries where a different price per each bundle is allowed.

An Item-Price Ascending Auction for Substitutes Valuations:

Initialization:

For every item $j \in M$, set $p_j \leftarrow 0$, and let $S_1 \leftarrow M$ and for every $i \neq 1$ let $S_i \leftarrow \emptyset$.

Repeat:

For each bidder i, let D_i be an ϵ-demand at prices $p_1 ... p_n$ relative to S_i that contains S_i.

If for all i, $S_i = D_i$ exit the loop; else, find i with $S_i \neq D_i$ and update:

- For every item $j \in D_i \setminus S_i$, set $p_j \leftarrow p_j + \epsilon$
- $S_i \leftarrow D_i$ and for every bidder $k \neq i$, $S_k \leftarrow S_k \setminus D_i$

Finally: Output the allocation $S_1, ..., S_n$.

Figure 2.1: A socially-efficient ascending auction for substitutes bidders.

repeatedly respond to the current prices by bidding on their most desired bundle of goods under the current prices. The auctioneer then repeatedly updates the prices by increasing some of them in some manner, until a level of prices is reached where the auctioneer can declare an allocation. (Intuitively, prices related to over-demanded items are increased until the demand equals supply.) There are several reasons for the popularity of ascending auctions, including their intuitiveness, the fact that private information is only partially revealed, that it is clear that they will terminate, and that they may increase the seller's revenue in some settings (see, e.g., [38]).

We will center on two representative families of ascending auctions. One auction uses a simple price scheme (item prices), and guarantees economic efficiency for a restricted class of bidder valuations. This discussion is based on the work of [82, 45, 70]. The second family, based on [124, 4], is socially efficient for every profile of valuations, but uses a more complex pricing scheme - prices for bundles – extending the demand queries defined earlier.

Definition 2.16. (Item/Bundle prices) *An auction uses* item prices (or linear prices)*, if, at each stage, the auctioneer presents a price p_j for each item j, and the price of every set S is additive: $p(S) = \sum_{j \in S} p_j$. We say that an auction uses* bundle prices (or non-linear prices) *if each bundle S may have a different price $p(S)$ (and we allow having a bundle S where $p(S) \neq \sum_{j \in S} p_j$).*

Ascending Item-price Auctions

Figure 2.1 describes an auction that is very natural from an economic point of view: increase prices until supply equals demand. The auction starts with zero item prices, iteratively collects the demands of the bidders at current prices, and increases the prices of over-demanded items by ϵ. Intuitively, when no item is demanded by more than a single bidder we are close to a Walrasian equilibrium which is socially optimal.

Of course, we know that a Walrasian equilibrium does not always exist in a combinatorial auction, so this can not always be true. The problem is that it does not suffice that the auction terminates when no item is over-demanded, since a Walrasian equilibrium requires that no item is under-demanded. Unfortunately, this may indeed happen: increasing the price of one item may reduce the demand for another item that is complementary to it, in this case it will not be possible to find a (near)-demand set D_i that contains the current set of items S_i. The following definition captures a class of valuations in which this cannot happen.

Definition 2.17. *A valuation v_i satisfies the substitutes (or gross-substitutes) property if for every pair of item-price vectors $\overrightarrow{q} \geq \overrightarrow{p}$ (coordinate-wise comparison), and for every $D \in D_i(\overrightarrow{p})$, there* *exists a bundle $A \in D_i(\overrightarrow{q})$ such that for all $j \in D$ with $p_j = q_j$ we have that also $j \in A$.*

I.e., in a substitutes valuation, increasing the price of certain items can not reduce the demand for items whose price has not changed. Simple sub-cases of substitutes valuations are additive valuations, unit-demand valuations. With such valuations, the auction maintains the property that every item is demanded by some bidder. The auction terminates when all the bidders receive their demanded bundles, and consequently, the auction converges to a (nearly) Walrasian equilibrium. We first formalize the notion of approximation that we use in the proof (and in the algorithm).

Definition 2.18. *We say that D_i is an ϵ-demand of v_i at prices $p_1...p_m$ relative to an existing bundle S_i, if $D_i \in D(p'_1...p'_n)$ where $p'_j = p_j + \epsilon$ for $j \notin S_i$ and $p'_j = p_j$ for $j \in S_i$. An allocation $S_1...S_n$ and a price vector $p_1...p_m$ are an ϵ-Walrasian equilibrium if for each i, S_i is an ϵ-demand of bidder i relative to itself.*

Theorem 2.5. *For bidders with substitutes valuations, the auction described in Figure 2.1 ends with an ϵ-Walrasian equilibrium. In particular, the allocation obtained is within $n\epsilon$ from the optimal social welfare. The total running time is polynomial in n, m and v_{max}/ϵ, where v_{max} is the maximum valuation.*

Proof. The theorem will follow from the following key claim:

Claim 2.1. *At every stage of the auction, for every bidder i there indeed exists an ϵ-demand D_i that contains S_i.*

First notice that this claim is certainly true at the beginning. Now lets see what an update step for some bidder i_0 causes. Clearly for i_0, S_{i_0} is an ϵ-demand by definition. For i's with $S_i \cap S_{i_0} = \emptyset$ only prices of items outside of it have been increased, so the demand is unchanged; S_i is also unchanged for such i's, thus it is still contained in i's ϵ-demand. For other i's we are exactly in a situation that fits the definition of the substitutes property, as the prices of items in the new S_i have not been raised.

Once we have this claim, it is directly clear that all items are always demanded (at every stage and for every i, the items in S_i will be contained in some S_j at the subsequent stage). Since the auction terminates only when all $D_i = S_i$ we get an ϵ-Walrasian equilibrium. The fact that an ϵ-Walrasian equilibrium is close to socially optimal is obtained just as in the proof of social-efficiency of Walrasian equilibria (see, e.g., [69, 102]).

For the running time analysis it is clear that the price of each item can be raised at most v_{max}/ϵ times. Each stage is clearly polynomial time, except that we need to verify that we can find the required D_i (that contains S_i) just with a demand query. This will be immediate when the demand set contains only a single bundle. Otherwise, each bidder should break ties by reporting an ϵ-demand that contains S_i, and answer consistently (keep reporting the same bundle when the price level is unchanged). □

It is useful to view this auction as implementing a *primal-dual* algorithm (see [75] for a survey on primal-dual algorithms). The auction starts with a feasible solution to the dual linear program (here, zero prices), and as long as the complementary-slackness conditions are unsatisfied proceeds by improving the solution to the dual program (i.e., increasing prices of over-demanded items).

A corollary from Theorem 2.5 is that there always exists a Walrasian equilibrium for environments with substitutes bidders, and it can be revealed by a simple, feasible ascending auction. The obvious question now is whether this can be obtained for wider classes of valuations. Unfortunately, the answer is negative: A Walrasian equilibrium does not necessarily exist for any wider class of valuations (see, e.g., [69, 102]). [70] show that no item-price ascending auction can determine the VCG prices, regardless of whether it reaches an equilibrium or not. [5] showed that the VCG prices can be computed by $n+1$ separate ascending trajectories of item prices.

The above results are obtained given that the bidders truthfully reveal their demand sets at each stage. Is this a reasonable assumption? If the valuation functions exhibit complementarities, then bidders will clearly have strong incentives not to report their true preferences in such auctions, due to a problem known as the *exposure problem*: Bidders who bid for a complementary bundle (e.g., a pair of shoes), are exposed to the risk that part of the bundle (the left shoe) may be taken from them later, and they are left liable for the price of the rest of the bundle (the right shoe) that is worthless for them.

However, even for substitutes preferences the incentive issues are not solved. The prices in Walrasian equilibria are not necessarily VCG prices, and therefore truthful bidding is not an ex-post equilibrium. The strategic weakness of Walrasian equilibria is that bidders may have the incentive to demand smaller bundles of items (*"demand reduction"*), in order to lower their payments (see the discussion in [54, 6]). When we further restrict the class of substitutes valuations such that each bidder desires at most one item ("unit-demand" valuations), then it is known that the auction reaches the lowest possible Walrasian-equilibrium prices that are also VCG prices, and hence these auctions are ex-post Nash incentive compatible ([96, 69]).

Ascending Bundle-price Auctions

We saw that a simple notion of a competitive equilibrium that uses item prices only exists under very restrictive assumptions. We will now extend the concept of a competitive equilibrium and see that such equilibria always exist and they can be reached by ascending auctions. However, such equilibria may use an exponential number of prices – a distinct price per each bundle – and also use a more expressive pricing method that uses personalized prices (or *non-anonymous* prices). That is, personalized bundle prices specify a price $p_i(S)$ per each bidder i and every bundle S.

Definition 2.19. *Personalized bundle prices $\overrightarrow{p}$ and an allocation $S = (S_1, ..., S_n)$ are called a* competitive equilibrium *if:*

- *For every bidder i, $S_i \in D_i(\overrightarrow{p_i})$, i.e., for any bundle $T \subseteq M$, $v_i(S_i) - p_i(S_i) \geq v_i(T) - p_i(T)$.*
- *The allocation S maximizes the* seller's revenue *under the current prices, i.e., for any other allocation $(T_1, .., T_n)$, $\sum_{i=1}^{n} p_i(S_i) \geq \sum_{i=1}^{n} p_i(T_i)$.*

It is easy to see that with personalized bundle prices, competitive equilibria always exist. An efficient allocation and the prices $p_i(S) = v_i(S)$ is a trivial example. It is known that both bundle prices and personalized prices are necessary to guarantee the existence of a competitive equilibrium. Even this weak notion of equilibrium, however, guarantees optimal social welfare:

Proposition 2.4. *In any competitive equilibrium $(\overrightarrow{p}, S)$ the allocation is socially efficient.*

Proof. Let $(\overrightarrow{p}, S)$ be a competitive equilibrium, and consider some allocation $T = (T_1, ..., T_n)$. Since $S_i \in D_i(\overrightarrow{p_i})$ for every bidder i, we have that $v_i(S_i) - p_i(S_i) \geq v_i(T_i) - p_i(T_i)$. By summing over all the bidders we get:

$$\sum_{i=1}^{n} v_i(S_i) - \sum_{i=1}^{n} p_i(S_i) \geq \sum_{i=1}^{n} v_i(T_i) - \sum_{i=1}^{n} p_i(T_i)$$

Incremental Auctions:
Initialization: For every player i and bundle S, $p_i(S) \leftarrow 0$.
Repeat:
• Let $\Delta_i = D_i(\overrightarrow{p_i} + \epsilon)$, i.e., all the bundles S that maximize $v_i(S) - (p_i(S) + \epsilon)$.
• Let $T_1, ..., T_n$ be a revenue-maximizing allocation with $p_i(T_i) > 0$ for every i, i.e., for every allocation $\{Y_i\}_{i \in N}$ we have $\sum_{i=1}^{n} p_i(T_i) \geq \sum_{i=1}^{n} p_i(Y_i)$.
• Terminate when $\Delta_i = \{\emptyset\}$ for every losing bidder i (i.e., where $T_i = \emptyset$)
• For every losing bidder i and for every non-empty bundle $S \in \Delta_i$, $p_i(S) \leftarrow p_i(S) + \epsilon$.

Figure 2.2: Incremental auctions end up with the socially-efficient allocation for any profile of bidders.

Since $\sum_{i=1}^{n} p_i(S_i) \geq \sum_{i=1}^{n} p_i(T_i)$, the welfare in the allocation S exceeds the welfare in T. □

Several iterative auctions are designed to end up with competitive equilibria. Figure 2.2 describes a family of such ascending-price auctions, we call here *Incremental Auctions*. Such auctions collect the demand of the bidders at each price level, then the auctioneer computes a tentative allocation, and all the losing bidders raise their bids on their demanded bundles. The auction maintains the property that the surplus of each bidder from all the bundles he demanded during the auction is equal (up to an ϵ), and it terminates when this surplus reaches zero for all the losing players. Note that the communication volume transmitted in Incremental Auctions is infeasible when auctioning more then a few items, since an exponential number of prices may be published at each stage, and each bidder may transmit an exponential number of bundles during the auction.

Again, since our auction uses discrete price increments, it will only reach an approximate competitive equilibrium. Following are the approximation definitions that we use.

Definition 2.20. *A bundle S is an* ϵ-demand *for a player i under the bundle prices $\overrightarrow{p_i}$ if for any other bundle T, $v_i(S) - p_i(S) \geq v_i(T) - p_i(T) - \epsilon$. An* ϵ-competitive equilibrium *is similar to a competitive equilibrium (Definition 2.19), except each bidder receives an ϵ-demand under the equilibrium prices.*

Theorem 2.6. *For any profile of valuations, Incremental Auctions (described in Figure 2.2) end with an ϵ-competitive equilibrium. In particular, the allocation obtained is within $n\epsilon$ from the optimal social welfare.*

Proof. We will show that such auctions terminate at an ϵ-competitive equilibrium. Similar arguments as in Proposition 2.4 show that an ϵ-competitive equilibrium is close up to an additive factor of ϵn to the optimal social welfare.

At each step of the auction, at least one price will be raised. Since a bundle price will clearly never exceed its value, the auction will terminate. What is left to be shown is that, upon termination, each bidder receives an ϵ-demand.

Losing bidders will clearly receive their ϵ-demand, the empty set, since by definition their surplus from all the other bundles is at most ϵ.

For each winning bidder, the surplus from every two bundles that were demanded (in Δ_i) during the auction will differ by at most ϵ. To see that, note that once a bundle S is in Δ_i gaining a surplus of x, its price may be raised by ϵ turning the surplus to x-ϵ. S must be demanded again before any bundle with a surplus smaller than x-ϵ. This way, a gap of at most ϵ is maintained between the surplus from S and the surplus from every bundle previously demanded. A winning bidder will receive only bundles that were once demanded (with non-zero prices), and it follows that each winning bidder will receive an ϵ-demand. □

Incremental Auctions will not necessarily terminate at VCG prices, and therefore myopic bidding (i.e., telling your true demand set at each stage) is not an ex-post Nash equilibrium. Overall, when comparing the pricing schemes of the two socially-efficient auctions – VCG and Incremental auctions – no strict conclusion can be reached in favor of one of them. See [4, 103] for a detailed comparison of VCG auctions and Incremental auctions (also called "Proxy auctions"). The VCG auction has the obvious advantage that truthful bidding is an ex-post Nash equilibrium. This also implies that the strategic and the computational burdens on the bidders will be low, as they simply report their true preferences.

The payment scheme calculated by Incremental Auctions, on the other hand, circumvents many flaws that make VCG mechanism impractical. For example, the revenue gained in VCG auctions is not monotone, in the sense that the revenue may decrease as the set of bidders expands. Also, the bidders in VCG auctions may have incentives to impersonate to several shill bidders or to plan collusive deviations with other players. The bidding rules in Incremental Auctions are simple and intuitive. Yet, they cannot overcome the inherent communication-complexity hardness. They do not possess a truthful ex-post equilibrium, what makes analyzing the strategic behavior of the bidders difficult. One weak positive equilibrium property is achieved when each bidder is committed in advance to a particular valuation ("proxy bidding"). Then, the auctions do admit ex-post Nash equilibria but these equilibria require the participants to possess considerable knowledge of the preferences of the other bidders.

Anonymous Ascending Auction for Super-Additive Valuations

Another restircted setting where a simpler ascending auction can compute the socially-efficient allocation is the case of super-additive valuations. These are valuations that admit "no-substitutabilities", i.e., the value of a bundle is always at least the sum of the values of its disjoint parts.

Definition 2.21. *A valuation v is called* super-additive*, if for any bundles S, T such that $S \cap T = \emptyset$ we have that $v(S \cup T) \geq v(S) + v(T)$*

An anonymous bundle-price auction can terminate with an optimal allocation for every profile of super-additive bidders (the iBundle(2) auction in [124, 123]). The optimal allocation cannot be found by such auctions for general valuation, where more complex non-anonymous auctions should be used [124, 4]. We give here a simple argument showing that anonymous auctions can determine the optimal allocation for super-additive valuations. A proof is given in Appendix A.1.

Proposition 2.5 ([123])**.** *An anonymous bundle-price auction can determine the optimal allocation for every profile of super-additive bidders.*

Chapter 3

Auctions with Severely Bounded Communication

3.1 Introduction

Recent years have seen the emergence of the Internet as a platform of multifaceted economic interaction, from the technical level of computer communication, routing, storage, and computing, to the level of electronic commerce in its many forms. Studying such interactions raises new questions in economics that have to do with the necessity of taking computational considerations into account. This chapter deals with one such question: how to design auctions optimally when we are restricted to use a very small amount of communication.

This chapter studies the effect of severely restricting the amount of communication allowed in a single-item auction. Each bidder privately knows his real-valued willingness to pay for the item, but is only allowed to send k possible messages to the auctioneer, who must then allocate the item and determine the price on the basis of the messages received. (For example, a bidder may only be able to send t bits of information, in which case $k = 2^t$). The simplest case is $k = 2$, i.e., each bidder sends a single bit of information. This is in contrast to the usual auction design formulation, in which bidders communicate real numbers.

While communicating a real number may not seem excessively burdensome, there are several motivations for studying auctions with such severe restrictions on the communication. First, if auctions are to be used for allocating low-level computing resources, they should use only a very small amount of computational effort. For example, an auction for routing a single packet on the Internet must require very little communication overhead, certainly not a whole real number. Ideally, one would like to "waste" only a bit or two on the bidding information, perhaps "piggy-backing" on some unused bits in the packet header of existing networking protocols (such as IP or TCP). Second, the amount of communication also measures the extent of information revelation by the bidders. Usually, bidders will be reluctant to reveal their exact private data (see, e.g., [130]). This work studies the tradeoff between the amount of revealed data and the optimality of the auctions. We show that auctions can be close to optimal even using a single Yes/No question per each bidder. Our results can also be applied to various environments where there is a need for discretize the bidding procedure; One example is determining the optimal bid increment in English auctions (see [97]). Finally, a restriction on communication may sometimes be viewed as a proxy for other simplicity considerations, such as simple user interface or small number of possible payments to facilitate their electronic handling. Chapter 4 show that the ideas illustrated in this work extend to general mechanism-design frameworks where the requirement for a small number of "actions" per each player are natural and intuitive, and the reader is referred to Chapter 4 for examples and references.

We examine the effect of severe communication bounds on both the problem of maximizing social welfare and that of maximizing the seller's expected profits (the latter under the restrictions of Bayesian incentive-compatibility and interim individual rationality of the bidders, and under a standard regularity condition on the distribution of bidders' valuations). We study both simultaneous mechanisms, in which the bidder send their bids without observing any actions of the other bidders, and sequential mechanisms where messages may depend on previous messages. We find that single-item auctions may be very close to fully optimal despite the severe communication constraints. This is in contrast to *combinatorial* auctions, in which exact or even approximate efficiency is known to require an exponential amount of communication in the number of goods [117].[1]

Both for welfare maximization and revenue maximization, we show that the optimal 2-bidder auction takes the simple form of a ''priority game" in which the player with the highest bid wins, but ties are broken asymmetrically among the players (i.e., some players have a pre-defined priority over the others when they send the same message). We show how to derive the optimal values for the parameters of the priority game. These optimal mechanisms are asymmetric by definition, although the players are a priory identical. Furthermore, we show that for any number of players, as the allowed number of messages grows, the loss due to bounded communication is in order of $O(\frac{1}{k^2})$. The bound is tight for some distributions of valuations (e.g., for the uniform distribution). In addition, we consider the case in which the number of players grows while each player has exactly two possible messages. We show that priority games are optimal for this case as well, and we also characterize the parameters for the optimal mechanisms and show that they can be generated from a simple recursive formula. We offer an asymptotic bound on the welfare and profit losses due to bounded communication as the number of players grows (it is $O(\frac{1}{n})$ for the uniform distribution).

Our analysis implies some expected as well as some unexpected results:

- **Low welfare and profit loss:** Even severe bounds on communication result in only a mild loss of efficiency. We present mechanisms in which the welfare loss and the profit loss decrease exponentially in the number of the communication bits (and quadratically in the number k of the allowed bids). For example, with two bidders whose valuations are uniformly distributed on $[0, 1]$, the optimal 1-bit auction brings expected welfare 0.648, compared to the first-best expected welfare 0.667.

- **Asymmetry helps:** Asymmetric auctions are better than symmetric ones with the same communication bounds. For example, with two bidders whose valuations are uniformly distributed in $[0, 1]$, symmetric 1-bit auctions only achieve expected welfare of 0.625, compared to 0.648 for asymmetric ones. We prove that both welfare- and profit-maximizing auctions must be discriminatory in both allocation and payments.

- **Dominant-strategy incentive-compatibility is achieved at no additional cost:** The

[1] There have been several other studies considering various computational considerations in auction design: timing (e.g., [90, 129]), unbounded supply (e.g., [60, 65, 11]), computational complexity in combinatorial auctions (see survey in [37]) and more.

auctions we design have dominant-strategy equilibria and are ex-post individually rational[2], yet are optimal even without any incentive constraints (for welfare maximization), or among all Bayesian-Nash incentive-compatible and interim individually rational auctions (for profit maximization). This generalizes well-known results for the case without any communication constraints.

- **Bidding using "mutually-centered" thresholds is optimal:** We show that in the optimal auctions with k messages, bidders simply partition the range of valuations into k intervals ranges and announce their interval. In 2-bidder mechanisms, each threshold will have the interesting property of being the average value of the other bidder in the respective interval. We denote such threshold vectors as "mutually-centered".

- **Sequential mechanisms can do better, but only up to a linear factor:** Allowing players to send messages sequentially rather than simultaneously can achieve a higher payoff than in simultaneous mechanisms. However, the payoff in any such multi-round mechanism among n players can be achieved by a simultaneous mechanism in which the players send messages which are longer only by a factor of n. This result is surprising in light of the fact that in general the restriction to simultaneous communication can increase communication complexity exponentially.

Although the welfare-maximizing mechanisms are asymmetric, we show that symmetric mechanisms may also be close to optimal when the allowed communication grows: We show that as the number k of possible messages grows, while the number of players is fixed, the loss in optimal symmetric mechanisms converges to zero at the same rate as the loss in efficient "priority games". However, the loss in optimal priority games is still smaller by a factor. On the other hand, when we fix the number of messages, and let the number n of players grow, we show that the loss in optimal asymmetric mechanisms converges *asymptotically faster* to zero than in optimal symmetric mechanisms ($O(\frac{1}{n})$ compared to $O(\frac{logn}{n})$, for the uniform distribution).

We now demonstrate the properties above with an example for the simplest case: a 2-player mechanism where each player has two possible bids (i.e., 1 bit) and the values are distributed uniformly.

Example 3.1. *Consider two players, Alice and Bob, with values uniformly distributed between* $[0,1]$*. A 1-bit auction among these players can be described by a 2x2 matrix, where Alice chooses a row, and Bob chooses a column. Each entry of the matrix specifies the allocation and payments given a bids' combination. The mechanism is allowed to toss coins to determine the allocations. Figure 3.1 describe an example for such a mechanism, and denote this mechanism as* g_1.

A strategy *defines how a player determines his bid according to his private value. We first note that in* g_1*, both players have dominant strategies, i.e., strategies that are optimal regardless of the actions of the other players: consider the following threshold strategy: "*bid 1 if your valuation is greater than $\frac{1}{3}$, else bid 0*". Clearly, this strategy is dominant for Alice in* g_1*: when her valuation is smaller than* $\frac{1}{3}$ *she will gain a negative utility if she bids "1"; When her valuation is greater than* $\frac{1}{3}$*, bidding "0" gives her a utility of zero, but she can get positive utility by bidding "1". Similarly, a threshold strategy with the threshold* $\frac{2}{3}$ *is dominant for Bob.*

[2]A mechanism is ex-post individually rational if a player never pays more than her value. Interim individual rationality is a weaker property, in which a player will not pay more than his value *on average*. Individual rationality constraints are essential for the study of revenue maximization (otherwise, the potential revenue is unbounded).

The social welfare *in a mechanism measures the total happiness of the players from the allocation, or in our case, the value of the player that receives the item. Now, the expected welfare in* g_1*, given that the players follow their dominant strategies, is easily calculated to be* $\frac{35}{54} = 0.648$*: Both player will bid "0" with probability* $\frac{1}{3}\frac{2}{3}$*, where the expected welfare equals the expected value of Bob,* $\frac{1}{2} \cdot \frac{2}{3}$*. Similar computations show that the expected welfare is indeed:*

$$\frac{1}{3}\frac{2}{3}\frac{\left(\frac{2}{3}\right)}{2} + \frac{1}{3}(1-\frac{2}{3})\frac{\left(1+\frac{2}{3}\right)}{2} + (1-\frac{1}{3})\frac{2}{3}\frac{\left(1+\frac{1}{3}\right)}{2} + (1-\frac{1}{3})(1-\frac{2}{3})\frac{\left(1+\frac{2}{3}\right)}{2} = \frac{35}{54}$$

We see that despite restricting the communication from an infinite number of bits to a single bit only, a relatively small welfare loss of $\frac{1}{54}$ *was incurred. Of course, a random allocation that can be implemented without communication at all will result in an expected welfare of* $\frac{1}{2}$*, and this may be regarded as our naive benchmark.*

It turns out that the mechanism described in Figure 3.1 maximizes the expected welfare: no other 1-bit mechanism achieves strictly higher expected welfare with any pair of bidders' strategies (that is, regardless of the concept of equilibrium we use). We note that the optimal mechanism is asymmetric (a "priority game") – ties are always broken in favor of Bob, and that this mechanism is optimal even when we allow randomized decisions. Note that the optimal symmetric *1-bit mechanism uses randomization, but only achieves an expected welfare of* 0.625 *(the mechanism is illustrated in Appendix B.2 and see also Footnote 17).*

Finally, we note that the optimal thresholds of the players are "mutually centered". That is, Alice's value $\frac{1}{3}$ *is the average value of Bob when he bids 0 and Bob's value* $\frac{2}{3}$ *is the average value of Alice when she bids 1. The intuition is simple: given that Bob bids "0", his average value is* $\frac{1}{2} \cdot \frac{2}{3} = \frac{1}{3}$*. For which values of Alice an efficient mechanism should give her the item? Clearly when her value is greater than the average value of Bob. Therefore, Alice should use the threshold* $\frac{1}{3}$*.*

The most closely related studies in the economic literature are by Harstad and Rothkopf [97], which considers similar questions in cases of restricting bid levels in oral auctions to discrete levels, and Wilson [148] and McAfee [99] who analyze the inefficiency caused by discrete priority classes of buyers. In particular, Wilson shows that as the number k of priority classes grows, the efficiency loss is asymptotically proportional to $\frac{1}{k^2}$. While in [148] the buyers' aggregate demand is known while supply is uncertain, in our model the demand is uncertain. Both [148, 97] restrict attention to symmetric mechanisms, while we show that creating endogenous asymmetry among ex ante identical buyers is beneficial. Another related work is by Bergemann and Pesendorfer [13], where the seller can decide on the accuracy by which bidders know their private values. This problem is different than ours, since the bidders in our model know their valuations.

The organization of the chapter is as follows: Section 3.2 presents our model definition and introduces our notations and Section 3.3 presents a characterization of the welfare- and profit-optimal 2-player auctions. Section 3.4 characterizes optimal mechanisms with arbitrary number of bidders, but 2 possible bids for each player. In Section 3.5 we give an asymptotic analysis of the minimal welfare and profit losses in the optimal mechanisms. Finally, Section 3.6 compares simultaneous and sequential mechanisms with bounded communication.

B / A	0	1
0	B wins and pays 0	B wins and pays 0
1	A wins and pays $\frac{1}{3}$	B wins and pays $\frac{2}{3}$

Figure 3.1: (g_1) A 2-bidder 1-bit game that achieves maximal expected welfare (efficiency). For example, when Alice (the rows bidder) bids "1" and Bob bids "0", Alice wins and pays $\frac{1}{3}$

3.2 The Model

3.2.1 The Bidders and the Mechanism

We consider single item, sealed bid auctions among n risk-neutral players. Player i has a private valuation for the object $v_i \in [\underline{a}, \overline{b}]$.[3] The valuations are independently drawn from cumulative probability functions F_i. In some parts of our analysis[4], we assume the existence of an always-positive probability density function f_i. We will sometime treat the seller as one of the bidders, numbered 0. The seller has a constant valuation v_0 for the item. We consider a normalized model, i.e., bidders' valuations for not having the item are $\underline{a}$.

The novelty in our model, compared to the standard mechanism-design settings, is that each bidder i can send a message of $t_i = \lg(k_i)$ **bits** to the mechanism, i.e., player i can choose one of possible k_i **bids** (or messages). Denote the possible set of bids for bidder i as $\beta_i = \{0, 1, 2, ..., k_i - 1\}$. In each auction, bidder i chooses a bid $b_i \in \beta_i$. A mechanism should determine the allocation and payments given a vector of bids $b = (b_1, ..., b_n)$:

Definition 3.1. *A mechanism g is composed of a pair of functions (a, p) where:*

- $a : (\beta_1 \times ... \times \beta_n) \to [0,1]^{n+1}$ *is the allocation scheme (not necessarily deterministic). We denote the i'th coordinate of $a(b)$ by $a_i(b)$, which is bidder i's probability for winning the item when the bidders bid b. Clearly, $\forall i \; \forall b \; a_i(b) \geq 0$ and $\forall b \; \sum_{i=0}^{n} a_i(b) = 1$. If $a_0(b) > 0$, the seller will keep the item with a positive probability.*

- $p : (\beta_1 \times ... \times \beta_n) \to \Re^n$ *is the payment scheme. $p_i(b)$ is the payment of the ith bidder given a bids' vector b.*[5]

Definition 3.2. *In a mechanism with k-possible bids, for every bidder i, $|\beta_i| = k_i = k$. We denote the set of all the mechanisms with k-possible bids among n bidders by $G_{n,k}$. We denote the set of all the n-bidder mechanisms in which $|\beta_i| = k_i$ for each bidder i, by $G_{n,(k_1,...,k_n)}$.*

A *strategy* s_i for bidder i in a game $g \in G_{n,(k_1,...,k_n)}$ describes how a bidder determines his bid according to his valuation, i.e., it is a function $s_i : [\underline{a}, \overline{b}] \to \{0, 1, ..., k_i - 1\}$. Let s_{-i} denote the strategies of the bidders except i, i.e., $s_{-i} = (s_1, ..., s_{i-1}, s_{i+1}, ..., s_n)$. We sometimes use the notation $s = (s_i, s_{-i})$.

[3]For simplicity, we use the range $[0, 1]$ in some parts of the chapter. Using the general interval will be required, though, for the characterization of the optimal mechanisms, mainly due to the reduction we use for maximizing the revenue that translates the original support to another interval.

[4]That is, in the characterization of the optimal mechanisms in Sections 3.3.2 and 3.4 and when using the concept of *virtual valuation* in Sections 3.3.3 and 3.5.2

[5]Note that we allow non-deterministic allocations, but we ignore non-deterministic payments (since we are interested in expected values, using lottery for the payments has no effect on our results).

Definition 3.3. *A real vector* $(t_0, t_1, ..., t_k)$ *is a vector of threshold values if* $t_0 \leq t_1 \leq ... \leq t_k$.

Definition 3.4. *A strategy* s_i *is a threshold strategy based on a vector of threshold values* $(t_0, t_1, ..., t_k)$*, if for every bid* $j \in \{0, ..., k_i - 1\}$ *and for every valuation* $v_i \in [t_j, t_{j+1})$*, bidder* i *bids* j *when his valuation is* v_i*, i.e.,* $s_i(v_i) = j$ *(and for every* v_i*,* $v_i \in [t_0, t_k]$*). We say that* s_i *is a threshold strategy, if there exists a vector* t *of threshold values such that* s_i *is a threshold strategy based on* c.

3.2.2 Optimality Criteria

The bidders aim to maximize their (quasi-linear) utilities. The utility of bidder i is 0 when he loses (and pay nothing), and $v_i - p_i$ when he wins and pay p_i. Let $u_i(g, s)$ denote the expected utility of bidder i from a game g when the bidders use the vector of strategies s (implicit here is that this utility depends on the value v_i).

Definition 3.5. *A strategy* s_i *for bidder* i *is dominant in a mechanism* $g \in G_{n,(k_1,...,k_n)}$ *if regardless of the other bidders' strategies* s_{-i}*,* i *cannot increase his expected utility by a deviation to another strategy, i.e.,*

$$\forall \widetilde{s_i} \;\; \forall s_{-i} \;\; u_i(g, (s_i, s_{-i})) \geq u_i(g, (\widetilde{s_i}, s_{-i}))$$

Definition 3.6. *A profile of strategies* $s = (s_1, ..., s_n)$ *forms a Bayesian-Nash equilibrium (BNE) in a mechanism* $g \in G_{n,(k_1,...,k_n)}$*, if for every bidder* i*,* s_i *is the best response for the strategies* s_{-i} *of the other bidders, i.e.,*

$$\forall i \;\; \forall \widetilde{s_i} \;\; u_i(g, (s_i, s_{-i})) \geq u_i(g, (\widetilde{s_i}, s_{-i}))$$

We use standard participation constraints definitions: We say that a profile of strategies $s = (s_1, ..., s_n)$ is *ex-post individually rational* in a mechanism g, if every bidder never pays more than his actual valuation (for any realization of the valuations). We say that a strategies profile $s = (s_1, ..., s_n)$ is *interim Individually Rational* in a mechanism g if every bidder i achieves a non-negative *expected* utility, given any valuation he might have, when the other bidders play with s_{-i}.

Our goal is to find optimal, communication-bounded mechanisms. As the mechanism designers, we will try to optimize "social" criteria such as *welfare* (efficiency) and the seller's *profit*.

The *expected welfare* from a mechanism g, when bidders use the strategies s, is the expected social surplus. Because the item is indivisible, the social surplus is actually the valuation of the bidder who receives the item. If the seller keeps the item, the social welfare is v_0.

Definition 3.7. *Let* $w(g, s)$ *denote the expected welfare (or expected efficiency) in the* n*-bidder game* g *when the bidders' strategies are* s*, i.e., the expected value of the player (possibly the seller) who receives the item in* g*. Let* $w^{opt}_{n,(k_1,...,k_n)}$ *denote the maximal possible expected welfare from any* n*-bidder game where each bidder* i *has* k_i *possible bids, with any vector of strategies allowed, i.e.,*

$$w^{opt}_{n,(k_1,...,k_n)} = \max_{g \in G_{n,(k_1,...,k_n)}, \; s} w(g, s)$$

When all bidders have k *possible bids we use the notation* $w^{opt}_{n,k} = w^{opt}_{n,(k,...,k)}$

Actually, the optimal welfare should have been defined as the maximum expected welfare that can be obtained *in equilibrium*. Since we later show that the optimal welfare without strategic considerations is dominant-strategy implementable, we use the above definition for simplicity.

Definition 3.8. *The seller's profit is the payment received from the winning bidder, or v_0 when the seller keeps the item.*[6] *Let $r(g,s)$ denote the expected profit in the n-bidder game g where the bidders' strategies are s. Let $r^{opt}_{n,k}$ denote the maximal expected profit from an n-bidder mechanism with k possible bids and some vector of interim individually-rational strategies s that forms a Bayesian-Nash equilibrium in g:*

$$r^{opt}_{n,k} = \max_{\substack{g \in G_{n,k} \\ s \text{ is interim IR and in BNE in } g}} r(g,s)$$

Note that we define the optimal welfare as the maximal welfare among all mechanisms and strategies, not necessarily in equilibria, and we define the optimal profit as the maximal profit achievable in interim-IR Bayesian-Nash equilibria in any mechanism. Yet, the optimal mechanisms (for both measures) that we present in this chapter implement these optimal values with dominant strategies and ex-post IR.[7]

Definition 3.9. *We say that a mechanism $g \in G_{n,k}$ achieves the optimal welfare **(resp. profit)**, if g has an interim-IR Bayesian-Nash equilibrium s for which the expected welfare **(resp. profit)** is $w(g,s) = w^{opt}_{n,k}$ (**resp.** $\mathbf{r(g,s) = r^{opt}_{n,k}}$).*

*We say that a mechanism $g \in G_{n,k}$ incurs a welfare loss **(resp. profit loss)** of L, if it achieves an expected welfare **(resp. profit)** which is additively smaller than the optimal welfare **(resp. profit)** with unbounded communication by L (the optimal results with unbounded communications are the best results achievable with interim-IR Bayesian-Nash equilibria).*

3.3 Optimal Mechanisms for Two Bidders

In this section we present 2-bidder mechanisms with bounded communication that achieve optimal welfare and profit. In Section 3.4 we will present the characterization of the welfare-optimal and profit-optimal n-bidder mechanisms with 2 possible bids for each bidder. The characterization of the optimal mechanisms in the most general case (n bidders and k possible bids) remains an open question. Anyway, our asymptotic analysis of the optimal welfare loss and the profit loss (in Section 3.5) holds for the general case, and shows *asymptotically* optimal mechanisms.

We first show that the allocation rules in efficient mechanisms have a certain structure we call *priority games*. The term priority game means that the allocation rule uses an asymmetric tie breaking rule: the winning player is the player with the highest priority among the bidders that bid the highest bid. One consequence is that the bidder with the lowest priority will win only when his bid is strictly higher than all other bids. Note that the term "priority game" refers to the asymmetry in the mechanism's allocation function, but additional asymmetry will also appear in the payment scheme. A *modified priority game* has a similar allocation, except the item is not allocated when all the bidders bid their lowest bid.[8]

[6] When $v_0 = 0$, the expected profit is equivalent to the seller's expected *revenue*.

[7] Note that ex-ante IR, i.e., when bidders do not know their type when choosing their strategies, is non-interesting in this model, since the auctioneer can then simply ask each bidder to pay her expected valuation.

[8] Modified priority games can be viewed as priority games that treat the seller as one of the bidders with the lowest "priority" (then, the seller always bids his second-lowest bid).

Definition 3.10. *A game is called a priority game if it allocates the item to the bidder i that bids the highest bid (i.e., when $b_i > b_j$ for all $j \neq i$, the allocation is $a_i(b) = 1$ and $a_j(b) = 0$ for $j \neq i$), with ties consistently broken according to a pre-defined order on the bidders.*

A game is called a modified priority game if it has an allocation as of a priority game, except when all bidders bid 0, the seller keeps the item.

It turns out to be useful to build the payment scheme of such mechanisms according to a given profile of threshold strategies:

Definition 3.11. *An n-bidder priority game based on a profile of threshold values' vectors $\overrightarrow{t} = (t^1, ..., t^n) \in \times_{i=1}^n \Re^{k+1}$ (where for every i, $t_0^i \leq t_1^i \leq ... \leq t_k^i$) is a mechanism whose allocation is a priority game and its payment scheme is as follows: when bidder j wins the item for a vector of bids b she pays the smallest valuation she might have and still win the item, given that she uses the threshold strategy s_j based on t^j, i.e, $p_j(b) = \min\{v_j | a_j(s_j(v_j), b_{-j}) = 1\}$. We denote this mechanism as $PG_k(\overrightarrow{t})$. A modified priority game with a similar payment rule is called a modified priority game based on a profile of threshold-value vectors, and is denoted by $MPG_k(\overrightarrow{t})$.*

For 2-bidder games, we may use the notations $PG_k(x, y)$, $MPG_k(x, y)$ (where x, y are some vectors of threshold values). The mechanisms $PG_k(x, y)$ and $MPG_k(x, y)$ are presented in Figure 3.2. Note $PG_k(x, y)$ and $MPG_k(x, y)$ differ only when bidder A bids "0" (i.e., the first line of the game's matrix).

We now observe that priority games and the modified priority games, with the payments schemes that were described above, have two desirable properties: they admit a dominant-strategy equilibrium, and they are ex-post individually rational when the players follow these dominant strategy.

As for the dominant strategies, a well known result in mechanism design (see, e.g., [25, 74] and the references within) states that for any monotone[9] allocation rule there is *some* transfer (i.e., payment) rule that would implement the desired allocation in dominant strategies. For deterministic auctions, to support this equilibrium, each winning bidder should pay the smallest valuation for which she still wins (fixing the behavior of the other bidders). The payments in Definition 3.11 are defined in this way, and therefore they support the dominant-strategy implementation. It follows that the threshold strategies based on the threshold values vector $\overrightarrow{t}$ are dominant in both $PG_k(\overrightarrow{t})$ and $MPG_k(\overrightarrow{t})$. It is clear from the definition of priority games and modified priority games that, when playing their dominant threshold strategies, winning players will never pay more than their value, and losing players will pay zero. Ex-post IR follows.

Actually, the observation about the payments that lead to dominant strategies is even more general. We observe that monotone mechanisms reveal enough information, despite the communication constraints, to find transfer rules that support the dominant-strategy implementation. Therefore, when characterizing the optimal mechanisms we can focus on defining monotone allocation schemes under the communication restrictions, and the transfers that lead to dominant-strategy equilibria can be concluded "for free". In other words, we can use the 2-stage approach that is widely used in the mechanism-design literature also for bounded-communication settings: first solve the optimal allocation rule, and then construct transfers that satisfy the desired incentive-compatibility and individual-rationality constraints.

[9] A mechanism is monotone if the probability that some bidder wins increases as he raises his bid, fixing the bids of the other bidders. See Definition 3.12 below for our model.

	0	1	...	k-2	k-1
0	$\mathbf{B}, y_0$	$\mathbf{B}, y_0$	...	$\mathbf{B}, y_0$	$\mathbf{B}, y_0$
1	A, x_1	$\mathbf{B}, y_1$	...	$\mathbf{B}, y_1$	$\mathbf{B}, y_1$
2	A, x_1	A, x_2	...	$\mathbf{B}, y_2$	$\mathbf{B}, y_2$
...	...	...	...	...	...
k-2	A, x_1	A, x_2	...	$\mathbf{B}, y_{k-2}$	$\mathbf{B}, y_{k-2}$
k-1	A, x_1	A, x_2	...	A, x_{k-1}	$\mathbf{B}, y_{k-1}$

	0	1	...	k-2	k-1
0	ϕ	$\mathbf{B}, y_1$	...	$\mathbf{B}, y_1$	$\mathbf{B}, y_1$
1	A, x_1	$\mathbf{B}, y_1$	...	$\mathbf{B}, y_1$	$\mathbf{B}, y_1$
2	A, x_1	A, x_2	...	$\mathbf{B}, y_2$	$\mathbf{B}, y_2$
...	...	...	...	...	...
k-2	A, x_1	A, x_2	...	$\mathbf{B}, y_{k-2}$	$\mathbf{B}, y_{k-2}$
k-1	A, x_1	A, x_2	...	A, x_{k-1}	$\mathbf{B}, y_{k-1}$

Figure 3.2: A priority game (left) and a modified priority game (right) both based on the threshold values vectors x, y. In each entry, the left argument denotes the winning bidder, and the right argument is the price she pays. The mechanisms differ in the allocation for all-zero bids, and the payments in the first row.

Remark 3.1. *This argument holds for more general environments: in environments in which each player has a one-dimensional private value and a quasi-linear utility, if a non-monetary allocation rule can be implemented in dominant strategies with some transfers, then any communication protocol*[10] *realizing this rule also reveals enough information to construct supporting transfers for the dominant strategies. To see this, recall that in direct-revelation mechanisms (i.e., with unbounded communication), if the allocation rule proves to be monotonic, there are transfers that support a dominant-strategy equilibrium. The transfers will be defined according to some allocation-dependent thresholds, e.g., for a deterministic allocation rule every bidder should pay the smallest valuation for which she still wins. By standard revelation-principle arguments, any monotonic allocation rule in bounded communication mechanisms, can be viewed as a monotonic direct-revelation mechanism with unbounded communication, and therefore such supporting transfers exist. The supporting transfers are determined by the changes in the allocation rule as the valuation of each bidder increases, so the transfers change as the allocation rule changes. Thus, with the same communication protocol that is used for determining the allocation, we can reveal the transfers that support a dominant-strategy implementation.*

3.3.1 The Efficiency of Priority Games

The characterization of the welfare-maximizing mechanism is done in two steps: we first show that the allocation scheme in 2-bidder priority games is optimal[11]. Afterwards, we will characterize the strategies of the players that lead to welfare maximization in priority games; this will complete the description of the outcome of the mechanism for every profile of bidder valuations. These two stages do not take strategic behavior of the bidders into account. Yet, as observed before, since the allocation scheme is proved to be monotone, there exists a payment scheme for which these strategies are dominant.

Definition 3.12. *A mechanism $g \in G_{n,k}$ is monotone if for any bids' vector b and for any bidder i, the probability that bidder i wins the item cannot decrease when only his bid increases, i.e.,*

$$\forall b \quad \forall i \quad \forall b_i^{'} > b_i \qquad a_i(b_i, b_{-i}) \leq a_i(b_i^{'}, b_{-i})$$

[10] Here we deal with simultaneous communication, i.e., where all bidders send their messages simultaneously. Our observation is not true for sequential mechanisms (see Section 3.6).

[11] We assume, w.l.o.g., throughout this chapter that in 2-bidder priority games $B \succ A$, i.e., the mechanism allocates the item to A if she bids a higher bid than B, and otherwise to B

In the following theorem we prove that priority games are welfare maximizing. The proof is composed of four steps: We first show that we can assume that the bidders in the optimal mechanisms use threshold strategies. Then, we show that the allocation in the optimal mechanisms is, w.l.o.g., monotone and deterministic. We then show that the optimal mechanisms do not "waste" communication, i.e., no two "rows" or two "columns" in the allocation matrix of the optimal mechanism are identical. Finally, we use these properties, together with several combinatorial arguments, to derive the optimality of priority games.

Theorem 3.1. ***(Priority games' efficiency)*** *For every pair of distribution functions of the bidders' valuations, and for every v_0, the optimal welfare (i.e., $w^{opt}_{2,k}$) is achieved in either a priority game or a modified priority game (with some pair of threshold strategies).*

Proof. We first prove the theorem given that the seller has a low reservation value, i.e., $v_0 \leq \underline{a}$. Recall that at this point we aim to find the welfare-maximizing allocation scheme, without taking the incentives of the bidders into account. The proof uses the following three claims. For a later use, Claims 3.1 and 3.2 are proved for n players.

Claim 3.1. ***(Optimality of threshold strategies)*** *Given any mechanism $g \in G_{n,(k_1,...,k_n)}$, there exists a vector of threshold strategies s that achieve the optimal welfare in g among all possible strategies, i.e., $w(g,s) = \max_{\widetilde{s}} w(g,\widetilde{s})$*

Proof. (sketch - a formal proof is given in Appendix B.1)

Given a profile of welfare-maximizing strategies in g, we can modify the strategy of each bidder (w.l.o.g., bidder 1) to be a threshold strategy maintaining at least the same expected welfare. The idea is that fixing the strategies s_{-1} of the other bidders, the expected welfare achieved when bidder 1 bids some bid b_1 is a linear function in bidder i's value v_1. The maximum of all these linear functions is a piecewise-linear function, and it specifies the optimal welfare as a function of v_1. Bidder 1 can use a threshold strategy according to the breaking points of this piecewise-linear function that choose the welfare-maximizing linear function at each segment. Clearly, there are at most $k-1$ breaking points. □

Claim 3.2. ***(Optimality of deterministic, monotone mechanisms)*** *For every n and $k_1, ..., k_n$, there exists a mechanism $g \in G_{n,(k_1,...,k_n)}$ with optimal welfare (i.e., there exists a profile s of strategies such that $w(g,s) = w^{opt}_{n,(k_1,...,k_n)}$) which is monotone, deterministic (i.e., the winner is fixed for each combination of bids) and in which the seller never keeps the item.*

Proof. Consider a mechanism $g \in G_{n,(k_1,...,k_n)}$ and a profile s of strategies that maximize the expected welfare, that is, $w(g,s) = w^{opt}_{n,(k_1,...,k_n)}$. A social planner, aiming to maximize the welfare, will always allocate the item to the bidder with the highest expected valuation. That is, for each bids' combination $b = (b_1, .., b_n)$ we will allocate the item (i.e., $a_i(b) = 1$) to a bidder i such that $i \in argmax_j(E(v_j|s_j(v_j) = b_j))$. The expected welfare clearly did not decrease. In addition, we always allocate the item (we assume that $v_0 \leq \underline{a}$), and the allocation is deterministic. Finally, we can assume, w.l.o.g., that for each bidder i the bids' names (i.e., "0","1" etc.) are ordered according to the expected value this bidder has. Then, the mechanism will also be monotone: if a winning bidder i increases his bid, his expected valuation will also increase, while the expected welfare of all the other bidders will not change. Thus, bidder i will still have the maximal expected valuation. □

Claim 3.3. ***(Additional bids strictly help)*** *Consider a deterministic, monotone mechanism $g \in G_{2,k}$ in which the seller never keeps the item. If g achieves the optimal expected welfare, then in the matrix representation of g no two rows (or columns) have an identical allocation scheme.*

Proof. The idea that an optimal protocol exploits all its communication resources is intuitive, although it does not hold in all settings (a trivial example is calculating the parity of two binary numbers, more involved examples can be found in [85]). We do not have a simple proof for this statement in our model, and the proof is based on Lemma B.1 in the appendix in the following way: Consider such an optimal mechanism $g \in G_{2,k}$ with two identical rows. This mechanism achieves the optimal welfare when the players use some profile of strategies s. g's monotonicity implies that the two identical rows are adjacent. Thus, there is a mechanism with $\widetilde{g} \in G_{2,(k-1,k)}$ with $k-1$ possible bids for the rows bidder that achieves exactly the same expected welfare as g (when the identical rows are united to one). This welfare is achieved with the same strategies s of the bidders, where the rows player bids the united row instead of the two identical rows. The claim will now follow from Lemma B.1 in the appendix; According to this lemma, the optimal welfare from a game where both bidders have k possible bids cannot be achieved when one of the bidders has only $k-1$ possible bids (i.e., $w^{opt}_{2,k} > w^{opt}_{2,(k-1,k)}$). □

Now, due to Claim 3.2, there is a deterministic, monotone game in which the item must be sold that achieves $w^{opt}_{2,k}$. In such games, the allocation scheme in some row i looks like $[A, ..., A, B...B]$. Due to Claim 3.3, in the matrix representation of this optimal game, there are no two rows with the same allocation scheme. There are $k+1$ possible monotone rows for the game matrix (with prefix of 0 to k A's), but our mechanism has only k rows. Similarly, we have k different columns (of possible $k+1$) in the mechanism. Assume that the row $[B, B, ..., B]$ is in g. Then, the column $[A, A..., A]$ is clearly not in g. Therefore, our game matrix consists of all the columns except $[A, A, ..., A]$, which compose the priority game where $B \succ A$. If the row $[B, B, ..., B]$ is not in g, then g is the priority game where $A \succ B$.

Next, we complete the proof for any seller's valuation v_0. Consider a mechanism $h \in G_{2,k}$ and a pair of threshold strategies based on some threshold-value vectors $\widetilde{x}, \widetilde{y}$ that achieve the optimal welfare among all mechanisms and strategies (due to Claim 3.1, such strategies exist). We will modify h, such that the expected welfare (with $\widetilde{x}, \widetilde{y}$) will not decrease. Let a be the smallest index such that $E(v_A | \widetilde{x_a} \leq v_A \leq \widetilde{x_{a+1}}) \geq v_0$. Let b be the smallest index such that $E(v_B | \widetilde{y_b} \leq v_B \leq \widetilde{y_{b+1}}) \geq v_0$. If $a = 0$ or $b = 0$, the item is never allocated to the seller, and the efficient mechanism is as if $v_0 \leq \underline{a}$.

When $a, b > 0$, consider some bids' vector (i, j). When $i < a$ and $j < b$, the expected valuations of both A and B are smaller than v_0 . Thus, the seller should keep the item for optimal welfare. When $i < a$ and $j \geq b$, the expected welfare of bidder B is above v_0, and A's expected welfare is below v_0, thus we can allocate the item to B and the welfare will not decrease. Similarly, we should allocate the item to A when $i \geq a$ and $j < b$. When $i < a$, the allocation is done regardless to i, thus we can assume that x_a is the first threshold (i.e., $a = 1$), and similarly $b = 1$.

Now, we show the optimal allocation for bids' combinations (i, j) such that $i \geq a$ and $j \geq b$. Here, the item will not be allocated to the seller, so we actually perform an auction with $k-1$ possible bids for each bidder, when the bidders' valuation are in the range $[\widetilde{x_1}, 1]$, $[\widetilde{y_1}, 1]$. Note that the proof (above) for the case of $v_0 \leq \underline{a}$ holds for such ranges, so the optimal welfare is achieved in a priority game. Altogether, the optimal mechanism turns out to be a modified priority game. □

3.3.2 Efficient 2-bidder Mechanisms with k Possible Bids

Now, we can finally characterize the efficient mechanisms in our model. It turns out that the optimal threshold values for priority games are *mutually centered*, i.e., each threshold is the expected valuation of *the other* bidder, given that the valuation of the other bidder lies between his two adjacent thresholds.

Definition 3.13. *The threshold values $x = (x_0, x_1, ..., x_{k-1}, x_k)$, $y = (y_0, y_1, ..., y_{k-1}, y_k)$ for bidders A, B respectively are mutually centered, if the following constraints hold:*

$$\forall 1 \le i \le k-1 \quad x_i = E(v_B\,|y_{i-1} \le v_B \le y_i) = \frac{\int_{y_{i-1}}^{y_i} f_B(v_B) \cdot v_B dv_B}{F_B(y_i) - F_B(y_{i-1})}$$
$$\forall 1 \le i \le k-1 \quad y_i = E(v_A\,|x_i \le v_A \le x_{i+1}) = \frac{\int_{x_i}^{x_{i+1}} f_A(v_A) \cdot v_A dv_A}{F_A(x_{i+1}) - F_A(x_i)}$$

It is easy to see that given any pair of distribution functions, a pair $\overrightarrow{x}, \overrightarrow{y}$ of mutually-centered vectors is uniquely defined (when $x_k = y_k$ and, w.l.o.g., $y_1 \ge x_1$). The basic idea is that if x_1 is known, we can clearly calculate y_1 (the smallest value that solves $x_1 = E_{v_B}(v_B\,|y_0 \le v_B \le y_1))$. Similarly, it is easy to see that all the variables x_i and y_i can be considered as continuous, monotone functions of x_1. Now, let z be the solution for the equation $y_{k-1} = E(v_A|x_{k-1} \le v_A \le z)$. For satisfying all the $2(k-1)$ equations, z must equal x_k. Since z is also a continuous monotone function of x_1, there is only a single value of x_1 for which all the equations hold.

The following intuition shows why the optimal thresholds in priority games must be mutually-centered: Assume that Alice bids i, that is, her value is in the range $[x_i, x_{i+1}]$. In a monotone mechanism, the mechanism designer has to decide what is the minimal value for which Bob wins when Alice bids i. If the value of Bob is at least the average value of Alice, given that she bids i, then Bob should clearly receive the item. Therefore, Bob's threshold will be exactly this expected value of Alice. The proof has to handle few subtleties for which the intuition above does not suffice (like the characterization of the first thresholds in the optimal modified priority games, see below), thus we will derive the mutually-centered condition from the solution of the optimization problem.

Let $x^w = (\underline{a} = x_0^w, x_1^w, ..., x_{k-1}^w, x_k^w = \overline{b})$ and $y^w = (\underline{a} = y_0^w, y_1^w, ..., y_{k-1}^w, y_k^w = \overline{b})$ be mutually-centered threshold values (w.l.o.g., $y_1^w \ge x_1^w$).
Let $\overline{x} = (\underline{a} = \overline{x_0}, \overline{x_1}, ..., \overline{x_{k-1}}, \overline{x_k} = \overline{b})$ and $\overline{y} = (\underline{a} = \overline{y_0}, \overline{y_1}, ..., \overline{y_{k-1}}, \overline{y_k} = \overline{b})$ be two threshold values vectors for which the following constraints hold:

- $(\overline{x_1}, ..., \overline{x_{k-1}}, \overline{b})$ and $(\overline{y_1}, ..., \overline{y_{k-1}}, \overline{b})$ are mutually-centered vectors[12].
- $\overline{x_1} = v_0$ *and* $\overline{y_1} = \frac{1}{F_A(\overline{x_2})} \cdot \left(v_0 F_A(v_0) + \int_{\overline{x_1}}^{\overline{x_2}} v_A f_A(v_A) dv_A\right)$

The following theorem says that if the valuation of the seller for the item (v_0) is small enough (e.g., $\underline{a}$), the efficient mechanism is a priority game based on x^w and y^w (which are mutually centered). Otherwise, the optimal welfare can be achieved in a modified priority game based on $\overline{x}$ and $\overline{y}$.

Theorem 3.2. *For any pair of distribution functions of the bidders' valuations, and for any seller's valuation v_0 for the item, the mechanism $PG_k(x^w, y^w)$* ***or*** *the mechanism $MPG_k(\overline{x}, \overline{y})$ achieves the optimal welfare (i.e., $w_{2,k}^{opt}$). In particular, $PG_k(x^w, y^w)$ achieves the optimal welfare when $v_0 = \underline{a}$.*

[12]Again, a unique solution exists when, w.l.o.g., $\overline{y_2} \ge \overline{x_2}$

We demonstrate the characterization given above by showing an explicit solution for the case of uniformly-distributed valuations in $[0, 1]$.

Corollary 3.1. *When the bidders' valuations are distributed uniformly on* $[0, 1]$ *and* $v_0 = 0$*, the mechanism* $PG_k(x, y)$ *achieves the optimal welfare where*

$$x = (0, \frac{1}{2k-1}, \frac{3}{2k-1}, ..., \frac{2k-3}{2k-1}, 1) \quad , \qquad y = (0, \frac{2}{2k-1}, \frac{4}{2k-1}, ..., \frac{2k-2}{2k-1}, 1)$$

Proof. According to Theorem 3.2 optimal welfare is achieved with $PG_k(x, y)$, when x, y are mutually centered. With uniform distributions, this derives the following constraints, for which the given vectors x, y are the unique solution: $\forall_{1\leq i\leq k-1} \quad x_i = \frac{y_{i-1}+y_i}{2} \quad y_i = \frac{x_i+x_{i+1}}{2}$ To see how the above constraints are implied, note that the conditional expectation of the second player's value, given that his value is uniformly distributed between y_{i-1} and y_i, is exactly $\frac{y_{i-1}+y_i}{2}$. □

For example, when $k = 2$ we have the constraints $x_1 = \frac{0+y_1}{2}$ and $y_1 = \frac{x_1+1}{2}$, implying that $x_1 = 1/3$ and $y_1 = 2/3$ as in the optimal 1-bit mechanism from Example 3.1. The optimal mutually-centered thresholds for $k = 4$ are, for instance, $x = (0, \frac{1}{7}, \frac{3}{7}, \frac{5}{7}, 1)$ and $y = (0, \frac{2}{7}, \frac{4}{7}, \frac{6}{7}, 1)$.

3.3.3 Profit-Optimal 2-bidder Mechanisms with k Possible Bids

Now, we present profit-maximizing 2-bidder mechanisms. Most results in the literature on profit-maximizing auctions, assume that the distribution functions of the bidders' valuations are *regular* (as defined below). When the valuations of all bidders are distributed with the same regular distribution function, it is well known that Vickrey's 2nd-price auction, with an appropriately chosen reservation price, is profit-optimal ([147, 108, 62]) with unbounded communication.

Definition 3.14. *([108]) Let f be a probability density function, and let F be its cumulative function. We say that f is regular, if the function*

$$\widetilde{v}(v) = v - \frac{1 - F(v)}{f(v)}$$

is monotone, strictly increasing function of v. We call the function $\widetilde{v}(\cdot)$ the virtual valuation of the bidder.

For example, when the bidders valuations are distributed uniformly on $[0, 1]$, a bidder with a valuation v has a virtual valuation of $\widetilde{v}(v) = 2v - 1$.

Definition 3.15. *The virtual surplus in a game is the virtual valuation of the bidder (including the seller[13]) who receives the item.*

The key observation of Myerson ([108]), which we also use, is that in a Bayesian-Nash equilibrium, *the expected profit equals the expected virtual-surplus* (in interim individually-rational equilibria where losing bidders are not getting any surplus). We use this property to reduce the profit-optimization problem to a welfare-optimization problem, for which we have already given a full solution. Myerson's observation was originally proved for direct-revelation mechanisms. We observe here that Myerson's observation also holds for auctions with bounded communication. That is, given a k-bid mechanism, the expected profit in every Bayesian-Nash equilibrium equals the expected virtual surplus.

[13]The seller's virtual valuation is defined to be his "original" valuation (v_0).

Proposition 3.1. *Let $g \in G_{n,k}$ be a mechanism with a Bayesian Nash equilibrium $s = (s_1, ..., s_n)$ and interim individual rationality. Then, the expected revenue achieved by s in g is equal to the expected virtual-surplus of s in g.*

Proof. Consider the following direct-revelation mechanism g_d: each player i bids her true valuation v_i. The mechanism calculates $s_i(v_i)$ for every i, and determines the allocation and payments according to g. An easy observation is that g_d is incentive-compatible (i.e. truthful bidding is a Bayesian-Nash equilibrium for the players) and interim individually rational. According to Myerson observation for direct revelation mechanisms, the expected revenue in g_d is equal to the expected virtual-surplus. However, for every combination of bids, both mechanism output identical allocations and payments. Thus, the expected revenue and the expected virtual-surplus are equal in both mechanisms. □

According to Theorem 3.2, the optimal *welfare* is achieved in either a priority game or a modified priority game. In a model where bidders consider their virtual valuations as their valuations, let $MPG(\widetilde{x}, \widetilde{y})$ or $PG(\overline{x}, \overline{y})$ be the mechanisms which are the candidates to achieve the optimal *welfare* (see Theorem 3.2 for a full characterization). Now, consider the same mechanisms, except each payment $\widetilde{c}$ in them is replaced by the respective "true" valuation $c = \widetilde{v}^{-1}(\widetilde{c})$ (i.e., $\widetilde{c} = \widetilde{v}(c)$). Denote these mechanisms by $PG_k(x^R, y^R)$, $MPG_k(x^r, y^r)$. These mechanisms achieve the optimal *profit* in our (original) model. Note that the distribution functions must be regular (but not necessarily identical) for this reduction to work.

Theorem 3.3. *When both bidders' valuations are distributed with regular distribution functions, the mechanism $MPG_k(x^r, y^r)$* ***or*** *the mechanism $PG_k(x^R, y^R)$ (see definitions above) achieve the optimal expected profit among all profits achievable in an interim-IR Bayesian-Nash equilibrium of a mechanism in $G_{2,k}$ (i.e., $r_{2,k}^{opt}$).*

Proof. Consider the threshold values vectors $(\widetilde{x}, \widetilde{y})$ and $(\overline{x}, \overline{y})$ defined above. The mechanism $MPG(\widetilde{x}, \widetilde{y})$ is efficient in the model where the bidders consider their virtual valuations as their valuation (the same proof holds if $PG(\overline{x}, \overline{y})$ is the efficient mechanism). The density function f is regular, and therefore the virtual valuation $\widetilde{v}(\cdot)$ is strictly increasing. Thus, $MPG_k(x^r, y^r)$ (when the bidders use their original valuations) will have exactly the same allocation for every bids' combination as $MPG_k(\widetilde{x}, \widetilde{y})$ (when the bidders consider their virtual-valuations as their valuations). We conclude that $MPG_k(x^r, y^r)$ achieves the optimal *expected virtual-surplus* and thus also the optimal profit. □

As in the case of welfare optimization, we give an explicit solution for the case of uniform distribution functions. This is a direct corollary of Theorem 3.3. Note that the optimal profit is achieved in a modified priority game. This holds since for the uniform distribution the bidders' expected virtual valuation is negative when they bid "0", so an efficient mechanism will not sell the item when all bidders bid "0".

Corollary 3.2. *When the bidders' valuations are distributed uniformly on $[0, 1]$ and $v_0 = 0$, the modified priority game $MPG_k(x, y)$ achieves the optimal expected profit among all the profits achievable in interim-IR Bayesian-Nash equilibria of mechanisms in $G_{2,k}$, where*

$$x = (0, \frac{1}{2}, \theta + \frac{1 \cdot (1-\theta)}{2k-3}, ..., \theta + \frac{(2k-5) \cdot (1-\theta)}{2k-3}, 1)$$

$$y=(0,\theta,\theta+\frac{2\cdot(1-\theta)}{2k-3},...,\theta+\frac{(2k-4)\cdot(1-\theta)}{2k-3},1)$$

and $\theta=\frac{-2\alpha+\sqrt{1+3\alpha}}{2(1-\alpha)}$ *for* $\alpha=\frac{1}{(2k-3)^2}$ ($\theta=\frac{5}{8}$ *when* k*=2*).

3.4 Optimal Mechanisms for n Bidders with Two Possible Bids

In this section we consider games among n bidders where each bidder has 2 possible bids (i.e., they can send only 1 bit to the mechanism). We give the characterization of the optimal mechanisms for general distribution functions. The characterization of the optimal n-bidder mechanism with k possible bids seems to be harder, and it remains an open question. The difficulty stems from the fact that the monotonicity of the allocation rule does not dictate the exact allocation rule in the general case. Rather, there are many possible allocation schemes that we cannot rule out before we know the strategies of the bidders.[14] Therefore, it seems that one should solve the involved combinatorial problem of finding the optimal allocation rule together with finding the optimal payments. Priority games with 2 possible bids per player can be interpreted as a sequence of take-it-or-leave-it-offer; the player with the highest priority in this interpretation is the first player to be offered, if he accepts the offer (i.e., bids "1") he will receive it. See [138] for the analysis of such take-it-or-leave-it mechanisms.

3.4.1 The Characterization of the Optimal Mechanisms

We first observe that priority games also maximize the welfare in n-bidder games with 2 possible bids. This is easier to see than in the k-possible-bids case. By Claim 3.1 in Theorem 3.1, the bidders will use threshold strategies. An efficient mechanism will allocate the item, for each bids' combination $\overrightarrow{b}$, to the bidder with the highest expected welfare when he bids b_i. Given that the distributions are i.i.d, if $x_i \geq x_j$ then $E\left(v_i \left| v_i \in [x_i, \bar{b}]\right.\right) \geq E\left(v_j \left| v_j \in [x_j, \bar{b}]\right.\right)$ and $E\left(v_i \left| v_i \in [\underline{a}, x_i]\right.\right) \geq E\left(v_j \left| v_j \in [\underline{a}, x_j]\right.\right)$. Therefore, ties will be broken according to the order of the thresholds. If the seller's reservation price v_0 is high enough, the efficient mechanism will be a modified priority game.[15]

We now show the characterization of the optimal thresholds for the priority games. We show that the optimal mechanisms use fully discriminatory payments: the bidder with the highest priority in the priority game pays the highest payment when she wins, and so forth. The optimal modified priority game is given by a simple recursive formula. When the seller allocates the item for when all bids are zero, the constraints become cyclic.

Let $\overrightarrow{x}=(x_1,...,x_n)$ and $\overrightarrow{y}=(y_1,...,y_n)$ be the profiles of threshold values for the n bidders

[14] Consider, for example, a 3-player 3-bid priority game, where the item is allocated to the player with the second-highest priority when all the players bid their highest bid. This mechanism is also monotone with no identical actions for the players.

[15] To see this, we must note that in an efficient mechanism the seller will never keep the item when one of the bidder bids 1 (then, a threshold higher than v_0 for this bidder will gain a higher welfare).

such that the following constraints hold:

$$\begin{aligned}
x_1 &= E\left(v\,|\underline{a} \le v \le x_n\right) && (3.1)\\
\forall_{1\le m\le n-2}\ x_{m+1} &= (1-F(x_m))\cdot E\left(v\,\middle|v\in[x_m,\overline{b}]\right)+F(x_m)\cdot x_m && (3.2)\\
x_n &= \frac{\sum_{i=1}^{n-1}(\prod_{j=i+1}^{n-1}F(x_j))\,(1-F(x_i))\,E\left(v\,\middle|v\in[x_i,\overline{b}]\right)}{1-\prod_{i=1}^{n-1}F(x_i)} && (3.3)\\
y_1 &= v_0 && (3.4)\\
\forall_{1\le m\le n-2}\ y_{m+1} &= (1-F(y_m))\cdot E\left(v\,\middle|v\in[y_m,\overline{b}]\right)+F(y_m)\cdot y_m && (3.5)
\end{aligned}$$

We will now prove that either the mechanism $PG_2(\overrightarrow{x})$ or the mechanism $MPG_2(\overrightarrow{y})$ achieve the optimal welfare. As the thresholds description shows, the thresholds for the modified priority game (i.e., when the seller keeps the item when all bids are zero) are defined by a simple, easy-to-compute recursive formula. The optimality of these thresholds can be shown by the following intuitive argument: Consider a new bidder i that joins a set of $i-1$ bidders. An efficient auction will allocate the item to bidder i if and only if his value is greater than the optimal welfare achievable from the first $i-1$ bidders. Therefore, the threshold for each bidder will equal the optimal welfare gained from the preceding bidders; and indeed, with probability of $1-F(y_{i-1})$, bidder i's valuation will be greater than the expected welfare attained from the other bidders (y_{i-1}) and his average contribution will be $E\left(v\,\middle|v\in[y_m,\overline{b}]\right)$); with probability of $F(y_{i-1})$ he will not contribute to the optimal welfare which remains y_{i-1}. This intuition shows why the revenue-maximizing thresholds above (Equations 3.4,3.5) are independent of the number of players.

Theorem 3.4. *When the bidders' valuations are distributed with the same distribution function, the mechanism $PG_2(\overrightarrow{x})$ or the mechanism $MPG_2(\overrightarrow{y})$ achieves the optimal expected welfare. In particular, when $v_0=\underline{a}$, $PG_2(\overrightarrow{x})$ is the efficient mechanism.*

Proof. We already observed that there exists a priority game that achieves the optimal welfare with threshold strategies. Consider a priority game among n bidders, indexed by their priorities (i.e., $1 \prec 2 \ldots \prec n$). Every bidder wins the item if he bids 1 and all the bidders *with higher priorities* bid 0. Thus, the probability that bidder i wins is $\left(\prod_{j=i+1}^{n}F(x_j)\right)\cdot(1-F(x_i))$. When all bidders bid 0, either bidder n wins or the seller keeps the item for herself. The expected welfare from this game, where the bidders use threshold strategies $x_1,\ldots,x_n$ is:

$$w(g,s) = \sum_{i=1}^{n}\left(\prod_{j=i+1}^{n}F(x_j)\right)(1-F(x_i))\,\frac{\int_{x_i}^{\overline{b}}f(v_i)v_i dv_i}{(1-F(x_i))}+\left(\prod_{i=1}^{n}F(x_i)\right)E_0$$

Where $E_0=E(v_n|v_n\in[\underline{a},x_n])$ in the priority game and $E_o=v_0$ in a modified priority game (the second term relates to the case when all the bidders bid 0). For maximum, the partial derivatives with respect to $x_1,\ldots,x_n$ should equal zero, resulting a characterization of the optimal solution.

For bidders $1\le m\le n-1$ we get (both in the priority game and in the modified priority game):

$$x_m = \sum_{i=1}^{m-1}\left(\prod_{j=i+1}^{m-1}F(x_j)\right)(1-F(x_i))\,E\left(v_i\,\middle|v_i\in[x_i,\overline{b}]\right)+\left(\prod_{i=1,i\ne m}^{m-1}F(x_i)\right)E\left(v_n\,|v_n\in[\underline{a},x_n]\right)$$

The recursive formula is reached by calculating $x_{m+1}-x_m$, from which Equations 3.2 and 3.5 follow. For bidder n in the priority game the first order conditions yield the constraint in Equation

3.3. When $m = 1$, we have $x_1 = E\left(v_n \left| \underline{a} \le v_n \le x_n\right.\right)$ (in the priority game) and $x_1 = v_0$ (in the modified priority game). □

As in Section 3.3, we characterize the profit optimal mechanism by a reduction to the welfare optimizing problem. Again, the reduction can be performed only for regular distributions. Consider the model where bidders take their virtual valuations as their valuations. Let $PG_2(\widetilde{u})$ or $MPG_2(\widetilde{z})$ be the mechanisms that achieve the optimal *welfare* in this model (see Theorem 3.4 above). Let $PG_2(u)$ and $MPG_2(z)$ be similar mechanisms respectively, except each payment $\widetilde{c}$ is replaced with its respective "original" valuation $c = \widetilde{v}^{-1}(\widetilde{c})$.

Theorem 3.5. *When the bidders' valuations are distributed with the same regular distribution function, the mechanism $PG_2(u)$ or the mechanism $MPG_2(z)$ achieves the optimal expected profit among all the profits achievable with a Bayesian-Nash equilibrium and interim IR.*

Proof. This is a corollary of Theorem 3.4. The reduction is done as in Theorem 3.3, and it is possible due to the regularity of the distribution function. □

Again, the optimal thresholds for the modified priority game can be given by a simple recursive formula with an intuitive meaning. The recursion is identical to the welfare optimizing formula (Equation 3.2), and the only difference is in the value of the first threshold that should hold $y_1 = \widetilde{v}^{-1}(v_0)$. The intuition is that given that the best revenue achievable from the first $i-1$ bidders is y_{i-1}, with probability $F(y_{i-1})$ a new player i will not be able to pay a higher price (due to the individual-rationality restriction) and therefore the optimal revenue remains y_{i-1}. When his value is greater than y_{i-1}, he cannot be charged more than his average value ($E(v|y_{i-1} \le v \le \overline{b})$).

Now, we give explicit solutions for the uniform distribution on the support $[0, 1]$. The following recursive constraints characterize the efficient and profit-optimal mechanisms – these are the constraints given in Theorems 3.4 and 3.5 for uniform distributions.

Let $(x_1, ..., x_n) \in [0, 1]^n$ be threshold values for which the following constraints hold:

$$x_1 = \frac{x_n}{2} \tag{3.6}$$

$$\forall m \in \{1, ..., n-2\} \quad x_{m+1} = \frac{1}{2} + \frac{x_m^2}{2} \tag{3.7}$$

$$x_n = \frac{\sum_{i=1}^{n-1}\left(\prod_{j=i+1}^{n-1} x_j\right)\left(1 - x_i^2\right)}{2\left(1 - \prod_{i=1}^{n-1} x_i\right)} \tag{3.8}$$

Let $y = (y_1, ..., y_n) \in [0, 1]^n$ be threshold values where $y_1 = \frac{1}{2}$ and:

$$\forall m \in \{1, ..., n-2\} \quad y_{m+1} = \frac{1}{2} + \frac{y_m^2}{2}$$

Corollary 3.3. *Consider the threshold values $\overrightarrow{x} = (x_1, ..., x_n)$ and $\overrightarrow{y} = (y_1, ..., y_n)$ defined above. When the bidders' valuations are distributed uniformly in $[0, 1]$ and $v_0 = 0$, $PG_2(\overrightarrow{x})$ achieves the optimal welfare and $MPG_2(\overrightarrow{y})$ achieves optimal profit.*

For example, when $n = 5$ we have $y = (0.5, 0.625, 0.695, 0.741, 0.775)$. We could not find a simple closed-form formula for the above optimal thresholds.

3.5 Asymptotic Analysis of the Welfare and Profit Losses

In this section, we measure the performance of the optimal mechanisms presented in earlier sections. Although we did not present a characterization of the optimal mechanisms in the general model of k possible bids and n bidders, we present here mechanisms for this general case that are *asymptotically* optimal. For simplicity, we assume that the valuations' range is $[0, 1]$ (all the results apply for a general range $[\underline{a}, \bar{b}]$ which only changes the constants in our analysis).

We analyze the welfare loss (Subsection 3.5.1), the profit loss (Subsection 3.5.2), and finally, in Subsection 3.5.3 we measure the profit loss and the welfare loss in 1-bit mechanism with n bidders. All the result are asymptotic with respect to the amount of the communication, except in Section 3.5.3 where it is with respect to the number of bidders.

3.5.1 Asymptotic Bounds on the Welfare Loss

The next theorem shows that no matter how the bidders' valuations are distributed, we can always construct mechanisms such that the welfare loss they incur diminishes quadratically in k. This is true for any number of bidders we fix (when $k > 2n$). In particular, the efficient mechanism presented in Theorem 3.2 incurs a welfare loss of $O(\frac{1}{k^2})$. The intuition behind the proof: given the distribution functions of the bidders, we construct a certain threshold strategy, which will be dominant for all bidders. When using this strategy, each bidder will bid any bid i with probability smaller than $\frac{1}{k}$. This way, the probability that a welfare loss may occur is $O(\frac{1}{k})$ (for two players, for instance, a welfare loss will be possible only on the diagonal of the game's matrix). The average welfare loss will also be $O(\frac{1}{k})$, resulting in a total expected loss of $O(\frac{1}{k^2})$. The proof appears in Appendix B.3.

Theorem 3.6. *For any (fixed) number of bidders n, and for any set of distribution functions of the bidders' valuations, there exist a set of mechanisms $g_k \in G_{n,k}$ ($k = 2n+1, 2n+2, ...$), that incur an expected welfare loss of $O(\frac{1}{k^2})$). These results are implemented in dominant strategies with ex-post individual rationality.*

The requirement that $k > 2n$ (here and in Proposition 3.2 below) is due to the construction of the symmetric mechanism in the following proof. These results hold even without this requirement, as shown by an asymmetric construction for a more general setting in Chapter 4.

Asymptotic quadratic bounds were also given by Wilson in [148], which studied similar settings regarding the effect of discrete priority classes of customers. In [148] the uncertainty was about the supply, while in this chapter the demand is uncertain as well. Both results are illustrations for the idea that the deadweight loss is second order in the price distortion. (The price distortion in our model is the maximum difference between the prices that different bidders are facing for the item given the others' bids, and it can be bounded above by $\frac{1}{k}$.) Indeed, a small price distortion ensures both that the probability of an inefficient allocation is small and that the inefficiency is small when it does occur.

Theorem 3.6 is related to proposition 4 in [117]. In [117], Nisan and Segal showed that discretizing an exactly efficient continuous protocol communicating d real numbers yields a "truly polynomial" approximation scheme that is proportional to d (i.e., for any $\epsilon > 0$ we can realize an approximation factor of $1 - \epsilon$ using a number of bits which is polynomial in $log(\epsilon^{-1})$). Here, we discretize a continuous efficient auction (e.g., first-price auction), where d is the number of bidders. Discretization then achieves an approximation error that is exponential in the (minus) number of bits sent per bidder, i.e., asymptotically proportional to $\frac{1}{k}$. However, here we care about average-case approximation which is even closer, because worst-case approximation within an error of ϵ ensures an average case approximation within ϵ^2 (the probability that an error is made is itself in the order of ϵ).

We now show that the asymptotic upper bound above is tight, i.e., for some distribution functions (and in particular, for the uniform distribution) the minimal welfare loss is exactly proportional to $\frac{1}{k^2}$. We show this for any constant number of bidders.

Theorem 3.7. *Assume that the bidders' valuations are uniformly distributed and that $v_0 = 0$. Then, the efficient 2-bidder mechanism $PG_k(x, y)$ described in Corollary 3.1 incurs a welfare loss of exactly $\frac{1}{6\cdot(2k-1)^2}$. Moreover, for any (fixed) number of bidders n and for any v_0, there exists a positive constant c such that **any** mechanism $g \in G_{n,k}$ incurs a welfare loss $\geq c \cdot \frac{1}{k^2}$.*

Proof. We first prove the first part of the theorem, regarding 2-bidder mechanisms. Note that the given mechanism can make non-optimal allocation only for bids' combinations that are on the diagonal or on the lower secondary diagonal in the matrix representation of the 2-bidder game (i.e., when $b_A = b_B$ or when $b_A = b_B + 1$). For such bids (i, j), the overlapping segment of $[x_i, x_{i+1}]$ and $[y_j, y_{j+1}]$ is of size $\frac{1}{2k-1}$. Given such bids' vector (i, j), if one of the valuations is not in this overlapping segment, the allocation is optimal (note that we allocate the item to B on the main diagonal, and to A on the secondary diagonal). The probability that both valuation are in this overlapping range is $\frac{1}{(2k-1)^2}$. The expected valuation in our priority game (when both valuation are in this overlapping segment) is exactly in the middle of this segment. The expected valuation in the optimal auction (with unbounded communications), restricted to this overlapping interval, will be in the $\frac{2}{3}$ point of this range. Thus, the welfare loss is $\frac{1}{6}$ of the segment, i.e., $\frac{1}{6} \cdot \frac{1}{2k-1}$. Thus, for every bids' vector on the main diagonal or on the secondary-diagonal the expected welfare loss is $\frac{1}{6}\frac{1}{(2k-1)^3}$. There are $(2k - 1)$ such bids' vector, thus the total welfare loss is exactly $\frac{1}{6(2k-1)^2}$.

A similar argument shows that even when the seller's valuation v_0 is non zero, the welfare loss is asymptotically greater than $\frac{1}{(2k-1)^2}$: let $z_1, ..., z_m$ be the sizes of the overlapping segments (only when the valuations of both bidders are greater than v_0). Clearly, $m \leq 2k - 1$ and $\sum_{i=1}^{m} z_i \leq 1$. Then, the welfare loss from the game is at least [16]:

$$(1 - v_0)^2 \cdot \sum_{i=1}^{m} z_i^2 \cdot \frac{z_i}{6} = \frac{(1 - v_0)^2}{6} \cdot \sum_{i=1}^{m} z_i^3 \geq \frac{(1 - v_0)^2}{6} \frac{2k - 1}{(2k - 1)^3} \geq \frac{(1 - v_0)^2}{6} \frac{1}{(2k - 1)^2}$$

The proof of the second statement is easily derived: Consider only the case where bidders 1 and 2 have valuations above $\frac{1}{2}$, and the rest of the bidders have valuations below $\frac{1}{2}$. This occurs with the constant probability of $\frac{1}{2^n}$. The best a mechanism can do is to always allocate the item to one of 1 or 2. But due to the first part of the theorem, in any 2-bidder mechanism a welfare loss of proportional to $\frac{1}{k^2}$ will be incurred (the fact that the valuation range is $[\frac{1}{2}, 1]$ and not $[0, 1]$ only changes the constant c). This will hold for any opportunity cost v_0 of the seller. Thus, any mechanism will incur a welfare loss of $\Omega(\frac{1}{k^2})$. □

Note that the same asymptotic results hold even if we restrict attention to symmetric mechanisms. Actually, we prove the upper bound in Theorem 3.6 by constructing a symmetric mechanism

[16] In the left inequality we use the fact that when $z = (z_1, ..., z_m)$ is in the m'th dimensional simplex, $\sum_{i=1}^{m} z_i^3 \geq \frac{1}{m^2}$.

(we can allocate the item to all the bidders who bid the highest bid with equal probabilities). However, asymmetric mechanisms do incur a strictly smaller welfare loss than symmetric mechanisms. For example, when the valuations are distributed uniformly, the optimal welfare loss is $\frac{1}{6(2k-1)^2}$ (by Theorem 3.7) compared with an optimal welfare loss of $\frac{1}{6k^2}$ attained by symmetric mechanisms[17] (i.e., the welfare loss in asymmetric mechanisms is about 4 times better). This observation is interesting in light of the results of Harstad and Rothkopf ([97]) and Wilson ([148]). [97] studied symmetric English auctions, and analyzed the optimal price-jumps in such auctions. Our results show that non-anonymous prices (i.e., different jumps for each bidder) can achieve better results than symmetric (or anonymous) jumps. We also characterize the optimal price-jumps for such auctions (mutually centered threshold values). [148] also studies only symmetric priority classes in his model, and also gives a convergence rate of $\frac{1}{n^2}$ for the efficiency loss (where n is the number of priority classes). We show that asymmetric mechanisms can incur smaller efficiency loss, although the asymptotic convergence rate is the same.

One obvious drawback of our characterization of the optimal mechanisms is that the design is not "detail-free" (as in Wilson's doctrine) – we must know the priors of the bidders for designing the mechanisms. But can we design a mechanism that regardless of the distribution functions, will always incur a low welfare loss? The answer is that we can, but they will not be as efficient as in the commonly-known priors case. We observe that a simple, symmetric mechanism that use equally spaced thresholds (i.e., $PG_k(x, ..., x)$, $x = (0, \frac{1}{k}, \frac{2}{k}, ..., \frac{k-1}{k}, 1)$), incurs a welfare loss not greater than $\frac{1}{k}$ for all possible distribution functions. We actually show (see Proposition B.1 in Appendix B.4) that this is actually the best possible "detail-free" mechanism: for any mechanism there exist distribution functions for which the expected welfare loss is at least in order of $\frac{1}{k}$. For severely low communication, the difference between the "detail-free" mechanisms and prior-aware mechanisms (with loss of $O(\frac{1}{k^2})$) may be substantial. Note that without communication constraints, socially-efficient results can be achieved by "detail-free" mechanisms – second-price auctions. In Proposition B.2 in Appendix B.4 we further dicuss detail-free mechanisms: we whow that when the density functions are bounded from above by a constant, then the simple detail-free mechanism above actually incurs a loss of $O(\frac{1}{k^2})$; in addition, given a lower bound on the values of the desnsity functions, we can show that there are density functions for which *every* mechanism incurs a loss if at least $\Omega(\frac{1}{k^2})$. Finally, we show (in Proposition B.3 in Appendix B.4) that the above simple mechanism is actually optimal even when we no longer assume statistical independence of the valuations: for every k, we can always find a joint distribution such that *every* k-bid mechanism incurs a welfare loss of $\Omega(\frac{1}{k^2})$.

3.5.2 Asymptotic Bounds on the Profit Loss

As done in Theorem 3.3, the profit optimization problem can be reduced to a welfare-optimization problem by maximizing the expected virtual surplus.

Proposition 3.2. *Assume that the bidders' valuations are distributed with regular distribution functions. Then, for any number of bidders n, there exist a set of mechanisms $g_k \in G_{n,k}$ ($k = 2n+1, 2n+2, ...$) that incur a profit loss of $O(\frac{1}{k^2})$. The profit loss is compared with the optimal, individually-rational mechanism that is unconstrained in communication.*

[17] It is easy to show that efficient symmetric mechanisms are similar to priority games, except the item is allocated with equal probabilities in cases of ties. The thresholds of the bidders simply divide the valuations' range to identical segments. Then, it is straightforward to show that the welfare loss is exactly $\frac{1}{6k^2}$.

Proof. Consider the model where bidders consider their virtual valuations $\widetilde{v}_i(v_i)$ as their valuations. As the range of the valuations in this model, we take the union of the ranges of all the bidders' virtual valuations. Denote this range as $[\alpha, \beta]$. Let $\widetilde{g} \in G_{n,k}$ be the mechanism that achieves the maximal *welfare* in this model. Due to Theorem 3.6, $\widetilde{g}$ incurs a welfare loss smaller than $c \cdot \frac{1}{k^2}$, for some positive constant c (the constant takes into account the size of the virtual valuations' range $\beta - \alpha$). Let g be the mechanism with the same allocation as in $\widetilde{g}$, only each payment $\widetilde{q}_i$ for bidder i in $\widetilde{g}$ is replaced with $q_i = \widetilde{v}_i^{-1}(\widetilde{q}_i)$ in g, i.e., $\widetilde{q}_i = \widetilde{v}_i(q_i)$. Since each $\widetilde{v}_i$ is non-decreasing (by their regularity), the allocation rules in g and $\widetilde{g}$ are identical for every bids' combination. Thus, g achieves the maximal expected virtual surplus, and the loss of expected virtual surplus is smaller than $c \cdot \frac{1}{k^2}$. The proposition follows. □

Again, this upper bound is asymptotically tight: with the uniform distribution, any mechanism incurs a profit loss of $\Omega(\frac{1}{k^2})$. This result is derived from Theorem 3.7 using similar arguments as in Proposition 3.2.

Proposition 3.3. *Assume that the bidders' valuations are distributed uniformly. Then, for any (fixed) number of bidders n, there exists a positive constant c such that* ***any*** *mechanism $g \in G_{n,k}$ incurs a profit loss $\geq c \cdot \frac{1}{k^2}$.*

So far, we assumed that the bidders' valuations are drawn from statistically independent distributions. We now point out that the relaxation to general joint distributions is non-interesting in our model. Specifically, we can show that the trivial priority game for which all the bidders use the threshold strategy based on the vector $x = (0, \frac{1}{k}, \frac{2}{k}, ..., \frac{k-1}{k}, 1)$ always incurs an expected welfare loss smaller than $\frac{1}{k}$, and no mechanism can do asymptotically better. In other words, there exists some joint distribution function for which *any* mechanism incurs a welfare loss proportional to $\frac{1}{k}$.

3.5.3 Asymptotic Bounds for a Growing Number of Bidders

In this subsection, we fix the size of communication allowed (to two possible bids), and we show asymptotic bounds as a function of the number of bidders rather than the amount of communication. Unfortunately, we have been able to prove such bounds only for the uniform distribution.

When we restrict our attention to *symmetric* mechanisms, the solution is simple. Using the threshold $x = n^{-\frac{1}{n-1}}$ (for all bidders) achieves the maximal expected welfare, and we have the exact formula showing that the optimal welfare loss is $O(\frac{logn}{n})$.[18]

We now show that optimal *asymmetric* mechanisms incur asymptotically smaller welfare and profit losses of $O(\frac{1}{n})$. These mechanisms fully discriminate between the agents.

Theorem 3.8. *Consider the mechanisms $PG_2(\overrightarrow{x})$ and $MPG_2(\overrightarrow{y})$ described in Corollary 3.3 (in Section 3.4.1). When the bidders' valuations are distributed uniformly, both the welfare loss in $PG_2(\overrightarrow{x})$ and the profit loss in $MPG_2(\overrightarrow{y})$ are smaller than $\leq \frac{9}{n}$.*

Proof. Let x be the *revenue*-optimizing thresholds from Corollary 3.3. We will bound the *welfare* loss in $PG_2(x)$, the efficient mechanism will incur even a smaller loss. We assume, w.l.o.g., that in

[18] The expected welfare then is given by: $x^n \cdot \frac{x}{2} + (1-x^n) \cdot \frac{1+x}{2}$. A maximum is achieved (first order conditions) with: $x = n^{-\frac{1}{n-1}}$. The welfare loss is thus: $\frac{n}{n+1} - \frac{1}{2} \quad 1 - n^{-\frac{1}{n-1}}(\frac{1}{n} - 1) \quad (\frac{n}{n+1}$ is the maximal welfare with unbounded communication). It is easy to see that if $1 - \quad \frac{1}{n} \quad ^{\frac{1}{n}}$ converges to $\frac{\log n}{n}$ then the welfare loss also converges to $\frac{\log n}{n}$. And indeed, $1 - \quad \frac{1}{n} \quad ^{\frac{1}{n}} = 1 - e^{-\frac{\log n}{n}} \approx \frac{\log n}{n}$ (since $1 - e^{-x} \approx x$ for small x's).

g, bidders are indexed according to their priorities (i.e., $1 \prec 2 ... \prec n$). When a bidder wins after bidding "1", the maximal welfare loss is $1 - x_i$. When all bidders bid "0", we use the trivial upper bound of 1 for the welfare loss . Therefore, we can bound the welfare loss with:

$$\sum_{i=1}^{n} \left(\prod_{j=i+1}^{n} x_j \right) (1 - x_i)(1 - x_i) + \prod_{i=1}^{n} x_i \tag{3.9}$$

The following two claims can be easily verified by induction:

Claim 3.4. $\forall_n \quad 1 - x_n \leq \frac{2}{n}$

Claim 3.5. $\forall_{n \geq 15} \quad x_n \leq \frac{2n-3}{2n}$

Now, we prove by induction on n that the first summand in Equation 3.9 is $\leq \frac{8}{n}$. Denote this first term by $\overline{wl_n}$. Note that $\overline{wl_{n+1}} = (1 - x_{n+1})^2 + x_{n+1}\overline{wl_n}$. Assuming that $\overline{wl_n} \leq \frac{8}{n}$, and using the two claims above, it is easy to prove that $\overline{wl_{n+1}} \leq \frac{8}{n+1}$ for $n > 14$. (the reader can verify that this also holds for $n \leq 14$.)

Next, we prove (again by induction on n) that the second expression is smaller than $\frac{1}{n}$. We assume $\prod_{i=1}^{n} x_i \leq \frac{1}{n}$ and prove that $\prod_{i=1}^{n+1} x_i \leq \frac{1}{n+1}$ (using Claim 3.5) :

$$\prod_{i=1}^{n+1} x_i = x_{n+1} \prod_{i=1}^{n} x_i \leq x_{n+1} \frac{1}{n} \leq \frac{2n-1}{2n+2} \cdot \frac{1}{n} < \frac{2n-1}{2n+2} \cdot \frac{1}{n} + \frac{1}{2n(n+1)} = \frac{1}{n+1}$$

Thus, the expected welfare loss is smaller than $\frac{8}{n} + \frac{1}{n} = \frac{9}{n}$

The statement about the profit loss can be derived from the result about the welfare loss (again, by reducing profit optimization to welfare optimization). Nevertheless, a direct proof is straightforward: with the same thresholds x from above, the profit loss is bounded from above by $\sum_{i=1}^{n} \left(\prod_{j=i+1}^{n} x_j \right) \cdot (1 - x_i) \cdot (1 - x_i)$ that was proved to be smaller than $\frac{8}{n}$.[19] □

3.6 Sequential Auctions

In sequential mechanisms, bidders split their bids into smaller messages and send them in an alternating order. In this section, we show that sequential mechanisms can achieve better results. However, the additional gain (in the amount of communication) is only up to a linear factor in the number of bidders.

A *sequential* mechanism is a mechanism in which each bidder sends several messages, in some order (not necessarily in a round-robin fashion). In each stage, each bidder knows what messages the other bidders have sent so far. After all the messages were sent, the mechanism determines the allocation and payments. The allocation scheme and the payment scheme are known to all bidders in advance. In addition, the sizes of the messages, their number and the order in which they are sent are also commonly known in advance. We measure the communication volume in a mechanism by the number of bits actually transmitted.

Definition 3.16. *The communication requirement of the mechanism is the maximal amount of bits which may be transmitted by the bidders in this mechanism.*

[19] In the priority games based on the thresholds $\overrightarrow{y}$, if bidder i wins the item, he pays y_i. Thus, the maximal profit loss when bidder i wins is $1 - y_i$.

B / A	0	1
0	$A, 0$	$B, \frac{1}{4}$
1	$A, \frac{1}{3}$	$B, \frac{3}{4}$

Figure 3.3: (h_1) This sequential game (when A bids first) attains a higher expected welfare than any simultaneous mechanism with the same communication requirement (2 bits). This outcome is achieved with Bayesian-Nash equilibrium.

A strategy for a bidder in a sequential mechanism is a *threshold strategy* if in each stage i of the game the bidder determines the message she sends by comparing her valuation to some threshold values $x_1, ...x_{\alpha_i}$ (where this bidder has $\alpha_i + 1$ possible bids in stage i).

Example 3.2. *The following sequential mechanism has a communication requirement of 2 (see Figure 3.3): Alice sends one bit to the mechanism first. Bob, knowing Alice's bid, also sends one bit. When Alice bids* 0*: Bob wins if he bids* 1 *and pays* $\frac{1}{4}$*; If he bids zero Alice wins and pays zero. When Alice bids* 1*: Bob also wins when he bids* 1*, but now he pays* $\frac{3}{4}$*; If he bids zero, Alice wins again, but now she pays* $\frac{1}{3}$*.*

It is easy to see that this mechanism has a Bayesian-Nash equilibrium[20] *that achieves an expected welfare of 0.653. We saw that the efficient* simultaneous *mechanism with a communication requirement of 2 bits is 0.648 (see Section 3.1). We conclude that sequential mechanisms can gain more efficiency than simultaneous mechanisms.*

Note that throughout the chapter we searched for optimal mechanisms among all the mechanisms with Bayesian-Nash equilibria, but we managed to implement this optimum in dominant strategy. In sequential mechanisms it is less likely to find dominant-strategy implementations, thus our above example uses Bayesian-Nash implementation. Our result below, however, do not assume any particular equilibrium concept in the sequential mechanisms.

How significant is the extra gain from sequential mechanisms over simultaneous mechanisms? The following theorem states for every sequential mechanism with a communication requirement of m there exists a simultaneous mechanism that achieves at least the same welfare with a communication requirement of nm (where n is the number of bidders)[21]. Note that in general (see, e.g., [85]), multi-round protocols can reduce the communication by an exponential factor. We observe that the gain from sequential mechanism is actually even smaller. In many environments, all messages are sent to a centralized authority (auctioneer); therefore, extra bits of communication will be required to inform the bidders about the previous messages of the other bidders. The following theorem holds for any order of transmission and any size of the sub-messages, even if these values are adaptively determined according to previous messages.

The goal of this section is to show that the gain from sequential auctions, compared to simultaneous auctions, is mild. We do not offer a comprehensive analysis of this case, not present welfare-maximizing and revenue-maximizing auctions. Several recent papers studied different aspects of sequential auctions with similar constraints. [138] analyze sequential auctions designed

[20]The following strategies are in Bayesian-Nash equilibrium: Alice uses the threshold $\frac{1}{2}$, and Bob uses the threshold $\frac{1}{4}$ when Alice bids "0" and $\frac{3}{4}$ when Alice bids 1.

[21]Note that in sequential mechanisms the bidders must be informed about the bits the other bidders sent (we do not take this into account in our analysis), so the total gain in communication can be very mild.

as sequences of take-it-or-leave-it offers. [83] studied sequential single-item auctions with discrete price increment, where information can be used in subsequent stages. [122] studied information elicitation in simultaneous and sequential auctions when the values are uncertain.

First, we observe that we can assume that the welfare-maximizing strategies of the bidders are threshold strategies. Again, we show that for each message chosen by bidder i, the welfare is a linear function in v_i. To show this we should use a backward-induction argument: in the last message, the bidders will clearly use thresholds. Therefore, in previous stages the welfare (as a function of v_i fixing the strategies of all the other bidders) is a linear combination of linear functions which is itself a linear function. The maximum over linear function is a piecewise linear function and the thresholds will be its crossing points.

Theorem 3.9. *Let h be an n-bidder sequential mechanism with a communication requirement m. Then, there exists a* simultaneous *mechanism g that achieves, with dominant strategies, at least the same expected welfare as h, with a communication requirement smaller than nm.*

Proof. Consider an n-bidder mechanism h with a Bayesian-Nash equilibrium, and with communication requirement m (for simplicity, assume n divides m, i.e., each bidder sends $\frac{m}{n}$ bits). There exists a profile $s = (s_1, ..., s_n)$ of threshold strategies that achieves the optimal welfare in h. First, we give an upper bound for the total number of thresholds each bidder uses in the game. For a bidder i, let $\alpha^i_1, ..., \alpha^i_{k_i}$ be the (positive) sizes of the k_i messages she sends in h. Let γ^i_j ($1 \leq j \leq k_i$) be the number of bits that were sent by all the bidders (including i), *before* bidder i sends his jth message. When choosing a message of size α^i_j, the bidder uses up to $2^{\alpha^i_j} - 1$ thresholds. In each stage, every bidder can use a different set of thresholds, for every possible history of the game. Thus, for sending her jth message she can use up to $2^{\gamma^i_j}\left(2^{\alpha^i_j} - 1\right)$ different thresholds. Summing up, bidder i uses no more than $T(i) = \sum_{j=1}^{k_i} 2^{\gamma^i_j}\left(2^{\alpha^i_j} - 1\right)$ thresholds. Now, assume w.l.o.g. that the bidders are numbered according to the order they send their *last* messages (i.e., $\gamma^1_{k_1} > \gamma^2_{k_2} > ... > \gamma^n_{k_n}$). Recall that the total number of bits sent by the bidders is m. When sending the last message, bidder 1 thus uses $2^{m-\alpha^1_{k_1}}\left(2^{\alpha^1_{k_1}} - 1\right) < 2^m$ different thresholds. Because all the messages have positive sizes, bidder 2 will have no more than $2^{m-1-\alpha^2_{k_2}}\left(2^{\alpha^2_{k_2}} - 1\right) < 2^{m-1}$ different thresholds for the last stage. Similarly, every bidder i can use at most 2^{m-i+1} thresholds for his last message. But therefore, for her before-last message bidder i uses at most 2^{m-i-1} different thresholds (the worst case occurs when one bidder sends one bit between bidder i's 2 last messages). It follows that the maximal number of different thresholds for bidder i is:

$$\begin{aligned} T(i) &= \sum_{j=1}^{k_i} 2^{\gamma^i_j}\left(2^{\alpha^i_j} - 1\right) < 2^{m-i+1} + 2^{m-i-1} + \sum_{j=1}^{k_i-2} 2^{\gamma^i_j}\left(2^{\alpha^i_j} - 1\right) \\ &< 2^{m-i+1} + 2^{m-i-1} + \sum_{j=1}^{m-i-2} 2^j < 2^{m-i+1} + 2^{m-i-1} + 2^{m-i-1} < 2^{m-i+2} \end{aligned}$$

Now, let g be a simultaneous mechanisms in which each bidder simply "informs" the mechanism between which of the thresholds he uses in h his valuation lies. Clearly, for every set of valuations of the bidders, this allocation in g and h is identical. Due to the inequality above, $m - i + 2$ bits suffice for bidder i to express this number. We conclude that the number of bits sent by all the bidders in g is smaller than: $\sum_{i=1}^{n}(m - i + 2) = nm - \frac{n(n-3)}{2}$.

Finally, we mention that we can set the allocation scheme and the payment scheme in g such that the threshold-strategies based on the thresholds in s will be an equilibrium and the expected welfare will not decrease. As shown in Section 3.3, we turn this mechanism to be monotone by allocating the item deterministically to the bidder with the highest expected value, in each bids' combination. A dominant-strategy equilibrium follows.

This analysis holds for any order and sizes of the bidder messages, even when they depend on the history of the messages, since counting the number of thresholds can be still done in the same way. □

3.7 Future Work

This chapter analyzed single-item auctions that are severely limited in their ability to elicit information from the bidders – few possible bids are available for each player although each player may have a continuum of types. We give a comprehensive analysis of such auctions, and present welfare- and revenue maximizing mechanisms under these restrictions, asymptotic analysis of the losses compared to auctions with unrestricted communication, and we also compare them to auction where bidder send their messages sequentially.

We leave several questions open. The most obvious problem is the characterization of the optimal mechanisms for arbitrary number of players and possible bids, and an asymptotic analysis of the welfare- and revenue loss as a function of both k and n (we provided a separate asymptotic analysis in these variables). In addition, it seems that the concepts and methods presented in this work extend to more general frameworks, like general single-parameter mechanism-design settings and settings with interdependent values (see Chapter 4). A broader view on some of these results will shed light on decision making under informational and incentive constraints.

An additional interesting question is regarding the gain from allocating the bits of communication non-uniformly among the agents. While in simple domains (like 2-bidder auctions) uniform distribution of the communication seems to be the best option, this is unclear, and probably untrue, in more general settings.

Finally, this work presented a partial study of sequential auctions with communication restrictions. This direction of research is very interesting and closer to many real-life auctions. We did not provide a characterization of the optimal sequential mechanisms, or a direct comparison of simultaneous and sequential mechanisms with the same communication requirement. Another interesting idea is to compare prior-aware sequential mechanisms and "detail-free" mechanisms (similar comparison for simultaneous mechanisms showed that detail-free mechanism can only achieve trivial results). It would also be interesting to take an integrated approach and study settings with partially-known priors.

Chapter 4

Implementation with a Restricted Action Space

4.1 Introduction

In standard mechanism-design settings, a social planner wishes to implement some social-choice rule that chooses an alternative based on the private information of the players. Since social planners cannot observe the private information of the players (their *types*), they design mechanisms that make decisions by observing the *actions* of the players. Each player determines his action in the mechanism according to his type in order to maximize his own utility. The challenge of the social planner is to elicit information that will allow him to implement system-wise goals although such goals may conflict with the objectives of the individual players.

Much of the literature on mechanism design restricts attention to *direct revelation* mechanisms, in which the action spaces of the players are identical to their type spaces. This focus is owing to the *revelation principle*[1], which asserts that every mechanism can be transformed into an equivalent incentive-compatible direct-revelation mechanism that implements the same social choice function.

Nonetheless, in most practical settings, direct-revelation mechanisms are not viable since the number of actions available to the players is significantly smaller than their preference space. Consider, for example, the screening model by [131], where an insurance firm wishes to sell different types of policies to different drivers based on their privately known caution levels. In this model, drivers may have a continuum of possible caution levels, but insurance companies offer only a small number of policies (e.g., a small number of deductible amounts in case of a claim) since it is probably infeasible to market and sell more then a few types of policies. Another example is the *signaling* model for the labor market by [146], where employees send signals about their skills to potential employers by the education level they acquire. Although there is a continuum of skill levels, it is unreasonable to expect more than a few education levels in practice (e.g., PhD, M.A., and B.A.).

Mechanisms with a small, manageable set of choices are widespread in practice, and the main reason for this phenomenon is probably their simplicity. This claim is also supported experimentally, e.g., by [78], who showed that a *choice overload* can hamper the willingness of the players to participate in the game, and can degrade their performance in a given transaction. Iyengar et al. compared decision making under a small set of choices and under larger choice sets (not unusually

[1] The work of [108], [66] and [42] discusses the foundations of the revelation principle.

large) and showed that such phenomena are significant even when the number of possible actions is increased from 6 to around 24 or 30. In fact, in many real-life mechanisms the players are required to map their complex preferences into discrete, often dichotomic, decisions. For instance, many mechanisms avoid negotiations and simply post prices for packages or services, and the players are left to decide whether they buy or not under the posted prices. In other settings, players decide whether they participate in or abstain from some transaction, vote for or against some issue in a referendum, and many other similar examples.

Additionally, there are clear evidences for the rare practical use of direct-revelation mechanisms, most prominently VCG mechanisms. One major reason for this fact relates to the *price discovery* process; players usually do not know their exact types and the discovery process may be prohibitively costly (hiring consultants, etc.) or even computationally intractable to compute (see, e.g., [88]). A well-designed mechanism with limited actions will guide the attention of the players to the information that is most relevant for the decision making. Another critical flaw of direct-revelation mechanisms is that players are typically unwilling to reveal their exact types, even if it is beneficial for them in the short run, worrying that this might harm them in future transactions. A small action space allows the players to preserve some degree of privacy. Papers by [130] and [7] provide more details on why VCG mechanisms are indeed rare.

Restrictions on the action space, for specific models, were studied in several earlier papers. [148] measured the effect of discrete "priority classes" of buyers on the efficiency of electricity markets and found that a few priority classes can realize most of the efficiency gains. In a related work, [99] showed that in matching and rationing problems at least half of the social value created by optimal complex schemes can be obtained using very coarse action schemes. [51] considered a simplified decision problem of a single agent searching for a low price with a limited memory; the memory restrictions force the player to divide the set of possible histories into a limited number of categories. It turns out that the optimal partition of the history is obtained, as in this chapter, by dividing the range of prices into disjoint intervals. Compared to the above work, this chapter incorporates incentives issues in general multi-player domains and also characterizes the exact effect of the expressiveness level allowed in the system. A similar result was obtained in a different setting, studied in [13]. There, a revenue-maximizing seller faces a set of bidders, who do not know their private types, and he needs to determine the accuracy level by which they learn their types. On the one hand, more information increases efficiency and thus the seller's revenue, but on the other hand, it increases the information rent of the bidders, thus decreases the seller's revenue. Once again, partitioning the information range into disjoint intervals is shown to maximize the seller's revenue. The work of [31] is the closest in spirit to our work. They studied single-item auctions with severely-restricted action space, and showed that nearly-optimal social welfare can be achieved even with very strict limitations on the action space. An earlier paper in a similar spirit is by [97] who analyzed discrete bid levels in English auctions.

We next present our framework and results.

4.1.1 Our Framework

We consider a general framework for the study of mechanism design in environments with a limited number of actions. We assume a Bayesian model where players have one-dimensional private types, independently distributed on real intervals, and a social planner who wishes to implement a *social-choice function* c that maps every profile of types to a chosen alternative. Due to the limited expressiveness that is implied by the restricted action space, the social planner will typically have uncertainty about the desired alternative. That is, for some realizations of the players' types, the decision of the social planner will unavoidably be incompatible with the social-choice function c. In order to quantify how well bounded-action mechanisms can approximate the original social-choice function, we assume that the social-choice function is derived from a *social-value* function g, which assigns a real value to every combination of alternative A and realization $\overrightarrow{\theta} = (\theta_1, .., \theta_n)$ of the players' types. The social-choice function c will maximize the social value, i.e., $c(\overrightarrow{\theta}) \in argmax_A\{g(\overrightarrow{\theta}, A)\}$.[2] Following are several simple examples of social-value functions:

- *Public goods.* A government wishes to build a bridge only if the sum of the benefits that players gain from it exceeds its construction cost C. There are two alternatives in this model: "*build*" and "*do not build*". The social value functions in a 2-player game is given by: $g(\theta_1, \theta_2,$"build"$)=\theta_1 + \theta_2 - C$, and $g(\theta_1, \theta_2,$"do not build"$)=0$.
- *Single-item auctions.* Consider a 2-bidder auction where the auctioneer wishes to allocate the item to the bidder with the highest value. The social-choice function is given by $g(\theta_1, \theta_2,$ "player 1 wins"$) = \theta_1$ and $g(\theta_1, \theta_2,$ "player 2 wins"$) = \theta_2$.
- *Message delivery over networks.* A message can be delivered over a network composed of two parallel edges. Each edge is owned by a selfish player that has a privately-known probability q_i of delivering the message successfully. A sender wishes to send his message through the network only if the probability of success is greater than, say, 90 percent - the known probability in an alternate network. That is, $g(q_1, q_2,$ "send over network"$)$=1-(1-q_1)·(1-q_2) and $g(q_1, q_2,$"send over the alternate network"$)$=0.9. Note that in this example the social-choice function is not welfare maximizing.

4.1.2 Our Contribution

This chapter centers on the following question: when the players are only allowed to use k actions, which mechanisms achieve the optimal expected social value, and how do they compare to optimal direct-revelation mechanisms? This question is actually composed of two questions.

1. An *information-theoretic* question: what is the optimal method to elicit information on the private information of the players when the players can only reveal information using k actions (recall that their type space may be continuous)?
2. A *game-theoretic* question: what is the best outcome achievable with k actions, given the additional constraint of implementation in dominant strategies?

These two questions raise the question about the "*price of implementation*": can the optimal information-theoretic result be always implemented in a dominant-strategy equilibrium, and to what extent does the dominant-strategy requirement degrade the optimal result?

Example 4.1. *Consider a public good model with two players whose types θ_1, θ_2 are uniformly distributed between $[0, 1]$. A social planner would like to build the bridge when $\theta_1 + \theta_2 > C$ where C is the construction cost of the bridge. It is well-known that if direct revelation is allowed, the VCG mechanism provides a socially-efficient solution. Assume now that only two actions are available*

[2] Observe that the social-value function is not necessarily the social *welfare* function – the social welfare function is a special case of g in which g is defined to be the sum of the players' valuations for the chosen alternative.

to the players: "No" and "Yes". Now, due to the inherent information-theoretic constraints, the social planner is no longer able to build the bridge exactly according to the objective function. What is an optimal 2-action mechanism? Consider the following allocation rule and strategies:

Allocation: the social planner always builds the bridge, unless both players report "No".

Strategies: both players use the following threshold strategy:
"Report "No" if $\theta_i \leq \frac{2}{3} \cdot C$*, otherwise report "Yes""*

As will be shown later in this chapter, the above solution is the best solution for the information-theoretic problem created by using only 2 actions (when $C \leq 1$*); there is no other allocation scheme and no other pair of 2-action strategies that together obtain a higher expected social value. The obvious question is whether this result can be obtained in equilibrium, and the answer is affirmative: it is easy to see that the following payment scheme implies that the above strategies are dominant for both players: "If a player is the only player to report "Yes" he should pay* $\frac{2}{3} \cdot C$*, and otherwise he pays zero". Consequently, the optimal information-theoretic solution can be supported with dominant strategies with no social-value loss!*

In the remainder of this section, we informally survey our three research questions and results.

Our first contribution presents a family of social-value functions for which solving the information-theoretic problem actually also solves the game-theoretic problem. The following theorem holds for any number of alternatives, any number of players, and any profile of distribution functions.

Theorem 1: *For all multilinear single-crossing social-value functions, the information-theoretically optimal social-choice rule is implementable in dominant strategies.*

The theorem assumes two properties of the social-value functions – multilinearity and single crossing. *Multilinear* social-value functions are polynomials where each variable has a degree of at most one in each monomial. They capture many important and well-studied models, and include, for instance, any social-*welfare* maximization setting where the valuations are linear in the types (like public-good and auction models), and other models like the above message-delivery example. *Single crossing* is a stronger property than monotonicity, where the latter is required to guarantee the dominant-strategy implementability of social-choice functions in the absence of restrictions on the actions. The reason for this stronger requirement is that action-bounded mechanisms will typically not be able to exactly implement the original social-choice function; therefore, the social value of all the alternatives should behave "monotonically," not only for those alternatives that are chosen by the desired social-choice function (and thus maximize the social value). A formal definition will be given in the next section.

For proving Theorem 1, we prove a useful lemma that presents an alternative characterization of social-choice functions whose "price of implementation" is zero. We show that for every social-choice function, the implementability of the best information-theoretic solution is equivalent to the property that the optimal expected social value is achieved with *threshold* (or *non-decreasing*) strategies.[3] This lemma actually implies that one can always implement in dominant strategies the optimal social-choice rule that is achievable with threshold strategies.

[3] The restriction to non-decreasing strategies is very common in the literature. One remarkable result by [3] shows that when a non-decreasing strategy is a best response for any other profile of non-decreasing strategies, a pure Bayesian-Nash equilibrium must exist. Another related result is by [51], who showed that the optimal way of an agent with limited memory to partition a given set of possible histories into a fixed number of categories is to use thresholds.

Our next result compares the expected social-value in k-action mechanisms to the optimal expected social value when the action space is unrestricted. For every number of players or alternatives, and for every profile of independent distribution functions, we construct mechanisms that are nearly optimal – up to an additive difference of $O(\frac{1}{k^2})$. This is the same asymptotic rate proved for specific environments by [148], [97] and [31]. Moreover, a better general upper bound cannot be obtained as the work of [31] shows that in some auction settings the optimal loss is exactly proportional to $\frac{1}{k^2}$. Note that there are social-choice functions that can be implemented with k actions with no loss at all (for example, the rule "always choose alternative A").

Our asymptotic result holds for any Lipschitz-continuous social-value function, i.e., functions for which the effect of local changes in the types on the social value is limited. In particular, all polynomials, including multilinear functions, are Lipschitz continuous.

Theorem 2: *For all single-crossing Lipschitz-continuous social-value functions, the optimal k-action mechanism incurs an expected social loss of $O(\frac{1}{k^2})$ compared with mechanisms with unrestricted action space.*

The proof for this theorem is constructive. We present mechanisms that never exceed this loss. Note that social planners can utilize this characterization to optimize the number of actions when this decision is under their control; that is, they should add an action only if its cost is smaller than the marginal contribution of the action to the expected social value. Due to the above result, we can bound the marginal contribution from an additional action by a value that is proportional to $\frac{1}{k^2} - \frac{1}{(k+1)^2}$, which is in the order of $\frac{1}{k^3}$.

Our final result concerns the problem of finding *the* mechanisms that maximize the expected social value. We fully characterize the optimal mechanisms in environments with two players and two alternatives for every number of actions k and every pair of distribution functions of the players' types. We present them in two parts: we first show that the optimal *allocation* scheme is "diagonal" in the sense that in its matrix representation one alternative will be chosen in, and only in, entries below one of the main diagonals. We then characterize the optimal *strategies* – strategies that are "mutually maximizers". Counter-intuitively (and in contrast to the results obtained in [31] in the context of auctions), the optimal "diagonal" mechanism may not utilize all the k available actions for some non-trivial social-value functions.

Theorem 3: *For all multilinear single-crossing social-value functions over two alternatives, the 2-player k-action mechanism that maximizes the social value is diagonal and it possesses dominant strategies that are mutually maximizers.*

Pinpointing the optimal action-bounded mechanism for multi-player or multi-alternative environments seems to be harder and remains an open question. The hardness stems from the fact that the number of diagonal mechanisms is growing exponentially in the number of players.

Finally, we present our results in the context of several natural applications. First, we provide an explicit solution for a public-good game with k-actions. We show that the optimum is achieved in symmetric mechanisms (in contrast to action-bounded auctions in [31]), and show how the optimal allocation scheme depends on the construction cost C. Then, we study the celebrated *signaling* model for the labor market, which is a natural application in our context since education levels are often discrete. Lastly, we present our results in the context of message delivery in networks. The latter example illustrates how our results apply to settings where the objective function of the social planner is other than welfare maximization.

The rest of the chapter is organized as follows: our model and notations are described in Section 4.2. We then describe our general results regarding implementation in multi-player and multi-alternative environments in Section 4.3. Asymptotic analysis of the social-value loss is given in Section 4.4. In Section 4.5 we fully characterize the optimal mechanisms in 2-player environments with two alternatives. Section 4.6 presents applications of our general results. Some of the proofs are deferred to the appendix.

4.2 Model and Preliminaries

We first describe a general mechanism-design model for players with one-dimensional types. Later, in Subsection 4.2.2, we impose limitations on the action space. Our model is a variant of existing models for environments with one-dimensional values that consider types that are drawn from a continuous support and a discrete set of alternatives. Consider n players and a set $\mathcal{A} = \{A_1, A_2, ..., A_m\}$ of m alternatives. Each player has a privately known type $\theta_i \in [\underline{\theta}_i, \overline{\theta}_i]$ (where $\underline{\theta}_i, \overline{\theta}_i \in \mathbb{R}$, $\underline{\theta}_i < \overline{\theta}_i$), and a type-dependent valuation function $v_i : [\underline{\theta}_i, \overline{\theta}_i] \times \mathcal{A} \to \mathbb{R}$. In other words, player i with type θ_i is willing to pay an amount of $v_i(\theta_i, A)$ for alternative A to be chosen. Each type θ_i is independently distributed according to a publicly known distribution F_i, with an always positive density function f_i. We denote the set of all possible type profiles by $\Theta = \times_{i=1}^{n} [\underline{\theta}_i, \overline{\theta}_i]$.

The social planner has a *social-choice function* $c : \Theta \to \mathcal{A}$, where the choice of alternatives is made in order to maximize a *social-value function* $g : \Theta \times \mathcal{A} \to \mathbb{R}$. That is, $c(\overrightarrow{\theta}) \in argmax_{A \in \mathcal{A}} \{g(\overrightarrow{\theta}, A)\}$

We assume that for every alternative $A \in \mathcal{A}$, the function $g(\cdot, A)$ is continuous and differentiable with respect to every type. The players reveal information about their types by choosing an *action*, from an action set B.[4]

A strategy of each player is a function $s_i : [\underline{\theta}_i, \overline{\theta}_i] \longrightarrow B$, mapping each possible type to an action. We denote a profile of strategies by $s = s_1, ..., s_n$ and the set of the strategies of all players except i by s_{-i}. We assume that players have quasi-linear utility functions. Thus, the utility of player i of type θ_i from alternative A under the payment p_i is $u_i = v_i(\theta_i, A) - p_i$.

4.2.1 Dominant-Strategy Implementation

Following is a standard definition of a mechanism. The action space B is usually implicit, but we mention it explicitly since we later examine limitations on B.

Definition 4.1. *A mechanism with an action set B is a pair (t, p) where:*

- $t : B^n \to \mathcal{A}$ *is the allocation rule.*[5]
- $p : B^n \to \mathbb{R}^n$ *is the payment scheme (i.e., $p_i(b)$ is the payment to the ith player given a vector of actions b).*

We say that a strategy s_i is *dominant* for player i in mechanism (t, p) if player i cannot increase his utility by reporting a different action than $s_i(\theta_i)$, regardless of the actions of the other players

[4]We assume that the action space is symmetric for all players, and this assumption can easily be relaxed (except for the characterization results in Section 4.5).

[5]We will show in the proof of Lemma 4.1 that, w.l.o.g., we can focus on deterministic allocation schemes.

b_{-i}. That is, for every type θ_i and action b'_i and b_{-i}, we have that

$$v_i\left(\theta_i,\ t(s_i(\theta_i), b_{-i})\right) - p_i(s_i(\theta_i), b_{-i}) \geq v_i\left(\theta_i,\ t(b'_i, b_{-i})\right) - p_i(b'_i, b_{-i})$$

Definition 4.2. *We say that a social-choice function h is* implementable with a set of actions B *if there exists a mechanism (t,p) with a dominant-strategy equilibrium $s_1, ..., s_n$ (where for each i, $s_i : [\underline{\theta}_i, \overline{\theta}_i] \longrightarrow B$) that always chooses an alternative according to h, i.e., for every θ we have $t(s_1(\theta_1), ..., s_n(\theta_n)) = h(\overrightarrow{\theta})$.*

Fundamental results in the mechanism-design literature state that under a "single-crossing" condition, *monotonicity* of the social-choice function is a sufficient and necessary condition for dominant-strategy implementability (in single-parameter environments). The single-crossing condition (in its different variants, like the Spence-Mirrlees condition, see [146] and [105], or the Milgrom-Shannon condition specified in [101]) appears, very often implicitly, in almost every paper on mechanism design in domains with one-dimensional types. Throughout this chapter, we assume that the preferences of each one of the players are single-crossing. Our definition of single-crossing valuations may be considered as a hybrid of the differential Spence-Mirrlees single-crossing property and the order-theoretic Single-Crossing property (discussions on these two variants can be found in the work of [101] and [53]). We find this variant convenient as it captures a multitude of models where the type space is a continuous real interval and the space of alternatives, for which players may have individual ordering, is discrete.

A valuation function for player i is single-crossing if for every two non-equivalent alternatives, the effect of an increment in θ_i is greater for one of these alternatives for every θ_i. The single-crossing condition actually defines an order on the alternatives for each one of the players. For example, if the value of player i for alternative A increases more rapidly than his value for alternative B, we can denote it by $A \succ_i B$. This definition rules out preferences where the value for an alternative increases more rapidly (compared to another alternative) on some parts of the support, and slower than the other alternative on different parts of the support. Later on, we will use these orders on the alternatives for defining monotonicity of social-choice functions.

Definition 4.3. *A valuation function $v_i : [\underline{\theta}_i, \overline{\theta}_i] \times \mathcal{A} \rightarrow \mathbb{R}$ is* single crossing *if there is a partial order $\succ_i$ on the alternatives, such that for every two alternatives $A_j \succ_i A_l$ we have that for every θ_i,*

$$\frac{\partial v_i(\theta_i, A_j)}{\partial \theta_i} > \frac{\partial v_i(\theta_i, A_l)}{\partial \theta_i}$$

and if neither $A_l \succ_i A_j$ nor $A_j \succ_i A_l$ (denoted by $A_j \sim_i A_l$) the functions are identical, i.e., $v_i(\theta_i, A_j) = v_i(\theta_i, A_l)$ for every θ_i. We also denote $A_j \succeq_i A_l$ if either $A_j \sim_i A_l$ or $A_j \succ_i A_l$.

Example 4.2. *Consider a single-item auction among 3 players, with 3 alternatives: A_1="1 wins", A_2="2 wins", and A_3="3 wins". For each player i, $v_i(A_i, \theta_i) = \theta_i$ and for $j \neq i$, $v_i(A_j, \theta_i) = 0$. Indeed, for player 1 the slope of $v_1(A_1, \theta_1)$ is greater than the slope of $v_1(A_2, \theta_1)$ and therefore $A_1 \succ_1 A_2$. The losing alternatives for player 1 gains her the same value therefore $A_2 \sim_1 A_3$.*

The definition of monotone social-choice functions requires an order on the actions as well. This order is implicit in most standard settings where, for example, it is defined by the order on the real numbers (e.g., in direct revelation mechanisms where each type is drawn from a real interval). When the action space is discrete, the order of the actions can be determined by the names of the actions, for example, "0", "1",...,"k-1" for k-action mechanisms. (We therefore describe this order with the standard relation on natural numbers $<, >$.)

Definition 4.4. *A deterministic mechanism is* monotone *if when player i raises his reported action, and fixing the actions of the other players, the mechanism never chooses an inferior alternative for i. That is, for every $b_{-i} \in \{0, ..., k-1\}^{n-1}$ if $b'_i > b_i$ then $t(b'_i, b_{-i}) \succeq_i t(b_i, b_{-i})$.*

Following is a classic result regarding the implementability of social-choice functions in single-parameter environments. The formal argument is given, for example, in [106] and [142].

Proposition 4.1. *Assume that the valuation functions $v_i(\theta_i, A)$ are single crossing and that the action space is unrestricted. A social-choice function c is dominant-strategy implementable if and only if c is monotone.*

4.2.2 Restrictions on the Action Space

We study environments where the action space B is restricted. We define a *k-action* mechanism to be a mechanism in which the number of possible actions for each player is k, i.e., $|B| = k$. In k-action mechanisms, the social planner typically cannot always choose an alternative according to the social-choice function c due to the informational constraints. Instead, we are interested in implementing a social-choice function that, with k actions, maximizes the *expected social value*: $E_{\overrightarrow{\theta}}\left[g\left(\overrightarrow{\theta}, t\left(s_1(\theta_1), ..., s_n(\theta_n)\right)\right)\right]$. We next define social-choice functions that can be achieved by k-action mechanisms. Note that, as opposed to Definition 4.2, this is an information-theoretic definition that does not involve strategic arguments.

Definition 4.5. *We say that a social-choice function $h : \Theta \rightarrow \mathcal{A}$ is* informationally achievable with a set of actions B *if there exists a profile of strategies $s_1, ..., s_n$ (where for each i, $s_i : [\underline{\theta}_i, \overline{\theta}_i] \longrightarrow B$), and an allocation rule $t : B^n \rightarrow \mathcal{A}$, such that t chooses the same alternative as h for every type profile, i.e., $t(s_1(\theta_1), ..., t(\theta_n)) = h(\overrightarrow{\theta})$. If $|B| = k$, we say that h is k-*action informationally achievable.

Example 4.3. *Consider an environment with two alternatives $\mathcal{A} = \{A, B\}$, and the following desired social-choice function: $\widetilde{c}(\theta_1, \theta_2) = A$ iff $\{\theta_1 > 1/2$ and $\theta_2 > 1/2\}$. $\widetilde{c}$ is informationally achievable with two actions: if both players report "1" when their value is greater than $1/2$ and "0" otherwise, then the allocation rule* "choose alternative A iff both players report 1" *results in exactly the desired allocation for every profile of types. Conversely, it is easy to see that the function $\hat{c}(\theta_1, \theta_2) = A$ iff $\theta_1 + \theta_2 > 1/2$ is not informationally achievable with two actions.*

Given a social-value function, we would like to determine mechanisms that maximize the expected social value, given the information-theoretic constraints.

Definition 4.6. *A social-choice function is k-*action informationally optimal with respect to the social-value function g, *if it is k-action informationally achievable, and it achieves the maximal expected social value among all the k-action informationally achievable social-choice functions.*[6]

As we will show later, it turns out that the monotonicity of the social-choice function will not suffice for ensuring the monotonicity of the k-action mechanisms. While monotonicity describes the structure of the choices that maximize the social value, mechanisms with discrete action spaces will

[6]By results shown later in the chapter, this maximum is attained and the optimal function is well defined. This holds since the optimal results are achieved by threshold strategies, hence every allocation scheme defines a compact set of social values, and there are finite number of different allocation schemes.

also take into account the social value obtained by other alternatives. Therefore, the social-value of all the alternatives should similarly be aligned with the preferences of the players. Therefore, we define a single-crossing property on the social-value function g which is stronger than monotonicity.

Definition 4.7. *Let $\succ_1, ..., \succ_n$ be the orders on the alternatives implied by the single-crossing condition on the valuations of the players. We say that the social-value function $g(\overrightarrow{\theta}, A)$ exhibits the single-crossing property if the following condition is met for every player i:*

for every two alternatives such that $A_j \succ_i A_l$ we have that for every $\overrightarrow{\theta} \in \Theta$,

$$\frac{\partial g(\overrightarrow{\theta}, A_j)}{\partial \theta_i} > \frac{\partial g(\overrightarrow{\theta}, A_l)}{\partial \theta_i}$$

and if $A_j \sim_i A_l$ then for every $\overrightarrow{\theta}$ we have $\frac{\partial g(\overrightarrow{\theta}, A_j)}{\partial \theta_i} = \frac{\partial g(\overrightarrow{\theta}, A_l)}{\partial \theta_i}$

Note that, unlike Definition 4.3, we do not require that the social value of equivalent alternatives will be identical, but we only require identical slopes.[7]

Example 4.4. *Consider the auction setting from Example 4.2. $A_1 \succ_1 A_2$, and indeed the social welfare in A_1 is v_1 and has a greater slope than the welfare in A_2, v_2, as a function of v_1. $A_2 \sim_1 A_3$, and indeed, v_2 and v_3 have identical slopes for all values of θ_1*

Finally, we call attention to a natural set of strategies – "non-decreasing" strategies, where each player reports a higher action as her type increases. Equivalently, such strategies are *threshold strategies* – strategies where each player divides his type support into intervals, and simply reports the interval in which her type lies.

Definition 4.8. *A real vector $x = (x_0, x_1, ..., x_k)$ is a vector of threshold values if $x_0 \leq x_1 \leq ... \leq x_k$.*

Definition 4.9. *A strategy s_i is a* threshold strategy based on a vector of threshold values $x = (x_0, x_1, ..., x_k)$*, if for every action j it holds that $s_i(\theta_i) = j$ iff $\theta_i \in [x_j, x_{j+1}]$. A strategy s_i is called a* threshold strategy *if there exists a vector x of threshold values such that s_i is a threshold strategy based on x.*

4.3 Implementation with a Limited Number of Actions

In this section, we study the general model of action-bounded mechanism design. Our first result is a lemma that provides a sufficient and necessary condition for the implementability of the optimal solution achievable with k actions: this condition says that the informationally optimal social-choice rule is achieved when all the players use non-decreasing strategies. The intuition behind it is that with non-decreasing strategies (i.e., threshold strategies) we can apply the single-crossing property to show that when a player raises his reported action, the expected value for his high-priority alternative increases faster; therefore, monotonicity must hold. The result holds for every number of players and alternatives, and for every profile of distribution functions on the players' types, as long as they are statistically independent [8].

[7]This difference can be demonstrated in multi-item auctions: two allocations in which player i receives the same bundle of items are clearly identical with respect to this player, but their social welfare may differ since the items may be allocated differently among the other players. However, the social welfare changes at the same rate as the value of player i increases.

[8]One can easily verify that this result does not hold if the players' types are dependent.

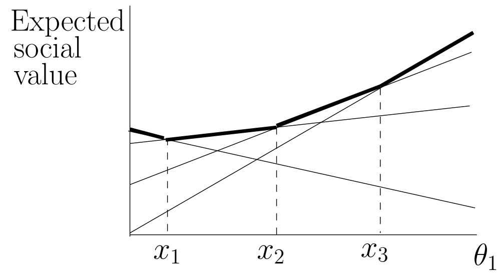

Figure 4.1: Each linear function in the diagram corresponds to the expected social value as a function of the type of Player 1 when she chooses a particular action. Since two linear functions cross at most once, the maximum of k linear functions has at most $k-1$ breaking points (e.g., x_1, x_2, x_3). Therefore, in order to maximize the social value, Player 1 uses a threshold strategy where those breaking points are the thresholds.

Lemma 4.1. *Consider a single-crossing social-value function g. The informationally optimal k-action social-choice function c^* (with respect to g) is implementable if and only if c^* achieves its optimum when the players use threshold strategies.*

The proof for this lemma can be found in the appendix. Theorem 4.1 below is based on one direction of the lemma (optimum with threshold strategies implies implementability); the other direction is given for completeness.[9]

Next, we show that for a wide family of social-value functions – multilinear functions – the information-theoretically optimal rule is dominant-strategy implementable. This family of functions captures many common settings from the literature.

Definition 4.10. *A multilinear function is a polynomial in which the degree of every variable in each monomial is at most 1.*[10] *We say that a social-value function g is* multilinear*, if $g(\cdot, A)$ is multilinear for every alternative $A \in \mathcal{A}$.*

The basic idea behind the proof of the following theorem is as follows: for every player, we show that the expected social welfare for every action he chooses (fixing the strategies of the other players) is a linear function of his type. This is a result of the multilinearity of the social-value function and of the linearity of expectation. The maximum over a set of linear functions is a piecewise-linear function, hence the optimal social value is achieved when the player uses threshold strategies (the thresholds are the breaking points of the piecewise linear function). Figure 4.1 graphically illustrates this argument. Since the optimum is achieved with threshold strategies, we can apply Lemma 4.1 to show the monotonicity of the social-choice rule. Note that in this argument we characterize the players' strategies that maximize the social value rather than the players' utilities.

Theorem 4.1. *If the social-value function is multilinear and single crossing, the informationally optimal k-action social-choice function is implementable.*

Proof. We will show that for every k-action mechanism, the optimal expected social value is achieved when all players use threshold strategies. This will be shown by proving that for every

[9] We currently do not have a concrete example for social choice functions that achieve an optimum with strategies other than threshold strategies and this remains an open problem.

[10] For example, $f(x, y, z) = xyz + 5xy + 7$.

player i and for every action b_i of player i, the expected welfare when she chooses the action b_i (fixing the strategies of the other players) is a linear function in player i's type θ_i. Then, it will follow from Lemma 4.1 that the social-choice function is implementable.

Let t be the allocation function of the mechanism, and let $s_{-i}(\theta_{-i})$ be the strategy profile of all players other than i. For a fixed action b_i of player i, let q_A denote the probability that alternative A is allocated, i.e.,

$$q_A = Pr_{\theta_{-i}} \left[t(s(\overrightarrow{\theta})) = A \,|s_i(\theta_i) = b_i \right]$$

Due to the linearity of expectation, the expected social value when player i with type θ_i reports b_i is:

$$E_{\theta_{-i}} \left[g(\theta_i, \theta_{-i}, t(b_i, s_{-i}(\theta_{-i}))) \right] = \sum_{A \in \mathcal{A}} q_A \, E_{\theta_{-i}} \left[\, g(\theta_i, \theta_{-i}, A) \mid t(b_i, s_{-i}(\theta_{-i})) = A \, \right] \tag{4.1}$$

$$= \sum_{A \in \mathcal{A}} q_A \int_{\theta_{-i}} g(\theta_i, \theta_{-i}, A) f_{-i}^A(\theta_{-i}) d(\theta_{-i}) \tag{4.2}$$

where $f_{-i}^A(\theta_{-i}) = \frac{\prod_{j \neq i} f_j(\theta_j)}{q_A}$ for type profiles θ_{-i} such that $t(b_i, s_{-i}(\theta_{-i})) = A$, and 0 otherwise.

Since g is multilinear, $g(\theta_i, \theta_{-i}, A)$ is a linear function in $\theta_i$for every alternative, where the coefficients depend on the values of θ_{-i}. Denote this function by $g(\theta_i, \theta_{-i}, A) = \lambda_{\theta_{-i}} \theta_i + \beta_{\theta_{-i}}$. Thus, we can write Equation 4.2 as:

$$\sum_{A \in \mathcal{A}} q_A \int_{\theta_{-i}} \left(\lambda_{\theta_{-i}} \theta_i + \beta_{\theta_{-i}} \right) f_{-i}^A(\theta_{-i}) d(\theta_{-i})$$

$$= \sum_{A \in \mathcal{A}} q_A \left(\theta_i \int_{\theta_{-i}} \lambda_{\theta_{-i}} f_{-i}^A(\theta_{-i}) d(\theta_{-i}) \; + \; \int_{\theta_{-i}} \beta_{\theta_{-i}} f_{-i}^A(\theta_{-i}) d(\theta_{-i}) \right)$$

In this expression, each integral is a constant independent of θ_i when the strategies of the other players are fixed. Therefore, each summand, thus the whole function, is a linear function in θ_i.

For achieving the optimal expected social value, the player must choose the action that maximizes the expected social value. A maximum of k linear functions is a piecewise-linear function with at most $k-1$ breaking points. These breaking points are the thresholds to be used by the player. For all types between subsequent thresholds, the optimum is clearly achieved by a single action; since linear functions are single-crossing, every action will be maximal in at most one interval.

The same argument applies to all the players, and therefore the optimal social value is obtained with threshold strategies.

Finally, we must handle one subtle issue. Showing that the informationally optimal k-action social-choice rule is monotone is actually not enough. We should also show that the same amount of actions also suffices for determining the *prices* that support the dominant-strategy implementation of this rule. This clearly holds in our setting. Formally, we can apply Proposition 1 from [144] that claims that in any simultaneous mechanism, the information that allows computing some implementable social-choice function is also sufficient for computing the supporting prices. □

Observe that the proof of Theorem 4.1 actually works for a more general setting. For proving that the information-theoretically optimal result is achieved with threshold strategies, it is sufficient to show that the social-choice function exhibits a *single-crossing condition in expectation*: given any allocation scheme, and fixing the behavior of the other players, the expected social value in any two actions (as a function of θ_i) should be single crossing. Theorem 4.1 shows that this requirement holds for multilinear functions, but we were not able to give an exact characterization of this general class of functions.

Also observe that if the valuation functions of the players are linear and single crossing, then the social-*welfare* function (i.e., the sum of the players' valuations) is multilinear and single-crossing. This holds since the single-crossing conditions on the valuations are defined with a similar order on the alternatives as in the social-value function. Therefore, an immediate conclusion from Theorem 4.1 is that the optimal social welfare, which is achievable with k actions, is implementable when the valuations are linear.

Corollary 4.1. *If the valuation functions $v_i(\cdot, A)$ are single crossing and linear in θ_i for every player i and for every alternative, then the informationally optimal k-action* social welfare *function is implementable.*

4.4 Asymptotic Analysis of the Social-Value Loss

In this section we prove an upper bound on the social-value loss as a function of the number of actions k. In particular, we show that the social value loss diminishes quadratically with the number of possible actions, k. This result holds for any social value function that is Lipschitz-continuous, and includes, among others, all the bounded-degree polynomials. The main challenge here, compared to earlier results, is in dealing with general social-value functions and any number of players and alternatives. In particular, the social-value function may be asymmetric with respect to the players' types and social-value loss may a-priori occur in every profile of actions.

The basic intuition for the proof is that we can construct mechanisms where the probability of having an allocation that is incompatible with the original social-choice function is $O(\frac{1}{k})$. This fact holds for all single-crossing social-value functions, even without the Lipschitz-continuous property. Then, Lipschitz-continuity implies that the social-value loss will always be $O(\frac{1}{k})$ in the mechanisms we construct. Taken together, the expected loss becomes $O(\frac{1}{k^2})$. We present an explicit construction for mechanisms that exhibit the desired loss in dominant strategies. The expected social-value loss clearly depends on the length of the support of the type space. Here, we assume that the type space is normalized to $[0, 1]$, that is, for every player i, $\underline{\theta}_i = 0$ and $\overline{\theta}_i = 1$.

Theorem 4.2. *Assume that the type spaces are normalized to* $[0, 1]$*. For every number of players and alternatives, and for every set of distribution functions of the players' types, if the social-value function g is single crossing and Lipschitz-continuous, then the informationally-optimal k-action social-choice function (with respect to g) incurs an expected social-value loss of $O(\frac{1}{k^2})$.*

Proof. Given a set of n players, we will define a k-action threshold strategy for each player where each action j is chosen with probability $O(\frac{1}{k})$, and the distance between each pair of consecutive thresholds is $O(\frac{1}{k})$. Using these strategies, we define a mechanism that achieves an $O(\frac{1}{k^2})$ loss. For simplicity, we assume that k is even.

<u>*Construction of the threshold strategies:*</u>
For each player i let $Y^i = \{y^i_0 = 0, y^i_1, ..., y^i_{\frac{k}{2}-1}, y^i_{\frac{k}{2}} = 1\}$ be a set of thresholds that divide the density function of player i to $\frac{k}{2}$ equi-mass intervals. That is, for every j we have $F_i(y^i_{j+1}) - F_i(y^i_j) = F_i(y^i_j) - F_i(y^i_{j-1}) = \frac{2}{k}$.

In addition, let $Z^i = \{z^i_0 = 0, z^i_1, ..., z^i_{\frac{k}{2}-1}, z^i_{\frac{k}{2}} = 1\}$ be a set of thresholds that divide the interval $[0,1]$ to $\frac{k}{2}$ equi-sized intervals. That is, for every j we have $y^i_{j+1} - y^i_j = y^i_j - y^i_{j-1} = \frac{2}{k}$.

Now, let $X^i = Y^i \cup Z^i$ be the set of thresholds for player i. That is, player i uses strategy s_i based on the thresholds X^i. Clearly, using a threshold strategy based on X^i (when the thresholds are ordered in an increasing order), player i chooses each action j with probability $O(\frac{1}{k})$, and the distance between each consecutive thresholds is $O(\frac{1}{k})$.

The allocation rule:

For each vector of actions b, the mechanism will choose an alternative that maximizes the expected social value when the players use the threshold strategies s based on the vectors X^i defined above. That is,

$$t(b) \in argmax_A \; E\left[g(\overrightarrow{\theta}, A) \middle| s(\overrightarrow{\theta}) = b \right]$$

Analysis:

We say that an action profile b is *decisive* if one alternative maximizes the social value for every profile of types (otherwise the profile is *indecisive*). Formally, an action profile b is *decisive* if there exists an alternative A for which $A \in argmax_B g(\theta_1, ..., \theta_n, B)$ for every profile $\overrightarrow{\theta}$, such that $s_i(\theta_i) = b_i$ for every player i. Similarly, the profile b is *decisive with respect to a pair of alternatives* A, B, if one of these alternatives is always superior to the other when the players choose the actions b.

We will prove that the above mechanism incurs an expected loss of $O(\frac{1}{k^2})$ using the following two claims. Claim 4.1 shows that the number of indecisive action profiles is $O(k^{n-1})$. Since the player chooses each action with probability $O(\frac{1}{k})$, each indecisive action profile is chosen with probability $O(\frac{1}{k^n})$, and therefore an indecisive profile will be chosen with probability of $O(k^{n-1} \cdot \frac{1}{k^n}) = O(\frac{1}{k})$. Claim 4.2 proves that the maximal possible social-value loss, compared to the optimal allocation with unrestricted actions, is $O(\frac{1}{k})$ for each indecisive action profile. Taken together, the expected social-value loss in the above k-action mechanism is $O(\frac{1}{k^2})$.

Claim 4.1. *For single-crossing social value functions, the number of indecisive action profiles is at most* $O(k^{n-1})$.

Proof. Consider a pair of players 1, 2 and a pair of alternatives A, B and fix the actions $b_{-\{1,2\}}$ of the other players. Let $(b_1, b_2, b_{-\{1,2\}})$ be an indecisive profile with respect to alternatives A and B (assume that $A \succ_1 B$ and $B \succ_2 A$, the other cases are treated analogously). Since the action profile is indecisive, there must be types θ_1, θ_2 for which $s_1(\theta_1) = b_1$ and $s_2(\theta_2) = b_2$, and also

$$E_{\theta_{-\{1,2\}}}[g(\theta_1, \theta_2, \theta_{-\{1,2\}}, A)] > E_{\theta_{-\{1,2\}}}[g(\theta_1, \theta_2, \theta_{-\{1,2\}}, B)]$$

Now consider an action profile b'_1, b'_2 such that $b'_1 > b_1$ and $b'_2 < b_2$. We will show that for any pair of types θ'_1, θ'_2 for which $s_1(\theta'_1) = b'_1$ and $s_2(\theta'_2) = b'_2$ we have:

$$E_{\theta_{-\{1,2\}}}[g(\theta'_1, \theta'_2, \theta_{-\{1,2\}}, A)] > E_{\theta_{-\{1,2\}}}[g(\theta'_1, \theta'_2, \theta_{-\{1,2\}}, B)]$$

The formal argument is proved similarly to the proof in Lemma 4.1, and it follows from the single-crossing condition: changing the types from θ_1, θ_2 to θ'_1, θ'_2 clearly increases the type of player 1 and decreases the type of player 2 – both changes increase the gap between the social value achieved with the alternative A and the alternative B. We conclude that if $b_1, b_2, b_{-\{1,2\}}$ is indecisive with respect to A, B, then any other indecisive action profile cannot include a smaller action for one of the players $1, 2$ and a higher action for the other. Thus, there are at most $2k - 1$ indecisive profiles for any profile $b_{-\{1,2\}}$ of the other players. Every indecisive action profile is clearly indecisive with respect to some pair of alternatives, thus the number of indecisive action profiles (given $b_{-\{1,2\}}$) is at most $\binom{|\mathcal{A}|}{2} \cdot (2k-1) = O(k)$. Therefore, for every pair of players (out of $\binom{n}{2}$ pairs), there are k^{n-2} different actions for the other players, each one allows at most a linear number of indecisive action profiles. The total number of indecisive action profiles will therefore be $O(k^{n-2}) \cdot O(k) = O(k^{n-1})$. □

Claim 4.2. *For all Lipschitz-continuous social-value functions, the social-value loss incurred when the players play an indecisive action profile is* $O(\frac{1}{k})$.

Proof. Consider an indecisive profile of actions b with respect to a pair of alternatives A, B. Given that the players choose the actions b, we show that the difference between the social value gained by choosing A and B is always at most $O(\frac{1}{k})$. It will follow immediately that the expected loss incurred given each action profile is also $O(\frac{1}{k})$.

Consider two profiles of types θ and θ' for which the profile of actions b is chosen by the players. We will prove that $g(\theta, A) - g(\theta', B) = O(\frac{1}{k})$.

Since the profile b is indecisive with respect to A, B, and since the social-value function is continuous, we know that there is a profile of types θ^* for which the players choose the actions in b and such that $g(\theta^*, A) = g(\theta^*, B)$. Consider some profile of types θ for which the profile of actions b is chosen. We will show that $|g(\theta, A) - g(\theta^*, A)|$ is at most $O(\frac{1}{k})$, and similarly one can show that $|g(\theta, B) - g(\theta^*, B)|$ is $O(\frac{1}{k})$ and the theorem will follow.

Since the social-value function is Lipschitz-continuous, there exists a non-negative constant α such that for every alternative A, $|g(\theta, A) - g(\theta^*, A)| \leq \alpha \cdot \sum_{i=1}^{n} |\theta_i - \theta_i^*|$. Since the same action profile is chosen for both θ and θ^*, our construction implies that for every i, $\theta_i - \theta_i^* < \frac{2}{k}$. The claim follows. □

This concludes the proof of the theorem. □

Moreover, as proved by [31], this bound is asymptotically tight in several environments. That is, there exist a set of distribution functions for the players and social-value functions (e.g., the uniform distribution in auctions and public-good settings) for which *every* mechanism incurs a social-value loss of at least an order of $\frac{1}{k^2}$. Obviously, this claim does not imply that the loss of *every* social-choice function will be proportional to $\frac{1}{k^2}$. For example, in the social-choice function that chooses the same alternative for every type profile, no loss will ever be incurred (even with 0 actions).

4.5 Optimal Mechanisms for Two Players and Two Alternatives

While in the previous section we presented k-action mechanisms that are asymptotically optimal, we will now consider the problem of finding *the* k-action mechanisms that maximize the expected social value. We will show a full solution for action-bounded environments with two players and two alternatives, when the social-value function is multilinear and single crossing, and for every pair of distribution functions and every number of actions. This solves the problem, for example, for 2-bidder auctions and 2-player public-good games. The characterization of the optimal mechanisms

	0	1	2	3
0	A	A	A	**B**
1	A	A	**B**	**B**
2	A	**B**	**B**	**B**
3	**B**	**B**	**B**	**B**

	0	1	2	3
0	A	A	A	A
1	A	A	A	**B**
2	A	A	**B**	**B**
3	A	**B**	**B**	**B**

	0	1	2	3
0	**B**	**B**	**B**	**B**
1	A	**B**	**B**	**B**
2	A	A	**B**	**B**
3	A	A	A	**B**

	0	1	2	3
0	A	A	A	**B**
1	A	A	**B**	**B**
2	A	**B**	**B**	**B**

Figure 4.2: The three left tables show all possible diagonal allocation scheme with 4 possible actions for each player. The rightmost table show an example for a diagonal allocation scheme where one of the player has only $k-1$ possible actions.

for arbitrary number of alternatives and players remains an open problem, and we will illustrate the intuition behind its hardness.

In this section, as in most parts of this chapter, we characterize monotone mechanisms by their *allocation* scheme and by a profile of *strategies* for the players. By doing this, we completely describe which alternative is chosen for every profile of types. It is well known that in monotone mechanisms for one-dimensional environments, the allocation scheme uniquely defines payments that support dominant-strategy implementation. We find this description, which does not explicitly mention the payments, simpler for presentation.

The characterization of the optimal k-action mechanisms is presented in two stages: we first illustrate the allocation scheme in the optimal mechanisms and prove that they must be "diagonal". We then define the optimal strategies in such mechanisms, and prove that they exhibit the "mutually-maximizers" property.

4.5.1 Diagonal Allocations

A key notion in our characterization of the optimal action-bounded mechanism is the notion of *non-degenerate* mechanisms. In a degenerate mechanism, there are two actions for one of the players that are identical in their allocation. Intuitively, a degenerate mechanism does not utilize all the action space it is allowed to use, and therefore one might infer that such a mechanism cannot be optimal. Using this property, we then define "diagonal" mechanisms that turns out to exactly characterize the optimal mechanisms.

Definition 4.11. *A mechanism is* degenerate with respect to player i *if there exist two actions* b_i, b_i' *for player i such that for all profiles* b_{-i} *of actions of the other players, the allocation scheme is identical whether player i reports* b_i *or* b_i' *(i.e.,* $\forall b_{-i}$, $t(b_i, b_{-i}) = t(b_i', b_{-i})$*).*

Consider a representation of the allocation scheme using a matrix, where each entry specifies the chosen alternative where the action of one player is choosing a row, and the action of the of the other player is choosing a column. Then, a 2-player mechanism is degenerate with respect to the row player, if there are two rows with identical allocation. We can now define diagonal allocation scheme.

Definition 4.12. *An allocation scheme for 2-player 2-alternative mechanisms with k-possible actions is called* diagonal *if it is monotone, and non-degenerate with respect to at least one of the players.*

The term "diagonal" originates from the matrix representation of these mechanisms, in which one of the diagonals determines the boundary between the choice of the two alternatives. Figure

4.2 depicts some diagonal 4-action allocation schemes. Simple combinatorial arguments show that diagonal mechanisms may come in very few forms. The direction of the diagonal is determined by whether the players have the same order $\succ_i$ on the alternatives (as in public-good games) or not (like in auctions).

Proposition 4.2. *Every diagonal 2-player mechanism has one of the following forms:*

1. *If both players favor the same alternative (w.l.o.g., $B \succ_i A$ for $i = 1, 2$) then either* $t(b_1, b_2) = B$ *iff* $b_1 + b_2 \geq k - 1$ *or* $t(b_1, b_2) = B$ *iff* $b_1 + b_2 \geq k$
2. *If the two players have conflicting preferences (w.l.o.g., $A \succ_1 B$ and $B \succ_2 A$) then either* $t(b_1, b_2) = B$ *iff* $b_1 \geq b_2$ *or* $t(b_1, b_2) = B$ *iff* $b_1 > b_2$
3. *One of the above mechanisms, when one of the fixed-allocation actions is removed for one of the players (i.e., we can subtract the action j of player i such that for any two actions b, b' of the other player we have $t(j, b) = t(j, b')$).*

Proof. Note that in a monotone allocation scheme, there are $k + 1$ possible columns with k alternatives (e.g., for $k = 3$, $[A, A, A]$, $[A, A, B]$, $[A, B, B]$, $[B, B, B]$). Assume that the mechanism is non-degenerate, for example, w.r.t. Player 2 (the column player). If the column $[A, ..., A]$ appears in the allocation matrix, then clearly the row $[B, ..., B]$ does not appear there, which leaves k possible distinct rows for the row player. Note that in this case, when we exclude the row $[A, ..., A]$ of the row player we are still left with k distinct actions for the column player (see Item 3 in the proposition). The actual matrix is defined by the orders on the alternatives, as shown in the proposition. □

We will show that the social-value is maximized in mechanisms with diagonal allocation scheme. Interestingly, one of the possible forms of diagonal mechanisms is degenerate with respect to one of the players (see Item 3 in Proposition 4.2); that is, it can be described as a mechanism with $k - 1$ actions for this player. For example, the rightmost allocation scheme in Figure 4.2 will maximize the social value for some 4-action environments, although the row player has only 3 actions. This auction can be viewed as the leftmost mechanism in Figure 4.2 when the bottom row has been removed.

4.5.2 Mutually Maximizer Threshold Strategies

In Section 4.5.1, we provided a characterization of the allocation scheme of the social-value maximizing mechanisms. Here, we complete the characterization of the optimal mechanisms by defining the optimal pricing rules – the pricing rules that support the optimal strategies. We define the notion of *mutually-maximizer* thresholds, and show that threshold strategies based on such thresholds are optimal. The intuition behind it is as follows. Consider some action b ("row" in the matrix representation) for Player 1. In a monotone mechanism, the allocation in such a row will be of the form $[A, A, ..., B, B]$ (assuming that $B \succ_2 A$). That is, alternative A will be chosen for low actions of Player 2, and alternative B will be chosen for higher actions of Player 2. By determining a threshold for Player 2 that will be used in his threshold strategy, the social planner actually determines the minimal type of Player 2 from which alternative B will be chosen when the row player chooses action b. For optimizing the expected social value, this type for Player 2 should clearly be the type for which the expected social value from A equals the expected social value from B (given that Player 1 chooses the action b); for greater values of Player 2, the single-crossing condition ensures that B will be preferred. The diagonal allocation scheme ensures that the value of each threshold follows from those arguments that concern only a single action of the other player.

Definition 4.13. *Consider a diagonal mechanism, where the players use threshold strategies based on the threshold vectors x, y.*[11] *We say that the threshold x_i of one player (w.l.o.g., Player 1) is a* maximizer *if*

$$E_{\theta_2}[\ g(x_i, \theta_2, A) \mid \theta_2 \in [y_j, y_{j+1}]\ ] = E_{\theta_2}[\ g(x_i, \theta_2, B) \mid \theta_2 \in [y_j, y_{j+1}]\ ]$$

where j is the action of player 2 for which the mechanism swaps the chosen alternative exactly when player 1 plays i, i.e., $t(i, j) \neq t(i-1, j)$.

The threshold vectors x, y are called mutually maximizers *if all their thresholds are maximizers (except the first and the last).*

Example 4.5. *Consider the public-good setting in Example 4.1. The types of the two players are uniformly distributed between $[0, 1]$, each player has 2 actions "0" and "1", and the mechanism builds the bridge unless both bidders choose the action "0" (this optimal mechanism is illustrated in the left table of Figure 4.3). Assume that Player 1 chooses the action "0", and uses a threshold strategy based on the threshold $\frac{2}{3} \cdot C$ (where C is the construction cost of the bridge). What is the minimal type of Player 2 for which the social planner will build the bridge? The expected value of Player 1, given that he chooses "0", is $\frac{C}{3}$. Therefore, the bridge should be built for any θ_2 such that $\frac{C}{3} + \theta_2 \geq C$, that is $\theta_2 \geq \frac{2}{3} \cdot C$. It follows that the threshold strategies based on $\frac{2}{3} \cdot C$ are mutually maximizers in this game. For further discussion on the public-good example, see Section 4.6.1.*

4.5.3 The Optimal 2-Action 2-Player Mechanisms

It turns out that in 2-player, 2-alternative environments, where the social-value rule is multilinear and single crossing, the optimal expected social value is achieved in diagonal mechanisms with mutually-maximizer strategies.

The proof centers on proving that the allocation scheme is non-degenerate with respect to one of the players. In non-trivial mechanisms, this, together with monotonicity, will also show that the other player will either have non-degenerate allocation, or slightly degenerate allocation (i.e., $k-1$ distinct actions). We actually show that in an optimal k-action allocation scheme one of the players will always have k distinct strategies, otherwise we can add a new action for this player and strictly increase the expected social welfare. The proof requires dealing with several sub-cases and is deferred to the appendix.

Theorem 4.3. *In non-trivial*[12] *environments with two alternatives and two players, if the social-value function is multilinear and single crossing, then the optimal k-action mechanism is diagonal, and the optimum is achieved with threshold strategies that are mutually maximizers.*

A corollary from the proof of Theorem 4.1 is that the optimal 2-player k-action mechanism may be, for some distribution functions, degenerate with respect to one of the players (that is,

[11] For simplicity of presentation, we assume that the mechanism is non-degenerate w.r.t. both players; otherwise, the definition is similar but requires adjusting the indices.

[12] Environments are non-trivial if the original social-choice function chooses each alternative with positive probability, and if for both players the single-crossing condition on the alternative is strict (i.e., $\succ_i$ and not $\succeq_i$). Otherwise, the solution is easy.

$C \leq 1$	0	1
0	No $p_1 = p_2 = 0$	**Yes** $p_1 = 0; p_2 = \frac{2C}{3}$
1	**Yes** $p_1 = \frac{2C}{3}; p_2 = 0$	**Yes** $p_1 = p_2 = 0$

$C \geq 1$	0	1
0	No $p_1 = p_2 = 0$	No $p_1 = p_2 = 0$
1	No $p_1 = p_2 = 0$	**Yes** $p_1 = p_2 = \frac{2}{3}C - \frac{1}{3}$

Figure 4.3: Optimal mechanisms in a 2-player, 2-alternative, 2-action public-goods game, when the types are uniformly distributed in $[0, 1]$. The mechanism on the left is optimal when $C \leq 1$ and the other is optimal when $C \geq 1$. The bridge is built in entries labeled as "Yes".

equivalent to a game where one of the players has only $k - 1$ different actions). Moreover, the proof also identifies the following sufficient condition under which the optimal mechanism will be non-degenerate with respect to both players: if the players have the same order on the alternatives (e.g., $B \succ_1 A$ and $B \succ_2 A$), then the optimal alternative must be identical under the profiles $(\underline{\theta}_1, \overline{\theta}_2)$ and $(\overline{\theta}_1, \underline{\theta}_2)$.[13] Similarly, if the players have the opposite order on the alternatives (e.g., $A \succ_1 B$ and $B \succ_2 A$), then the alternative with the higher social value must be identical for $(\underline{\theta}_1, \underline{\theta}_2)$ and $(\overline{\theta}_1, \overline{\theta}_2)$. This condition clearly holds in the public-good model presented in Section 4.6.1 and in auctions.

The full characterization of the optimal mechanisms in multi-player and multi-alternative environments is still an open question. The hardness stems from the fact that the necessary conditions we specified before for the optimality of the mechanisms (i.e., non-degenerate and monotone allocations) are not restrictive enough for the general model. In other words, the number of monotone and non-degenerate mechanisms rapidly increases as the number of players n grows (it can be shown to grow exponentially in n, a proof is given by [24]). Unlike the 2-player 2-alternative case, it seems that pinpointing the best allocation scheme cannot be done independently of finding the optimal strategies, causing a considerable growth in the complexity of determining the solution.

4.6 Applications

In this section, we demonstrate the applicability of our results to public-good models, signaling games and message delivery in networks.

4.6.1 Application 1: Public Goods

This section will discuss in more details the public-good model which was discussed above. The model deals with a social planner who needs to decide whether to supply a public good, such as building a bridge. Let Yes and No denote the respective alternatives of building and not building the bridge. $v = v_1, \ldots, v_n$ is the vector of the players' types – the values they gain from using the bridge drawn from the interval $[0, 1]$. The decision that maximizes the social welfare is to build the bridge if and only if $\sum_i v_i$ is greater than its cost, and this cost is denoted by C. If the bridge is built, the social welfare is $\sum_i v_i - C$, and zero otherwise; thus, $g(v, Yes) = \sum_i v_i - C$, and $g(v, No) = 0$. The utility of player i under payment p_i is $u_i = v_i - p_i$ if the bridge is built, and 0 otherwise. It is

[13] More precisely, the condition for non-degeneracy when $B \succ_1 A$ and $B \succ_2 A$ is that $sign(g(\underline{\theta}_i, \overline{\theta}_i, A) - g(\underline{\theta}_i, \overline{\theta}_i, B)) = sign(g(\overline{\theta}_i, \underline{\theta}_i, A) - g(\overline{\theta}_i, \underline{\theta}_i, B))$ (when $sign(0)$ is considered both negative and positive).

well-known that under no restriction on the action space, it is possible to induce truthful revelation by VCG mechanisms, therefore full efficiency can be achieved. Obviously, when the action set is limited to k actions, we cannot achieve full efficiency due to the informational constraints. Yet, since $g(v, Yes)$ and $g(v, No)$ are multilinear and single crossing, we can directly apply Theorem 4.1. Hence, the information-theoretically optimal k-action mechanism is implementable in dominant strategies.

Corollary 4.2. *The k-action informationally-optimal social welfare in the n-player public-good game is implementable in dominant strategies.*

Moreover, as Theorem 4.3 suggests, in the public-good game we can fully characterize the optimal k-action 2-player mechanisms. As mentioned in Section 4.5.3, when $g(\overline{\theta}_1, \underline{\theta}_2, A) = g(\overline{\theta}_1, \underline{\theta}_2, B)$ the mechanism is non-degenerate with respect to both players. This condition clearly holds here $(1 + 0 - C = 0 + 1 - C)$, therefore the optimal mechanisms will use all k actions. An immediate corollary from Theorem 4.3 is a full characterization of the optimal mechanisms in this setting:

Corollary 4.3. *The optimal expected welfare in a 2-player k-action public-good game is achieved with one of the following mechanisms (where $x_0 = y_0 = 0$ and $x_k = y_k = 1$):*

1. *Allocation: Build the bridge iff $b_1 + b_2 \geq k$.*

   Strategies: Threshold strategies based on the vectors $\overrightarrow{x}, \overrightarrow{y}$ where for every $1 \leq i \leq k$-1,

   $x_i = C - E[v_2 | v_2 \in [y_{k-i}, y_{k-i+1}]]$

   $y_i = C - E[v_1 | v_1 \in [x_{k-i}, x_{k-i+1}]]$

2. *Allocation: Build the bridge iff $b_1 + b_2 \geq k - 1$.*

   Strategies: Threshold strategies based on the vectors $\overrightarrow{x}, \overrightarrow{y}$ where for every $1 \leq i \leq k$-1:

   $x_i = C - E[v_2 | v_2 \in [y_{k-i-1}, y_{k-i}]]$

   $y_i = C - E[v_1 | v_1 \in [x_{k-i-1}, x_{k-i}]]$

The construction cost determines which of the two mechanisms above obtains a better result. Recall that we define the optimal mechanisms by their (monotone) allocation scheme and by the optimal strategies for the players. The payments that support dominant-strategy implementation satisfy the rule that a player pays his lowest value for which the bridge is built, when the action of the other player is fixed. Therefore, the payments for players 1 and 2 reporting the actions b_1 and b_2 are as follows: in mechanism 1 from Proposition 4.3, $p_1 = x_{b_2}$ and $p_2 = y_{b_1}$; in mechanism 2 from Proposition 4.3, $p_1 = x_{b_2 - 1}$ and $p_2 = y_{b_1 - 1}$.

We now apply Corollary 4.3 for the specific case of the uniform distribution. The example shows how the optimal mechanism is determined by the cost C: a mechanism of type 1 is optimal for construction costs smaller than 1, while a mechanism of type 2 is optimal for higher costs. Note that the optimal mechanisms are symmetric, unlike the solution for auctions in [31].

Example 4.6. *Suppose that the types of both players are uniformly distributed on* $[0, 1]$*. Figure 4.3 illustrates the optimal mechanisms for $k = 2$, and how they depend on the construction cost C. For every number of actions k, the welfare-maximizing mechanisms are:*

- *If the cost of building is at least* 1*:*

  Allocation: Build iff $b_1 + b_2 \geq k$

  Strategies: The thresholds of both players are (for $i = \{1, \ldots, k-1\}$), $x_i = \frac{2(k-i) \cdot C}{2k-1} - \frac{2k-4i+1}{2k-1}$

- *If the cost of building is smaller than* 1*:*

  Allocation: Build iff $b_1 + b_2 \geq k - 1$

  Strategies: The thresholds of both players are (for $i = \{1, \ldots, k-1\}$*),* $x_i = \frac{2iC}{2k-1}$

4.6.2 Application 2: Signaling

We now study a signaling model in labor markets. In this model, the type of each worker, $\theta_i \in [\underline{\theta}, \overline{\theta}]$, describes the worker's productivity level. The firm wishes to make her hiring decisions according to a decision function $f(\overrightarrow{\theta})$. For example, the firm may want to hire the most productive worker (like the auction model), or hire a group of workers only if their sum of productivity levels is greater than some threshold (similar to the public-good model). However, the worker's productivity is invisible to the firm; the firm only observes the worker's education level e that should convey signals about her productivity level. The standard assumption here is that acquiring education, at any level, does not affect the productivity of the worker, but only signals about the worker's skills.

A main component in this model, is the fact that as the worker is more productive, it is easier for him to acquire high-level education. In addition, the cost of acquiring education increases with the education level. More formally, a continuous function $C(e, \theta)$ describes the cost to a worker from acquiring each education level as a function of his productivity. The standard assumptions about the cost function are: $\frac{\partial C}{\partial e} > 0$, $\frac{\partial C}{\partial \theta} < 0$, $\frac{\partial C}{\partial e \partial \theta} < 0$, where the latter requirement is exactly equivalent to the single-crossing property. The utility of a worker is determined according to the education level he chooses and the wage $w(e)$ attached to this education level, that is, $u_i(e, \theta_i) = -C(\theta_i, e) + w(e)$.

An action for a worker in this game is the education level he chooses to acquire. In standard models, this action space is continuous, and then a *fully separating equilibrium* exists (under the single-crossing conditions on the cost function). That is, there exists an equilibrium in which every type is mapped into a different education level; thus, the firm can induce the exact productivity levels of the workers by this signaling mechanism. However, a continuum of education levels is somewhat unrealistic. It is usually the case that there are only several discrete education levels (e.g., BSc, MSc, PhD).

With k education levels, the firm may not be able to exactly follow the decision function f. For achieving the best result in k actions, the firm may want the workers to play according to specific threshold strategies. It is easy to verify that the standard condition, the single-crossing condition on the cost function, suffices for ensuring that these threshold strategies will be dominant for the players. We can now apply Theorem 4.2, and show that if the decision function f of the firm is Lipschitz-continuous (i.e., the decisions are made to maximize a set of Lipschitz-continuous functions), then the firm can design the education system such that the expected loss will be $O(\frac{1}{k^2})$, with a dominant-strategy equilibrium. Note that while in the classic example of the job market it is unreasonable for each firm to select the education level, in other reasonable applications the social planners may be able to determine the thresholds, e.g., by fixing the levels of qualifying exams or other means for the players to demonstrate their skills.

Corollary 4.4. *Consider a Lipschitz-continuous decision function* f *and a single-crossing cost function for the players. With* k *education levels, the firm can implement in dominant strategies a decision function that incurs a loss of* $O(\frac{1}{k^2})$ *compared with the original decision function* f*.*

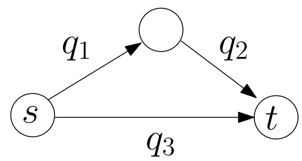

Figure 4.4: An example for a parallel-path network, where each link has a probability q_i for transmission success. We show that the overall probability of success in such networks is multilinear in q_i, and thus the optimal k-action social-choice function is dominant-strategy implementable.

4.6.3 Application 3: Message Delivery in Networks

Lastly, we show the applicability of our results to settings where messages should be delivered over lossy communication networks. Different parts of the networks, i.e., edges in their graph, are owned by rational players who possess a privately known probability of successfully delivering a message (or completing another task) over this edge. Each player owns at most one edge. A sender knows the topology of the networks, and has to devise a mechanism for deciding which network has the highest success probability. It is natural to assume that the players (i.e., links) may not be able to report (or to figure out) their accurate probabilities of success, but only, e.g., whether these are "low", "intermediate", or "high".

Consider a set of networks, where each network is composed of multiple parallel paths from a given source to a given destination. An example for such a network appears in Figure 4.4; in this example, the probability that a message will be transmitted successfully in the upper path, for instance, is $q_1 \cdot q_2$. The sender wishes to send the message through the network with the highest success probability.

In this example we assume that every player has a single-crossing valuation function over the alternatives. That is, each player wishes that the message will be sent through his network, and his benefit is positively correlated with his private data (e.g., the valuation of player i for delivering the task may be exactly q_i). We would like to emphasize that the social planner in this example (the sender) does not aim to maximize the social welfare. That is, the social value is neither the sum of the players' types nor any weighted sum of the types ("affine maximizer").

The success probability of sending a message through a parallel-path network is multilinear, since it can be expressed by the following multilinear formula (where $\mathcal{P}$ denotes the set of all paths between the source and the sink):

$$f(\vec{q}) = 1 - \prod_{P \in \mathcal{P}} (1 - \prod_{j \in P} q_j) \tag{4.3}$$

For example, in the network presented in figure 4.4, the probability of success is given by

$$f(\vec{q}) = 1 - (1 - q_1 q_2)(1 - q_3)$$

Thus, if all the candidate networks are parallel-path networks, the social-value function is multilinear.[14] We also note that for every edge i, the partial derivative in q_i of the success probability written in Equation 4.3 is positive where in all the other networks, that do not contain link i, the

[14]The results obtained here hold for all directed networks with no cycles (also known as DAG – directed acyclic graphs).

partial derivative is clearly zero. Therefore, the social-value function is single crossing and our general results apply.

Corollary 4.5. *For all social-choice functions that maximize the success probability over parallel-path networks, the informationally-optimal k-action social-choice function is implementable in dominant strategies (for every k).*

Chapter 5

Informational Limitations of Ascending Combinatorial Auctions

5.1 Introduction

Combinatorial auctions are a general name given to auctions in which multiple heterogeneous items are concurrently sold and in which bidders may place bids on *combinations* of items rather than just on single items. Such *combinatorial bidding* is desired whenever items sold are complements or substitutes of each other, at least for some of the bidders. In such cases, the combinatorial bidding allows the bidders to better express their complex preferences, allowing the auction to achieve higher social welfare, and often (but not necessarily) higher profit as well. Combinatorial auctions have been used in many settings such as truckload transportation ([92, 145]), airport slot allocation ([126, 39]), industrial procurement ([16]), and, prominently, radio spectrum auctions ([40, 56]). Additionally, combinatorial auctions serve as a common abstraction for many resource allocation problems in decentralized computerized systems such as the Internet, and may serve as a central building block of future electronic commerce systems.

The design of combinatorial auctions faces multiple types of complexities: informational, cognitive, computational, and strategic. Indeed, the design of combinatorial auctions is still part art and part science. While many aspects have been analyzed mathematically or empirically, many other aspects remain an art form. In many cases the design is ad-hoc for a given application, and it is usually not clear how well the existing design performs relative to the other non-implemented alternatives. Indeed, when the US Federal Communications Commission held a series of workshops addressing the intended design of their multi-billion dollar combinatorial auctions for radio spectrum (see, e.g., [57]), there has been very little agreement among the participants. We refer the reader to the recent tomes ([37],[103]) that elaborate on various aspects, applications and suggestions for combinatorial auctions.

This chapter concerns a large class of combinatorial auction designs which contains the vast majority of implemented or suggested ones: ascending auctions. In this class of auctions, the auctioneer publishes prices, initially set to zero (or some other minimum prices), and the bidders repeatedly respond to the current prices by bidding on their most desired bundle of goods under the current prices. The auctioneer then repeatedly updates the prices by increasing some of them in some manner, until a level of prices is reached where the auctioneer can declare an allocation. (Intuitively, prices related to over-demanded items are increased until the demand equals supply.)

There are several reasons for the popularity of ascending auctions, including their intuitiveness, the fact that private information need only be partially revealed, that they increase the trust in the auctioneer as bidders see the prices gradually emerging, that it is clear that they will terminate, that they may sometimes reduce the winner's curse and increase the seller's revenue ([104]). See [38] for a survey of the advantages and disadvantages of ascending auctions.

Ascending auctions may vary from each other in the bidding rules, in the price update scheme, in the termination condition, etc. The most notable difference is in the types of prices used. Some auctions attach a price to each item, and the price of each bundle of items is the sum of the item prices. Such auctions are termed *item-price* auctions or linear-price auctions. A more general class of auctions maintains a separate arbitrary price for each bundle of items. These are called *bundle-price* auctions or non-linear price auctions. Some auctions present the same set of prices to all bidders – these are called *anonymous-price* auctions. Others maintain a separate set of prices for each bidder – these are called *personalized* price auctions (or *non-anonymous* price auctions). It is clear that item-price auctions are preferable to bundle-price ones in terms of simplicity, and similarly that anonymous-price ones are simpler than personalized-prices ones. This simplicity is important in many respects, including the cognitive, computational, and communication burden placed on the bidders and on the auctioneer. In particular, such auctions tend to be simpler to bid on, will run faster, and will require less communication and computation and thus will be feasible for a larger number of items. The question is really whether the more complex types of auctions' added expressiveness offers benefits that overcome the cost in complexity. Indeed, presentations at the FCC's 2003 conference ([57]) reveal an interesting debate along these lines between the suggestions of David Porter, Stephen Rassenti and Vernon Smith (on the simplicity side) and of Larry Ausubel, Peter Cramton and Paul Milgrom (on the complexity side).

Two families of ascending auctions have received most of the mathematical analysis so far. The first family of item-price auctions is an extension of the "Deferred-Acceptance Mechanisms" from the literature on matching (see the survey by [128]). This family includes work by [82], [45] and [69], and the main idea is quite simple: increase the price of over-demanded items until every item is demanded by at most one bidder. The basic theorem shows that if all bidders have (gross) substitutes valuations, then this converges to a competitive (Walrasian) equilibrium and thus leads to social efficiency. The restriction to having (gross) substitutes valuations is critical; for example, [69] show that for any bidder whose preferences fail the substitutes condition we can add a set of unit demand bidders such that the resulting economy has no Walrasian equilibrium.

The second family of auctions uses personalized bundle prices and includes those of [124] and [4]. The main idea here is that the auctioneer computes, at each stage, an optimal tentative allocation, and then losers in this tentative allocation are allowed to increase their bids (i.e., essentially losers' personalized prices are increased). The basic theorem states that when no loser wants to increase his bid, then an optimal allocation has been reached. This holds for arbitrary bidder valuations.

The fundamental question that we address is whether the added complexities of bundle prices and of personalized prices are indeed necessary. We present a strong affirmative answer on both counts. We prove that no ascending item price auction (using anonymous or personalized prices) can always reach a socially-efficient allocation among arbitrary bidder valuations. Similarly, we prove that no anonymous-price auction (using either item prices or bundle prices) can always reach the socially optimal allocation. Our basic theorems are proved by analyzing two very simple scenarios in which we show that the appropriate type of auction can simply not gather sufficient information from the bidders.

We then prove several stronger variants of our theorems showing that our impossibility results are very robust is several senses. We show that not only is it impossible for an ascending item-price auction to obtain the social optimum, but even if we allow multiple, sub-exponentially many, "ascending paths" (e.g., as used in [5]), then the impossibility remains. Finally, we also show that the loss of welfare is extreme both for item price auctions and for anonymous-price auctions, and that only a vanishingly small fraction of the social welfare may be captured[1]. This last pair of results is proved using a sophisticated combinatorial construction of valuations that are "hard to elicit" by these restricted types. We also show that our examples are not "unusual" by showing that for any set of substitutes bidders, it is possible to add a single extra bidder making it impossible to find the social optimum by item-price auctions. Recall that environments with substitutes preferences are the most general setting where item-price ascending auctions are known to be able to determine the efficient allocation.

All of our results are in a very general setting: they do not rely on any incentive constraints and hold even if bidders simply bid "as told". As long as their response at every stage is just of function of the desired bundle at the current prices, the impossibility holds. Namely, even if one managed to motivate the players to truthfully respond to the queries in the auction, such ascending auction would obtain poor result. In particular, the impossibilities do not rely on any inter-dependencies between the bidders' valuations and hold for simple private values. Our impossibility results do not assume that any particular type of equilibrium will be achieved upon termination, and hold whether or not any equilibrium is achieved – they allow taking into account the whole amount of information obtained during the auction. The results do not rely on any computational limitations or limitations on the amount of communication that is transmitted, and hold even if unbounded (and unrealistic) computation and communication capabilities are available to the auctioneer and bidders.

While most previous work on combinatorial auctions has actually studied specific types of auctions, a few other impossibility results have been shown that should be compared to ours. First, are the known theorems (see, e.g., [141, 82, 19, 102, 44]) that for general, non-substitutes valuations, certain types of competitive equilibria cannot be found without personalized bundle prices. Our results, on the other hand, do not assume that any type of equilibrium is reached. To strengthen the contrast, note that item-price *non*-ascending auctions *can* obtain the social optimum, despite the lack of any equilibrium (see Chapter 7 and [27]). Other related results are the communication lower bounds proved in [117] showing that exponential communication is required by any type of combinatorial auction for obtaining the optimum. These results are quantitative and are not delicate enough to qualitatively distinguish between different types of auctions, as we do here. Additionally, such lower bounds on the amount of the transmitted communication cannot be applied in our setting, as we show in the chapter's body that an amount of information that is exponentially greater than the number of items can be elicited by ascending auctions, even with item prices.

The closest result to ours, in spirit, is by [70] who showed that ascending anonymous item-price auctions can not come up with VCG prices even for (gross)-substitutes valuations, despite the fact that the social optimum can be achieved in such cases. In contrast, our impossibility is for just finding the optimum, or even a reasonable approximation, rather than calculating a particular set of prices. Additionally, in contrast to our results, the impossibility in [70] is very delicate, even in the generalized version we prove in Chapter 6 for non-anonymous auctions, and stops holding if

[1] Formally, we show that no better than a $4/\sqrt{m}$ fraction of welfare may be captured by each auction type, where m is the number of items.

multiple ascending rounds are allowed, as in [5].

The bottom line of this chapter is a formal analysis showing that simple combinatorial auction schemes that use only item prices or that use only anonymous prices do have severe informational limitations. This will not allow them to match the performance guarantees of the more complex schemes. The exact tradeoff between these limitations and the significant costs of the more complex scheme remains part of the "art" of combinatorial auction construction.

The structure of the rest of the chapter is as follows: in section 2 we formally present our model and definitions. Section 3 gives the impossibility results for item-price auctions, while section 4 gives the impossibility results for anonymous-price auctions. In the body of the chapter we provide the full (and simple) proofs of the basic impossibility theorems; the proofs of the stronger variants are postponed to the appendix. Appendix D.1 contains some definitions to be used in proofs that appear later in Appendices D.2 and D.3.

5.2 The Model

A seller wishes to sell a set M of m heterogeneous indivisible items to a set of n bidders. Each bidder i has a valuation function $v_i : 2^M \rightarrow \mathbb{R}_+$ that attaches a non-negative real value $v_i(S)$ for any bundle $S \subseteq M$. We assume two conventional assumptions on the preferences: (i) Free disposal (monotonicity), i.e., if $S \subset T$ then $v_i(S) \leq v_i(T)$. (ii) Normalization, i.e., $v_i(\emptyset) = 0$ for every bidder i.

The goal of the auctioneer is to find an *efficient allocation* of the items, that is, to find a partition $S_1, ..., S_n$ of the items that maximizes the *social welfare*, $\sum_{i=1}^{n} v_i(S_i)$. We do not study revenue maximization in this chapter.

In this work, we concentrate on iterative auctions where, at each stage, the auctioneer publishes a set of prices p for the bundles, and each bidder responds with her *demand* given the published prices, that is, a bundle S that maximizes her (quasi linear) utility $u_i(S, p) = v_i(S) - p(S)$, where $p(S)$ denotes the price of S under the price level p.[2] The stages of the auction are ordered by time, and at each stage, a single set of prices is presented to each bidder. The prices can be presented in different ways. For example, the seller can explicitly publish a price for each bundle, or use a succinct representation for the prices (e.g., by only publishing item prices). We touch several common representations below.

The specific *auction* is determined by the method that the auctioneer determines which prices will be presented to the bidders at each stage. The seller can determine the prices adaptively, i.e., as a function of the history of the published prices and responses. The seller can also use information gained from the responses of one bidder for determining the future prices for other bidders. At the end of the auction, the auctioneer analyzes the information received during the auction, and determines the final allocation accordingly. That is, the data that is available to the auctioneer at the end of the auction is exactly $\{(p_i^t, S_i^t) |$ for every bidder i and every stage $t\}$, where S_i^t denotes the demand of bidder i at stage t under the price vector p_i^t. To strengthen our results, and as opposed to most of the existing literature, we consider a general model where the allocation can be determined by all the information gathered during the auction, and not only

[2] All our results hold for any consistent tie-breaking rule by the bidders or by the auctioneer. Moreover, our result will also hold if every bidder i reports, at each stage, all the bundles that maximize her utility, i.e., her whole demand set $\{S \subseteq M \mid v_i(S) - p(S) \geq v_i(T) - p(T) \text{ for every } T \subseteq M\}$. An equivalent model allows the bidders to raise their "bids" on their desired bundles.

according the demands at the final stage of the auction. Note that, to strengthen our results, we do not assume any limitations on the power of the participants, except for information limitations. In particular, the auctioneer may be computationally unbounded (including, e.g., the ability to solve hard problems classified as "NP-hard" in the computer-science terminology).

This chapter centers on auctions with non-decreasing prices:

Definition 5.1. (Ascending auctions) *In an* ascending auction*, each bidder responds with his demand under every price level presented to him, and prices presented to the same bidder can only increase in time. Formally, let p be a set of prices presented to bidder i, and q be the prices for bidder i at a later stage in the protocol. Then, for all sets $S \subseteq M$, we have $q(S) \geq p(S)$.*

Two highly important factors in the design of ascending combinatorial auctions concern the representation of the prices. First, the seller might choose to present only prices for the individual items, or, with greater expressiveness, publish a price per every possible bundle. Another pricing decision is whether to present personalized prices for each bidder, or present every price level to all bidders.

Definition 5.2. (Item/Bundle prices) *An auction uses* item prices (or linear prices)*, if, at each stage, the auctioneer presents a price p_j for each item j, and the price of a set S is additive: $p(S) = \sum_{j \in S} p_j$. We say that an auction uses* bundle prices (or non-linear prices) *if each bundle S may have a different price $p(S)$ (which is not necessarily equal to the sum of the prices of the items in S).*

Definition 5.3. (Anonymous/Non-Anonymous prices) *An auction uses* anonymous *prices, if the prices seen by the bidders at any stage in the auction are the same, i.e., whenever a set of prices is presented to some bidder, the same set of prices is also presented to all other bidders. In auctions with non-anonymous (personalized) prices, each bidder i is presented with personalized prices for the bundles denoted by $p_i(S)$.*[3]

Observe that with bundle prices, the number of distinct prices presented by the seller in each stage may be exponentially greater than the number of items (a price per every subset of items). Consequently, such auctions may be practically infeasible when selling more than few items.

5.3 Limitations of Item-Price Ascending Auctions

Before describing their limitations, we would like to demonstrate that item-price ascending auctions are not trivial in their power. The most prominent example is their ability to end up with a Walrasian equilibrium (which is, in particular, efficient) for environments with (gross) substitutes valuations, see [82] and [69].

We would also like to point out that despite using a linear number of item prices, ascending auctions may elicit a significantly larger amount of information from the bidders. Namely, if small enough increments are allowed, such auctions can elicit an amount of information that exceeds the number of items by an exponential factor. This is shown in Example D.1 in Appendix D.2.

Without restricting the prices to be ascending, analyzing the demand of the bidders under different price levels enables the auctioneer to easily determine the efficient allocation in any combinatorial auction (see Chapter 7 and [27]). However, as we show in this section, this is no longer

[3]Note that a non-anonymous auction can clearly be simulated by n parallel anonymous auctions.

	$\mathbf{v(ab)}$	$\mathbf{v(a)}$	$\mathbf{v(b)}$
Bidder 1	2	$\alpha \in (0,1)$	$\beta \in (0,1)$
Bidder 2	2	2	2

Figure 5.1: This example shows that no item-price ascending auction can always determine the optimal allocation: no such auction can tell whether α is greater than β or vice versa.

true when the prices are restricted to be ascending, even for settings with only two items and two bidders. After proving this negative result, we strengthen it in several directions: in Theorem 1a, we show that the number of ascending trajectories of prices that are required for finding the efficient allocation is exponentially larger than the number of items; Then, in Theorem 1b, we show that a single item-price ascending auction can only guarantee a small fraction of the optimal welfare, a fraction that diminishes with the number of items. Finally, Theorem 1c indicates that these impossibilities may rise for every profile of bidders with substitutes preferences following an addition of a single bidder.

An example for a combinatorial auction that cannot be solved by any ascending auction is given in Figure 5.1. In this example, for determining the efficient allocation, the auctioneer has to know which one of the two singleton bundles has a greater value for Bidder 1. However, an ascending auction can only elicit information about one of the singletons, so the efficient outcome cannot be obtained. The basic idea is that in order to gain any information about one of the singletons, the price of the *other* item must be increased significantly, otherwise the bidder will continue demanding the whole bundle. Since the prices cannot decrease, it follows that the demand of Bidder 1 will be independent of his value for the latter item.

Theorem 5.1. *No item-price ascending auction can determine the efficient allocation for all profiles of bidder valuations.*

Proof. Consider the two valuations described in Figure 5.1. All the values are known to the auctioneer, except for the values α and β (between $(0,1)$) that Bidder 1 attaches to the singletons a and b, respectively. For such preferences, the only way to achieve a welfare greater than 2 is to allocate one singleton to Bidder 1 and the other to Bidder 2. Therefore, the identity of the efficient allocation depends on which of the two singletons gains a greater value for Bidder 1. We prove that a single ascending trajectory of item prices can reveal information only on one of these values. We first claim that no information is elicited as long as both prices are low.

Claim 5.1. *As long as p_a and p_b are both below 1, Bidder 1 demands the whole bundle $\{ab\}$.*

Proof. For every price level p in which both prices are smaller than 1, Bidder 1's utility from the bundle ab will be strictly greater than the utility from either a or b separately. For example, we show that $u_1(ab,p) > u_1(a,p)$ (the same statement for the singleton b can be similarly shown):

$$\begin{aligned} u_1(ab,p) &= 2-(p_a+p_b) && (5.1)\\ &= 1-p_a+1-p_b && (5.2)\\ &> v_A(a)-p_a+1-p_b && (5.3)\\ &> u_1(a,p) && (5.4) \end{aligned}$$

Where Equation 5.1 is due to the linearity of the prices, Inequality 5.3 holds since $v_A(a)$ is smaller than 1, and Inequality 5.3 follows from the assumption that p_b is smaller than 1. □

Thus, in order to gain *any* information about the unknown values α and β, the auctioneer must arbitrarily (i.e., without any new information) choose one of the items (w.l.o.g., a) and increase its price to be greater than 1. But then, since the prices are ascending, the singleton a will not be demanded by Bidder 1 throughout the auction, thus no information at all will be gained about α. Hence, the auctioneer will not be able to identify the efficient allocation.

Since the valuation of one of the bidders is fully known in advance to the auctioneer, the theorem holds even for *non-anonymous* item-price ascending auctions. □

The proof of Theorem 5.1 describes a profile of preferences for which no ascending trajectory of prices can elicit enough information for determining the optimal allocation. This would hold even if the auctioneer had some exogenous information (or a "good guess") telling him what is the "right" way to increase the prices.[4] Similar arguments show that this hardness result also holds for the similar family of *descending-price* auctions (otherwise, the "reversed" price trajectory would be an ascending auction that finds the optimal allocation).

While Theorem 5.1 proved that a single ascending trajectory of prices cannot guarantee finding the efficient allocation, it does not rule out the possibility that a small number of trajectories will achieve this goal. For instance, a similar question was studied regarding the number of ascending auctions that are required for calculating VCG prices for bidders with substitutes preferences: A negative result by [70] showed that the VCG payments for substitutes valuations cannot be found by a single ascending-price trajectory; However, [5] presented an $(n+1)$-trajectory ascending auction that achieved this task. Below, we extend the result presented in Theorem 5.1 and show that for guaranteeing that an efficient allocation will be discovered, for any profile of valuations, an exponential number (in the number of items) of ascending-price trajectories is required.

We define a *k-trajectory ascending auction* as an auction in which the price vectors presented to the bidders at the different stages of the auction can be partitioned into up to k sets, ordered according to the time they were published, where the prices published within each set only increase in time (for a formal definition, see Definition D.4 in Appendix D.2). Note that we use a general definition; It allows the trajectories to run in parallel or sequentially, and to use information elicited in some trajectories for determining the future queries in other trajectories.

The theorem is proved by presenting preferences for two bidders, where the efficient allocation depends on the identity of a particular $\frac{m}{2}$-sized bundle that gains one of the bidders a high value. For eliciting information about the value of some $\frac{m}{2}$-sized bundle S, the prices of all the items that are not in S should be very high, otherwise a larger bundle would be demanded. Therefore, every ascending auction can only reveal information on a *single* $\frac{m}{2}$-sized bundle. Since an exponential number of such bundles exists, the theorem follows. The proof can be found in Appendix D.2.

Theorem 5.1.a. *The number of ascending item-price trajectories needed for revealing the efficient allocation, for every profile of bidder valuations, must be exponentially greater than the number of items.*

The next question is whether item-price ascending auctions can find an allocation with a close-to-optimal welfare. Again, we present a strong negative answer to this question. We show that no item-price ascending auction can guarantee more than a vanishing fraction the optimal welfare. Formally, we show that such auctions cannot guarantee a fraction of the efficient welfare that is

[4] Protocols that allow the usage of an exogenous data are often named "non-deterministic" protocols in the computer-science literature.

greater than $\max\{\frac{4}{n}, \frac{4}{\sqrt{m}}\}$.[5] We emphasize that this result even holds for non-anonymous item-price ascending auctions, that is, auctions with a personalized ascending trajectory of prices per each bidder.

A sketch of the proof: we create a profile of valuations for the n bidders with certain combinatorial properties that make them hard to be elicited by any ascending auction. This is done by defining a set of bundles that form a special combinatorial structure: we divide the items to several partitions; Every two bundles from different partitions intersect (*"mutually-intersecting partitions"*), and therefore achieving the optimal allocation is possible only by partitioning the items according to one of these partitions. The values that each bidder attaches for these bundles are unknown to the auctioneer and are either 0 or 1. To gain any information about one of these bundles, the prices of *every* bundle from all the other partitions must exceed 1 (since the bidders have a value of 2 for some larger bundles). It follows that the bundles from the other partitions will not be demanded any more during the ascending-price auction. This way, the auctioneer can elicit information about bundles from at most one partition for each bidder. This is shown to be insufficient for achieving a reasonable approximation for the social welfare. The proof appears in Appendix D.2.

Theorem 5.1.b. *No item-price ascending auction can guarantee better than a fraction of* $\max\{\frac{4}{n}, \frac{4}{\sqrt{m}}\}$ *of the efficient welfare for all profiles of bidder valuations.*

Our final result regarding item-price ascending auctions illustrates how our impossibility results hold even for preferences that are slightly away from having the substitutes property. Substitutes preferences are, informally, preferences with the property that when a bidder demands a certain bundle, and some of the prices in this bundle increase, then the bidder will still demand the other items in this bundle (an exact definition is presented in Definition D.3 in Appendix D.2). As mentioned, it is well known that item-price ascending auctions can determine the efficient allocation for substitutes valuations. We show that for every profile of players with substitutes valuations, the efficient outcome cannot be found after an addition of a single player. The proof takes advantage of the fact that the aggregate demand of n substitutes valuations also has the substitutes property. Therefore, the marginal contributions of bundles to the welfare of the n players must have complementarities. We construct a valuation for the new player that obtains a greater value than the marginal values for some of the bundles. Due to the presence of complementarities, we argue that an ascending auction will not be able to determine which bundle obtains the highest additional gain.

This result applies for any profile of substitutes valuations, except for the degenerate case where the aggregation of these players forms an additive valuation (i.e., where for every two disjoint bundles S, T, the aggregate valuation exhibits exactly $v(S) + v(T) = v(S \cup T)$).[6]

Theorem 5.1.c. *For every n, and for every profile of n substitutes valuations that their aggregation is not an additive valuation, there exists an additional bidder such that no item-price ascending auction can determine the efficient allocation among the $n+1$ bidders.*

[5] From a computer-science perspective, this is a strong indication that a "non-trivial" approximation guarantee cannot be achieved by item-price ascending auctions.

[6] A valuation w is called the aggregation of the valuations $v_1, ..., v_n$ if for every bundle S, $w(S)$ equals the optimal welfare achieved by allocating the items in S over the n players.

Bidder 1	$v_1(ac) = 2$	$v_1(bd) = 2$	$v_1(cd) = \alpha \in (0, 1)$
Bidder 2	$v_2(ab) = 2$	$v_2(cd) = 2$	$v_2(bd) = \beta \in (0, 1)$

Figure 5.2: This example shows that anonymous ascending auctions cannot always determine the efficient allocation. The value of every bundle that is not explicitly specified equals to the maximal value of a bundle it contains.

5.4 Limitations of Anonymous Ascending Auctions

All the ascending auctions in the literature that are proved to find the optimal allocation for unrestricted valuations are non-anonymous bundle-price auctions (e.g., iBundle(3) by [124] and the "Proxy Auction" by [4]). Yet, several *anonymous* ascending auctions with bundle prices have been suggested (e.g., AkBA by [149], the PAUSE auction by [81], and iBundle(2) by [124]). The power of such anonymous auctions is not trivial, as they can reach an efficient outcome for super-additive preferences ([123]). We first show that no anonymous ascending auction can always find the efficient solution for general valuations, even for environments with only two bidders and four items, and even if it is allowed to use bundle prices. Later in this section, we extend this negative result and show that such auctions can only guarantee a diminishing fraction of the social welfare.

In Figure 5.2, we present a class of valuations for which the efficient allocation cannot be determined by any anonymous bundle-price ascending auction. The basic idea: In the example, Bidder 1 and Bidder 2 have unknown values for some bundles S_1 and S_2, respectively. However, Bidder 1 also has a high value for S_2 and bidder 2 has a high value for S_1. Therefore, in order to reveal information about $v_1(S_1)$, the price of S_2 must be increased significantly, and thus "hide" the value $v_2(S_2)$. Similarly, for gaining information about $v_2(S_2)$ the price of S_1 must increase, "hiding" the value $v_1(S_1)$. This stems from the anonymity of the auction – the bidders face the same ascending trajectory of prices. Consequently, the auctioneer will only be able to attain information about both values, what will prevent him from identifying the efficient allocation.

Theorem 5.2. *No anonymous bundle-price ascending auction can determine the efficient allocation for all profiles of bidder valuations.*

Proof. Consider the pair of valuations described in Figure 5.2. Each bidder has a value of 2 for two 2-item bundles, and some unknown value, between 0 and 1, for a third 2-item bundle. The values of the other bundles equals the maximal value of a bundle that they contain. For finding the optimal allocation the auctioneer must know whether α is greater than β or vice versa: If $\alpha > \beta$, the optimal allocation will allocate *cd* to Bidder 1 and *ab* to bidder 2. Otherwise, it should allocate *bd* to bidder 2 and *ac* to Bidder 1. Notice that since each item can be allocated only once, at most one bidder can gain a value of 2.

In an anonymous ascending auction, however, one can only elicit information on one of the values α and β: as long as the prices of both *cd* and *bd* are below 1, both bidders will clearly demand their high-valued bundles (that gain them utilities greater than 1). Therefore, in order to elicit any information, the auctioneer must raise one of these prices to be greater than 1, w.l.o.g., the price of *bd*. Thus, since the prices cannot decrease, no information will be gained about β. □

We now strengthen the impossibility result above by showing that anonymous auctions, even with bundle prices, cannot guarantee more than a vanishing fraction of the social welfare, namely, at most a $\frac{4}{\sqrt{m}}$-fraction of the efficient welfare. This result may indicate that using anonymous bundle prices is wasteful; such auctions may potentially use an exponential number of prices (a price for each subset of items), although a similar fraction of the optimal welfare, $O(\frac{1}{\sqrt{m}})$, can be achieved using a significantly smaller amount of prices (that is, with polynomial-sized communication - see, e.g., Chapter 7).

For proving the limitations of anonymous auctions, we build a profile of valuations that, due to certain combinatorial properties, cannot be solved by anonymous ascending auctions. These preferences are different than those used in Theorem 1b. Nevertheless, we use the same combinatorial construction of *mutually-intersecting partitions* that was introduced in the proof for Theorem 1b. Recall that *mutually-intersecting partitions* are a set of partitions of the items with the property that every two bundles from different partitions have at least one item in common. We show that for the class of valuations that we build, before the auctioneer elicits any information, the prices of *all the bundles* from some partition should exceed 1. Since all the unknown values are below 1, an anonymous ascending auction will gain no information about the values that the bidders have for the bundles in this partition. Allocating bundles from this partition to different bidders may form an efficient allocation, but the auctioneer will not have enough information to correctly match those bundles to the bidders. We refer the reader to the full proof in Appendix D.3.

Theorem 5.2.a. *No anonymous ascending auction can guarantee better than a fraction of* $\max\{\frac{4}{n}, \frac{4}{\sqrt{m}}\}$ *of the efficient welfare for all profiles of bidder valuations, even when it uses bundle prices.*

5.5 Conclusions

This chapter considered ascending-price auctions for combinatorial auctions. It presented several impossibility results, providing insights about the power of different pricing models for such auctions. The chapter showed that both bundle prices and personalized prices are necessary in order to achieve efficient, or even approximately efficient, outcomes by ascending combinatorial auctions. Proposals for other kinds of ascending auctions carry the burden of proof for showing that good results can occur in their particular settings.

Chapter 6

More on the Power of Ascending Combinatorial Auctions

6.1 Introduction

Most of the suggested mechanisms for combinatorial auctions, most prominently in the design of spectrum auctions, maintain an ascending-price property: they start with a low price level, and increase prices while observing the demands of the players at each stage, until an allocation is announced. Chapter 2.3.2 of this dissertation surveys several families of ascending auctions, and Chapter 5 gives a systematic analysis of the information that may be elicited by every type of ascending auctions. In this chapter, we continue to characterize the power of ascending combinatorial auctions, with an emphasis on item-price auctions. The discussion also handles incentive issues of the bidders. This chapter uses the model and notations as in Chapters 2.3.1 and 5.

We begin with a discussion, in Section 6.2, of an important and well-studied class of valuations: *substitutes* (or gross-substitutes) valuations. These valuations, for settings with discrete goods, were defined in [82]. Remarkably, a simple ascending item-price auction ends up with the optimal allocation for such valuations for any number of bidders [82, 45, 70]. We first study a variant of the classic "integrability problem" from the economic literature (see, e.g., [98]) with respect to substitutes valuations: can we discover the preference of a single bidder that possesses a substitutes valuation according to his demand along an ascending path of prices? We present a negative answer: the preferences of such a player cannot be learned by an ascending auction, or even by $m/2$ separate ascending auctions (where m is the number of items). This may come as a surprise since the optimal allocation for such bidders can be computed by an ascending auctions for arbitrary large number of bidders.

We then examine whether ascending auctions for bidders with substitutes bidders can also compute VCG payments; this would result in an ex-post Nash equilibrium in these auctions. Gul and Stacchetti [70] proved that this cannot happen using a single path of ascending prices, but Ausubel [5] presented a neat auction that computes VCG prices using $n+1$ ascending paths (n denotes the number of players). Our work strengthen the results of [70] and justifies the multi-trajectory auction of Ausubel, as we show that even *non-anonymous* ascending auctions, which are composed from n separate price trajectories (one for each bidder), cannot compute VCG prices. This solves an open problem from [44]. The above negative results do not assume any computational constraints or rational behavior on behalf of the players and the auctioneer. As in Chapter 5, these results are only derived from analyzing the information requirements in such protocols.

Actually, we prove stronger versions of the above two negative results. They are proved for a sub-class of substitutes valuations that includes all the valuations that aggregate the demand of unit-demand bidders (unit-demand bidders are interested in at most one item). We also show, for the first time, that every valuation from this class (denoted as "OXS" valuations in [93]) can be succinctly represented by at most m^2 values. This holds although each valuation may be the aggregation of an unlimited number of unit-demand valuations. Whether substitutes valuations have a succinct representation (e.g., by sub-exponential number of values) has been an open problem for a long time, and this is one main reason for our interest in this question. We believe that the proving that aggregations of unit-demand valuations has a succinct representation is one step towards the solution for the substitutes case.

The design of ascending auctions involves many technical decisions about the auction procedure. For example, the auction may use anonymous prices or non-anonymous prices, update the prices adaptively or obliviously, and run sequentially or simultaneously between the bidders. We present several results that separate the power of different variants of ascending auctions. In Section 6.3 we present classes of valuations that can be solved by one type of auction, but not by others. Most of these results are non-surprising, but require non-trivial constructions and proofs. Since such decisions have other implications, like their effect on the strategic behavior of the bidders, creating a clear hierarchy that will separate the power of the different models may assist the designers of such protocols. For example, we show combinatorial-auction settings where the optimal welfare can be found by ascending auctions, but cannot be solved by any descending auction. Another example is that non-anonymous ascending auction that run simultaneously between the bidders, and can share the information between them, can do better than simultaneous non-anonymous ascending auctions.

Finally, in Section 6.4, we present several positive results on ascending auctions. We show that several natural "pieces of information" can be disclosed by ascending auctions. These pieces of information are modeled as "queries" (see Chapters 2.3.1 and 7). For showing that ascending auctions can simulate the above queries, we prove a useful lemma showing that auctions that change prices continuously over time (or by ϵ at each stage) can reveal the values of *all* the bundles that were demanded at any stage during the auction. An interesting open question is whether socially-efficient item-price ascending auctions exist for wider classes of valuations than substitutes valuations. In the spirit of this question, we describe an item-price *descending* auction that achieves *at least half* of the optimal welfare for bidders with submodular valuations. This auction is a variant of an algorithm suggested in [93].

6.2 Ascending Item-Price Auctions for Substitutes Valuations

Substitutes (also known as gross-substitutes) valuations play a central role in the analysis of combinatorial auctions. Intuitively, in a substitutes valuation, increasing the price of certain items can not reduce the demand for items whose price has not changed.

Definition 6.1. *A valuation v_i satisfies the substitutes (or gross-substitutes) property if for every pair of item-price vectors $\overrightarrow{q} \geq \overrightarrow{p}$ (coordinate-wise comparison), and for every $D \in D_i(\overrightarrow{p})$, there exists a bundle $A \in D_i(\overrightarrow{q})$ such that for all $j \in D$ with $p_j = q_j$ we have that also $j \in A$.*

Given bidders with substitutes valuations, simple item-price ascending auctions can be used for determining the socially-efficient allocation, since a *tatonnement* procedure (see [82, 45, 69] and Chapter 2.3.2) converges to a Walrasian equilibrium. Informally speaking, such a procedure starts from low item prices, and raises the prices of over-demanded items at each stage until the demand equals supply.

6.2.1 Learning Substitute Valuations

One important unsolved question concerns the complexity of describing substitutes valuations. In particular, whether there exist a "succinct" representation for such valuations. In other words, can a bidder disclose the exact details on his valuation without conveying an exceptionally-large amount of information. Recall that a naive representation of a valuation in a combinatorial auction requires 2^m-1 values – a value per each non-empty subset of items. While the above question remains unsolved, we will present two related results.

First, we show that while there exists a simple ascending auction that computes the optimal allocation for an unlimited number of bidders, exactly *learning* the valuation of a single bidder by an ascending auction is impossible. We say that a protocol learns the valuations of a certain bidder, if it collects sufficient information for determining the value this bidder assigns to every subset of the items. This is a variant of the well-studied *integrability problem* in economics: can one recover the preferences of a player by observing his demand? (see, e.g., [77]).

Theorem 6.1. *No ascending auction can exactly learn every substitutes valuation. Moreover, substitute valuations cannot be exactly learned even by $\frac{m}{2}$ ascending-price trajectories ($m > 3$).*

We actually prove a stronger result, by proving this theorem for a sub-class of substitutes valuations – aggregations of unit-demand valuations (also called "OXS valuations" in [93]). A bidder has a unit-demand valuation if he is only interested in singletons:

Definition 6.2. *A valuation v is called a* unit-demand *valuation if $v(S) = max_{j \in S}\ v(\{j\})$ for all S. A valuation v is an* aggregation of unit-demand valuations *if there is a set of unit-demand valuation $v_1, ..., v_k$ such that $v(S) = max_{S_1,...,S_k} \sum_{i=1}^{k} v_i(S_k)$ where the maximum is taken over all allocations $S_1, ..., S_k$.*

A unit-demand valuation can be represented by the values it attaches to the singletons. An aggregation of unit-demand players can hence be represented by separately representing all the unit-demand valuations that it aggregates.

We will first illustrate the intuition for the proof by a simple 2-bidder 4-item example, and then turn to formally prove the theorem.

Example 6.1. *Consider 4 items denoted by a_1, a_2, b_1, b_2 and the two unit-demand valuations v_1 and v_2 described in Figure 6.1 where α and β are unknown to the auctioneer and drawn from $(0, 1)$. The aggregation of this pair of valuations, denoted by v, cannot be learned by a single item-price ascending auction. Specifically, one ascending auction cannot elicit information from this bidder both on α and on β. To see this, note that α and β only affect the value gained from the bundles a_1b_2 and a_2b_1, respectively. Consider now a price level p in which a_1b_2 is demanded by the bidder. At this price level, a_1b_2 will be preferred over the bundle b_2, thus $v(a_1b_2) - p_{a_1} - p_{a_2} \geq v(b_2) - p_{b_2}$. Since $v(a_1b_2) = 3 + \beta$ and $v(b_2) = 3$ we get that $p_{a_1} < 1$. Since a_1b_2 is also preferred over a_1a_2 at p, similar computations show that $p_{a_2} > 2$. Symmetrically, if a_2b_1 is demanded at a price level q then necessarily, $q_{a_2} > 2$ and $q_{a_1} < 1$. Thus, these two bundles cannot be demanded on the same ascending trajectory of prices and information will not be revealed on the two unknown values.*

	$v(a_1)$	$v(a_2)$	$v(b_1)$	$v(b_2)$
v_1	3	0	$\alpha \in (0,1)$	3
v_2	0	3	3	$\beta \in (0,1)$
v_3	3	0	0	0
v_4	0	3	0	0
v_5	0	0	3	0
v_6	0	0	0	3

Figure 6.1: We use these valuations to show that substitutes valuations cannot be learned by item-price ascending auctions, and that non-anonymous ascending auctions cannot always reveal the VCG prices.

Proof. **(Of Theorem 6.1)** Denote the m goods by $a_1, ..., a_{\frac{m}{2}}$ and $b_1, ..., b_{\frac{m}{2}}$ (assume $m \geq 4$). We will first define a set of $\frac{m}{2}$ unit-demand valuations: for every i $1 \leq i \leq \frac{m}{2}$, let v_i denote the unit-demand valuation in which $v_i(a_i)$ is 3, the value of any singleton b_j is 3 for $j \neq i$, and $v_i(b_i)$ has an unknown value denoted by $\underline{b_i} \in (0,1)$; the values of all the other singletons are zero. All the unknown values $\underline{b_i}$ are informationally independent – having information on the realization of the value of some of them does not reveal any information on the values of the others.

Consider now the aggregation v of these $\frac{m}{2}$ unit-demand valuations. We will show that an ascending auction cannot learn this valuation.

Claim 6.1. *For every $1 \leq i \leq \frac{m}{2}$, the value of the bundle $\{b_i, a_1, ..., a_{i-1}, a_{i+1}, ..., a_m\}$ (denoted this bundle as $a_{-i}b_i$) depends on the realization of $\underline{b_i}$. The values of all the other bundles are independent of $\underline{b_i}$.*

Proof. Let B be some bundle in which the item b_i contributes $\underline{b_i}$ to the value. All the unit-demand valuations must contribute to the value, otherwise b_i could have contributed 3 in the unused unit-demand valuation. If there is a set of unit-demand valuations, except the ith valuation, in which the value is calculated according to the values of the b's, then a permutation of the b's among the terms must achieve a contribution of 3 for any of them. Thus, $|B| = \frac{m}{2}$, and it contains b_i and all a_j for $j \neq i$. □

Thus, in order to learn any information on the value $\underline{b_i}$, bidder i must have the opportunity to demand the bundle $a_{-i}b_i$ at least once along the ascending trajectory. Let $\overrightarrow{p} = (p_1, ..., p_m)$ be the vector of prices for which $a_{-i}b_i$ is demanded. Thus, the utility from this bundle is not smaller than the utility from the bundle $a_{-i-j}b_i$ (i.e., when we remove some item $j \neq i$ from $a_{-i}b_i$). Therefore, due to the linarity of prices, $\underline{b_i} + 3 - p_{b_i} - p_{a_j} \geq 3 - p_{b_i}$. We conclude that $p_{a_j} \leq \underline{b_i} < 1$. $a_{-i}b_i$ will also be preferred over the bundle $a_{-i}a_i$ (i.e., when replacing b_i with a_i). Thus, $\underline{b_i} - p_{b_i} \geq 3 - p_{a_i}$ and we get $p_{a_i} \geq 3 - \underline{b_i} + p_{b_i} \geq 3 - \underline{b_i} > 2$. We see that when $a_{-i}b_i$ is demanded, $p_{a_i} > 2$ and $p_{a_j} < 1$ for any $j \neq i$.

Now, let $\overrightarrow{q} = (q_1, .., q_m)$ be a price vector for which the bundle $a_{-j}b_j$ is demanded (for some $j \neq i$). From symmetry arguments, we have that $p_{a_j} > 2$ and $p_{a_i} < 1$. Therefore, p and q cannot be on the same ascending price trajectory, and only one of them could be demanded in an ascending auction. Since i, j were chosen arbitrarily, it follows that in a single ascending trajectory only the value of one of the b_i's will be disclosed. Since Claim 6.1 argued that the valuation v depends on the values of all the b_i's, we would need at least $\frac{m}{2}$ ascending trajectories to exactly learn v. □

6.2.2 Computing VCG Prices for Substitutes Valuations

Valuations that are aggregations of unit-demand valuations also exhibit the substitutes property [93]. Therefore, a Walrasian equilibrium may be found using an item-price ascending auction assuming that the bidders respond with their true demand at each price level. For the more restricted case of players with unit-demand valuations, such ascending auctions reach the lowest possible Walrasian-equilibrium prices that are also VCG prices, and hence these auctions are ex-post Nash incentive compatible. However, [70] showed that no ascending auction can compute the VCG prices for substitiute valuation. They showed that revealing some of the VCG payments requires high price levels that "hide" the other VCG payments. [70] prove their result for an ascending auction with a single ascending trajectory of prices. Ausubel [5], on the other hand, showed an auction with $n + 1$ ascending trajectories that computes the VCG payments. We strengthen the result of [70], and show that the added complexity in Ausubel's auction is required, since the VCG payments cannot be computed even by *non-anonymous* ascending auctions, i.e., where there are n separate ascending paths of prices, one per each bidder. Doing that, we answer an open question from [44]. Again, we prove a stronger claim by proving the claim even for aggregations of unit-demand valuations.

Theorem 6.2. *No item-price ascending auction can compute VCG payments for every profile of bidders with substitutes valuations, even with non-anonymous prices.*

Proof. Consider the items a_1, a_2, b_1, b_2 and the following three aggregations of unit-demand valuations: w_1 aggregates the valuations v_1 and v_2 in Figure 6.1 (as in Example 6.1); w_2 aggregates the valuations v_3 and v_6 from Figure 6.1; w_3 aggregates the valuations v_4 and v_5 from the same figure. These valuations clearly have the substitutes property, as they are aggregatations of unit-demand valuations.

The optimal welfare is 12. For calculating the VCG payment of the bidder whose valuation is w_3 we must find the optimal welfare when w_3 is excluded. This welfare is $9+\alpha$. When w_2 is excluded, the optimal welfare equals $9+\beta$. Therefore, for calculating the VCG payments for those 3 bidders one must exactly know the values of both α and β. But as argued in Example 6.1 (and proved in Theorem 6.1), no ascending auction can reveal both values. Since both unkown values are held by one player, even non-anonymous ascending auctions cannot complete this task. □

As mentioned, an intriguing open question is whether substitutes valuations have a succinct represetnation. While we were not able to answer this question, we do show that the sub-class of valuations that are aggregations of unit-demand valuations can be succinctly represented. Although such valuations may be defined by an aggregation of unlimited number of unit-demand players, we show that at most m unit-demand bidders really play a role when defining a particular valuation of this kind. Therefore, every valuation of this kind can be described by at most m^2 values – m values for each unit-demand valuation. The intuition is that at most m unit-demand valuations will contribute to the aggreagate value of the whole bundle M, and we can prove by induction that only those unit-demand players will contribute to the value of any subset of the items.

Lemma 6.1. *An aggregation of any number of unit-demand bidders can equivalently be defined as an aggregation of at most m unit-demand bidders. It follows that such valuations can always be represented by at most m^2 values.*

Proof. Consider a valuation v which is an aggregation of l unit-demand valuations $v_1, ..., v_l$, denoted as the "atomic" valuations. The value of some bundle S is defined by allocating each item to at most one of these atomic valuations, and in this case we say that these atomic valuations *contribute* to the value of S (or that the paricular item is *valuated* by the atomic valuation v_i).

We will prove that for any bundle S and any item $a \in S$, the set of atomic valuations that contribute value to $v(S \setminus a)$ is a subset of the atomic valuations that contribute to $v(S)$. Since clearly at most m atomic valuations contribute to the value of M, it follows (inductively) that a subset of these atomic valuations will contribute to the value of every bundle, and the lemma follows.

Consider some bundle of k items $S = (a_1, ..., a_k)$, and denote, w.l.o.g., $v_1, ..., v_k$ as the atomic valuations that contribute to its value, respectively. When valuating the bundle $S \setminus a_1$, if no item is valuated in the atomic valuation v_1, then the other items $a_2, ..., a_k$ are valuated by exactly the same atomic valuations as in valuating S (otherwise the value for S could have increased). If v_1 contributes to the value of $S \setminus a_1$, w.l.o.g by valuating the item a_2, then we check if the atomic valuation v_2 valuates one of the items $a_3, ..., a_k$, and we proceed (by induction), until reaching the atomic valuation v_j that does not contribute value to any item (and then $a_{j+1}, ..., a_k$ are valuated using by the same atomic valutations as in $v(S)$), or until the last item a_k is valuated (by the valuation v_{k-1}). (We assume, w.l.o.g., that the items are indexed in the order derived by the proof.) Therefore, in any case, every item will be valuated by an atomic valuation from $\{v_1, ..., v_k\}$. $\square$

6.3 Separation Results for Different Types of Ascending Auctions

The design of iterative combinatorial auctions involves many decision on the auction format. Every decision has a significant effect on the strategic behavior of the players, and also on the information that these auctions can elicit on the bidders' preferences. In this section, we would like to separate the power of different variants of item-price ascending auctions, and show that some variants are strictly stronger than others and some variants are incomparable in their power. The following theorem present several separation results regarding the auctions' ability to determine the socially-efficient allocation. Some of the auction types mentioned in the theorem (items 1 and 2) were previously defined and the other types (in items 3-5) are described right after the theorem. The proofs are given by separate theorems in Appendix E.1. The proof for each separation result constructs a class of valuations from which the valuations are drawn and that derive the separations.

Theorem 6.3. *Consider item-price ascending auctions. There exist classes of valuations such that:*

1. *The efficient allocation can be determined by an* ascending *auction but not by a* descending *auction, and* vice versa.

2. *The efficient allocation can be determined by a* non-anonymous *ascending auction but not by an* anonymous *ascending auction.*

3. *The efficient allocation can be found by a* non-deterministic *ascending auction, but not by a* determinisitic *ascending auction.*

4. *The efficient allocation can be determined by a* simultaneous *non-anonymous ascending auction but not by a* sequential *non-anonymous ascending auction.*

5. *The efficient allocation can be determined by an* adaptive *ascending auction but not by an* oblivious *ascending auction.*

Deterministic vs. Non-Deterministic Auctions: Non-deterministic ascending auctions can be viewed as auctions where some benevolent teacher that has complete information guides the auctioneer on how she should raise the prices. That is, the optimal allocation can be obtained by a non-deterministic ascending auction, if there is *some* ascending trajectory that elicits enough information for determining the optimal allocation (and verifying that it is indeed optimal). The concept of non-deterministic protocols is central in computer-science theory (see, e.g., [120]), and also appears in different ways in the economic literature (like the concept of "competitive equilibria", or, e.g., [133]).

Sequential vs. Simultaneous Non-Anonymous Auctions: A non-anonymous auction is called *simultaneous* if at each stage, the price of some item is raised by ϵ for *every* bidder. The auctioneer can use the information gathered until each stage, in all the personalized trajectories, to determine the next queries.

A non-anonymous auction is called *sequential* if the auctioneer performs an auction for each bidder separately, in sequential order. The auctioneer can still determine the next query based on the information gathered in the trajectories completed so far and on the history of the current trajectory.

Adaptive vs. Oblivious Auctions: If the auctioneer determines the queries regardless of the bidders' responses (i.e., the queries are predefined) we say that the auction is *oblivious*. Otherwise, the auction is *adaptive*. We prove that an adaptive behaviour of the auctioneer may be beneficial.

6.4 Some Positive Results on the Power of Ascending Auctions

In this section, we illustrate several examples of the capabilities of item-price ascending auctions and their close variants.

Simulating Queries by Ascending auctions

In Chapter 7 we will show how several natural types of "queries" in combinatorial auctions can be simulated by a polynomial number of demand queries (polynomial in n, m and in the number of bits required to represent the values of the bundles). One example is a "value query", that inquires for the bidder's value for a given bundle S. Here, we show that these queries can be simulated by series of demand queries with *ascending* prices. The queries are defined in Section 7.2 of Chapter 7, and the proposition is proved in Appendix E.2.

Proposition 6.1. *Every value query, marginal-value query, indirect-utility query and relative-demand query can be simulated by a single ascending trajectory of item-price demand queries. The number of queries required is polynomial in* n, m *and* L/δ*, where* L *is an upper bound on the values for the bundles, and all the values are multiples of* δ*.*

The proof that a single ascending auction can simulate any value query (in the proof of Proposition 6.1 in Appendix E.2) actually proves a stronger, useful result regarding the information elicited by iterative auctions. This result says that in *any* iterative auction in which the changes of prices

A descending auction for bidders with submodular valuations: **Initialization:** set all item prices to L. Let X_i be the current items allocated to bidder i, and for each bidder initialize $X_i \leftarrow \emptyset$. **Repeat:** For all items $i = 1, ..., m$ (the items are arbitrarily ordered), decrease the price p_i of item i by $\epsilon = \delta$. Allocate the item to the first bidder j that demands his current bundle X_j together with item i (i.e., $X_j \leftarrow X_j \cup \{i\}$).

Figure 6.2: This item-price descending auction guarantees at least $\frac{1}{2}$ of the optimal welfare for submodular valuations. We do not know if there is an ascending auction achieving the same approximation ratio.

are small enough at each stage ("ϵ-continuous" auctions), the value of *all* the bundles that were demanded, even once, during the auction can be computed. The basic idea is that when the bidder moves from demanding some bundle T_i to demanding another bundle T_{i+1}, there is a point in which she is indifferent between these two bundles. Thus, knowing eventually the value of some demanded bundle (e.g., the empty set) enables computing the values of all other bundles that were demanded.

We say that an auction is *"ϵ-continuous"*, if it only uses demand queries, and at each step, the price of at most one item is changed by ϵ (for some $\epsilon \in (0, \delta]$, where, again, all the values are multiples of δ) with respect to the previous query. Note that the definition does not require that the prices will be ascending.

Proposition 6.2. *Consider any ϵ-continuous auction (not necessarily ascending) in which bidder i demands the empty set at least once during the auction. Then, the value of every bundle demanded by bidder i throughout the auction can be calculated at the end of the auction up to an error of ϵ.*

A Descending Auction for Submodular Valuations

As mentioned, item-price ascending auctions are well known for their ability to compute a socially-efficient Walrasian equilibrium for substitute valuations. The immediate question is whether a similar auction can be designed for the well-studied super-class of submodular valuations.

Definition 6.3. *A valuation v is* submodular *if for every two bundles S, T we have that*

$$v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$$

An equivalent, yet more intuitive, definition of submodular valuations says that a valuation is submodular if and only if it exhibits non-increasing marginal values: for every two bundles S and T such that $S \subseteq T$, and for every item x, we have that $v(S \cup x) - v(S) \geq v(T \cup x) - v(T)$. Namely, the marginal contribution of an item to a bundle decreases as the bundle expands. While an ascending-price auction can compute the optimal allocation for submodular bidders remains an open question, we are able to show that at least half of the optimal allocation can always be found by a *descending*-price auction. This is done by an adaptation of the greedy algorithm described in [93]. The descending auction is illustrated in Figure 6.2 and a proof is given in Appendix E.2. Whether such an approximation can be obtained by an *ascending* auction remains an open problem.

Proposition 6.3. *For any profile of sub-modular valuations, the descending auction described in Figure 6.2 achieves at least $\frac{1}{2}$ of the social welfare.*

Chapter 7

On the Power of Iterative Auctions: Demand Queries

7.1 Introduction

Combinatorial auctions have recently received a lot of attention. In a combinatorial auction, a set M of m non-identical items are sold in a single auction to n competing bidders. The bidders have preferences regarding the *bundles of items* that they may receive. The preferences of bidder i are specified by a valuation function $v_i : 2^M \rightarrow R^+$, where $v_i(S)$ denotes the value that bidder i attaches to winning the bundle of items S. We assume "free disposal", i.e., that the v_i's are monotone non-decreasing. The usual goal of the auctioneer is to optimize the social welfare $\sum_i v_i(S_i)$, where the allocation $S_1 ... S_n$ must be a partition of the items. Applications include many complex resource allocation problems and, in fact, combinatorial auctions may be viewed as *the* common abstraction of many complex resource allocation problems. Combinatorial auctions face both economic and computational difficulties and are a central problem in the recently active border of economic theory and computer science. A recent book [37] addresses many of the issues involved in the design and implementation of combinatorial auctions.

The design of a combinatorial auction involves many considerations. In this chapter we focus on just one central issue: the communication between bidders and the allocation mechanism – "preference elicitation". Transferring all information about bidders' preferences requires an infeasible (exponential in m) amount of communication. Thus, "direct revelation" auctions in which bidders simply declare their preferences to the mechanism are only practical for very small auction sizes or for very limited families of bidder preferences. We have therefore seen a multitude of suggested "iterative auctions" in which the auction protocol repeatedly interacts with the different bidders, aiming to adaptively elicit enough information about the bidders' preferences as to be able to find a good (optimal or close to optimal) allocation.

Most of the suggested iterative auctions proceed by maintaining temporary prices for the bundles of items and repeatedly querying the bidders as to their preferences between the bundles under the current set of prices, and then updating the set of bundle prices according to the replies received (e.g., [82, 45, 70, 124, 4]). Effectively, such an iterative auction accesses the bidders' preferences by repeatedly making the following type of *demand query* to bidders: "Query to bidder i: a vector of bundle prices $p = \{p(S)\}_{S \subseteq M}$; Answer: a bundle of items $S \subseteq M$ that maximizes $v_i(S) - p(S)$.". These types of queries are very natural in an economic setting as they capture the "revealed preferences" of the bidders. Some auctions, called *item-price* or *linear-price* auctions, specify a price p_i for each *item*, and the price of any given bundle S is always linear, $p(S) = \sum_{i \in S} p_i$. Other auctions, called *bundle-price* auctions, allow specifying arbitrary (non-linear) prices $p(S)$ for bundles.

In this chapter, we embark on a systematic analysis of the computational power of iterative auctions that are based on demand queries. We do not aim to present auctions for practical use but rather to understand the limitations and possibilities of these kinds of auctions. Our main question is what can be done using a polynomial number of these types of queries? That is, polynomial in the main parameters of the problem: n, m and the number of bits t needed for representing a single value $v_i(S)$. Note that from an algorithmic point of view we are talking about sub-linear time algorithms: the input size here is really $n(2^m - 1)$ numbers – the descriptions of the valuation functions of all bidders. There are two aspects to computational efficiency in these settings: the first is the communication with the bidders, i.e., the number of queries made, and the second is the "usual" computational tractability. Our lower bounds will depend only on the number of queries – and hold independently of any computational assumptions like $P \neq NP$. Our upper bounds will always be computationally efficient both in terms of the number of queries and in terms of regular computation. As mentioned, this chapter concentrates on the single aspect of preference elicitation and on its computational consequences and does not address issues of incentives. This strengthens our lower bounds, but means that the upper bounds require evaluation from this perspective also before being used in any real combinatorial auction.[1]

In Chapters 5 and 6 we studied similar questions for the more restricted natural case of *ascending-price* combinatorial auctions.

7.1.1 Extant Work

Many iterative combinatorial auction mechanisms rely on demand queries (see the survey in [122]). For our purposes, two families of these auctions serve as the main motivating starting points: the first is the ascending item-price auctions of [82, 70] that with computational efficiency find an optimal allocation among "(gross) substitutes" valuations[2], and the second is the ascending bundle-price auctions of [124, 4] that find an optimal allocation among general valuations – but not necessarily with computational efficiency. The main lower bound in this area, due to [117], states that indeed, due to inherent communication requirements, it is not possible for any iterative auction to find the optimal allocation among general valuations with sub-exponentially many queries. A similar exponential lower bound was shown by [117] also for even approximating the optimal allocation to within a factor of $m^{1/2-\epsilon}$. Several lower bounds and upper bounds for approximation are known for some natural classes of valuations – these are summarized in Figure 7.1.

In [117], the universal generality of demand queries is also shown: any *non-deterministic* communication protocol for finding an allocation that optimizes the social welfare can be converted into one that only uses demand queries (with bundle prices). In [143] this was generalized also to non-deterministic protocols for finding allocations that satisfy other natural types of economic

[1] We do observe however that some weak incentive property comes for free in demand-query auctions since "myopic" players will answer all demand queries truthfully. We also note that in some cases (but not always!) the incentives issues can be handled orthogonally to the preference elicitation issues, e.g., by using Vickrey-Clarke-Groves (VCG) prices (e.g., [5, 121]).

[2] Informally, the substitutes property means that the bidder will continue to demand an item when the prices of some *other* items were raised. See exact definition in [82, 70].

Valuation family	Upper bound	Reference	Lower bound	Reference
General	$min(n, O(\sqrt{m}))$	Section 7.5	$min(n, m^{1/2-\epsilon})$	[117]
Substitutes	1	[117]		
Submodular	1.5818	[59]	$1+\frac{1}{2m}$	[117]
Subadditive	2	[58]	2-ϵ	[47]
k-duplicates	$O(m^{1/k+1})$	[33]	$O(m^{1/k+1})$	[49]
Procurement	$\ln m$	[117]	$(\log m)/2$	[112, 117]

Figure 7.1: The best approximation factors currently achievable by computationally-efficient combinatorial auctions, for several classes of valuations. All lower bounds in the table apply to all iterative auctions; all upper bounds in the table are achieved with item-price demand queries.

requirements (e.g., approximate social efficiency, envy-freeness). However, in [117] it was demonstrated that this "completeness" of demand queries holds only in the nondeterministic setting, while in the usual deterministic setting, demand queries (even with bundle prices) may be exponentially weaker than general communication.

Bundle-price auctions are a generalization of (the more natural and intuitive) item-price auctions. It is known that indeed item-price auctions may be exponentially weaker: a nice example is the case of valuations that are an XOR of k bundles[3], where k is small (say, polynomial). Lahaie and Parkes [87] show an economically-efficient bundle-price auction that uses a polynomial number of queries whenever k is polynomial. In contrast, [21] show that there exist valuations that are XORs of $k = \sqrt{m}$ bundles such that any item-price auction that finds an optimal allocation between them requires exponentially many queries.

The organization of the rest of the chapter is as follows: First, in Section 7.2, we present an informal exposition that describes our new results and their context. Section 7.3 describes our model. In Section 7.4 we discuss the power of different types of queries, and Section 7.5 studies the approximability of the social welfare with a polynomial number of queries. In Section 7.6, we show how demand queries enable solving linear programs for winner determination problems. Finally, Section 7.7 studies the representation of bundle-price demand queries.

7.2 A Survey of Our Results

7.2.1 Comparison of Query Types

We first ask what other natural types of queries could we imagine iterative auctions using? Here is a list of such queries that are either natural, have been used in the literature, or that we found useful.

1. *Value query*: The auctioneer presents a bundle S, the bidder reports his value $v(S)$ for this bundle.

2. *Marginal-value query*: The auctioneer presents a bundle A and an item j, the bidder reports how much he is willing to pay for j, given that he already owns A, i.e., $v(j|A) = v(A \cup \{j\}) - v(A)$.

[3]These are valuations where bidders have values for k specific packages, and the value of each bundle is the maximal value of one of these packages that it contains.

3. *Demand query (with item prices)*: The auctioneer presents a vector of item prices $p_1 ... p_m$; the bidder reports his demand under these prices, i.e., some set S that maximizes $v(S) - \sum_{i \in S} p_i$.[4]

4. *Indirect-utility query*: The auctioneer presents a set of item prices $p_1 ... p_m$, and the bidder responds with his "indirect-utility" under these prices, that is, the highest utility he can achieve from a bundle under these prices: $max_{S \subseteq M}(v(S) - \sum_{i \in S} p_i)$. [5] We apply this query, for example, when describing our welfare-approximation algorithm in Section 7.5.

5. *Relative-demand query*: the auctioneer presents a set of non-zero prices $p_1 ... p_m$, and the bidder reports the bundle that maximizes his value per unit of money, i.e., some set that maximizes $\frac{v(S)}{\sum_{i \in S} p_i}$.[6]

Theorem: Each of these queries can be efficiently (i.e., in time polynomial in n, m, and the number of bits of precision t needed to represent a single value $v_i(S)$) simulated by a sequence of demand queries with item prices.

In particular, this shows that demand queries can elicit all information about a valuation by simulating all $2^m - 1$ value queries. We also observe that value queries and marginal-value queries can simulate each other in polynomial time and that demand queries and indirect-utility queries can also simulate each other in polynomial time. We prove that exponentially many value queries may be needed in order to simulate a single demand query.[7]

7.2.2 Welfare Approximation

The next question that we ask is how well can a computationally-efficient auction that uses only demand queries *approximate* the optimal allocation? Two separate obstacles are known: In [117], a lower bound of $min(n, m^{1/2-\epsilon})$, for any fixed $\epsilon > 0$, was shown for the approximation factor obtained using any polynomial amount of communication. A computational bound with the same value applies even for the case of single-minded bidders, but under the assumption of $NP \neq ZPP$ [136]. As noted in [117], the computationally-efficient greedy algorithm of [95] can be adapted to become a polynomial-time iterative auction that achieves a nearly matching approximation factor of $min(n, O(\sqrt{m}))$. This iterative auction may be implemented with bundle-price demand queries but, as far as we see, not as one with item prices. Since in a single bundle-price demand query an exponential number of prices can be presented, this algorithm can have an exponential communication cost. In Section 7.5, we describe a different item-price auction that achieves, for the first time, the same approximation factor with a polynomial number of demand queries (and thus polynomial communication).

Theorem: There exists a computationally-efficient iterative auction with item-price *demand queries* that finds an allocation that approximates the optimal welfare between arbitrary valuations to within a factor of $min(n, O(\sqrt{m}))$.

[4] A tie breaking rule should be specified. All of our results apply for any fixed tie breaking rule.

[5] This is exactly the utility achieved by the bundle which would be returned in a demand query with the same prices. This notion relates to the Indirect-utility function studied in the Microeconomic literature (see, e.g., [98]).

[6] Note that when all the prices are 1, the bidder actually reports the bundle with the highest per-item price. We found this type of query useful, for example, in the design of the approximation algorithm described in Figure 7.4 in Section 7.5.

[7] It is interesting to note that for the restricted class of substitutes valuations, demand queries may be simulated by polynomial number of value queries [14].

Query type	Upper bound	Reference	Lower bound	Reference
General Communication	$min(n, O(m^{1/2}))$	[95]	$min(n, m^{1/2-\epsilon})$	[117]
Demand Queries	$min(n, O(m^{1/2}))$	**new**	$min(n, m^{1/2-\epsilon})$	[117]
Value Queries	$O(\frac{m}{\sqrt{\log m}})$	[76])	$O(\frac{m}{\log m})$	**new**

Figure 7.2: Achievable approximation factors for the social welfare using polynomially many value queries, demand queries (with item prices), and general queries (communication).

One may then attempt obtaining such an approximation factor using iterative auctions that use only the weaker value queries. However, we show that this is impossible:

Theorem: Any iterative auction that uses a polynomial (in n and m) number of *value queries* can not achieve an approximation factor that is better than $O(\frac{m}{\log m})$.[8]

Note however that auctions with only value queries are not completely trivial in power: the bundling auctions of [76] can easily be implemented by a polynomial number of value queries and can achieve an approximation factor of $O(\frac{m}{\sqrt{\log m}})$ by using $O(\log m)$ equi-sized bundles. We do not know how to close the (tiny) gap between this upper bound and our lower bound. Figure 7.2 summarizes these upper and lower bounds.

7.2.3 Demand Queries and Linear Programs

The winner determination problem in combinatorial auctions may be formulated as an integer program. In many cases solving the linear-program relaxation of this integer program is useful: for some restricted classes of valuations it finds the optimum of the integer program (e.g., substitute valuations [82, 70]) or helps approximating the optimum (e.g., by randomized rounding [47, 50, 22]). However, the linear program has an exponential number of variables. Nisan and Segal [117] observed the surprising fact that despite the exponential number of variables, this linear program may be solved within polynomial communication. The basic idea is to solve the dual program using the Ellipsoid method (see, e.g., [80]). The dual program has a polynomial number of variables, but an exponential number of constraints. The Ellipsoid algorithm runs in polynomial time even on such programs, provided that a "separation oracle" is given for the set of constraints. Surprisingly, such a separation oracle can be implemented using a single demand query (with item prices) to each of the bidders.

The treatment of [117] was somewhat ad-hoc to the problem at hand (the case of substitute valuations). Here we give a somewhat more general form of this important observation. Let us call the following class of linear programs "generalized-winner-determination-relaxation (GWDR)

[8] This was also proven independently by Shahar Dobzinski and Michael Schapira.

LPs":

$$\begin{aligned}
\textbf{Maximize} \sum_{i \in N, S \subseteq M} & w_i \, x_{i,S} \, v_i(S) & \\
\textbf{s.t.} \sum_{i \in N,\ S | j \in S} & x_{i,S} \leq q_j & \forall j \in M \\
\sum_{S \subseteq M} & x_{i,S} \leq d_i & \forall i \in N \\
& x_{i,S} \geq 0 & \forall i \in N, S \subseteq M
\end{aligned}$$

The case where $w_i = 1, d_i = 1, q_j = 1$ (for every i, j) is the usual linear relaxation of the winner determination problem. More generally, w_i may be viewed as the weight given to bidder i's welfare, q_j may be viewed as the quantity of units of good j, and d_i may be viewed as duplicity of the number of bidders of type i.

Theorem: Any GWDR linear program may be solved in polynomial time (in n, m, and the number of bits of precision t) using only demand queries with item prices.[9]

7.2.4 Representing Bundle Prices

Finally, we deal with a critical issue with bundle-price auctions that was side-stepped by our model, as well as by all previous works that used bundle-price auctions: how are the bundle prices represented? For item-price auctions this is not an issue since a query needs only to specify a small number, m, of prices. In bundle-price auctions that situation is more difficult since there are exponentially many bundles that require pricing. Our basic model (like all previous work that used bundle prices, e.g., [124, 121, 4]), ignores this issue, and only requires that the prices be determined, *somehow*, by the protocol. A finer model would fix a specific *language* for denoting bundle prices, force the protocol to represent the bundle-prices in this language, and require that the *representations of the bundle-prices* also be polynomial.

What could such a language for denoting prices for all bundles look like? First note that specifying a price for each bundle is equivalent to specifying a *valuation*. Second, as noted in [113], most of the proposed *bidding languages* are really just languages for representing valuations, i.e., a syntactic representation of valuations – thus we could use any of them. This point of view opens up the general issue of *which* language should be used in bundle-price auctions and what are the implications of this choice.

Here we initiate this line of investigation. We consider bundle-price auctions where the prices must be given as a XOR-bid, i.e., the protocol must explicitly indicate the price of every bundle whose value is different than that of all of its proper subsets. Note that all bundle-price auctions that do not explicitly specify a bidding language must implicitly use this language or a weaker one, since without a specific language one would need to list prices for all bundles, perhaps except for trivial ones (those with value 0, or more generally, those with a value that is determined by one of their proper subsets.) We show that once the representation length of bundle prices is taken into account (using the XOR-language), bundle-price auctions are no more strictly stronger than item-price auctions. Our proof relies on the sophisticated known lower bounds for constant depth circuits due to Hastad [72]. We were not able to find an elementary proof.

[9] The produced optimal solution will have polynomial support and thus can be listed fully.

Define the *cost* of an iterative auction as the total length of the queries and answers used throughout the auction (in the worst case).

Theorem: For some class of valuations, bundle price auctions that use the XOR-language require an exponential cost for finding the optimal allocation. In contrast, item-price auctions can find the optimal allocation for this class within polynomial cost.

This puts doubts on the applicability of bundle-price auctions like [4, 124], and it may justify the use of "hybrid" pricing methods such as Ausubel, Cramton and Milgrom's Clock-Proxy auction ([36]).

7.3 The Model

A single auctioneer is selling m indivisible non-homogeneous items in a single auction, and let M be the set of these items and N be the set of bidders. Each one of the n bidders in the auction has a valuation function $v_i : 2^M \to \{0, 1, ..., L\}$, where for every bundle of items $S \subseteq M$, $v_i(S)$ denotes the value of bidder i for the bundle S and is an integer in the range $0...L$. We will sometimes denote the number of bits needed to represent an integer in the range $0...L$ by $t = \log L$. We assume free disposal, i.e., $S \subseteq T$ implies $v_i(S) \leq v_i(T)$ and that $v_i(\emptyset) = 0$ for all bidders.

A valuation is called a *k-bundle XOR* if it can be represented as a XOR combination of at most k atomic bids [111], i.e., if there are at most k bundles S_i and prices p_i such that for all S, $v(S) = max_{i|S \supseteq S_i} p_i$.[10]

7.3.1 Iterative Auctions

The auctioneer sets up a protocol (equivalently an "algorithm"), where at each stage of the protocol some information q – termed the "query" – is sent to some bidder i, and then bidder i replies with some reply that depends on the query as well as on his own valuation. In this chapter, we assume that we have complete control over the bidders' behavior, and thus the protocol also defines a reply function $r_i(q, v_i)$ that specifies bidder i's reply to query q. The protocol may be adaptive: the query value as well as the queried bidder may depend on the replies received for past queries. At the end of the protocol, an *allocation* $S_1...S_n$ must be declared, where $S_i \cap S_j = \emptyset$ for $i \neq j$.

We say that the auction finds an *optimal allocation* if it finds the allocation that maximizes the social welfare $\sum_i v_i(S_i)$. We say that it finds a c-approximation if $\sum_i v_i(S_i) \geq \sum_i v_i(T_i)/c$ where $T_1...T_n$ is an optimal allocation. The running time of the auction on a given instance of the bidders' valuations is the total number of queries made on this instance. The running time of a protocol is the worst case cost over all instances. Note that we impose no computational limitations on the protocol or on the players.[11] This of course only strengthens our hardness results. Yet, our positive results will not use this power and will be efficient also in the usual computational sense.

Our goal will be to design computationally-efficient protocols. We will deem a protocol computationally-efficient if its cost is polynomial in the relevant parameters: the number of bidders n, the number of items m, and $t = \log L$, where L is the largest possible value of a bundle. Note that

[10] For example, consider a bidder with values of 5,3,4 for the atomic bundles $abcd, ac, b$, respectively. For this valuation, $v(ac) = 3$, $v(dcb) = 4$ but $v(abcd) = 5$.

[11] The running time really measures communication costs and not computational running time.

all of our results give concrete bounds, where the dependence on the parameters is given explicitly; we use the standard big-Oh notation just as a shorthand.

7.3.2 Demand Queries

Most of the chapter will be concerned with a common special case of iterative auctions that we term "demand auctions". In such auctions, the queries that are sent to bidders are demand queries: the query specifies a price $p(S) \in \Re^+$ for each bundle S. The reply of bidder i is simply the set most desired – "demanded" – under these prices. Formally, a set S that maximizes $v_i(S) - p(S)$. It may happen that more than one set S maximizes this value. In which case, ties are broken according to some fixed tie-breaking rule, e.g., the lexicographically first such set is returned. All of our results hold for any fixed tie-breaking rule.

Note that even though in our model valuations are integral, we allow the demand query to use arbitrary real numbers. A practical issue here is how will the query be specified: in the general case, an exponential number of prices needs to be sent in a single query. Formally, this is not a problem as the model does not limit the length of queries in any way – the protocol must simply define what the prices are in terms of the replies received for previous queries. We look into this issue further in Section 7.7.

Many auctions in the literature restrict the prices' representation to item prices (or linear prices):

Definition 7.1. Item Prices: *The prices in each query are given by prices p_j for each item j. The price of a set S is additive: $p(S) = \sum_{j \in S} p_j$.*

7.4 The Power of Different Types of Queries

In this section we compare the power of the various types of queries defined in the introduction. We will present computationally-efficient simulations of these query types using item-price demand queries. In Chapter 6, we show that these simulations can also be done using item-price *ascending* auctions. The opposite, however, is false: we show that an exponential number of some of these queries may be needed for simulating demand queries. Figure 7.3 summarizes the relations between the different query types. Some parts of the following lemmas are elementary, and some are harder. These lemmas will be used in the analysis in the rest of this chapter. All missing proofs can be found in Appendix F.1.

Lemma 7.1. *A value query can be simulated by m marginal-value queries. A marginal-value query can be simulated by two value queries.*

Lemma 7.2. *A value query can be simulated by mt demand queries (where $t = \log L$ is the number of bits needed to represent a single bundle value).*[12]

As a direct corollary we get that demand auctions can always fully elicit the bidders' valuations by simulating all possible $2^m - 1$ queries and thus elicit enough information for determining the optimal allocation. Note, however, that this elicitation may be computationally inefficient.

The next lemma shows that demand queries can be exponentially more powerful than value queries.

[12] Note that t *bundle-price* demand queries can easily simulate a value query by setting the prices of all the bundles except S (the bundle with the unknown value) to be L, and performing a binary search on the price of S.

	Value	Mar-value	Demand	Ind-util	Rel-demand
Value query	1	2	exp	exp	exp
Marginal-value query	m	1	exp	exp	exp
Demand query	mt	poly	1	mt+1	poly
Indirect-utility query	1	2	m+1	1	poly
Relative-demand query	-	-	-	-	1

Figure 7.3: Each entry in the table specifies how many queries of the relevant row are needed to simulate a query from the relevant column.

Lemma 7.3. *An exponential number of value queries may be required for simulating a single demand query.*

Proof. We will actually show an example where a single demand query suffices for finding the optimal allocation, but an exponential number of value queries are required for that. Consider a player with a valuation of $2|S|$ for any bundle S, except for some "hidden" bundle H of size $\frac{m}{2}$ with a valuation of $2|S|+2$, and a second player with a known valuation of $2|S|+1$ for every bundle S. The only optimal allocation gives the hidden set H to the first bidder. In a demand query with a price of $2+\epsilon$ for every item, the first bidder demands his "hidden" set, and thus reveals the optimal allocation.

However, consider any algorithm that uses only value queries. An adversary will answer each value query $v(S)$ to the first bidder with $v(S) = 2|S|$. As long as two sets S of size $\frac{m}{2}$ have not been queried any of them can be the hidden set H and the optimal allocation can not be determined. Thus, $\Omega(2^m)$ value queries will be needed in the worst case. □

Indirect utility queries are, however, equivalent in power to demand queries:

Lemma 7.4. *An indirect-utility query can be simulated by $mt+1$ demand queries. A demand query can be simulated by $m+1$ indirect-utility queries.*

Demand queries can also simulate relative-demand queries:[13] According to our definition of relative-demand queries, they clearly cannot simulate even value queries.

Lemma 7.5. *Relative-demand queries can be simulated by a polynomial number of demand queries.*

7.5 Approximating the Social Welfare with Value and Demand Queries

We know from [117] that iterative combinatorial auctions that only use a polynomial number of queries can not find an optimal allocation among general valuations and in fact can not even approximate it to within a factor better than $min\{n, m^{1/2-\epsilon}\}$. In this section we ask how well can this approximation be done using demand queries with item prices, or using the weaker value

[13] Note: although in our model values are integral, we allow the query prices to be arbitrary real numbers, thus we may have bundles with arbitrarily close relative demands. In this sense the simulation above is only up to any given ϵ (and the number of queries is $O(\log L + \log \frac{1}{\epsilon})$). When the relative-demand query prices are given as rational numbers, exact simulation is implied when $\log \epsilon$ is linear in the input length.

An Approximation Algorithm:

Initialization: Let $T \leftarrow M$ be the current items for sale.
Let $K \leftarrow N$ be the currently participating bidders.
Let $S_1^* \leftarrow \emptyset, ..., S_n^* \leftarrow \emptyset$ be the provisional allocation.
Repeat until $T = \emptyset$ or $K = \emptyset$:
Ask each bidder i in K for the bundle S_i that maximizes her per-item value, i.e., $S_i \in argmax_{S \subseteq T} \frac{v_i(S)}{|S|}$.
Let i be the bidder with the maximal per-item value, i.e., $i \in argmax_{i \in K} \frac{v_i(S_i)}{|S_i|}$, and set: $s_i^* = s_i, \quad K = K \setminus i, \quad M = M \setminus S_i$
Finally: Ask the bidders for their values $v_i(M)$ for the grand bundle.
If allocating all the items to some bidder i improves the social welfare achieved so far (i.e., $\exists i \in N$ such that $v_i(M) > \sum_{i \in N} v_i(S_i^*)$), then allocate all items to this bidder i.

Figure 7.4: This algorithm achieves a $\min\{n, 4\sqrt{m}\}$-approximation for the social welfare, which is asymptotically the best worst-case approximation possible with polynomial communication. This algorithm can be implemented with a polynomial number of demand queries.

queries. We show that, using demand queries, the lower bound can be matched, while value queries can only do much worse.

Figure 7.4 describes a polynomial-time algorithm that achieves a $min(n, O(\sqrt{m}))$ approximation ratio. This algorithm greedily picks the bundles that maximize the bidders' per-item value (using "relative-demand" queries, see Section 7.4). As a final step, it allocates all the items to a single bidder if it improves the social welfare (this can be checked using value queries). Since both value queries and relative-demand queries can be simulated by a polynomial number of demand queries with item prices (Lemmas 7.2 and 7.5), this algorithm can be implemented by a polynomial number of demand queries with item prices.[14]

Theorem 7.1. *The auction described in Figure 7.4 can be implemented by a polynomial number of demand queries and achieves a* $\min\{n, 4\sqrt{m}\}$*-approximation for the social welfare.*

Proof. We first observe that the algorithm can be implemented by a polynomial number of value queries and relative demand queries: querying a bidder for the bundle that maximizes his per-item value is a relative-demand query when all the item prices are 1. Querying a bidder for his value for the grand bundle can be done by a value query. In Section 7.4 we show that any value query and any relative-demand query can be implemented by a polynomial number of demand queries. Each bidder is asked at most m relative demand queries, and exactly one value query, thus a polynomial number of demand queries can implement this algorithm.

Next, we prove that the algorithm achieves an approximation ratio of at least $\min\{n, 4\sqrt{m}\}$. The algorithm will clearly achieve a $\frac{1}{n}$-approximation since we allocate the whole bundle M to the bidder with the highest valuation if it improves the welfare achieved. Next, we prove that the algorithm achieves at least $\frac{1}{4\sqrt{m}}$ of the optimal welfare.

Let $OPT = \{T_1, ..., T_k, Q_1, ..., Q_l\}$ be an optimal allocation where for every $i \in \{1, ..., k\}$ $|T_i| \leq \sqrt{m}$ and for every $j \in \{1, ..., l\}$ $|Q_j| > \sqrt{m}$ $(l, k \in \{0, ..., n\})$. Let ALG be the allocation found by the algorithm, and let $v(OPT)$ and $v(ALG)$ be the optimal welfare and the welfare achieved by

[14] This algorithm can be also implemented by two descending-price auctions (where we allow removing items during the auction), see [27].

the algorithm, respectively. First, we analyze cases where "large" bundles contribute most of the optimal welfare, i.e., $\sum_{i=1}^{l} v_i(Q_i) \geq \sum_{i=1}^{k} v_i(T_i)$. Then,

$$v(OPT) \leq 2\sum_{i=1}^{l} v_i(Q_i) \leq 2\sum_{i=1}^{l} v(ALG) = 2l \cdot v(ALG) \leq 2\sqrt{m} \cdot v(ALG)$$

Where the first inequality holds since $v(OPT) = \sum_{i=1}^{l} v_i(Q_i) + \sum_{i=1}^{k} v_i(T_i)$ and the second holds since the last stage of the algorithm verifies that the welfare achieved by the algorithm is at least the valuation of every player for the whole bundle M. The last inequality holds since there are no more than $\sqrt{m}$ bundles of size of at least $\sqrt{m}$.

The analysis of the case where "small" bundles contribute most of the optimal welfare (i.e., $\sum_{i=1}^{l} v_i(Q_i) < \sum_{i=1}^{k} v_i(T_i)$) is more involved. Let $I \subseteq \{1, ..., k\}$ be the set of bidders that receives a "small" bundle (i.e., bundles in $\{T_1, ..., T_k\}$) in OPT that does not intersect any bundle in ALG. Consider the following sum:

$$\sum_{i=1}^{k} \frac{v_i(T_i)}{|T_i|} = \sum_{i\in I} \frac{v_i(T_i)}{|T_i|} + \sum_{i\in\{1,...,k\}\setminus I} \frac{v_i(T_i)}{|T_i|} \tag{7.1}$$

In the two claims below, we show that each of the summands in the right side of Equation 7.1 is not greater than $v(ALG)$. This immediately derives that $\sum_{i=1}^{k} \frac{v_i(T_i)}{|T_i|} \leq 2 \cdot v(ALG)$. And since for every $i \in 1, ..., k,\ |T_i| \leq \sqrt{m}$, we have: $\sum_{i=1}^{k} v_i(T_i) \leq 2\sqrt{m} \cdot v(ALG)$. Most of the optimal welfare is contributed by "small" bundles, hence:

$$v(OPT) \leq 2\sum_{i=1}^{k} v_i(T_i) \leq 4\sqrt{m} \cdot v(ALG)$$

What is left to be proved is that both summands in Equation 7.1 are not greater than $v(ALG)$:

Claim 7.1. $\sum_{i\in I} \frac{v_i(T_i)}{|T_i|} \leq v(ALG)$

Proof. Consider a bidder i that receives a small bundle T_i in OPT such that T_i is disjoint to all bundles in ALG. We observe that this bidder surely receives a non-empty bundle S_i in ALG. This holds since the items in T_i are not allocated at the end of the algorithm (they are not in any bundle in ALG), but player i has a non-zero value for T_i.

Since the algorithm picked some $S_i \in ALG$ and not T_i for bidder i, $\frac{v_i(T_i)}{|T_i|} \leq \frac{v_i(S_i)}{|S_i|}$. Therefore,

$$\sum_{i\in I} \frac{v_i(T_i)}{|T_i|} \leq \sum_{i\in I} \frac{v_i(S_i)}{|S_i|} \leq \sum_{i\in I} v_i(S_i) \leq \sum_{i=1}^{n} v_i(S_i) = v(ALG)$$

□

Claim 7.2. $\sum_{i\in\{1,...,k\}\setminus I} \frac{v_i(T_i)}{|T_i|} \leq v(ALG)$

Proof. For every bidder $i \in \{1, ..., k\} \setminus I$, T_i intersects at least one bundle from ALG, and let $F(i)$ be the first bidder for which the algorithm allocates a bundle that intersects T_i. Then,

$$\sum_{i\in\{1,...,k\}\setminus I} \frac{v_i(T_i)}{|T_i|} \leq \sum_{j=1}^{n} \sum_{i|F(i)=j} \frac{v_i(T_i)}{|T_i|} \leq \sum_{j=1}^{n} \sum_{i|F(i)=j} \frac{v_i(S_j)}{|S_j|} \leq \sum_{j=1}^{n} |S_j| \frac{v_i(S_j)}{|S_j|} \leq \sum_{j\in ALG} v_j(S_j)$$

Where the second leftmost inequality holds since bidder $j = F(i)$ demands $S_j \in ALG$ when all the items in T_i are still on sale and the third inequality holds since each S_j intersects at most $|S_j|$ bundles from $\{T_1, ..., T_k\}$ (all T_i's are disjoint). □

We showed before how the theorem follows from these two claims. □

We now ask how well can the optimal welfare be approximated by a polynomial number of *value queries*. First we note that value queries are not completely powerless: In [76] it is shown that if the m items are split into k fixed bundles of size m/k each, and these fixed bundles are auctioned as though each was indivisible, then the social welfare generated by such an auction is at least $\frac{m}{\sqrt{k}}$-approximation of that possible in the original auction. Notice that such an auction can be implemented by $2^k - 1$ value queries to each bidder – querying the value of each bundle of the fixed bundles. Thus, if we choose $k = \log m$ bundles we get an $\frac{m}{\sqrt{\log m}}$-approximation while still using a polynomial number of queries.

We show that not much more is possible using value queries:

Lemma 7.6. *Any iterative auction that uses only value queries and distinguishes between k-tuples of 0/1 valuations where the optimal allocation has value 1, and those where the optimal allocation has value k requires at least $2^{\frac{m}{k}}$ queries.*

We conclude that a polynomial time protocol that uses only value queries cannot obtain a better than $O(\frac{m}{\log m})$ approximation of the welfare. This can be immediately derived from Lemma 7.6: achieving any approximation ratio k which is asymptotically greater than $\frac{m}{\log m}$ needs an exponential number of value queries.

Theorem 7.2. *An iterative auction that uses a polynomial number of value queries cannot achieve better than $O(\frac{m}{\log m})$-approximation for the social welfare.*

7.6 Demand Queries and Linear Programming

In this section, we show that the linear-programming relaxation of the combinatorial-auction problem can be optimally solved using demand queries. This observation turns to be useful for the design of approximation algorithms for combinatorial auctions and other related resource-allocation problems (see, e.g., [47, 23, 58]).

Consider the following linear-programming relaxation for the generalized winner-determination problem in combinatorial auctions (the "primal" program):

$$\begin{aligned}
\textbf{Maximize} \quad & \sum_{i \in N, S \subseteq M} w_i \, x_{i,S} \, v_i(S) & \\
\textbf{s.t.} \quad & \sum_{i \in N,\, S | j \in S} x_{i,S} \leq q_j & \forall j \in M \\
& \sum_{S \subseteq M} x_{i,S} \leq d_i & \forall i \in N \\
& x_{i,S} \geq 0 & \forall i \in N, S \subseteq M
\end{aligned}$$

Note that the primal program has an exponential number of variables. Yet, we will be able to solve it in polynomial time using demand queries to the bidders. The solution will have a polynomial size support (non-zero values for $x_{i,S}$), and thus we will be able to describe it in polynomial time.

Here is its dual:

$$\begin{aligned}
\textbf{Minimize} \quad & \sum_{j \in M} q_j p_j + \sum_{i \in N} d_i u_i \\
\textbf{s.t.} \quad & u_i + \sum_{j \in S} p_j \geq w_i v_i(S) && \forall i \in N, \;\; S \subseteq M \\
& p_i \geq 0, \;\; u_j \geq 0 && \forall i \in M, \; j \in N
\end{aligned}$$

Notice that the dual problem has exactly $n + m$ variables but an exponential number of constraints. Thus, the dual can be solved using the Ellipsoid method in polynomial time – *if* a "separation oracle" can be implemented in polynomial time. Recall that a separation oracle, when given a possible solution, either confirms that it is a feasible solution, or responds with a constraint that is violated by the possible solution.

We construct a separation oracle for solving the *dual* program, using a single demand query to each of the bidders. Consider a possible solution $(\overline{u}, \overline{p})$ for the dual program. We can re-write the constraints of the dual program as:

$$u_i / w_i \geq v_i(S) - \sum_{j \in S} p_j / w_i$$

Now a demand query to bidder i with prices p_j / w_i reveals exactly the set S that maximizes the RHS of the previous inequality. Thus, in order to check whether $(\overline{u}, \overline{p})$ is feasible it suffices to (1) query each bidder i for his demand D_i under the prices p_j / w_i; (2) check only the n constraints $u_i + \sum_{j \in D_i} p_j \geq w_i v_i(D_i)$ (where $v_i(D_i)$ can be simulated using a polynomial sequence of demand queries as shown in Lemma 7.2). If none of these is violated then we are assured that $(\overline{u}, \overline{p})$ is feasible; otherwise we get a violated constraint.

What is left to be shown is how the *primal* program can be solved. (Recall that the primal program has an exponential number of variables.) Since the Ellipsoid algorithm runs in polynomial time, it encounters only a polynomial number of constraints during its operation. Clearly, if all other constraints were removed from the dual program, it would still have the same solution (adding constraints can only decrease the space of feasible solutions). Now take the "reduced dual" where only the constraints encountered exist, and look at its dual. It will have the same solution as the original dual and hence of the original primal. However, look at the form of this "dual of the reduced dual". It is just a version of the primal program with a polynomial number of variables – those corresponding to constraints that remained in the reduced dual. Thus, it can be solved in polynomial time, and this solution clearly solves the original primal program, setting all other variables to zero.

7.7 The Representation of Bundle Prices

In this section we explicitly fix the language in which bundle prices are presented to the bidders in bundle-price auctions. This language requires the algorithm to explicitly list the price of every bundle with a non-trivial price. "Trivial" in this context is a price that is equal to that of one of its proper subsets (which was listed explicitly). This representation is equivalent to the XOR-language for expressing valuations. Formally, each query q is given by an expression: $q = (S_1 : p_1) \oplus (S_2 : p_2) \oplus ... \oplus (S_l : p_l)$. In this representation, the price demanded for every set S is simply $p(S) = max_{\{k=1...l | S_k \subseteq S\}} p_k$.

Definition 7.2. *The* length *of the query* $q = (S_1 : p_1) \oplus (S_2 : p_2) \oplus ... \oplus (S_l : p_l)$ *is* l. *The* cost *of an algorithm is the sum of the lengths of the queries asked during the operation of the algorithm on the worst case input.*

Note that under this definition, bundle-price auctions are not necessarily stronger than item-price auctions. An item-price query that prices each item for 1, is translated to an exponentially long bundle-price query that needs to specify the price $|S|$ for each bundle S. But perhaps bundle-price auctions can still find optimal allocations whenever item-price auction can, without directly simulating such queries? We show that this is not the case: indeed, when the representation length is taken into account, bundle price auctions are sometimes seriously inferior to item price auctions.

Consider the following family of valuations: Each item is valued at 3, except that for some single set S, its value is a bit more: $3|S| + b$, where $b \in \{0, 1, 2\}$. Note that an item price auction can easily find the optimal allocation between any two such valuations: Set the prices of each item to $3 + \epsilon$; if the demand sets of the two players are both empty, then $b = 0$ for both valuations, and an arbitrary allocation is fine. If one of them is empty and the other non-empty, allocate the non-empty demand set to its bidder, and the rest to the other. If both demand sets are non-empty then, unless they form an exact partition, we need to see which b is larger, which we can do by increasing the price of a single item in each demand set.

We will show that any bundle-price auction that uses only the XOR-language to describe bundle prices requires an exponential cost (which includes the sum of all description lengths of prices) to find an optimal allocation between any two such valuations.

The complication in the proof stems from the fact that using XOR-expressions, the length of the price description depends on the number of bundles whose price is strictly larger than each of their subsets – this may be significantly smaller than the number of bundles that have a non-zero price. (The proof becomes easy if we require the protocol to explicitly name every bundle with non-zero price.) We do not know of any elementary proof for this lemma (although we believe that one can be found). Instead we reduce the problem to a well known lower bound in boolean circuit complexity [72] stating that boolean circuits of depth 3 that compute the majority function on m variables require $2^{\Omega(\sqrt{m})}$ size.

Lemma 7.7. *Every bundle-price auction that uses XOR-expressions to denote bundle prices requires* $2^{\Omega(\sqrt{m})}$ *cost in order to find the optimal allocation among two valuations from the above family.*

Chapter 8

Conclusions

Designing economic markets is a complex task, mainly due to asymmetric information held by the participants in these markets. The designers usually aim to optimize some system-wise objective functions that often depend on the secret data of the players. Therefore, deciding on the protocols by which information is exchanged is a major part of the design process. This dissertation studies scenarios where the mechanisms use several natural and common patterns of communication. In particular, bidders cannot directly disclose their privately known data so the celebrated "revelation principle" cannot hold.

In some settings, a mere information-theoretic analysis provides interesting and even surprising description of the power of mechanisms restricted to certain communication patterns. For example, ascending item-price auction cannot reveal optimal, or even nearly optimal, allocations in combinatorial auctions (Chapter 5). A positive example shows that a very small number of bids in single-item auctions incurs a very mild loss in social efficiency (Chapter 3).

In other settings, we showed that the optimal results under some communication pattern can be implemented in dominant strategies without any additional informational requirement. This holds, for example, in single-parameter mechanism-design domains where the system-wise target functions are multilinear in the players' types (Chapter 4).

Finally, we also presented settings where additional expressiveness is required from a mechanism for achieving an incentive-compatible implementation of the outcomes. One example is the impossibility to compute VCG payments for substitutes valuations by item-price auctions illustrated in Chapter 6 (even with non-anonymous prices) .

Future work should identify settings where direct-revelation mechanisms are unreasonable, due to various reasons, and try to measure the performance of mechanisms in these settings under such "communication filters". Interesting properties of these environments can be achieved by a pure information theoretic analysis, or by measuring the cost of incentive-compatible implementation. The high level goal should be to characterize tradeoffs and dependencies between the various components of the markets: information, incentives, economic properties and computational complexity. This characterization is especially important since in many natural settings these properties are mutually exclusive. The main theme of this dissertation is that the first component - the information effects - is an important factor that should be explicitly taken into account and it cannot be discarded by "revelation principle" arguments.

This dissertation leaves many open questions. The main open question in the first part of this dissertation is whether the information-theoretically optimal results, under restrictions on the number of actions, can be always implemented in dominant strategies with no additional communication. We prove such a result given that all the social-value functions are multilinear, but failed to either prove it or disprove it for the general case. The second part of this dissertation studied the power of different types of iterative combinatorial auctions. Open questions that emerged in this part include whether strong negative results with respect to ascending auctions can also be shown for natural restricted sub-classes of valuations, e.g., for submodular valuations or any other application-dependent class of valuations, and whether a cost is incurred when implementing these results in an equilibrium. Finally, this dissertation studied specific toy models for economic interactions. Such questions should be addressed for more sophisticated information models, like models with interdependent or common types, models with externalities, and environments where monetary transfers are not allowed.

Appendix A

Preliminaries: Mechanism Design

A.1 Combinatorial Auctions - Missing Proofs

Proof of Proposition 2.5

Proof. Consider the following simple protocol that uses anonymous bundle prices: it starts with zero prices, collects the demand of all the bidders at each step (bidder report all the bundles in their demand set at each stage). Given the demands at a particular price level, it raises the prices of every bundle that is demanded by at least one bidder until all the bidders demand the empty set (actually, we can stop at an earlier stage at a competitive equilibrium, exactly as done in Incremental auctions [124, 4] only with anonymous prices). The auction outputs the optimal allocation among the allocations that allocate for each bidder a bundle in his demand set at the final stage.

What is left to be shown is that when the bidders have super-additive preferences, all the bundles in an optimal allocation will be demanded at the final stage of the auction. Let $S_1, .., S_n$ be some optimal allocation. We first note that bidder i has a higher valuation for the bundle S_i than all the other bidders. Assume that for some bidder j, $v_j(S_i) > v_i(S_i)$. Then, $v_j(S_j \vee S_i) \geq v_j(s_j)+v_j(s_i) > v_j(s_j) + v_i(s_i)$. Thus, we can increase the efficiency by allocating s_i to j, contradiction to the optimality of $S_1, .., S_n$.

Since the prices are increased only for demanded bundles, the price of a bundle S_i of the optimal allocation cannot increase above $v_i(S_i)$. Thus, it will be demanded by i at some stage of the auction. □

Appendix B

Auctions with Severely Bounded Communication

B.1 Efficient Mechanisms: Missing Proofs

Proof of Claim 3.1 in Theorem 3.1

Proof. Given a vector of strategies s^* which achieve optimal welfare in g (i.e., $\max_{\widetilde{s}\in\times_{i=1}^{k}\varphi_{k_i}} w(g,\widetilde{s})$), we will show that for every player i we can modify s_i^* to be a threshold strategy, and the welfare will not decrease.

Assume s_i^* is not a threshold strategy. Therefore, there must be $\alpha, \beta, \gamma \in [\underline{a}, \overline{b}]$, $\alpha < \beta < \gamma$ such that $s_i^*(\alpha) = s_i^*(\gamma) = m$ but $s_i^*(\beta) \neq m$ (where m is some bid of player i). We will show that a strategy vector s identical to s^*, except that for every such β $s_i(\beta) = m$, we have that $w(g,s) \geq w(g,s^*)$.

Denote the probability that all players except i bids b_{-i} as $Pr(b_{-i})$. Thus, the expected welfare from a game g given that bidder i with valuation v_i bids m and that the other players use strategies s_{-i}^* is:

$$\sum_{b_{-i}} Pr(b_{-i}) \left(a_i(m, b_{-i}) \cdot v_i + \sum_{j\neq i} a_j(m, b_{-i}) \cdot E(v_j \left| s_j^*(v_j) = b_j \right.) \right)$$

Note that this expected welfare is a linear function of v_i, and we denote it by $h(m) \cdot v_i + t(m)$ (the constants $h(m)$ and $t(m)$ depend on the bid m).

We know that s^* achieve optimal welfare in g and that $s_i^*(\alpha) = m$. Therefore, there is no other bid l such that if $s_i^*(\alpha) = l$, the expected welfare will increase, i.e.:

$$\forall l \neq m \quad h(m) \cdot \alpha + t(m) \geq h(l) \cdot \alpha + t(l) \tag{B.1}$$

Similarly, because $s_i^*(\gamma) = m$:

$$\forall l \neq m \quad h(m) \cdot \gamma + t(m) \geq h(l) \cdot \gamma + t(l) \tag{B.2}$$

Because β is a convex combination of α and γ, and due to Equations B.1 and B.2:

$$\forall l \neq m \;\; h(m) \cdot \beta + t(m) \geq h(l) \cdot \beta + t(l)$$

Thus, the expected welfare for player i, given $v_i = \beta$, is maximal when she bids m. Therefore, when modifying s_i^* such that $s_i^*(\beta) = m$ the total expected welfare will not decrease. We can repeat this process until s_i^* becomes a threshold strategy.[1] □

Lemma B.1. $w^{opt}_{2,(k,k)} > w^{opt}_{2,(k-1,k)}$ *for every* $k > 1$.

Proof. Let $g \in G_{2,(k-1,k)}$ be a deterministic, monotone mechanism that achieves the optimal welfare with threshold strategies based on the vectors (x, y). Each row in g is of the form $[A, ..., A, B, ..., B]$, and let $l_i \in \{0, ..., k\}$ be the first index in row i in which B wins. We will modify g to $\widetilde{g} \in G_{2,(k,k)}$ by adding some missing row, and change the threshold strategy x to $\widetilde{x} \in \Re^{k+1}$, such that the welfare strictly improves. We assume, w.l.o.g., that the thresholds are unique (i.e., $0 < x_1 < ... < x_{k-1} < 1$, $0 < y_1 < ... < 1$).

Case 1. *The row* $[B, ..., B]$ *is in the game's matrix.*

Let $x_1' = \frac{E_{v_B}(v_B | 0 \le v_B \le y_1)}{2}$, and let $\widetilde{x} = (0, x_1', x_1, x_2, ..., x_{k-2}, 1)$. We will create a new game $\widetilde{g}$ by adding the line $[B, ..., B]$ as the first line. It is easy to see that the allocation in g and $\widetilde{g}$ is identical in all rows except the new one. When $v_A \in [0, x_1']$ and $v_B \in [0, y_1]$ $\widetilde{g}$ allocates the item to B where g allocated the item to A. The distribution functions are always positive, hence this will occur with a positive probability. Since $E(v_A | v_A \in [0, x_1']) < x_1' < E_{v_B}(v_B | 0 \le v_B \le y_1)$, the expected welfare has strictly increased. For higher bids of bidder B, the allocation in the first row will clearly be efficient now, therefore no welfare loss was incurred.

Case 2. *The row* $[B, ..., B]$ *does not appear in* g*'s game matrix.*

Due to the monotonicity, g must have two rows i and $i+1$ and two columns j and $j+1$ such that we allocate the item to B when the bids are $(i, j), (i, j+1)$ and to A when the bids are $(i+1, j), (i+1, j+1)$. We will create a mechanism $\widetilde{g}$ by adding a row i' identical to row $i+1$ except that B wins in index $j+1$. The new threshold is constructed as follows:

If $\mathbf{E(v_B | y_j \le v_B \le y_{j+1}) < x_{i+1}}$:
Let $x_{i+1}' = E(v_B | y_j \le v_B \le y_{j+1})$, and let $\widetilde{x} = (0, x_1, ..., x_i, x_i', x_{i+1}, ..., 1)$. As in previous cases, the welfare in all entries hasn't changed, except a strictly positive improvement in the (i', j) entry.

If $\mathbf{E(v_B | y_j \le v_B \le y_{j+1}) \ge x_{i+1}}$:
Let $x_i' = E(v_B | y_{j+1} \le v_B \le y_{j+2})$ and let $\widetilde{x} = (0, x_1, ..., x_i, x_i', x_{i+1}, ..., 1)$. We show that since g is efficient $x_{i+1} < x_i' < x_{i+2}$: First, $E(v_B | y_{j+1} \le v_B \le y_{j+2}) > E(v_B | y_j \le v_B \le y_{j+1}) \ge x_{i+1}$; Also, since A wins for the bids $(i+1, j+1)$, we have $E(v_B | y_{j+1} \le v_B \le y_{j+2}) \le E(v_A | x_{i+1} \le v_A \le x_{i+2}) < x_{i+2}$. It follows that the expected welfare has strictly increased in the entry $(i', j+1)$, and has not decreased in all other entries. □

Proof of Theorem 3.2

Proof. First, we prove that $PG_k(x^w, y^w)$ is optimal when $v_0 = \underline{a}$. According to Theorem 3.1 there is a pair of threshold values' vectors $x = (x_0, x_1, ..., x_k)$,$y = (y_0, y_1, ..., y_k)$ such that $PG_k(x, y)$ achieves the optimal welfare. Note that $x_0 = y_0 = \underline{a}$ and $x_k = y_k = \overline{b}$, so we have $2(k-1)$ variables to optimize.

[1] See analysis of similar problems, e.g., by Athey in [3].

We will calculate the total expected welfare by summing first the expected welfare in the entries of the game matrix where B wins the item, then summing the entries where A is the winner.

$$
\begin{aligned}
w(g,s) &= \sum_{i=1}^{k} (F_B(y_i) - F_B(y_{i-1})) \cdot (F_A(x_i) - F_A(x_0)) \cdot \frac{\int_{y_{i-1}}^{y_i} f_B(v_B) v_B dv_B}{F_B(y_i) - F_B(y_{i-1})} \\
&+ \sum_{i=2}^{k} (F_A(x_i) - F_A(x_{i-1})) \cdot (F_B(y_{i-1}) - F_B(y_0)) \cdot \frac{\int_{x_{i-1}}^{x_i} f_A(v_A) v_A dv_A}{F_A(x_i) - F_A(x_{i-1})} \\
&= \sum_{i=1}^{k} F_A(x_i) \cdot \int_{y_{i-1}}^{y_i} f_B(v_B) v_B dv_B + \sum_{i=2}^{k} F_B(y_{i-1}) \cdot \int_{x_{i-1}}^{x_i} f_A(v_A) v_A dv_A
\end{aligned}
$$

We assume here that a probability density function exists for each bidder. Thus, we can express the partial derivatives with respect to all variables:

$$
(w(g,s))'_{x_i} = \left(\int_{y_{i-1}}^{y_i} f_B(v_B) v_B dv_B \right) \cdot f_A(x_i) + f_A(x_i) \cdot x_i \cdot F_B(y_{i-1}) - f_A(x_i) \cdot x_i \cdot F_B(y_i) = 0
$$

$$
(w(g,s))'_{y_i} = \left(\int_{x_i}^{x_{i+1}} f_A(v_A) v_A dv_A \right) \cdot f_B(y_i) + f_B(y_i) \cdot y_i \cdot F_A(x_i) - f_B(y_i) \cdot y_i \cdot F_A(x_{i+1}) = 0
$$

Rearranging the terms derives that $y_i = E_{v_A}(v_A \,| x_i \leq v_A \leq x_{i+1})$ and that $x_i = E_{v_B}(v_B \,| y_{i-1} \leq v_B \leq y_i)$ and therefore, x, y should be mutually centered for optimal efficiency.

Now, we no longer assume $v_0 = \underline{a}$: According to Theorem 3.1, if the optimal welfare is not achieved in the priority game above, it will be achieved in a modified priority game. For some threshold values' vectors x, y, the expected welfare in $MPG_k(x, y)$ is given by the formula:

$$
\begin{aligned}
& F_A(x_1) \cdot F_B(y_1) \cdot v_0 + F_A(x_1) \int_{y_1}^{\bar{b}} v_B f_B(v_B) dv_B + F_B(y_1) \int_{x_1}^{\bar{b}} v_A f_A(v_A) dv_A \\
+ & \sum_{i=2}^{k} (F_A(x_i) - F_A(x_1)) \int_{y_{i-1}}^{y_i} v_B f_B(v_B) dv_B \quad + \quad \sum_{i=3}^{k} (F_B(y_{i-1}) - F_B(y_1)) \int_{x_{i-1}}^{x_i} v_A f_A(v_A) dv_A
\end{aligned}
$$

First-order condition similarly derive the constraints on x_1 and y_1 given in the above definition of $\overline{x}, \overline{y}$, and that $(x_1, ..., x_{k-1}, x_k)$ and $(y_1, ..., y_{k-1}, y_k)$ should be mutually-centered[2].

□

B.2 Optimal Symmetric 1-bit Mechanisms

Following are the optimal 1-bit 2-player mechanisms assuming independent uniform distributions for all values. The socially-efficient symmetric 1-bit mechanism achieves an expected welfare of 0.625 compared to 0.648 that is achieved in an asymmetric 1-bit mechanism and 2/3 that is achieved with unrestricted communication. Similarly, the revenue-maximizing symmetric 1-bit mechanism below achieves an expected profit of 0.385 compared to 0.39 with 1-bit symmetric mechanisms and

[2] The results are not surprising, since except for the case when one of the bidders bids 0, we have a priority game's allocation for which the optimal threshold values must be mutually centered (due to the first part of the proof).

$5/12 = 0.417$ with unrestricted communication (obtained by second-price auction with a reserve price).

The following mechanism acheives the optimal welfare among all the symmetric 1-bit mechanisms:

	0	1
0	w.p. $\frac{1}{2}$ A wins, pays 0 w.p. $\frac{1}{2}$ B wins, pays 0	B wins and pays $\frac{1}{4}$
1	A wins and pays $\frac{1}{4}$	w.p. $\frac{1}{2}$ A wins, pays $\frac{1}{2}$ w.p. $\frac{1}{2}$ B wins, pays $\frac{1}{2}$

Proving the social efficiency of the mechanism can be done by the following idea: First note that a symmetric, efficient mechanism will clearly allocate the item to the player that bids 1 when the other player bids 0, and allocate with equal probabilities of $\frac{1}{2}$ when the bids are equal. With threshold strategies (x, y) the expected welfare is:

$w(x,y) = x \cdot y \cdot \left(\frac{1}{2} \cdot \frac{x}{2} + \frac{1}{2} \cdot \frac{y}{2}\right) + x \cdot (1-y) \cdot \frac{(1+y)}{2} + (1-x) \cdot y \cdot \frac{(1+x)}{2} + (1-x) \cdot (1-y) \cdot \left(\frac{1}{2} \cdot \frac{(1+x)}{2a} + \frac{1}{2} \cdot \frac{(1+y)}{2}\right)$

Maximum is achieved when $(x, y) = (\frac{1}{2}, \frac{1}{2})$.

The mechanism below is the revenue-maximizing mechanism:

	0	1
0	No allocation	B wins and pays $\frac{1}{\sqrt{3}}$
1	A wins, pays $\frac{1}{\sqrt{3}}$	w.p. $\frac{1}{2}$ A wins, pays $\frac{1}{\sqrt{3}}$ w.p. $\frac{1}{2}$ B wins, pays $\frac{1}{\sqrt{3}}$

The idea behind the optimality of the above mechanism over all the 1-bit symmetric mechanisms: in the profit-maximizing symmetric mechanism if a player bids 0 and the other bids 1, the latter wins and pays x. When both players bid 1, they will pay $\overline{x}$ with equal probabilities. It is easy to see that under the ex-post IR assumption, $x = \overline{x}$. The expected profit is thus: $r(x) = x(1-x)x + (1-x)xx + (1-x)(1-x)(\frac{1}{2}x + \frac{1}{2}x)$. Maximum is achieved ($x \in [0, 1]$) when $x = \frac{1}{\sqrt{3}}$.

B.3 Asymptotic Analysis - Missing Proofs

Proof of Theorem 3.6

Proof. The proof's idea: we construct a priority game in which all bidders have the same dominant threshold strategy, such that the probability for a bidder to bid each bid is smaller than $\frac{n}{k}$. This is done by dividing the density functions of all the bidders to $\frac{k}{n}$ intervals with equal mass, then combining these thresholds to a vector of k threshold values. Because the bidders use the same threshold strategy, a welfare loss is possible only when more than one bidder bids the highest bid. This observation leads to the upper bound.

Let $\alpha_1, ..., \alpha_n$ be integers such that $\sum_{i=1}^{n} \alpha_i = k - 2$, and for every i, $\alpha_i \geq \lfloor \frac{k}{n} \rfloor - 1$ (clearly such numbers exist). For every bidder i, let $Y^i = (y_1^i, ..., y_{\alpha_i}^i)$ be a set of threshold values that divide her distribution function f_i to $\alpha_i + 1$ segments with the same mass (when $y_0^i = 0, y_{\alpha_i+1}^i = 1$), i.e., for every bid j, $F_i(y_{j+1}) - F_i(y_j) = \frac{1}{\alpha_i+1}$.

Let $X = \{\bigcup_{i=1}^{n} Y^i\} \bigcup \{v_0\}$, $|X| = k - 1$, be the union of all the threshold values (we add arbitrary threshold values if the size of X is smaller than $k - 1$). Let $x = (0, x_1, ..., x_{k-1}, 1)$ be a threshold-value vector created by ordering the threshold values in X from smallest to largest. Now, consider the n-bidder mechanism $MPG_k(\widetilde{t})$ where $\widetilde{t} = (x, .., x)$. The threshold strategy based on x is dominant for all the bidders, with ex-post IR. By the construction of the sets $Y^1, ..., Y^n$, every bidder will bid any particular bid w.p. $\leq \frac{2n}{k}$.[3]

Next, we will bound the welfare loss. We divide the possible cases according to the number of bidders who bid the highest bid. Since all the bidders use the same threshold strategy, if only one bidder bids the highest bid, no welfare loss is incurred (he will definitely have the highest valuation). If more than 1 bidder bid the highest bid i, the expected welfare loss will not exceed $x_{i+1} - x_i$. For a set of bidders $T \subseteq N$, denote the probability that all the bidders in T bid i by $Pr(T = i)$, and the probability that all the bidders not in T have bids smaller than i by $Pr(N \setminus T < i)$. Thus, the expected welfare loss is smaller than (when $2n < k$):

$$\begin{aligned}
& \sum_{j=2}^{n} \sum_{T\subseteq N,\ |T|=j} \sum_{i=1}^{k} Pr(T = i) Pr(N \setminus T < i)\,(x_{i+1} - x_i) \\
\leq\ & \sum_{j=2}^{n} \sum_{T\subseteq N,\ |T|=j} \sum_{i=1}^{k} Pr(T = i)\,(x_{i+1} - x_i) \\
\leq\ & \sum_{j=2}^{n} \sum_{1}^{\binom{n}{j}} \sum_{i=1}^{k} \left(\frac{2n}{k}\right)^{j} (x_{i+1} - x_i) \quad = \quad \sum_{j=2}^{n} \binom{n}{j} \left(\frac{2n}{k}\right)^{j} \quad < \quad 2^n \cdot 4n^2 \cdot \frac{1}{k^2}
\end{aligned}$$

When the valuations of all the bidders are smaller than v_0, there is no welfare loss (it is easy to see that we can assume, w.l.o.g., that $x_1 = v_0$). Note that despite the coefficient of $\frac{1}{k^2}$ is exponential in n, we consider it as a constant because n is fixed. For Example, when $n = 2$ a similar proof shows a welfare loss smaller than $\frac{8}{k^2}$ (when $k > 3$). □

B.4 More Asymptotic Results

Proposition B.1. *The n-player mechanism $PG_k(x, ..., x)$, $x = (0, \frac{1}{k}, \frac{2}{k}, ..., \frac{k-1}{k}, 1)$, incurs an expected welfare loss $\leq \frac{1}{k}$ for any set of distribution functions of the players' valuations. Moreover, for any mechanism g there exists a set of distribution functions for which the expected welfare loss in g is greater than $\frac{n-1}{n^2} \cdot \frac{1}{k}$ (i.e., $\Omega(\frac{1}{k})$).*

Proof. When all players use the same threshold strategy in priority games, non-optimal allocation is possible only when more than one bidder bid the highest bid. Since the difference between subsequent thresholds is $\frac{1}{k}$, the expected welfare loss is clearly not greater than $\frac{1}{k}$.

For proving the lower bound, consider a mechanism $g \in G_{n,k}$ with an equilibrium $s_1, ..., s_n$. We can prove, similarly to the proof of Claim 3.1 in Theorem 3.1, that every mechanism with a Bayesian-Nash equilibrium, has an equilibrium of threshold strategies. Thus, we can assume that $s_1, ..., s_n$ are threshold strategies based on some threshold-value vectors $x^1, ..., x^n$. Observe that there are no more than nk bids' combinations $b = (b_1, ..., b_n)$ with overlapping valuations for *all* players, i.e., for every pair of players i, j: $[x^i_{b_i}, x^i_{b_i+1}] \bigcap [x^j_{b_j}, x^j_{b_j+1}] \neq \emptyset$. (The maximal number of different thresholds for all players is $(k-1)n$, and every two subsequent thresholds define such an

[3] For every bidder i, and every bid j, $F_i(x_{j+1}) - F_i(x_j) \leq \frac{1}{\lfloor \frac{k}{n} \rfloor} \leq \frac{2n}{k}$

"overlapping segment".) In addition, the sum of the sizes of these overlapping segments is 1. Thus, there must be a bids' combination $\tilde{b} = (\tilde{b_1}, ..., \tilde{b_n})$ with an overlapping segment with size of at least $\frac{1}{nk}$. Denote this segment as $[\underline{m}, \overline{m}]$ ($\overline{m} - \underline{m} \geq \frac{1}{nk}$). Assume w.l.o.g that for the bids vector $\tilde{b}$, player n wins the item with probability not greater than $\frac{1}{n}$ (such a player must exist). Now assume that the players' valuations are distributed such that player n has the constant valuation $\overline{m}$ and all the other players have the constant valuation $\underline{m}$. Then, the allocation will not be optimal (i.e., player n will not win) with probability of at least $\frac{n-1}{n}$, and the welfare loss is at least $\overline{m} - \underline{m} \geq \frac{1}{nk}$. The total welfare loss will therefore be at least $\frac{n-1}{n^2} \cdot \frac{1}{k}$. Thus, the expected welfare loss is bounded from below by a proportion of $\frac{1}{k}$. □

If we assume that the distribution functions of the players are bounded from above or from below, we can get even stronger results for this simple mechanism:

Definition B.1. *We say that a probability density function f is bounded from above (resp. below) if for every x in its domain, $f(x) \leq c$ (resp. $f(x) \geq c$) for some constant $c > 0$.*

Proposition B.2. *For every set of probability density functions of the players' valuations which are bounded from above, the mechanism $PG_k(x, ..., x) \in G_{n,k}$, where $x = (0, \frac{1}{k}, \frac{2}{k}, ..., \frac{k-1}{k}, 1)$, incurs an expected welfare loss $\leq c_1 \cdot \frac{1}{k^2}$ for some positive constant c_1 (i.e., $O(\frac{1}{k^2})$ in the CS notations). For every set of probability density functions which are bounded from below, every mechanism incurs an expected welfare loss $\geq c_2 \cdot \frac{1}{k^2}$ for some positive constant c_2 (i.e., $\Omega(\frac{1}{k^2})$).*

Proof. For proving the first statement, say that the distribution function is bounded from above by $\overline{q}$. When the players use the same threshold strategy, a welfare loss is only possible when more than one player bid the higher bid. Every subset of players can be the set of players that bids the highest bid, and this bid can be any bid in $1, ..., k-1$. (The welfare added when all player bid "0" is negligible.) The maximal welfare loss is $\frac{1}{k}$, thus the expected welfare loss is smaller than:

$$\begin{aligned}
& \sum_{T \subseteq N \ |T| \geq 2} \sum_{i=2}^{k} \left(\overline{q} \cdot \frac{1}{k}\right)^{|T|} \cdot \left(\overline{q} \cdot \frac{i-1}{k}\right)^{n-|T|} \frac{1}{k} \\
\leq \ & \overline{q}^n \sum_{T \subseteq N, \ |T| \geq 2} \sum_{i=2}^{k} \frac{1}{k^{n+1}} \cdot (i-1)^{n-|T|} \\
\leq \ & \overline{q}^n \sum_{T \subseteq N, \ |T| \geq 2} \sum_{i=2}^{k} \frac{1}{k^{n+1}} \cdot (k)^{n-|T|} \\
\leq \ & \overline{q}^n \sum_{T \subseteq N, \ |T| \geq 2} \frac{1}{k^n} \cdot (k)^{n-|T|} \\
\leq \ & \overline{q}^n \sum_{T \subseteq N, \ |T| \geq 2} \frac{1}{k^2} \ \leq \ (2\overline{q})^n \frac{1}{k^2}
\end{aligned}$$

As for the second part of the theorem, we first prove it for 2 players (n=2). Assume that the distribution functions of the players are bounded from below by $\underline{q}$, and that the players A, B use the threshold strategies based on $x = (x_0, ..., x_k)$ and $y = (y_0, ..., y_k)$. We say that the bids i, j for players A, B (respectively) are overlapping, if $[x_i, x_{i+1}] \cap [y_j, y_{j+1}] \neq \emptyset$. Consider a mechanism $g \in G_{2,k}$ such that the threshold strategies based on x, y are in Bayesian-Nash equilibrium. Let m

be the number of overlapping pairs of bids. It is easy to see that $m \le 2k-1$. For each overlapping pair i ($i = 1, ..., m$), let z_i be the size of the overlapping segment i. Clearly, $\sum_{i=1}^{m} z_i = 1$. Given that the 2 players' valuations are in the ith segment, the maximal welfare loss is z_i. Thus, the expected welfare loss is greater than[4].

$$\sum_{i=1}^{m} (\underline{q} z_i) \cdot (\underline{q} z_i) \cdot z_i \;=\; \underline{q}^2 \sum_{i=1}^{m} z_i^3 \;\ge\; \underline{q}^2 \cdot \left(\frac{1}{m}\right)^2 \;\ge\; \underline{q}^2 \cdot \frac{1}{(2k-1)^2}$$

The proof for n players is straightforward now: consider only the case where players 1 and 2 have valuations above $\frac{1}{2}$, and the other players will have valuations below $\frac{1}{2}$. This will occur with probability not smaller than the constant $\frac{\underline{q}^n}{2^n}$. We saw that any 2-player mechanism incurs a loss which is bounded from below by a proportion of $\frac{1}{k^2}$ (only with a different coefficient)[5]. □

One can interpret Proposition B.2 as a contest between "nature" and the mechanism designers: when they choose a mechanism first, and then "nature" chooses the distribution functions, the designers can ensure a welfare loss of no more than a proportion of $\frac{1}{k^2}$. When "nature" chooses the distribution function first, and then we choose the mechanism, "nature" ensures that the welfare loss will be at least proportional to $\frac{1}{k^2}$.

So far, we assumed that the players' valuations are drawn from statistically-independent distributions. Next, we relax this assumption and deal with general joint distributions of the valuations. For this case, we show that a trivial mechanism is asymptotically optimal. In particular, it derives an asymptotically tight upper bound of $\frac{1}{k}$ for the efficiency loss in n-player games.

Proposition B.3. *The mechanism $PG_k(x, ..., x) \in G_{n,k}$ where $x = (0, \frac{1}{k}, \frac{2}{k}, ..., \frac{k-1}{k}, 1)$ incurs an expected welfare loss $\le \frac{1}{k}$ for any joint distribution ϕ on the players' valuations.*

Moreover, for every k there is a joint distribution function ϕ_k such that any mechanism $g \in G_{n,k}$ incurs a welfare loss $\ge c \cdot \frac{1}{k}$ (where c is a positive constant independent of k).

Proof. The straightforward proof of the first statement is identical to the case of independent valuations (see Proposition B.1).
We first prove the second statement for $n = 2$. For every k, we construct a joint distribution that incurs, for any mechanism $g \in G_{2,k}$, an efficiency loss which is greater than $\frac{1}{16k}$. Consider the following joint distribution: v_A is distributed uniformly in the range $[\frac{1}{4k}, 1 - \frac{1}{4k}]$, and v_B is $v_A + \frac{1}{4k}$ or $v_A - \frac{1}{4k}$ with equal probabilities. We say that v_A is *dominated*, if there is a threshold y_j of player B, such that $|v_A - y_j| \le \frac{1}{4k}$. B's thresholds 0 and 1 clearly cannot dominate any $v_A \in [\frac{1}{4k}, 1 - \frac{1}{4k}]$. Each one of the other $k-1$ thresholds of B dominates a range of size $\frac{1}{2k}$, so the total range of dominated values for A is at most $\frac{k-1}{2k}$. The probability that a random v_A will be dominated is thus at most $\frac{\frac{k-1}{2k}}{1-\frac{1}{2k}} < \frac{1}{2}$. When v_A is not dominated, $v_A, v_A + \frac{1}{4k}$ and $v_A - \frac{1}{4k}$ will lie within the same entry in the matrix representation of g. This will therefore happen with probability $> \frac{1}{2}$ (the probability that v_A is not dominated). v_B is determined randomly (and uniformly), thus whatever allocation is made in this entry, welfare loss will be incurred with probability $\ge \frac{1}{2}$. The welfare loss (if incurred) will clearly be of at least $\frac{1}{4k}$. Thus, the expected efficiency loss will be greater than $\frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{4k} = \frac{1}{16k}$. The generalization for n players is easy now (see e.g., Proposition B.2). □

[4]Again, we used the fact that when $z = (z_1, ..., z_m)$ is in the m'th dimensional simplex, $\sum_{i=1}^{m} z_i^3 \ge \frac{1}{m^2}$
[5]see Theorem 3.7 for similar analysis

Appendix C

Implementation with a Restricted Action Space

C.1 Missing Proofs from Section 4.3

Proof of Lemma 4.1:

Proof. We first observe that for every social-value function there exists an informationally optimal k-action mechanism with a *deterministic* allocation scheme. This observation is general and does not require the use of threshold strategies or single-crossing conditions. Consider an optimal k-action mechanism that achieves the optimal result with some set of strategies $s = s_1, ..., s_n$. At least the same expected social value will clearly be achieved by the following deterministic allocation scheme: for each profile of actions b, the mechanism chooses an alternative that maximizes the expected social value, i.e., $t(b) \in argmax_{A'} E_{\overrightarrow{\theta}}\left[g(\overrightarrow{\theta}, A') \mid \forall i \; s_i(\theta_i) = b_i\right]$. Of course, this procedure may ruin incentive-compatibility properties of the mechanisms, but we will handle the incentive considerations separately.

With this observation in hand, we now turn to prove the two directions of the lemma. By Proposition 4.1, it is sufficient to show that the optimum is achieved with threshold strategies if and only if the optimal k-action mechanism is monotone. $\underline{\Longleftarrow}$:

Denote the thresholds used by player i by $x^i_0, x^i_1, ..., x^i_k$. Namely, when player i reports an action b_i and uses a threshold strategy, her type lies between $[x^i_{b_i}, x^i_{b_i+1}]$. Consider a deterministic choice rule as described above, and consider an action profile $b = (b_1, ..., b_n)$. Let A and B be two alternatives such that $A \succeq_i B$ (as determined by the single-crossing property). Now consider another action vector $b' = (b'_i, b_{-i})$, where $b'_i > b_i$. An optimal mechanism chooses for each profile of bids the alternative that maximizes the expected social value. Let A be the alternative chosen by the mechanism under action profile b. For For proving monotonicity, it suffices to show that if A gains a higher pxpected social value than B for the action profile b, this will also hold for the action profile b'. That is, if

$$E_{\overrightarrow{\theta}}\left[g(\overrightarrow{\theta}, A) \mid s(\overrightarrow{\theta}) = b\right] \;\; \geq \;\; E_{\overrightarrow{\theta}}\left[g(\overrightarrow{\theta}, B) \mid s(\overrightarrow{\theta}) = b\right]$$

then

$$E_{\overrightarrow{\theta}}\left[g(\overrightarrow{\theta}, A) \mid s(\overrightarrow{\theta}) = b'\right] \;\; \geq \;\; E_{\overrightarrow{\theta}}\left[g(\overrightarrow{\theta}, B) \mid s(\overrightarrow{\theta}) = b'\right]$$

This will be an immediate conclusion from the following intuitive statement: fixing θ_{-i}, the expected difference in social value between alternatives A and B is greater for b' than for b.[1] Formally,

$$E_{\theta_i}\left[g(\overrightarrow{\theta},A) \;-\; g(\overrightarrow{\theta},B)\;\middle|\; s_i(\theta_i)=b'_i\right] \tag{C.1}$$

$$= \frac{1}{F_i(x^i_{b'_i+1})-F_i(x^i_{b'_i})}\int_{x^i_{b'_i}}^{x^i_{b'_i+1}}\left(g(\overrightarrow{\theta},A) \;-\; g(\overrightarrow{\theta},B)\right)f_i(\theta_i)d\theta_i \tag{C.2}$$

$$\geq \frac{1}{F_i(x^i_{b'_i+1})-F_i(x^i_{b'_i})}\int_{x^i_{b'_i}}^{x^i_{b'_i+1}}\left(g(x^i_{b_i+1},\theta_{-i},A) \;-\; g(x^i_{b_i+1},\theta_{-i},B)\right)f_i(\theta_i)d\theta_i \tag{C.3}$$

$$= \frac{1}{F_i(x^i_{b_i+1})-F_i(x^i_{b_i})}\int_{x^i_{b_i}}^{x^i_{b_i+1}}\left(g(x^i_{b_i+1},\theta_{-i},A) \;-\; g(x^i_{b_i+1},\theta_{-i},B)\right)f_i(\theta_i)d\theta_i \tag{C.4}$$

$$\geq \frac{1}{F_i(x^i_{b_i+1})-F_i(x^i_{b_i})}\int_{x^i_{b_i}}^{x^i_{b_i+1}}\left(g(\overrightarrow{\theta},A) \;-\; g(\overrightarrow{\theta},B)\right)f_i(\theta_i)d\theta_i \tag{C.5}$$

$$= E_{\theta_i}\left[g(\overrightarrow{\theta},A) \;-\; g(\overrightarrow{\theta},B)\;\middle|\; s_i(\theta_i)=b_i\right] \tag{C.6}$$

Where inequalities C.3 and C.5 are due to the single-crossing property of the social-value function, and Equations C.3 and C.4 are equal since they are the expected value of the same constant value.

$\underline{\Longrightarrow}$: We now assume that a mechanism possesses a monotone allocation scheme, and prove that the optimum is achieved with threshold strategies.

The basic idea: we consider the expected social value of some player as a function of her type θ_i when she chooses a particular action. We show that such functions, for every two actions $b_i < b'_i$, cross at most once; that is, if for some θ^*_i the expected social value is equal when player i chooses either b_i or b'_i, then for any $\theta_i > \theta^*_i$ the expected social value when choosing b_i is at most the expected social value when choosing b'_i. The optimality of threshold strategies for this mechanism will be derived directly from this weak single-crossing property.

Consider two actions $b'_i > b_i$ for player i. Let θ^*_i be a type for which the expected social value is equal either when player chooses b'_i or b_i, that is (we denote the actions of the players except i when their types are θ_{-i} by $s_{-i}(\theta_{-i})$):

$$E_{\theta_{-i}}\left[\, g(\theta^*_i,\ \theta_{-i},\ t(b_i,s_{-i}(\theta_{-i})\ )\right] = E_{\theta_{-i}}\left[\ g(\theta^*_i,\ \theta_{-i},\ t(b'_i,s_{-i}(\theta_{-i})\ )\right] \tag{C.7}$$

We will show that for every $\theta_i > \theta^*_i$, the expected social value when player i chooses b_i is at most the expected social value in b'_i .

Monotonicity implies that $t(b'_i,b_{-i}) \succeq_i t(b_i,b_{-i})$ for every b_{-i}. Consider some profile of actions of the other players b_{-i}, and denote $t(b_i,b_{-i})=A$ and $t(b'_i,b_{-i})=B$. Since the social value function

[1] Note that due to the linearity of expectation,

$$E_{\overrightarrow{\theta}}\left[g(\overrightarrow{\theta},A)\quad s(\overrightarrow{\theta})=b\right] \;-\; E_{\overrightarrow{\theta}}\left[g(\overrightarrow{\theta},B)\quad s(\overrightarrow{\theta})=b\right]$$
$$= \; E_{\theta_{-i}}\left[\ E_{\theta_i}\left[g(\overrightarrow{\theta},A) \;-\; g(\overrightarrow{\theta},B)\quad s_i(\theta_i)=b_i\right]\quad s_{-i}(\theta_{-i})=b_{-i}\right]$$

is single crossing, the change in the expected social value when alternative B is chosen must be greater, that is:

$$\begin{aligned} & E_{\theta_{-i}}\left[g(\theta_i,\theta_{-i},B) \mid s_{-i}(\theta_{-i})=b_{-i}\right] - E_{\theta_{-i}}\left[g(\theta_i^*,\theta_{-i},B) \mid s_{-i}(\theta_{-i})=b_{-i}\right] \\ \geq\ & E_{\theta_{-i}}\left[g(\theta_i,\theta_{-i},A) \mid s_{-i}(\theta_{-i})=b_{-i}\right] - E_{\theta_{-i}}\left[g(\theta_i^*,\theta_{-i},A) \mid s_{-i}(\theta_{-i})=b_{-i}\right] \end{aligned}$$

Now, considering all possible b_{-i} and using the linearity of expectation, we get:

$$E_{\theta_{-i}}\left[g\left(\theta_i,\theta_{-i},t(b_i',s_{-i}(\theta_{-i}))\right)\right] - E_{\theta_{-i}}\left[g(\theta_i^*,\theta_{-i},t(b_i',s_{-i}(\theta_{-i})))\right] \quad \text{(C.8)}$$

$$\geq\ E_{\theta_{-i}}\left[g(\theta_i,\theta_{-i},t(b_i,s_{-i}(\theta_{-i})))\right] - E_{\theta_{-i}}\left[g(\theta_i^*,\theta_{-i},t(b_i,s_{-i}(\theta_{-i})))\right] \quad \text{(C.9)}$$

Due to Equation C.7 and Inequality C.8-C.9, indeed for any $\theta_i > \theta_i^*$ the expected social value in b_i is at most the social welfare in b_i':

$$E_{\theta_{-i}}\left[g(\theta_i,\theta_{-i},t(b_i',s_{-i}(\theta_{-i})))\right] \geq E_{\theta_{-i}}\left[g(\theta_i,\theta_{-i},t(b_i,s_{-i}(\theta_{-i})))\right]$$

Finally, we conclude that the optimal social value can be achieved with threshold strategies for k-action games; each player should choose, for every type θ_i, the action that maximizes the expected social value. The maximum over k pairwise single-crossing functions have at most $k-1$ switching points between the functions, therefore the social value is maximized using a threshold strategy that always chooses the action with the highest social value. (A similar argument is given for maximum of linear functions, that are also single-crossing in this since, in Theorem 4.1.) □

C.2 Missing Proofs from Section 4.5

Proof of Theorem 4.3:

Proof. Since the social-value function is multilinear and single crossing, the optimal expected social value is achieved by threshold strategies and therefore in a monotone mechanism (Lemma 4.1 and Theorem 4.1). To show that the mechanism is diagonal, we should also show that the allocation scheme is non-degenerate with respect to one of the players.

We prove the theorem for the case where the preferences $\succ_i$ of the player are conflicting, and the proof for correlated preferences is similar. We assume, w.l.o.g., that $A \succ_1 B$ and $B \succ_2 A$ and that $g(\underline{\theta}_1,\underline{\theta}_2,A) \geq g(\underline{\theta}_1,\underline{\theta}_2,B)$. For such preferences, we show that the optimal mechanism will be non-degenerate with respect to Player 2. In other words, in the matrix representation of the optimal mechanism there will be no identical columns. Showing this will suffice, as it is easy to see that in a monotone allocation scheme where the column player has k distinct columns, the row player clearly has either $k-1$ or k players.[2]

If Player 2 has two identical columns, then monotonicity derives that these columns will be adjacent, so in an equivalent allocation scheme this player will actually have $k-1$ possible actions. We will prove that a mechanism where Player 2 has $k-1$ possible actions cannot be optimal, since

[2] Assuming that $g(\underline{\theta}_1,\underline{\theta}_2,A) \geq g(\underline{\theta}_1,\underline{\theta}_2,B)$, we can show that the optimal allocation scheme is non-degenerate with respect to Player 2. If the converse is true, we can show in the same way that the optimal allocation scheme is non-degenerate w.r.t. Player 1. Similar arguments also prove that when $g(\overline{\theta}_1,\overline{\theta}_2,A) \geq g(\overline{\theta}_1,\overline{\theta}_2,B)$ the optimal allocation will be non-degenerate w.r.t. Player 2 (otherwise, w.r.t. Player 1). Therefore, a sufficient condition for having an optimal allocation scheme that is non-degenerate w.r.t. both players is having both $g(\underline{\theta}_1,\underline{\theta}_2,A) \geq g(\underline{\theta}_1,\underline{\theta}_2,B)$ and $g(\overline{\theta}_1,\overline{\theta}_2,A) \leq g(\overline{\theta}_1,\overline{\theta}_2,B)$, or when both inequalities are in the opposite direction.

we can add a new column and strictly increase the expected social value. We therefore assume that the optimal k-action social value is achieved when Player 1 uses the threshold vector $x_0, ..., x_k$ and Player 2 has $k-1$ possible actions and uses the threshold vector $y_0, ..., y_{k-1}$.

Case 1: The column $[A, A, ..., A]$ does not appear in the allocation matrix.

We will add this column to the game as the first column (action "0"), and add an additional threshold y' such that the expected social value strictly improves in the new mechanism when Player 2 uses the threshold vector $y_0, y', y_1, ..., y_{k-1}$. Consider the expected difference between the social value of the two alternatives when both players report 0, as a function of the second threshold of Player 2:

$$diff(y) \quad = \quad E_{\theta_1,\theta_2}\left[g(\theta_1,\theta_2,A) \; - \; g(\theta_1,\theta_2,B) \;\mid\; \theta_1 \in [x_0,x_1],\; \theta_2 \in [y_0,y]\,\right]$$

We know that $diff(y_0) > 0$ (since we assumed that $g(\underline{\theta}_1, \underline{\theta}_2, A) \geq g(\underline{\theta}_1, \underline{\theta}_2, B)$ and due to the single-crossing property). We also know that $diff(y_1) < 0$, otherwise alternative A would be preferred in this entry and the column $[A, ..., A]$ would have existed (monotonicity). Due to the Intermediate-Value theorem, there must be some $y^* \in (y_0, y_1)$ for which $diff(y^*) = 0$ ($diff(\cdot)$ is clearly continuous since each both $g(\theta_1, \theta_2, A)$ and $g(\theta_1, \theta_2, B)$ are continuous w.r.t. θ_2). Setting y' to be, for example, $\frac{y_0+y^*}{2}$ ensures that when θ_2 is between $[y_0, y']$ and when Player 1 reports "0", the expected social value strictly increases. The allocation in all other cases remains unchanged.

Case 2: when the column $[A, A, ..., A]$ exists.

Since there are $k+1$ possible columns of the form $[B, B, ..., A, A]$ and only $k-1$ columns in the allocation matrix, it must be the case that some "internal" column is missing, hence, there are actions $i, i+1$ for Player 1 and $j, j+1$ for Player 2 such that $t(i,j) = t(i+1,j) = A$ and $t(i,j+1) = t(i+1,j+1) = B$. We will show that adding an action (column) j' for Player 2, between actions j and j' in the order on the actions, that is identical to the allocation in column j except $t(i,j') = B$, will strictly increase the expected social value. For the exact construction, we have to consider two different sub-cases: if the expected social value when Player 1 reports 0 and Player 2's type is y_{j+1} is greater for alternative A than for B, then we will define a new threshold which is greater than y_{j+1}; Otherwise, the threshold will be smaller than y_{j+1}:

Case 2.1.: $E\,[\,g(\theta_1, y_{j+1}, A) \mid \theta_1 \in [x_i, x_{i+1}]\,] \geq E\,[\,g(\theta_1, y_{j+1}, B) \mid \theta_1 \in [x_i, x_{i+1}]\,]$.

Due to the (strict) single-crossing condition, clearly

$$E\,[\,g(\theta_1, y_{j+1}, A) \mid \theta_1 \in [x_{i+1}, x_{i+2}]\,] > E\,[\,g(\theta_1, y_{j+1}, B) \mid \theta_1 \in [x_{i+1}, x_{i+2}]\,]$$

Therefore, due to similar intermediate-value considerations, there must be some threshold $y^* > y_{j+1}$ for which

$$E\,[\,g(\theta_1, y_{j+1}, A) \mid \theta_1 \in [x_{i+1}, x_{i+2}]\,] = E\,[\,g(\theta_1, y_{j+1}, B) \mid \theta_1 \in [x_{i+1}, x_{i+2}]\,]$$

Now, let Player 2 use the threshold strategy based on the vector $y_0, ..., y_{j+1}, y', .., y_{k-1}$, for example, $y' = \frac{y_{j+1}+y^*}{2}$. The expected social value strictly increases when $\theta_1 \in [x_i, x_{i+1}], \theta_2 \in [y_{j+1}, y']$), while the allocation in all other cases remains unchanged.

Case 2.2.: $E\left(g(\theta_1, y_{j+1}, A \mid \theta_1 \in [x_i, x_{i+1}]\right) < E\left(g(\theta_1, y_{j+1}, B \mid \theta_1 \in [x_i, x_{i+1}]\right)$

Let y^* be again the value for which

$$E\left(g(\theta_1, y^*, A) \mid;\ \theta_1 \in [x_i, x_{i+1}]\right) = E\left(g(\theta_1, y^*, B) \mid \theta_1 \in [x_i, x_{i+1}]\right)$$

Clearly, now $y^* < y_{j+1}$. Similar arguments show that adding a new threshold $y' = \frac{y^*+y_{j+1}}{2}$ yields a higher expected social surplus.

Given that the mechanism is diagonal, it is clear that each threshold of a player affects the decision that is made only for one action of the other player. Therefore, it is easy to see that each threshold must be a maximizer, based on the arguments given in Section 4.5.2. □

Appendix D

Informational Limitations of Ascending Combinatorial Auctions

D.1 Critical Price Levels

In this subsection we give a simple, formal argument, to be used in the proofs of the impossibility results, saying that if an auction does not give an opportunity for a bidder to demand some bundle S, by presenting relevant levels of prices ("critical price levels"), then the auction reveals no information at all about the value of S.

Some notations that describe the uncertainty of the auctioneer regarding the bidders: Denote the set of all the possible valuations for bidder i by V_i. Also denote the set of all possible values for the bundle S in V_i by $Q_i(S) = \{v_i(S) \mid v_i \in V_i\}$. Finally, denote the set of the possible values for the bundle S, given that the realization of the value of some other bundle T is c_T, by $Q_i(S \mid v_i(T){=}c_T) = \{v_i(S) \mid v_i \in V \text{ and } v_i(T){=}c_T\}$.

First we define informationally-independent classes of valuations – valuations where obtaining information regarding any set of bundles adds no new information about the possible values of other bundles.

Definition D.1. *We say that a set V_i of valuations for bidder i is* informationally independent*, if for any bundle S, and any realization of the values of the other bundles $\{c_T\}_{T \neq S}$, the set of possible values for S remains unchanged. Namely, for every $S \subseteq M$,*

$$Q_i(S) \quad = \quad Q_i\left(S \mid v_i(T) = c_T \text{ for every } T \neq S\right)$$

Definition D.2. *Denote the class of all possible valuations of bidder i by V_i. We say that the price level p is* critical for Bidder i with respect to the bundle S*, if for some $v_i \in V_i$, Bidder i demands the bundle S under the price level p.*

The next easy proposition implies that if no critical price vector is presented to a bidder regarding some bundle S, then no information at all will be elicited on the value of this bundle. The proposition also holds for non-ascending auctions, and for all pricing schemes.

Proposition D.1. *Consider a bidder i, with an informationally-independent set of possible valuations V_i. If an auction reaches no critical price level for Bidder i with respect to a bundle S, then, at the end of the auction, no information is revealed on the value of S, that is, the set of possible values for S remains $Q_i(S)$.*

Proof. The proof is straightforward: Since no critical price level with respect to the bundle S is presented to Bidder i, then the data accumulated throughout the auction is completely independent of the value $v_i(S)$. Since the demands of the other bidders are also unchanged, and these demands are the only data that is available to the auctioneer, the auctioneer will not be able to differentiate between different values of $v_i(S)$. Therefore, no value of $v_i(S)$ can be ruled out. □

D.2 Limitations of Item-Price Ascending Auctions

Example D.1. *This example shows that a single item-price auction can elicit an exponential amount of information. Consider two bidders in a combinatorial auction with preferences of the following type:* $v(S) = 1$ *for every bundle* S *with more than* $\frac{m}{2}$ *items,* $v(S) = 0$ *if* $|S| < \frac{m}{2}$ *and every* S *such that* $|S| = \frac{m}{2}$ *has an unknown value from of either* 0 *or* 1. *As proved by [117], for determining the efficient allocation, the bidders may be required to communicate an amount of information which is exponentially larger than the number of items. However, using small enough increments, it is easy to determine the values of all the bundles of size* $\frac{m}{2}$ *by an ascending auction.*[1] *This information clearly suffices for determining the optimal allocation.*

Definition D.3. *([82]) A valuation* v *is said to satisfy the substitutes (or gross-substitutes) property if for every pair of item-price vectors* $\overrightarrow{q} \geq \overrightarrow{p}$ *(coordinate-wise comparison), if* $S = \{j \in M | p_j = q_j\}$ *and* A *maximizes the bidder's utility under the price vector* $\overrightarrow{p}$*, then there exists a bundle* B *that maximizes the bidder's utility under the price vector* $\overrightarrow{q}$ *such that* $S \cap A \subseteq B$*.*

Definition D.4. ***(k-trajectory ascending auctions)*** *Consider an auction* $\mathcal{A}$*, and denote the set of all the price vectors presented to bidder* i *in* $\mathcal{A}$ *by* $\mathcal{P}_i$*.*[2] *We say that* $\mathcal{A}$ *is a* k-trajectory ascending auction *if for every bidder* i*, the set* $\mathcal{P}_i$ *can be divided into* k *ascending trajectories of prices* $\mathcal{P}_i(1), ..., \mathcal{P}_i(k)$*. Formally,* $\cup_{j=1}^{k} \mathcal{P}_i(j) = \mathcal{P}_i$ *and for every* $j \in \{1, .., k\}$*, and for every two price vectors* $p, q \in \mathcal{P}_i(j)$ *such that* q *was presented to bidder* i *at a later stage in* $\mathcal{A}$ *than* p*, and for every bundle* $S \subseteq M$*, we have that* $q(S) \geq p(S)$*.*

Proof of Theorem 1a:

Proof. Consider a single agent with a valuation with the following properties: For every bundle S such that $|S| > \frac{m}{2}$ we have $v(S) = 2$, and for every $|S| \leq \frac{m}{2}$ we have $v(S) = 0$, except for a single unknown bundle T of size $\frac{m}{2}$ that either has a value of $1 - \delta$ (for some small $\delta > 0$) or 0. We first show that finding the hidden bundle T requires an exponential number of ascending item-price *trajectories*, even if the auctioneer knows these properties of the valuations.

Recall that under a "critical" price level with respect to the bundle S, the player demands S for some realization of his valuation (see Definition D.2 in Appendix D.1). We first prove the following claim:

Claim D.1. *In an ascending auction, if the bidder is presented with a critical price vector for some bundle* S *of size* $\frac{m}{2}$*, then no critical price vector will be published at later stages of the ascending auction with respect to any other* $\frac{m}{2}$*-sized bundle.*

[1] This can be done by enumerating on all the different bundles of size $\frac{m}{2}$, and for each bundle S set the prices of the items in S to some value λ and set the prices of the items not in S to $\lambda + \epsilon$ for sufficiently small ϵ. Clearly, the bundle S will be demanded if and only if $v_i(S) = 1$. Using exponentially small increments, we can construct such vectors of prices during a a single ascending path of prices.

[2] Recall that each price vector p specifies a price $p(S)$ for every bundle $S \subseteq M$.

Proof. Let $\overrightarrow{p}$ be a critical price vector presented to the bidder with respect to some bundle S, $|S| = \frac{m}{2}$. Thus, for some possible value of $v(S)$ and for any item $x \in M \setminus S$, the bidder (weakly) prefers the bundle S over the bundle $\{S \cup x\}$, i.e., $v(S \cup x) - p(S \cup x) \le v(S) - p(S)$. Since the prices are linear, and since $v(S)$ is always smaller than 1, it follows that: $p_x \ge v(S \cup x) - v(S) > 2 - v(S) > 1$. Thus, the price of *any item* in $M \setminus S$ is strictly greater than 1. Since the prices are ascending, it follows that the bidder will not demand any bundle of size $\frac{m}{2}$ containing an item from $M \setminus S$ at later stages of the auction. (Clearly, the only bundle of size $\frac{m}{2}$ that does not contain any item from $M \setminus S$ is S.) □

Due to Claim D.1, an ascending path of prices can only contain critical price levels with respect to one of the $\frac{m}{2}$-sized bundles. Therefore, this ascending trajectory will be independent of the values of all the other $\frac{m}{2}$-sized bundles, and no new information will be elicited on them (this holds since the valuations are informationally independent – see Proposition D.1). It follows that in each ascending trajectory, the auctioneer has to arbitrarily decide which $\frac{m}{2}$-sized bundle will be checked. An adversary (or "nature") may choose a valuation such that the last (or before last) bundle to be checked is the bundle T. Since the number of $\frac{m}{2}$-item bundles is exponential in m,[3] an exponential number of ascending trajectories is required for finding the hidden bundle.

Now, consider a second bidder that has a value of 2 for every bundle of size $\frac{m}{2}$ or more. The optimal allocation will clearly allocate the bundle T to Bidder 1, and the other $\frac{m}{2}$ items to the second bidder. Finding the efficient allocation for these two bidders is equivalent to finding the bundle T. The theorem follows. □

Proof of Theorem 1b:

Proof. Consider n bidders and n^2 items for sale, and assume that n is prime.[4] We construct a total of n^2 distinct bundles with the following properties: for each bidder i ($1 \le i \le n$), we define a partition $S^i = (S^i_1, ..., S^i_n)$ of the n^2 items to n bundles, such that any two bundles from different partitions intersect (i.e., for every two bidders $i \ne j$, and every k, l we have $S^i_k \cap S^j_l \ne \emptyset$). We call this combinatorial structure *mutually-intersecting partitions*. In Appendix D.4, we show an explicit construction of mutually-intersecting partitions using the properties of linear functions over finite fields. The rest of the proof is independent of the specific construction.

We now build a set of valuations for the bidders, and prove that they are hard to elicit by item-price ascending auctions. Each bidder i will have a value of 2 for every bundle that contains a union of two bundles from different partitions, and an unknown value of either 0 or $1 - \delta$ (for some small $\delta > 0$) for bundles that contain only a single bundle from a partition (henceforth, the "low-valued" bundles). More formally, each bidder will have the following valuation (the value of any other bundle is the maximal value of a bundle that it contains):

- A value of 2 for the bundle $S^{j'}_k \cup S^j_l$, for every k, l and every $j' \ne j$.
- A value of either 0 or $1 - \delta$ (unknown to the seller) for the bundle S^j_k, for every j, k.

[3] According to Stirling's formula, the number of distinct bundles of size $\frac{m}{2}$, out of m distinct items, is approximately $\sqrt{\frac{2}{\pi m}} \cdot 2^m$.

[4] Due to the celebrated Bertrand Conjecture from 1845 (proved by Chebyshev in 1850), for every natural number n there exists at least one prime number between n and $2n$. Therefore, we can assume that n is prime, where the number of items is at most twice the original number. This will result in an additional factor of 2 in our approximation result.

Note that at most one bidder can gain a value of 2, since every two 2-valued bundles contain bundles from different partitions and thus must intersect. Therefore, for achieving more than a welfare of 2, we must allocate low-valued bundles. However, as the following claim shows, the demand of a bidder during a single ascending auction can only reveal information about his values for bundles from a single partition.

Claim D.2. *If a bidder is presented with a critical price vector with respect to a bundle from one partition, no critical price levels will be presented to this bidder with respect to bundles from other partitions at later stages of the ascending auction.*

Proof. Let p be a critical price level for Bidder i with respect to his low-valued bundle S_k^j. Then, for every bundle S_k^l from a different partition (i.e., $l \neq k$), we have:

$$v(S_k^j) - p(S_k^j) \quad \geq \quad v(S_k^j \cup S_k^l) - p(S_k^j \cup S_k^l)$$

Since the prices are linear, it follows that:

$$p(S_k^l) \quad \geq \quad p(S_k^j \cup S_k^l) - p(S_k^j) \quad \geq \quad v(S_k^j \cup S_k^l) - v(S_k^j) \quad > \quad 1$$

where the final inequality holds since $v(S_k^j) < 1$. Hence, the bundle S_k^l will not be demanded before the auction concludes. □

It follows from the claim above that every ascending trajectory of prices will be independent of the values of every bidder to bundles from all the partitions, except at most one partition. Hence, for each bidder, the auctioneer will gain information about at most one partition of the n partitions. Therefore, for every ascending auction, there must exist a partition j (i.e., $S_1^j, ..., S_n^j$) for which *at most* one bidder revealed some information. An adversary ("nature") can set the values of the bundles in all the other partitions such that any way of allocating them will result in a total value of at most 2. In addition, the total value of the bidders to bundles in partition j may be arbitrary close to n (that is, $n - n\delta$) – each bidder will have a value of $1 - \delta$ for one distinct bundle from this partition. The auctioneer does not have any information on the values that the bidders (except, maybe, one) have for bundles in this partition, and therefore the auctioneer will not be able to correctly match the bundles in this partition to the bidders; The auctioneer can only guarantee a value of 2 by allocating all items to a single bidder, as opposed to the optimal welfare that can be arbitrarily close to n (and here, $n = \sqrt{m}$). The theorem follows (as mentioned, we lose an additional factor of 2 since we assumed that n is prime). □

Proof of Theorem 1c:

Proof. Let w be the valuation that aggregates the preferences of the n original players. Since all the original valuations hold the substitutes property, then their aggregation, w, also has the substitutes property ([93]). Substitutes valuations are, in particular, "complement-free" – that is, for every two bundles S, T we have that $w(S) + w(T) \geq w(S \cup T)$. Due to the assumption that the w is not additive, there are two bundles S and T for which the inequality is strict,

$$w(S) + w(T) > w(S \cup T) \tag{D.1}$$

Substitute valuations are also submodular, and thus exhibit diminishing marginal valuations (see, e.g., [93]). Therefore, the marginal contribution of $M \setminus (S \cup T)$ in Inequality D.1 is greater for T than for $S \cup T$, thus,

$$w(S) + w(M \setminus S) > w(M) \tag{D.2}$$

Denote $\epsilon = w(S) + w(M \setminus S) - w(M)$. Now, consider the "dual" valuation to w denoted by $\overline{w}$, i.e., for every bundle X, $\overline{w}(X) = w(M) - w(M \setminus X)$. The dual valuation specifies the contribution of the bundle X to the welfare of the n players, given that they already hold the other items. Clearly, if an additional player has a value for S that exceeds $\overline{w}(S)$, allocating this bundle to her will increase the total welfare. Using Inequality D.2, we thus have that the bundles S and $M \setminus S$ are complements with respect to $\overline{w}$,

$$\begin{aligned} & \overline{w}(S) + \overline{w}(M \setminus S) & \text{(D.3)} \\ = \; & w(M) - w(M \setminus S) + w(M) - w(S) & \text{(D.4)} \\ = \; & w(M) - (w(M \setminus S) + w(S) - w(M)) & \text{(D.5)} \\ = \; & \overline{w}(M) - \epsilon & \text{(D.6)} \end{aligned}$$

We define an additional bidder k with the valuation $v_k(\cdot)$ for which $v_k(M) = \overline{w}(M)$ (which also equals $w(M)$), and the values $v_k(S)$ and $v_k(M \setminus S)$ are unknown to the auctioneer and may take the following values: $v_k(S) \in \{\ \overline{w}(S),\ \ \overline{w}(S) + \frac{\epsilon}{6},\ \ \overline{w}(S) + \frac{\epsilon}{3}\ \}$ and
$v_k(M \setminus S) \in \{\ \overline{w}(M \setminus S),\ \ \overline{w}(M \setminus S) + \frac{\epsilon}{6},\ \ \overline{w}(M \setminus S) + \frac{\epsilon}{3}\ \}$.

The values of all the other bundles is the maximal value of a bundle, from the above bundles, that they contain.

An efficient auction clearly has to determine which of the bundles S and $M \setminus S$ adds more value for the new bidder with respect to $\overline{w}$. We will show that an ascending item-price auction will not be able to find this bundle using the following claim. (The concept of critical price levels is defined in Definition D.2 in Appendix D.1.)

Claim D.3. *If a critical price level is presented to player k with respect to the bundle S, no critical price levels will be presented with respect to the bundle $M \setminus S$ at later stages of the ascending auction.*

Proof. Let p be a critical price level with respect to the bundle S. Then for some value of $v_k(S)$ the player will prefer this bundle to the whole bundle: $v_k(S) - p(S) \geq v_k(M) - p(M)$. Due to the linearity of the prices and the definition of $v_k(\cdot)$ it follows that:

$$\begin{aligned} & p(M \setminus S) \quad \geq \quad v_k(M) - v_k(S) \quad \geq \quad \overline{w}(M) - \overline{w}(S) - \frac{\epsilon}{3} & \text{(D.7)} \\ > \; & \overline{w}(M \setminus S) + \epsilon - \frac{\epsilon}{3} \quad > \quad v_k(M \setminus S) & \text{(D.8)} \end{aligned}$$

Where Inequality D.8 follows from Equation D.6. The price of the bundle $M \setminus S$ is greater than all its possible values, and this bundle will not be demanded at future stages since the prices are ascending. □

Similarly, we can also show that if a critical price is presented with respect to $M \setminus S$, then all future price levels will be independent of the value of S. Therefore, the auctioneer will be able to elicit information only on one of the bundles S and $M \setminus S$, and the optimal allocation will remain unknown. □

D.3 Limitations of Anonymous Ascending Auctions

Proof of Theorem 2a:

Proof. Consider n bidders and n^2 items, and assume that n is prime.[5] Consider n^2 distinct bundles defined by mutually-intersecting partitions (see Theorem 1b), that is, for each bidder, we define a partition $S^i = (S_1^i, ..., S_n^i)$ of the n^2 items to n bundles, such that any two bundles from different partitions intersect. (As mentioned, an explicit construction is given in Appendix D.4.)

Using these n^2 bundles we construct the following valuations. We will define the values that the bidders have for each one of these n^2 bundles, and again, the value of any other bundle is the maximal value of a bundle that it contains. A bidder i has a value of 2 for any bundle S_j^i in his partition (i.e., the i'th partition). For all the bundles in the other partitions, he has a value of either 0 or of $1 - \delta$ (for some small $\delta > 0$), and these values are unknown to the auctioneer. Since every pair of bundles from different partitions intersect, at most one bidder can receive a bundle with a value of 2. Nonetheless, for some realizations of the bidders' preferences, we may allocate the bundles of a particular partition, one bundle per each bidder, such that one bidder gains a value of 2 and all the others receive a value of $1 - \delta$.

Consider the valuations described above. In every anonymous ascending auction, a bidder will not demand one of his low-valued bundle as long as the price of at least one of his high-valued bundles is below 1 (which gains him a utility greater than 1 for this bundle). Therefore, for eliciting any information about low-valued bundles, the auctioneer should first arbitrarily choose a bidder (w.l.o.g., Bidder 1) and raise the prices of *all* the bundles $S_1^1, ..., S_n^1$ to be greater than 1. Since the prices cannot decrease, no critical price level (see Definition D.2) will be presented with respect to any of these bundles at later stages of the auction for any bidder. Since the valuations are informationally independent, no information at all will be gained by the auctioneer on the values of these bundles (see Definition D.1 and Proposition D.1). It might happen that the low values of all the bidders for the bundles not in Bidder 1's partition are zero (i.e., $v_i(S_j^k) = 0$ for every bidder i and any partition $k \neq 1$ and every bundle j in it). However, allocating each bidder a different bundle from Bidder 1's partition, might achieve a welfare of $n + 1 - (n-1)\delta$ (Bidder 1's valuation is 2, and $1 - \delta$ for all other bidders); The auctioneer has no information on the values that the other bidders have for these bundles. Therefore, for every decision the auctioneer makes about the allocation, an adversary ("nature") may choose a profile of valuations for which no more than a welfare of 2 is achieved (2 for Bidder 1's high-valued bundle, 0 for all other bidders). We conclude that no anonymous bundle-price ascending auction can guarantee a welfare greater than 2 for this class, where the optimal welfare can be arbitrarily close to $n + 1$. The theorem follows. □

D.4 Constructing Mutually-Intersecting Partitions

We now present an explicit construction for the combinatorial structure used in Theorem 1b.We also use this combinatorial structure when we prove the inefficiency of anonymous bundle-price ascending auctions in Theorem 2a. We assume that there are n bidders and n^2 items (n is prime). For every bidder i, we define a partition $S^i = (S_1^i, ..., S_n^i)$ of the n^2 items to n bundles of size n, such that any two bundles from different partitions intersect (i.e., $S_j^i \cap S_l^k \neq \emptyset$ for every $i \neq k$ and every l, j). Figure D.1 describes such a construction for 3 bidders and 9 items.

[5] We can assume this and lose a factor of two in the approximation ratio. See the proof of Theorem 1b.

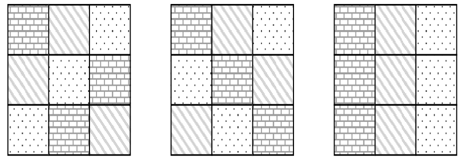

Figure D.1: Mutually-Intersecting partitions for 3 bidders and 9 items. Each partition is defined by parallel linear functions over the relevant field. Indeed, bundles from different partitions intersect.

We use the properties of linear functions over finite fields (for that, we denote the bidders by $0, ..., n-1$):

Recall that $Z_n = \{0, ..., n-1\}$ is a field if (and only if) n is prime. Denote the n^2 items for sale by pairs of numbers in Z_n. Each linear function $ax + b$ over the finite field Z_n denotes an n-item bundle (a total of n^2 bundles where $a, b \in Z_n$). The items in each bundle are the pairs $(x, ax + b)$ for every $x \in Z_n$. The bundles assigned to Bidder i are the n bundles $ix + b$ where $b \in Z_n$ (that is, all the parallel linear functions with a slope i). We need to show that the bundles assigned to Bidder i form a partition, and indeed the functions $ix + b_1$ and $ix + b_2$ cannot intersect when $b_1 \neq b_2$. It is also easy to see that every two bundles that are assigned to different bidders do intersect: consider the functions $ix + b_1$ and $jx + b_2$. Since z_n is a field, clearly an x exist such that $x(j - i) = (b_1 - b_2)$ when $j \neq i$ for any b_1, b_2. The jth bundle of Bidder i is therefore, $S_j^i = \{(0, i \cdot 0 + j), ..., (n - 1, i \cdot (n - 1) + j)\}$.

Appendix E

More on the Power of Ascending Combinatorial Auctions

E.1 Separation Results for Different Types of Ascending Auctions

The classes of valuations used in the proofs in this section are represented by "XOR" valuations (defined in Sub-section 2.3.1).

Proposition E.1. *There are classes of valuations for which the efficient allocation can be determined by a* descending *auction but not by an* ascending *auction.*

Proof. Consider a class of XOR valuations with 3 terms over 3 items $\{a, b, c\}$ of the following form:

$$v = (abc : 2) \oplus (x : 1) \oplus (y : \alpha)$$

where α is an unknown value between $(0, \frac{1}{2})$ and x, y are some unknown consecutive items (i.e., $(x, y) \in \{(a, b), (b, c), (c, a)\}$).

We first show that the following descending auction can fully elicit this class of valuations: start with a price of 1 for all items. Then, the bidder will demand $\{x\}$. The identity of x also reveals item y, thus we can decrease p_y until $\{y\}$ is demanded, thus α is revealed. (Note that the bidder will not demand the bundle $\{abc\}$, since its price stays above 2.)

Next, we show that no ascending auction can learn this class of valuation: starting from zero prices, a change in the demand can occur only if either $p_x + p_w \geq 1$ or $p_y + p_w \geq 1$ (where w is the 3rd item besides x, y). Thus, information will be gained only after the auctioneer arbitrarily increases the price of one of the items above $\frac{1}{2}$ (without any input received until this point in the auction). If this item is item y, the bidder will clearly never demand the bundle $\{y\}$, and thus the value of α will not be revealed.

We now prove similar result for the elicitation of this class of valuations. Consider the class of valuations described above. We first show that a descending auction can always find the optimal allocation for any number of bidders (we assume $n > 2$, otherwise allocating all the items to any bidder is optimal). We start with a price of 1 for all items. Under these prices, each bidder will demand the bundle with a value of 1 (i.e., $\{x\}$). If we have three bidders that demand three different items, we allocate each of the item to the bidder that demands it, and it is clearly the optimal allocation. If all the bidders demand the same item x, then the optimal allocation is achieved by allocating all items to one of the bidders. If the bidders demand two different items, these items must be consecutive, w.l.o.g. a and b. An optimal allocation will allocate a, b to bidders that demands them, and c to the bidder with the highest value for it. This bidder is the first bidder to demand c when we decrease p_c. (The price of the whole bundle will still be greater than 2, so no bidder will demand it.)

Next, we show that no item-price ascending auction can determine the optimal allocation for the above class. Consider three bidders drawn from the class of valuations above. One bidder have a valuation of 1 for some item x_1, and the two other bidders have a valuation of 1 for the subsequent item x_2, and let x_3 be the 3rd item ($x_1, x_2, x_3 \in \{a, b, c\}$). The optimal allocation will allocate x_1 to the first bidder, x_3 to the bidder with the highest value for it (i.e., with the highest α), and x_2 to the remaining bidder. For any set of item prices which are all smaller than $\frac{1}{2}$, there will be no change in the demand of the bidders and no information about the identity of these items will be extracted. Thus, in order to elicit any information, the auctioneer must arbitrarily raise one of the prices of the items above $\frac{1}{2}$. However, if this item turns out to be x_3, then no bidder will demand the bundle $\{x_3\}$, and the auctioneer cannot know who is the bidder with the highest valuation for this bundle. Therefore, the allocation may not be optimal. □

Proposition E.2. *There are classes of valuations for which the efficient allocation can be determined by an* ascending *auction but not by a* descending *auction.*

Proof. Consider 4 items $\{x_1, x_2, x_3, x_4\}$, and a class of valuations of the form $v = (x_{i-1}x_ix_{i+1} : 2.5) \oplus (x_{i-1}x_i : 2) \oplus (x_i : \alpha)$ *or* of the form $v = (x_{i-1}x_ix_{i+1} : 2.5) \oplus (x_ix_{i+1} : 2) \oplus (x_i : \alpha)$, where α is an unknown value between $(0, 1)$, the index i is unknown and the indices are cyclic.

An ascending auction can learn such valuations: for zero prices, the bidder demands $x_{i-1}x_ix_{i+1}$, revealing the index i.[1]

After increasing the price of x_{i-1} and x_{i+1} to $\frac{1}{2}$, the bidder demands the 2-item bundle. Now, we raise the prices of all items except x_i to L. Finally, we raise the price of x_i and the price where the bidder stops demanding $\{x_i\}$ is α.

Next, we show that no item price descending auction can fully elicit this valuation. No information can be elicited as long as all prices are greater or equal to 1, since no bundle will be demanded (except, maybe, the 2-item bundle from which the identity of the item x_i is still unknown). Let x_j be the first item for which $p_j < 1$. The elicitor must arbitrarily choose such item for eliciting some information about x_i. However, it might happen that x_j is the second item (besides x_i) in the 2-item atomic bundle. In this case, we claim that the bundle $\{x_i\}$ will never be demanded since:

$$u(x_ix_j) = 2 - (p_{x_i} + p_{x_j}) = 1 - p_{x_i} + 1 - p_{x_j} > 1 - p_{x_i} > v(x_i) - p_{x_i} = u(x_i)$$

Thus, no information about $\alpha = v(x_i)$ will be revealed.

Now, we present a class which cannot be elicited by a descending auction, but is elicitable by an ascending auction. Consider two bidders from the class described above, with an extra information that the singletons for which they have non-zero valuations are not consecutive and are not the same item. In addition, there is a third bidder with the valuation:

$$v = (x_1x_2x_3 : 2.5) \oplus (x_1x_2x_4 : 2.5) \oplus (x_1x_3x_4 : 2.5) \oplus (x_2x_3x_4 : 2.5)$$

[1]If the bidder's tie-breaking rule favors the grand bundle, we can raise the price of every item in turn, with a small increment, until $x_{i-1}x_ix_{i+1}$ is demanded.

First, we show that a descending auction cannot find the optimal allocation for every realization of the valuations. For determining the optimal allocation, we must know which bidder has the greatest value for a singleton, i.e., we must find the value of the x_i for the two players with the unknown valuation. Even if we knew one of these values, we would still need to know whether the other value is smaller or greater. However, exactly the same proof as above shows that a descending auction cannot guarantee to extract any information about this unknown value.

An ascending auction can find the optimal allocation: under zero prices, both players demand the 3-item bundle with the valuation of 2.5. Thus, the singletons x_i and x_j with the non-zero valuations are revealed. We raise the prices of the other items (except x_i and x_j) to L. Since x_i and x_j are not adjacent, the utility from all the bundles, except these singletons, will be negative. Raising the prices of these two items reveals the unknown α's. □

Proposition E.3. *There are classes of valuations for which the efficient allocation can be found by a* non-anonymous *item-price ascending auction, but not by an* anonymous *item-price ascending auction.*

Proof. Consider three players with the following valuations:

$$v_1 = (ab : 2) \oplus (a : \alpha)$$

$$v_2 = (bc : 2) \oplus (b : \beta)$$

$$v_3 = (ca : 2) \oplus (c : \gamma)$$

Where α, β, γ are unknown values between $(0, 1)$. The optimal allocation should allocate a singleton to the bidder with the highest singleton valuation, and give the other items to the player that has a valuation of 2 for them.

A non-anonymous item-price ascending auction can easily find the optimal allocation, by raising, for each bidder, the price of the item not in his singleton (e.g., p_b for bidder 1), until each bidder demands his singleton and thus revealing his unknown value.

Any anonymous auction must raise the price of some item above 1, before it encounters any change in the demands of the bidders or gaining any other information about the unknown values. No information will be elicited about the value of this item for the player that has a non-zero value for it. □

Proposition E.4. *There are classes of valuations for which the efficient allocation can be found by a* non-deterministic *ascending auction, but not by a* deterministic *ascending auction.*

Proof. Consider bidder 1 with one of the following XOR valuations:

$$v = (ab : 3) \oplus (a : \alpha)$$

$$v = (ab : 3) \oplus (b : \beta)$$

Where $\alpha, \beta \in \{0.4, 0.6\}$ are unknown to the auctioneer. Clearly, a non-deterministic algorithm can guess the singleton, raise the other item until the singleton is demanded, and then increase the price of the singleton until the value is discovered.

No deterministic algorithm, however, can learn the valuation. For zero prices, the bidder will clearly demand the bundle ab. The bidder's demand could change only if the price of the bundle

ab is greater than 2, i.e., when either p_a or p_b are greater than 1. In this case, this singleton term will not be demanded, and the valuation will not be fully elicited.

Consider bidder 1 described above, and two other bidders with the XOR valuations $v_2 = (ac : 2.5)$ and $v_3 = (bc : 2.5)$.

The optimal allocation clearly depends on the value of the singleton of bidder 1, but we saw that bidder 1's valuation cannot be learned with a deterministic ascending auction. (Note that even non-anonymous deterministic ascending auctions cannot elicit these valuations.) □

Proposition E.5. *There are classes of valuations for which the efficient allocation can be found by a* simultaneous *non-anonymous ascending auction, but not by an* sequential *non-anonymous ascending auction.*

Proof. Consider three bidders with the following valuations:

$$v_1 = (abc : 2) \oplus (x : \frac{1}{3})$$

$$v_2 = (abc : 2) \oplus (y : \beta)$$

$$v_3 = (a : 2) \oplus (b : 2) \oplus (c : 2)$$

Where $x, y \in \{a, b, c\}$ are unknown items, and β is an unknown number between $(0, \frac{1}{2})$.

If $x = y$, the optimal allocation allocates x to bidder 1 if $\beta < \frac{1}{3}$, or otherwise it allocates y to bidder 2 (in both cases, bidder 3 receives the other items). If x and y are distinct, each bidder should receive one of these items, and the third item goes to bidder 3. Therefore, for determining the optimal allocation, the auctioneer has to reveal the identity of x, the identity of y and its value.

First, we show that a simultaneous auction cannot find the optimal allocation. Bidder 1 definitely demands the whole bundle $\{abc\}$ when its price is below $\frac{3}{2}$. Since in a simultaneous auction the sum of the prices in all trajectories is equal at every stage, one of the prices for bidder 2 must exceed $\frac{1}{2}$ at this point of time. If this item turns out to be y, player 2 will never demand this singleton, thus the value of β will never be revealed.

A sequential auction, however, can find the optimal allocation. We first raise bidder 1's price of some item to $\frac{5}{3}$ and thus find x[2]. Then, we raise the price bidder 2's price for an item different than x. If some singleton is demanded, we found y and its value. If no singleton is demanded, it follows that y and x are distinct items, thus the optimal allocation do not depend on the value of β. □

Proposition E.6. *There are classes of valuations for which the efficient allocation can be found by a* adaptive *item-price ascending auction, but not by an* oblivious *item-price ascending auction.*

Proof. Consider the following class of XOR valuations over the three items a, b, c:

$$v = \{(xy : 3) \oplus (yz : 3) \oplus (x : \alpha) \oplus (z : \beta)\}$$

Where α, β are unknown values between (0,1) (may be different between bidders) and $x, y, z \in \{a, b, c\}$ are distinct items.

A simple adaptive algorithm that fully elicits this class of valuations is the following: Raise the prices of the items in the bundle demanded under zero prices by $\epsilon = \frac{\delta}{2}$ (then the bidder clearly

[2] If no singleton is demanded, this also reveals x.

demands the other 2-item atomic bundle). Now, raise the price of the item in the intersection of the two bundles demanded so far ('y'), until the bidder demands some singleton. Then raise the price of this singleton to L, and continue raising the price of y until the other singleton is demanded. Concluding the values from the responses is then straightforward.

An oblivious algorithm cannot know in advance what is the item in the intersection of the two 2-item bundles: A necessary condition for the bidder to demand a singleton is that $p(xy) > 2$ and $p(yz) > 2$. Assume w.l.o.g. that a is the first item for which the oblivious algorithm increases its price above 1. Then, for a valuation with the terms $(ab : 3) \oplus (bc : 3) \oplus (a : \alpha) \oplus (b : \beta)$ the value α will not be learned.

Consider a bidder 1 whose valuation is drawn from the class described above. and a bidder 2 with the XOR valuation $v = (ab : 3) \oplus (bc : 3) \oplus (ac : 3)$. The optimal allocation will allocate the singleton with the highest value to bidder 1, and the other items to bidder 2. We saw that an oblivious auction cannot learn the first valuation, but an adaptive auction can. Since the second valuation is fully known, the claim about the elicitation follows. (Note that the theorem also holds for non-anonymous auctions.) □

E.2 More Positive Results on Ascending Auctions

Proof of Proposition 6.1

Proof. **Simulating value queries:**
The following ascending auction learns the value of a given bundle S:
<u>Initialization</u>: *start with a zero price for every item in S, and price of L for every item in $M \setminus S$.*
<u>Repeat</u>: *raise the price of each item in S by $\epsilon = \delta$ in turn, in a round-robin fashion.*
<u>Finally</u>: *Terminate when the bidder demands the empty set.*

We claim that from the information elicited by this ascending auction we can calculate $v(S)$. In the initial stage, the bidder demands S or another bundle with the same value (due to the free-disposal assumption). Let $T_1, ..., T_k$ be the bundles demanded by the bidder in the order they were demanded (bundles might repeat). We know that $T_1 = S$ and $T_k = \emptyset$. We prove that if we know the value of some bundle T_{i+1} we can calculate $v(T_i)$. Thus, since we know that $v(\emptyset) = 0$, $v(S)$ can be calculated (by induction).

If we could raise the prices continuously, the proof would be very easy. Since prices are increased in a discrete manner, we should be more careful. In particular, we assume that we know the tie breaking rule of the bidder (i.e., which bundle he would demand if he had few bundles with the highest utility).

Let $\overrightarrow{p}$ be the smallest vector of prices in which the bidders demands T_{i+1}, and this happened after we raised the price of some item k by ϵ. If T_i has a priority over T_{i+1} in the bidder's tie breaking rule, then we know that his utility from T_{i+1} under the prices $\overrightarrow{p}$ is ϵ-higher than his utility from T_i (clearly, $k \in T_i$ but $k \notin T_{i+1}$, otherwise the demand change wouldn't happen). Thus,

$$v(T_{i+1}) - p(T_{i+1}) = v(T_i) - p(T_i) + \epsilon$$

Since the prices are known, we can calculate $v(T_i)$ from $v(T_{i+1})$. Similarly, if the bidder's tie breaking rule favors T_{i+1}, then the utilities at this point should be equal, i.e.,

$$v(T_{i+1}) - p(T_{i+1}) = v(T_i) - p(T_i)$$

and we can similarly calculate $v(T_i)$ from $v(T_{i+1})$. The total running time is at most $\frac{m \cdot L}{\delta}$.

Simulating marginal-value queries:
Simulating $v(j|S)$: start with zero prices, and increase the prices of all the items in $M \setminus \{S \cup \{j\}\}$ to be L. Then, gradually increase p_j by δ and stop when the bidder stops demanding $M \setminus \{S \cup \{j\}\}$- and this price for the item j is $v(j|S)$. (A δ may be added to this value, as derived from the tie breaking rules of the bidder).

Simulating indirect-utility queries:
For simulating $IU(\overrightarrow{p})$ using demand queries, we first ask the bidder for his desired bundle under these prices $S = D_i(\overrightarrow{p})$. Then, we calculate $v(S)$ according to the procedure described for the above simulation of value queries and 6.2.
Simulating relative-demand queries:
We simulate $RD(\overrightarrow{p})$ by the following ascending auction:
Initialization: start with a price vector $\epsilon \overrightarrow{p}$ ($\epsilon > 0$).
Repeat: for the vector of prices $\overrightarrow{q}$, if the bidder demands a non empty set, raise prices to $\overrightarrow{q} + \epsilon \overrightarrow{p}$.
Finally: If the bidder demands the empty set at stage $t+1$, terminate the auction, and return the bundle S demanded at stage t as the answer.

Now we show that for the price vector $\overrightarrow{p}$, every other bundle T has a smaller relative weight than S (up to ϵ), i.e.,

$$\frac{v(S)}{p(S)} \geq \frac{v(T)}{p(T)} - \epsilon. \tag{E.1}$$

At time t, the bundle S was demanded, therefore $v(S) - \epsilon t p(S) \geq 0$. Thus, $\frac{v(S)}{p(S)} \geq \epsilon t$. Assume that Inequality E.1 does not hold, then it follows that $\frac{v(T)}{p(T)} - \epsilon > \epsilon t$, or $v(T) > \epsilon(t+1)p(T)$. But in time $t+1$ no bundle achieved a positive utility since the empty set was demanded. Contradiction. □

Proof of Proposition 6.3

Proof. The algorithm by [93] arbitrarily orders the items and allocates each item in turn to the bidder with the highest marginal valuation for it (given the items already allocated to him). The descending auction decreases the price of each item i, until a bidder demands it together with the bundle S he already owns. At this stage, up to the ϵ used, his utility from the bundle $S \cup \{i\}$ is zero, thus his value for this bundle equals its current price, i.e. $v(S \cup \{i\}) = \sum_{j \in S \cup \{i\}} p_j$. Similarly, this bidder valuates S by $v(S) = \sum_{j \in S} p_j$. Thus, $p_i = v(S \cup \{i\}) - v(S)$ is, by definition, the marginal valuation of this bidder for item i. By decreasing the price of item i, we exactly find the bidder with the highest marginal valuation for it. □

Appendix F

On the Power of Iterative Auctions: Demand Queries

F.1 Simulating Queries by Demand Queries

Proof of Lemma 7.1:

Proof. By definition, value queries can simulate marginal-value queries: $v(j|S) = v(S \cup \{j\}) - v(S)$. The simulation of a value query S by $|S| \leq m$ marginal-value queries is given by the equation $v(S) = \sum_{j \in S} v(j|\{j' \in S|j' < j\})$. □

Proof of Lemma 7.2:

Proof. We will show that demand queries can simulate any marginal-value query $v(j|S)$ using t queries, and then invoke the previous lemma. Set the prices of all the items in S to zero, and the prices of all other items (except j) to ∞. Then, we perform a binary search on p_j to find its lowest value for which the bidder demands $v(S)$. It is straightforward to see that this price is indeed the marginal value of item j: at this price, the utilities from the bundles S and $S \cup \{j\}$ are equal, thus $v(S) - 0 = v(S \cup \{j\}) - p_j$ and the claim follows.

A binary search makes t demand queries, and m marginal-value queries are needed to simulate a single value query thus $v(S)$ can be simulated by mt demand queries. □

Proof of Lemma 7.4:

Proof. An indirect-utility query with prices $\overrightarrow{p}$ can be answered by first querying for the demand D under these prices and then simulating the value query $v(D)$.

The following algorithm uses $m+1$ indirect-utility queries to simulate a demand query with some price vector $\overrightarrow{p}$:
Initialization: start with the price vector $\overrightarrow{p}$ for which the player answers some utility x.
Repeat: for every item $i = 1, ..., m$, raise the price of item i by some $\epsilon \in (0, 1)$. If the answer to the indirect-utility query now is other than x, we decrease its price back by ϵ in all future queries. If the answer was x, we use the new price for i in all future queries.
Finally: After all the $m+1$ indirect-utility queries are done, return the bundle of all items for which the answer was changed when we increased their prices.

In the algorithm above, if we raised the price of some item i, and the reported maximal-utility did not change, then there would clearly be utility-maximizing bundles that do not contain i, thus we can ignore this item. If the maximal-utility changed, then any utility-maximizing bundle under the current prices clearly contains i, thus we include it in our answer. Leaving the price of item i (of the first kind) at $p_i + \epsilon$, ensures that any bundle that contains it will not be output (but we are guaranteed to have other utility-maximizing bundles). □

Proof of Lemma 7.5:

Proof. For any $\epsilon > 0$, we simulate $RD(\overrightarrow{p})$ by the following binary search (up to an ϵ, see below):
Initialization: start with a price vector $c\overrightarrow{p}$ $(c > 0)$.
Binary search: find with a binary search the value $c^* \in \Re^+$ for which the bidder has a non-empty demand for the price vector $c^* \cdot \overrightarrow{p}$ and the bidder demands the empty set for $(c^* + \epsilon) \cdot \overrightarrow{p}$.
Finally: return the bundle S demanded under the price vector $c^* \cdot \overrightarrow{p}$.

Now we show that for the price vector $\overrightarrow{p}$, every other bundle T has a smaller weight than S (up to ϵ), i.e.,

$$\frac{v(S)}{p(S)} \geq \frac{v(T)}{p(T)} - \epsilon. \tag{F.1}$$

Denote $c^* = \epsilon t$ for some $t \in \Re^+$. The bundle S was demanded under the prices $\epsilon t \cdot \overrightarrow{p}$, therefore $v(S) - \epsilon t p(S) \geq 0$. Thus, $\frac{v(S)}{p(S)} \geq \epsilon t$. Assume that Inequality F.1 does not hold, then it follows that $\frac{v(T)}{p(T)} - \epsilon > \epsilon t$, or $v(T) > \epsilon(t+1)p(T)$. But for the price vector $(c^* + \epsilon)\overrightarrow{p} = \epsilon(t+1)\overrightarrow{p}$ no bundle achieved a positive utility. Contradiction. □

F.2 Missing Proofs

Proof of Lemma 7.6

Proof. Consider the following family of valuations: for every S, such that $|S| > m/2$, $v(S) = 1$, and there exists a single set T, such that for $|S| \leq m/2$, $v(S) = 1$ iff $T \subseteq S$ and $v(S) = 0$ otherwise. Now look at the behavior of the protocol when all valuations v_i have $T = \{1...m\}$. Clearly in this case the value of the best allocation is 1 since no set of size $\frac{m}{2}$ or lower has non-zero value for any player. Fix the sequence of queries and answers received on this k-tuple of valuations.

Now consider the k-tuple of valuations chosen at random as follows: a partition of the m items into k sets $T_1...T_k$ each of size $\frac{m}{k}$ each is chosen uniformly at random among all such partitions. Now consider the k-tuple of valuations from our family that correspond to this partition – clearly T_i can be allocated to i, for each i, getting a total value of k. Now look at the protocol when running on these valuations and compare its behavior to the original case. Note that the answer to a query S to player i differs between the case of T_i and the original case of $T = \{1...m\}$ only if $|S| \leq \frac{m}{2}$ and $T_i \subseteq S$. Since T_i is distributed uniformly among all sets of size exactly $\frac{m}{k}$, we have that for any fixed query S to player i, where $|S| \leq \frac{m}{2}$,

$$Pr[T_i \subseteq S] \leq \left(\frac{|S|}{m}\right)^{|T_i|} \leq 2^{-\frac{m}{k}}$$

Using the union-bound, if the original sequence of queries was of length less than $2^{\frac{m}{k}}$, then with positive probability none of the queries in the sequence would receive a different answer than for the original input tuple. This is forbidden since the protocol must distinguish between this case and the original case – which cannot happen if all queries receive the same answer. Hence there must have been at least $2^{\frac{m}{k}}$ queries for the original tuple of valuations. □

Proof of Lemma 7.7

Proof. Consider the protocol running on the following two valuations: the first has $b = 0$ (i.e. is simply additive), and the second has $b = 1$ for the set S of all items. In this case the outcome must be to allocate all to the second bidder. Let $e_1 ... e_t$ be the queries made on this input, where each $e_i = E_i^1 \oplus E_i^2 \oplus ... \oplus E_i^{l_i}$. Now consider what happens when the first valuation is changed so that for some S of size exactly $m/2$, be get a bonus $b = 2$ – clearly the allocation must change so that this S is allocated to the first player – hence one of the queries $e_1 ... e_t$ must change its answer. We will see that the fact that this is true for every such S implies that $\sum_{i=1}^{t} l_i$ is exponential.

First note that if in e_i there exists some set of size $m/2 + 1$ that has price zero, then the answer will not change as this set will give a surplus of at least $3m/2 + 3$ as opposed to at most $3m/2 + 2$ that S gives. Let us focus at an e_i that does not have such a set. We build a boolean DNF formula from this expression as follows: the variable set will be $x_1 ... x_m$ – a variable for each item. Consider a term (atomic bid) $E_i^j = (B_i^j, p_i^j)$ in e_i. We call this term essential if there exists some bundle of size exactly $m/2+1$ whose price in e_i is exactly p_i^j. For every *essential* term (B_i^j, p_i^j) in e_i we build a conjunction of the variables in it (ignoring the price for this bundle). We then take the disjunction of all of these conjunctions. First notice that this DNF must accept all inputs with more than $m/2$ 1's in the input – since otherwise consider a set that is not accepted by this expression, and the value of this set in e_i must be zero.

Now notice that if an input with 1's in exactly the set S of size exactly $m/2$ is accepted by this formula, then the answer to query e_i will not change. The reason is that an accepted set S contains some essential bundle B_i^j, and thus its price in e_i would be at least p_i^j. However, since the bundle is essential, there exists some set of size $m/2+1$ that is priced at exactly p_i^j – this set would clearly be preferable to S – the only set whose value has changed. Since for *every* set S of size exactly $m/2$ the answer to one of the queries must change, at least one of the formulas constructed must reject the input with 1's exactly in S.

We now take the conjunction of all boolean expressions built for all i. This formula accepts all inputs with exactly $m/2 + 1$ 1's, and rejects all inputs with exactly $m/2$ 1's. Note that this formula is a conjunction of disjunctions of conjunctions of variables – a, so called, monotone depth 3 formula. Since it is a monotone formula, it computes the majority function. Its size is clearly bounded from above by the total length of all expressions e_i. We are now ready to invoke the well known lower bound by Hastad [72] that states that a depth 3 formula for majority must have size at least $2^{\Omega(\sqrt{m})}$. □

Bibliography

[1] amazon. Web Page: http://www.amazon.com.

[2] Kenneth Arrow. *Social Choice and Individual Values.* Wiley, 1951.

[3] Susan Athey. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–89, 2001.

[4] L. M. Ausubel and P. R. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1:1–42, 2002.

[5] Lawrence Ausubel. An efficient dynamic auction for heterogeneous commodities. *American Economic Review*, 96(3):602–629, 2006.

[6] Lawrence M. Ausubel and Peter C. Cramton. Demand reduction and inefficiency in multi-unit auctions, 2002. Working paper, University of Maryland.

[7] Lawrence M. Ausubel and Paul Milgrom. *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 1. The Lovely but Lonely Vickrey Auction.* MIT Press., 2006.

[8] B. Awerbuch, Y. Azar, and A. Meyerson. Reducing truth-telling online mechanisms to online optimization. In *The 35th ACM Symposium on Theory of Computing*, pages 503–510, 2003.

[9] Moshe Babaioff and Liad Blumrosen. Computationally-feasible auctions for convex bundles. In *7th. International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX). LNCS Vol. 3122.*, pages 27–38, 2004.

[10] Moshe Babaioff and Liad Blumrosen. Computationally-feasible auctions for convex bundles. *Games and Economic Behavior, to appear*, 2006.

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

[12] Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi unit combinatorial auctions. In *Proceedings of the 9th conference on Theoretical aspects of rationality and knowledge*, pages 72–87, 2003.

[13] Dirk Bergemann and Martin Pesendorfer. Information structures in optimal auctions. Technical Report 2991, C.E.P.R. Discussion Papers, October 2001.

[14] Alejandro Bertelsen. Substitutes valuations and m$^\natural$-concavity". M.Sc. Thesis, The Hebrew University of Jerusalem, 2005.

[15] H. Bhargava, J. Feng, , and D. Pennock. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS Journal on Computing, to appear.*, 2006.

[16] Martin Bichler, Andrew Davenport, Gail Hohner, and Jayant Kalagnanam. *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 22. Industrial Procurement Auctions.* MIT Press., 2006.

[17] S. Bikhchandani, S. Chatterji, R. Lavi, A. Mu'alem, N. Nisan, and A. Sen. Weak monotonicity characterizes deterministic dominant strategy implementation. *Forthcoming in Econometrica*, 2006.

[18] Sushil Bikhchandani and John W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of economic theory*, 74:385 – 413, 1997.

[19] Sushil Bikhchandani and John W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of economic theory*, 74:385 – 413, 1997.

[20] Avrim Blum, Jeffrey C. Jackson, Tuomas Sandholm, and Martin A. Zinkevich. Preference elicitation and query learning. In *The 16th Annual Conference on Learning Theory*, 2003.

[21] Avrim Blum, Jeffrey C. Jackson, Tuomas Sandholm, and Martin A. Zinkevich. Preference elicitation and query learning. *Journal of Machine Learning Research*, 5:649–667, 2004.

[22] Liad Blumrosen and Shahar Dobzinski. Welfare maximization in congestion games. manuscript, The Hebrew University, July 2005.

[23] Liad Blumrosen and Shahar Dobzinski. Welfare maximization in congestion games. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 52–61, 2006.

[24] Liad Blumrosen and Michal Feldman. Implementation with a bounded action space. Discussion paper no. 412, the center for the study of rationality, The Hebrew University, November 2005.

[25] Liad Blumrosen and Michal Feldman. Implementation with a bounded action space. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 62–71, 2006.

[26] Liad Blumrosen and Noam Nisan. Auctions with severely bounded communications. 43th Annual Symposium on Foundations of Computer Science (FOCS 2002), 2002.

[27] Liad Blumrosen and Noam Nisan. On the computational power of iterative auctions I: demand queries. Discussion paper no. 381, The Center for the Study of Rationality, The Hebrew University., 2005. An extended abstract in EC'05 contained preliminary results.

[28] Liad Blumrosen and Noam Nisan. On the computational power of iterative auctions II: ascending auctions. Discussion paper no. 382, The Center for the Study of Rationality, The Hebrew University., 2005. An extended abstract in EC'05 contained preliminary results.

[29] Liad Blumrosen and Noam Nisan. Informational limitations of ascending combinatorial auctions. Working paper, The Hebrew University., 2006.

[30] Liad Blumrosen, Noam Nisan, and Ilya Segal. Multi-player and multi-round auctions with severely bounded communication. 11th Annual European Symposium on Algorithms (ESA), 2003.

[31] Liad Blumrosen, Noam Nisan, and Ilya Segal. Auctions with severely bounded communications. *Journal of Artificial Intelligence Research, to appear.*, 2006.

[32] Paul J. Brewer and Charles R. Plott. A binary conflict ascending price (bicap) mechanism for the decentralized allocation of the right to use railroad tracks. *International Journal of Industrial Organization*, 14(6):857–886, 1996.

[33] Patrick Briest, Piotr Krysta, and Berthold Vocking. Approximation techniques for utilitarian mechanism design. In *the 37th ACM symposium on Theory of computing*, pages 39–48, 2005.

[34] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, pages 17–33, 1971.

[35] T.H. Cormen, C.L. Leiserson, R.L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press,MA, second edition edition, 2001.

[36] P. Cramton, L.M. Ausubel, and P.R. Milgrom. *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 5. The Clock-Proxy Auction: A Practical Combinatorial Auction Design*. MIT Press., 2006.

[37] P. Cramton, Y. Shoham, and R. Steinberg (Editors). *Combinatorial Auctions*. MIT Press., 2006.

[38] Peter Cramton. Ascending auctions. *European Economic Review*, 42(3–5):745–756, 1998.

[39] Peter Cramton. Comments on alternative policy options for managing capacity and mitigating congestion and delay at LaGuardia airport. U.S. Department of Transportation. FAA-2001-9852, FAA-2001-9854, 2002.

[40] Peter Cramton. *In Martin Cave, Sumit Majumdar, and Ingo Vogelsang, eds., Handbook of Telecommunications Economics. Chapter 14: Spectrum auctions*. Elsevier Science B.V., 2002.

[41] Peter Cramton. *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 5. Simultaneous Ascending Auctions*. MIT Press., 2006.

[42] P.S. Dasgupta, P.J. Hammond, and E.S. Maskin. The implementation of social choice rules: Some general results on incentive compatibility. *Review of Economic Studies*, 46:185 – 216, 1979.

[43] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 71–78, 2006.

[44] S. de Vries, J. Schummer, and R. V. Vohra. On ascending Vickrey auctions for heterogeneous objects. *Forthcoming in J. Econ. Theory*, 2006.

[45] G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94:863–872, 1986.

[46] Nikhil R. Devanur and Vijay V. Vazirani. The spending constraint model for market equilibrium: algorithmic, existence and uniqueness results. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 519–528, 2004.

[47] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *The 37th ACM symposium on theory of computing (STOC).*, pages 610–618, 2005.

[48] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In *STOC*, 2006.

[49] Shahar Dobzinski and Michael Schapira. Optimal upper and lower approximation bounds for k-duplicates combinatorial auctions. Working paper, the Hebrew University.

[50] Shahar Dobzinski and Michael Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1064–1073, 2006.

[51] James Dow. Search decisions with limited memory. *Review of Economic Studies*, 58(1):1–14, 1991.

[52] eBay. Web page:. http://www.ebay.com.

[53] A. S. Edlin and C. Shannon. Strict single crossing and the strict Spence-Mirrlees condition: a comment on monotone comparative statics. *Econometrica*, 68(6):1417–1425, 1998.

[54] Richard Engelbrecht-Wiggans and Charles M. Kahn. Multi-unit auctions with uniform prices. *Economic Theory*, 12:227 – 258, 1998.

[55] Alex Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. In *FOCS 04*, pages 604–612, 2004.

[56] FCC auctions home page. http://wireless.fcc.gov/auctions, 2006.

[57] Combinatorial bidding conference. See Conferences page in http://wireless.fcc.gov/auctions, 2003.

[58] Uriel Feige. On maximizing welfare where the utility functions are subadditive. In *38th ACM Symposium on Theory of Computing*, pages 41–50, 2006.

[59] Uriel Feige and Jan Vondrak. Approximation algorithms for allocation problems: Improving the factor of 1-1/e. In *47th Annual IEEE Symposium on Foundations of Computer Science*, 2006.

[60] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of muliticast transmissions. *Journal of Computer and System Sciences.*, 63:21–41, 2001.

[61] D. Fudenberg and J. Tirole. *Game Theory.* MIT Press, 1991.

[62] Riley J. G. and Samuelson W. F. Optimal auctions. *American Economic Review*, 71(3):381–392, 1981.

[63] Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41:587–601, 1973.

[64] A. V. Goldberg, J. D. Hartline, A. R. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242 – 269, 2006.

[65] Andrew V. Goldberg, Jason D. Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Symposium on Discrete Algorithms*, pages 735–744, 2001.

[66] J. Green and J.J. Laffont. Characterization of satisfactory mechanism for the revelation of preferences for public goods. *Econometrica*, pages 427–438, 1977.

[67] T. Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.

[68] Hongwei Gui, Rudolf Muller, and Rakesh V. Vohra. Dominant strategy mechanisms with multidimensional types. Technical report, Northwestern University, 2004.

[69] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95 – 124, 1999.

[70] Faruk Gul and Ennio Stacchetti. The english auction with differentiated commodities. *Journal of Economic Theory*, 92(3):66 – 95, 2000.

[71] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and David C. Parkes. Adaptive limited-supply online auctions. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 71–80, 2004.

[72] J. Hastad. Almost optimal lower bounds for small depth circuits. In *18th STOC*, pages 6–20, 1986.

[73] Johan Hastad. Clique is hard to approximate to within $n^{1-\epsilon}$. *Acta Mathematica*, 182, 1999.

[74] Benjamin E. Hermalin. Lecture notes in economics, 2005. University of California, Berkeley.

[75] Dorit S. Hochbaum. *Approximation Algorithms for NP-Hard Problems*. PWS publishing company, 1997.

[76] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatrial auctions. *Games and Economic Behavior*, 47:104–123, 2004.

[77] L. Hurwicz and H. Uzawa. *In J. Chipman, L. Hurwicz and H. Sonnenschein (editors). Preferences, Utility and Demand. Chapter 6: On the integrability of demand functions.* Harcourt Brace, Jovanovich, 1971.

[78] Sheena S. Iyengar and Mark R. Lepper. When choice is demotivating : Can one desire too much of a good thing? *Journal of personality and social psychology*, 79(6):995–1006, 2000.

[79] Matthew O. Jackson. A crash course in implementation theory. *Social Choice and Welfare*, 18(4):655 – 708, 2003.

[80] H. Karloff. *Linear Programming.* Birkhäuser Verlag, 1991.

[81] Frank Kelly and Richard Steinberg. A combinatorial auction with multiple winners for universal service. *Management Science*, 46:586–596, 2000.

[82] A.S. Kelso and V.P. Crawford. Job matching, coalition formation, and gross substitute. *Econometrica*, 50:1483–1504, 1982.

[83] A. Kress and C. Boutilier. A study of limited precision, incremental elicitation in auctions. In *AAMAS 04*.

[84] V. Krishna. *Auction theory.* Academic press, 2002.

[85] Eyal Kushilevitz and Noam Nisan. *Communication Complexity.* Cambridge University Press, 1997.

[86] A. Kwasnica, J. Ledyard, D. Porter, and C. DeMartini. A new and improved design for multi-object iterative auctions. *Management Science*, 51(3):419 – 434, 2005.

[87] Sebastien Lahaie and David C. Parkes. Applying learning algorithms to preference elicitation. In *The 5th ACM Conference on Electronic Commerce*, pages 180–188, 2004.

[88] Kate Larson and Tuomas Sandholm. Costly valuation computation in auctions. In *Proceedings of the 8th conference on Theoretical aspects of rationality and knowledge*, pages 169–182, 2001.

[89] Ron Lavi, Ahuva Mua'lem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *FOCS 03*, pages 574–583, 2003.

[90] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. *Theoretical Computer Science. A preliminary version appeared in ACM Conference on Electronic Commerce, 2000.*, 310(1):159–180, 2004.

[91] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In *46th Annual IEEE Symposium on Foundations of Computer Science*, pages 595–604, 2005.

[92] John O. Ledyard, Mark Olson, David Porter, Joseph A. Swanson, and David P. Torma. The first use of a combined-value auction for transportation services. *Interfaces*, 32(5):4 – 12, 2002.

[93] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behaviour. (A preliminary version appeared in EC'01)*, 55(2):270–296, 2006.

[94] Daniel Lehmann, Rudolf Muller, and Tuomas Sandholm. *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 12. The Winner Determination Problem.* MIT Press., 2006.

[95] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. In *JACM 49(5)*, pages 577–602, Sept. 2002.

[96] H. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91:1–36, 1983.

[97] Harstad R. M. and Rothkopf M. H. On the role of discrete bid levels in oral auctions. *European Journal of Operations Research*, pages 572–581, 1994.

[98] A. Mas-Collel, W. Whinston, and J. Green. *Microeconomic Theory*. Oxford university press, 1995.

[99] P. McAfee. Coarse matching. *Econometrica*, 70(5):2025–2034, 2002.

[100] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized on-line matching. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 264–273, 2005.

[101] P. Milgrom and C. Shannon. Monotone comparative statics. *Econometrica*, 62:157–180, 1994.

[102] Paul Milgrom. Putting auction theory to work: the simultaneous ascending auction. *Journal of Political Economy*, 108(2):245 – 272, 2000.

[103] Paul Milgrom. *Putting auction theory to work*. Cambridge University Press, 2004.

[104] Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982.

[105] J. Mirrlees. An exploration in the theory of optimum income taxation. *Review of Economic Studies*, 38:175–208, 1971.

[106] Dilip Mookherjee and Stefan Reichelstein. Dominant strategy implementation of bayesian incentive compatible allocation rules. *Journal of Economic Theory*, 56(2):378–399, 1992.

[107] Rudolf Muller. *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 13. Tractable Cases of the Winner Determination Problem*. MIT Press., 2006.

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

[109] R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265 – 281, 1983.

[110] John F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.

[111] Noam Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of the 2nd ACM conference on Electronic commerce*, pages 1–12, 2000.

[112] Noam Nisan. The communication complexity of approximate set packing and covering. In *Proceedings of the 29th International Colloquium on Automata, Languages and Programming*, pages 868–875, 2002.

[113] Noam Nisan. *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 1. Bidding Languages*. MIT Press., 2006.

[114] Noam Nisan. *In Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay Vazirani (Editors), Algorithmic Game Theory. Chapter 10. Combinatorial Auctions.* Cambridge University Press. To appear., 2007.

[115] Noam Nisan. *In Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay Vazirani (Editors), Algorithmic Game Theory. Chapter 9. Introduction to Mechanism Design (for Computer Scientists).* Cambridge University Press. To appear., 2007.

[116] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 35:166 – 196, 2001. A preliminary version appeared in STOC 1999.

[117] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192 – 224, 2006.

[118] M. J. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT press, 1994.

[119] M. Ostrovsky, B. Edelman, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Technical report, Stanford Univeristy, 2005.

[120] Christos H. Papadimitriou. *Computational Complexity.* Addison-Wesley Publishing, 1994.

[121] D. C. Parkes and L. H. Ungar. An ascending-price generalized Vickrey auction. Presented at the Stanford Institute for Theoretical Economics (SITE) Summer Workshop, June 2002.

[122] David Parkes. *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 3. Iterative Combinatorial Auctions.* MIT Press., 2006.

[123] David C. Parkes. Iterative combinatorial auctions: Achieving economic and computational efficiency. Ph.D. Thesis, Department of Computer and Information Science, University of Pennsylvania., 2001.

[124] David C. Parkes and Lyle H. Ungar. Iterative combinatorial auctions: Theory and practice. In *AAAI/IAAI*, pages 74–81, 2000.

[125] David Porter, Stephen Rassenti, Anil Roopnarine, and Vernon Smith. Combinatorial auction design. *Proceedings of the National Academy of Sciences*, 100:11153–11157, 2003.

[126] S.J. Rassenti, V.L. Smith, and R.L. Bulfin. A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics*, 13(2):402–417, 1982.

[127] Kevin Roberts. The characterization of implementable choise rules. *In Jean-Jacques Laffont, editor, Aggregation and Revelation of Preferences.*, pages 321–349, 1979.

[128] A. E. Roth and M. Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis.* Cambridge University Press, 1992.

[129] Alvin E. Roth and Axel Ockenfels. Late-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon on the internet. *American Economic Review*, 92(4):1093–1103, 2002.

[130] Michael H Rothkopf, Thomas J Teisberg, and Edward P Kahn. Why are Vickrey auctions rare? *Journal of Political Economy*, 98(1):94–109, 1990.

[131] Michael Rothschild and Joseph E Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4):630–49, 1976.

[132] Tim Roughgarden. *Selfish Routing and the Price of Anarchy.* MIT Press, 2005.

[133] Ariel Rubinstein. Why are certain properties of binary relations relatively more common in natural languages. *Econometrica*, 64:343–356, 1996.

[134] Michael Saks and Lan Yu. Weak monotonicity suffices for truthfulness on convex domains. In *Proceedings of the 6th ACM conference on Electronic commerce*, 2005.

[135] T. Sandholm. Expressive commerce and its application to sourcing. In *Innovative Applications of Artificial Intelligence (IAAI)*, 2006.

[136] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. In *Artificial Intelligence. Preliminary version appeared in IJCAI-99*, volume 135, pages 1–54, 2002.

[137] Tuomas Sandholm. *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 14. Optimal Winner Determination Algorithms.* MIT Press., 2006.

[138] Tuomas Sandholm and Andrew Gilpin. Sequences of take-it-or-leave-it offers: Near-optimal auctions without full valuation revelation. In *Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, 2006.

[139] P. Santi, V. Conitzer, and T. Sandholm. Towards a characterization of polynomial preference elicitation with value queries in combinatorial auctions. In *The 17th Annual Conference on Learning Theory*, 2004.

[140] Mark Allen Satterthwaite. Strategy-proofness and arrow's condition: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, pages 187–217, 1975.

[141] Herbert Scarf. Some examples of global instability of the competitive equilibrium. *International Economic Review*, 1(3):157 – 172, 1960.

[142] Ilya Segal. Optimal pricing mechanisms with unknown demand. *American Economic Review*, 93(3):509–529, 2003.

[143] Ilya Segal. The communication requirements of social choice rules and supporting budget sets. *Journal of Economic Theory, to appear.*, 2004.

[144] Ilya Segal and Ronald Fadel. The communication cost of selfishness, 2006.

[145] Yossi Sheffi. Combinatorial auctions in the procurement of transportation services. *Interfaces*, 34(4):245 – 252, 2004.

[146] Michael Spence. Job market signalin. *The Quarterly Journal of Economics,*, 87(3):355 – 374, 1973.

[147] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.

[148] R. Wilson. Efficient and competitive rationing. *Econometrica*, 57:1–40, 1989.

[149] P.R. Wurman and M.P. Wellman. Akba: A progressive, anonymous-price combinatorial auction. In *Second ACM Conference on Electronic Commerce*, pages 21 – 29, 2000.

[150] Martin A. Zinkevich, Avrim Blum, and Tuomas Sandholm. On polynomial-time preference elicitation with value queries. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 176–185, 2003.

[151] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 681–690, 2006.