Optimization in the Private Value Model: Competitive Analysis Applied to Auction Design

Jason D. Hartline

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Washington

2003

Program Authorized to Offer Degree: Computer Science and Engineering

University of Washington

Abstract

Optimization in the Private Value Model:
Competitive Analysis Applied to Auction Design

by Jason D. Hartline

Chair of Supervisory Committee:

Professor Anna R. Karlin
Computer Science and Engineering

We consider the study of a class of optimization problems with applications towards profit maximization. One feature of the classical treatment of optimization problems is that the space over which the optimization is being performed, i.e., the input description of the problem, is assumed to be publicly known to the optimizer. This assumption does not always accurately represent the situation in practical applications. Recently, with the advent of the Internet as one of the most important arenas for resource sharing between parties with diverse and selfish interests, this distinction has become more readily apparent. The inputs to many optimizations being performed are not publicly known in advance. Instead they must be solicited from companies, computerized agents, individuals, etc. that may act selfishly to promote their own self-interests. In particular, they may lie about their values or may not adhere to specified protocols if it benefits them.

An auction is one of the simplest applications where the classical (a.k.a. public value) optimization approach fails to work as expected in the presence of selfish agents with private data. We consider casting profit optimization problems into the game theoretic framework of mechanism design and consider the design of auction mechanisms to maximize the profit of the auctioneer. We show how a competitive analysis can be used to gauge the performance of profit maximizing mechanisms. We develop a number of techniques for designing auctions and show how they can be extend to more complex profit maximization problems.

TABLE OF CONTENTS

List of Figures **iv**

I The Basic Auction Problem **1**

Chapter 1: Introduction **2**

1.1 Problems Studied 5

1.2 Related Work 7

1.3 Contributions 9

Chapter 2: Definitions **12**

2.1 The Basic Auction Problem 12

2.2 The Game Theoretic Model and Truthful Mechanism Design 13

2.3 The Vickrey Auction 15

2.4 Bid-Independent Auctions 16

2.5 More Truthful Auctions 19

2.6 Profit Extraction and the Basic Auction Decision Problem 20

Chapter 3: Analysis Framework **23**

3.1 Competitive Auction Framework 24

3.2 Mass Markets 27

3.3 Concentration 28

Chapter 4: Competitive Auctions **30**

4.1 Existing Methods 31

4.2 The Bid-independent Optimal Price Auction 33

4.3 Deterministic Impossibility . . . 34
4.4 Random Sampling Auctions . . . 36
4.5 Randomized Consensus of Revenue Estimates . . . 43
4.6 Randomization via Weighted Pairing . . . 50
4.7 Analysis of RSOP . . . 54

Chapter 5: Limited Supply **59**

Chapter 6: Envy-free **61**

6.1 Definitions . . . 62
6.2 Lower Bound . . . 63
6.3 Envy-free with High Probability . . . 66
6.4 Truthful with high probability . . . 67

Chapter 7: On the Competitive Ratio **71**

7.1 Asymmetry . . . 71
7.2 Auctions for two items . . . 72
7.3 Auctions for three items . . . 72
7.4 Worst Case Verse Mass Market Analysis . . . 74
7.5 A Lower Bound on the Competitive Ratio . . . 76

Chapter 8: Monotonicity **81**

8.1 Hard-coded Auctions . . . 82
8.2 Monotonicity . . . 83
8.3 Analysis of Auctions . . . 84
8.4 Non-monotone variant of $\mathrm{BI}_{\mathrm{opt}}$ . . . 88

II Extensions **89**

Chapter 9: The Multi-item Auction Problem **90**

9.1 Definitions 91
9.2 Fixed Price Auction and Optimal Prices 92
9.3 The Random Sampling Auction 93
9.4 Concluding Remarks 98

Chapter 10: Auction Problems with Production Costs and Market Segmentations 100

10.1 Introduction 100
10.2 Competitive Framework 104
10.3 The Fixed Cost Basic Auction 105
10.4 Applications of Cancellable Auctions 109
10.5 An Upper Bound on the Profit of Truthful Mechanisms 111
10.6 Conclusions 115

Chapter 11: The Double Auction Problem 116

11.1 Introduction 116
11.2 Preliminaries 117
11.3 Reducing Competitive Double Auctions to Competitive Basic Auctions 120
11.4 The General Reduction to the Private Value Decision Problem 123
11.5 Consensus and Revenue Estimate for Double Auctions 127
11.6 Monotonicity 133
11.7 Conclusions 135

Chapter 12: Conclusions and Future Research Directions 137

Bibliography 141

LIST OF FIGURES

9.1 Two item auction with regions R_1 and R_2 95
9.2 Signatures in a two item auction 96

ACKNOWLEDGMENTS

I will forever be in debt to my thesis advisor, Anna Karlin. Throughout all aspects of my time as a graduate student, she is without a doubt the best possible advisor I could have had. I would also single out Andrew Goldberg for introducing me to the fascinating field that is the subject of this dissertation and for years of fruitful collaboration, the results of which appear here.

Part I

The Basic Auction Problem

Chapter 1

INTRODUCTION

In classic algorithm design, the standard unspoken assumption is that the algorithm, before processing, is actually given its input. This assumption does not hold for many practical applications of algorithms. Consider, as an example, the following problem. A real estate agency is selling houses to a set of home buyers. Each house can only be sold to one home buyer and each home buyer only wishes to purchase one house. In general, the houses may be different and thus each home buyer may have a different valuation for the worth of each house. The real estate agency has the problem of selling the houses to the home buyers so as to maximize the agency's profit.

An instance of this problem can be modeled as a weighted bipartite graph with home buyers on one side and houses on the other. The weight on the edge between a home buyer and a house corresponds to the valuation that that home buyer has for the worth of the house. The classical algorithmic approach gives us the natural solution of the real estate agency's problem as that of computing the maximum weighted matching in the aforementioned weighted bipartite graph. A fundamental flaw in this solution is the assumption that the real estate agency can obtain the home buyers' true valuations so as to be able to run the maximum weighted matching algorithm on the correct input. As a trivial example, if there is one home buyer and one house a natural strategy for the home buyer is to report a valuation of one cent for the house. The maximum weighted matching, of course, allocates the house to this buyer at a price of one cent. In general, understanding how the buyers should report their values requires a game theoretic analysis.

In this thesis, we consider a game theoretic approach to algorithm design whereby we look for algorithms for optimization problems such as the one above that have provably good performance even in the presence of potential gaming of the participants. Recently, with the advent of the Internet as one of the most important arenas for resource sharing between parties with diverse and selfish interests and as a platform for electronic commerce, the need for a game theoretic treatment of algorithms has become more readily apparent. The inputs to many optimizations being performed are not publicly known in advance. Instead they must be solicited from companies, computerized agents, individuals, etc. that may act selfishly to promote their own self-interests. In particular, they may lie about their values or may not adhere to specified protocols if it benefits them.

We will be referring to the classical algorithmic case where the input is readily available to the algorithm as the *public value model* and the case where the algorithms input must be obtained from selfish agents as the *private value model*. In general, it is possible to recast many traditional public value optimization problems into the private value model. Our goal is to get an understanding for optimization in the private value model and develop new design and analysis techniques for this type of optimization problem.

In general the optimal strategy of a player in a game depends on the game being played and the strategies of the other players involved. To prove bounds on the performance of the mechanism (game/algorithm) it is necessary to have a model for the strategies of the agents (players). A popular approach to this problem is to design mechanisms that are *truthful*. In a truthful mechanism there is no incentive for any agent to falsely report their true valuations. In a truthful mechanism it is assumed that all agents do in fact play by their optimal strategy of reporting their true valuation. It is easy to see that the maximum bipartite matching problem does not result in a truthful mechanism for the real estate problem.

We assume that each agent has a utility function that represents the agent's value for any outcome of the mechanism in monetary units. Reasoning about how agents determine their utility function is a very interesting area of research that is orthogonal to that of this thesis. For our purposes, we will assume that the agents each know their utility function. An easy way to think about this is to consider the example where the agent is attempting to buy a house; however, they intend on immediately reselling it for some amount, e.g., a million dollars. If the outcome of the house selling mechanism is to allocate the house to this agent then their utility is a million dollars. Otherwise, their utility for the outcome is zero. Typically, if the selling mechanism allocates an item to an agent it also requires that the agent make some payment. We assume that the agents goal is to maximize their profit, the difference between their utility of the outcome and the payment they are required to make.

An interesting class of mechanisms is one where the mechanism also has a utility function over possible outcomes and it is attempting to maximize its own profit, the sum of its utility function for the chosen outcome and the payments that it receives from the agents. This is the class of problems we will focus on in this thesis.

Prior to this work, profit maximization in mechanism design was considered in a Bayesian framework. In such a framework, it is assumed that the agents' valuations are drawn from some probability distribution. It is then assumed that the mechanism designer has knowledge of this prior distribution. The goal is to design the *Bayesian optimal mechanism* for the given prior distribution. The obvious drawback of this approach is that the mechanism designer must know this prior distribution. This is exactly the problem we were hoping to solve: that of performing well without knowing what the input is in advance. In particular our goal is a single mechanism that works well for all inputs and does not have to be tuned to reflect changes in the preferences of the agents.

Our mechanism design problem is the problem of designing the truthful mechanism that maximizes its profit for the case where nothing is known about the agents private values in advance. This is made difficult by the fact that for any non-trivial profit maximization problem there is no truthful mechanism that obtains the highest possible profit on all inputs. In particular, for any truthful mechanism, $\mathcal{M}$, there exists an input, $\mathbf{I}$, and alternative mechanism, $\mathcal{M}'$, such that $\mathcal{M}'$ on $\mathbf{I}$ obtains a higher profit than that of $\mathcal{M}$ on $\mathbf{I}$. This type of problem usually arises when an obstacle prevents a perfect optimization. The obstacle we are faced with is that we do not know the true values of the agents in advance of running the mechanism.

The fields of *approximation algorithms* and *online algorithms* overcome similar obstacles. Approximation algorithms are useful for obtaining polynomial time algorithms for problems that are impossible to solve exactly in polynomial time. The performance of such an algorithm is gauged by measuring either the additive or multiplicative approximation factor that the algorithm obtains in comparison to that of the true optimal on the worst possible input. An online algorithm is one that responds to input as it is received and must make performance affecting decisions without knowledge of what future inputs it might receive. In online algorithms the obstacle the algorithm faces is that of not knowing the future inputs. To gauge performance of an online algorithm, its performance is compared to the *optimal offline* algorithm's performance, i.e., the performance of an algorithm that is endowed with knowledge of all the inputs in advance. In the field of online algorithms, this technique is known as *competitive analysis*. The algorithm design goal in a competitive analysis framework is to obtain the algorithm with the best *competitive ratio*.

We adopt this general approach as our means for evaluating profit maximizing mechanisms for private value problems. In this analysis we are interested in finding the mechanism that obtains a profit within the best multiplicative ratio of an "optimal" mechanism that is endowed with perfect knowledge of the private values of the agents. We refer to such a mechanism interchangeably as the *optimal omniscient mechanism* or the *optimal public value mechanism*.

1.1 Problems Studied

The approach we will take to understand optimization in the private value model will be to consider a simple mechanism design problem and obtain results that we can then extend to more complex problems. The simple mechanism design problem we will consider throughout the majority of this thesis is the *basic auction problem*. The basic auction problem is a special case of the aforementioned house selling problem where all houses are identical and there is an *unlimited supply* of these houses. By unlimited supply, we mean that there are as many houses as there are buyers. Such a scenario may be relevant for the sale of a digital good as all copies are identical to the original, the seller can potentially make a copy for each consumer, and each consumer only has use for one copy. We will show later that the limited supply version of the problem, when there are only a fixed number of identical items for sale, reduces to the unlimited supply version. As such, this is perhaps the simplest problem where the fact that the agents' valuations are private is an important obstacle for profit maximization.

One generalization of the basic auction problem is to the case where there are multiple items in unlimited supply available for sale, yet each consumer desires at most one of these items. This *multi-item auction problem* is essentially, the unlimited supply variant of the real estate problem discussed previously. Like in the real estate example, bidders may have different valuations for the different items. As an example application, consider an airplane flight where passengers have individual movie screens and can choose to view one out of a dozen movies that are broadcast simultaneously. The flight is only long enough for one movie to be seen. The airline wants to price movies to maximize its profit.

Another generalization of the basic auction problem is the *multicast pricing problem* of Feigenbaum et al. [18]. In this problem a content provider is attempting to sell and distribute a digital good via multicast over a computer network such as the Internet. It is assumed that the consumers are located at various places in the network and that there is an implicit cost to the content provider for sending the good across a network link. The link costs are assumed to be publicly available. The content provider collects payments from consumers selected to receive the good but must pay for each link in a tree connecting themselves to all of these consumers.[1] We further generalize the multicast pricing problem in two ways. First we assume that the bidders are divided into markets. The auction mechanism may distinguish bidders in different markets and differential pricing based on the markets segmentation is possible. Second we assume that there is a cost to the seller that is a function of to which markets the good is supplied. In the multicast pricing problem the markets correspond to nodes in the multicast tree. Of course the basic auction is a special case of this problem where there is only a single market containing all bidders and the cost function is identical zero.

Finally, we consider the generalization where the auctioneer is actually a broker mediating the exchange of items between buyers and sellers. In this exchange the auctioneer can pocket the the difference between what is paid by the buyers and what is paid to the sellers. In this *double auction problem* the goal is to design the auction mechanism to maximize the auctioneer's profit. We assume that each buyer only desires one item and that each seller

[1] this is a private value version of a *prize collecting Steiner tree problem* [28]. The prizes at nodes in the Steiner tree problem correspond to private values of bidders in the multicast problem.

only has one item to sell. Naturally, the basic auction is a special case of the double auction problem in which all sellers are willing to sell their items for price zero.

1.2 Related Work

There is an extensive amount of Economics literature on auctions. For a survey, see [30].

Unlike the approach taken in this thesis, the standard approach to profit maximization is to look for the *Bayesian optimal mechanism.* This approach assumes that the private values come from some known probability distribution. The goal then is to design the mechanism that obtains the highest expected profit on said input distribution. For i.i.d. *prior distribution*'s the Bayesian optimal auction is given in [38, 11, 42]. We will discuss these results more in Chapter 2. For arbitrarily correlated prior distributions, Ronen considers the computational issues of implementing the Bayesian optimal auction for selling a single item. Ronen shows that it is computational infeasible (i.e., NP-hard) to implement the optimal auction but instead an approximately optimal solution can be implemented [43]. In [34], the Bayesian approach to the optimal auction was extended to the multicast pricing problem that we discuss in Chapter 10.

Recently, Segal has considered the uninformed auction problem in a setting similar to the Bayesian setting [48]. He shows that auction mechanisms related to the ones presented later in this thesis are asymptotically optimal when the bidders valuations are drawn independently from an unknown or partially unknown prior distribution.

A problem related to the basic auction is the *online auction problem.* In the online auction problem, it is assumed that bidders arrive one at a time and that the auction must compute a price for them before seeing the next bidder's bid. The analysis framework that we have developed in this thesis was originally extended to this problem by Bar-Yossef et al. [9]. An improved solution to the online auction problem, based on online learning, was given by Blum et al. [10]. This improvement also includes analysis of the *online posted price problem* where the mechanism must offer each consumer a price and the consumer may choose to accept or reject the offered price. The mechanism does not learn anything about the bidder's value except whether it was above or below the price the mechanism offered.

This problem was further studied by Klienberg and Leighton [29]; however, additive loss is considered instead of a multiplicative competitive ratio. Finally, Awerbuch et al. consider a general approach for adding incentive properties to online algorithms. They show how to convert an online algorithm into a truthful online mechanism with only logarithmic additive loss to the competitive ratio [6].

There has been substantial recent interest private value problems in Computer Science. This work has been primarily focused in two directions, (a) looking at the computational complexity and communication complexity of implementing traditional economic objectives, and (b) looking at designing mechanisms for objectives not previously considered by economists.

In the first direction, the most notable example is work looking at the intractability and approximation of the *common welfare maximization* objective for the *combinatorial auction* problem (See, e.g., [40, 32, 2]). The fundamental observation made is that the Vickrey-Clarke-Groves (VCG) [51, 13, 27] mechanism that implements common welfare maximization, when applied to the combinatorial auction problem, is required to solve an NP-hard optimization problem. Further, use of an approximately optimal solution in place of the true optimal solution causes the VCG mechanism to lose important properties [40]. In a network design setting, Feigenbaum et al.'s consider the communication complexity of the distributed implementation of the multicast pricing problem both with the objective of *budget-balanced*, i.e., zero profit and zero deficit, and the objective of common welfare maximization. Finally, looking beyond these traditional objectives, [39, 3] consider the private value scheduling problem where the objective is to minimize the *makespan* (the completion time of the last job).

One interesting observation about the objective of common welfare maximization is that the profit or loss of the mechanism is not at all optimized. As an example, the basic auction (n items and n bidders) with the goal of common welfare maximization would result in a mechanism that sells an item to each bidder at price zero. Further, in the multicast pricing problem with the goal of common welfare maximization, the mechanism profit is always non-positive. Archer and Tardos have studied the private value *shortest path* problem where the goal is to design a mechanism to buy the cheapest path [4]. They note that for this problem the VCG mechanism, the mechanism that maximizes the common welfare, may end up spending a lot more than the cost of the cheapest path. It turns out that all mechanisms for the private value shortest path share this problem [4, 17]. In contrast, for matriod optimization problems, e.g., minimum spanning tree, the VCG mechanism does find a solution at close to the cost of the optimal solution [50]. We note that using the VCG mechanism for the shortest path and minimum spanning tree problems was first proposed by Nisan and Ronen [39].

Other notable work in the intersection of Computer Science, Game Theory, and Economics includes the study of the *price of anarchy* [31] and *market equilibria* [14]. The price of anarchy considers the difference between the social welfare of an optimal allocation of resources and the social welfare of a allocation obtained when a Nash equilibrium is found by selfish agents competing for the resources. Along these lines, Roughgardin et al. have studied the problem of routing traffic in a communication network [46, 45]. The market equilibrium question looks at taking a market with agents and goods and computing prices for the goods such that they can be reallocated among the agents in an optimal way such that there is neither a surplus or a shortage of any good [14, 16].

1.3 Contributions

The work presented in this thesis is the result of collaborations between the author and Yossi Azar, Kostub Deshmukh, Amos Fiat, Andrew Goldberg, Anna Karlin, Claire Kenyon, Mike Saks, and Andrew Wright. Many of the results presented have been published with subsets of the above as coauthors [26, 22, 20, 15, 23, 24] and some are in manuscript form as of this writing [25, 8].

One main contribution of this work is the introduction of competitive analysis to economic mechanisms [26, 20, 24]. This type of analysis allows for design and analysis of profit maximizing mechanisms that have no prior knowledge about the nature of their input. Before this work, profit maximizing mechanisms were only studied in a Bayesian setting where the mechanism was assumed to be endowed with knowledge of the distribution from which the agents valuations are drawn.

This work can also be viewed as introducing game theoretic considerations into standard profit optimization problems studied in Computer Science. Prior work in Computer Science for profit maximization assumes the public value model. As already argued, this model is not appropriate for many interesting and potentially practical optimization problems. We show how techniques from game theory, mechanism design, and economics can be combined with a standard computer science approach to design and analysis of algorithms.

For the basic auction problem, we develop a number of auctions with worst case performance guarantees [26, 20] the best of which achieves a worst case competitive ratio of 3.39 [23]. Further, we show that no mechanism can achieve a competitive ratio better than 2.42 [25]. This worst case arises on inputs where the optimal auction only sells a small number of items. Intuitively, if the sale of some item gives a large portion of the profit of the mechanism then it is much harder to perform well in worst case.[2] There is less margin for error and if a mistake is made for one bidder it could affect the auction performance by a large amount. If, on the other hand, we are willing to assume that the number of items sold by an optimal auction is large, i.e., each item's sale only contributes a small amount to the profit of the mechanism, then it is possible to develop auctions that are $(1+\epsilon)$-competitive with an optimal omniscient mechanism [25, 24].

An unfortunate property of the competitive auctions we develop is that they all fail the following natural fairness criteria: the outcome of the basic auction should be a single sale price such that all bidders bidding above this price are allocated the good. In other words, these auctions are not *envy-free* in the sense that some bidder might prefer the outcome of another bidder. We show that there is no basic auction that simultaneously is envy-free and truthful and also performs well in worst case. Faced with this impossibility we consider an alternative solution concept to truthful mechanism design, that of truthfulness with with high probability, and give an envy-free auction that performs well in worst case and is truthful with high probability [24].

In our study of the basic auction we develop a number of techniques that can be used in the solution of more complex profit maximizing mechanism design problems. We generalize

[2]In such markets, it would be better to use some sort of market analysis to determine how to price the item instead of applying a worst case mechanism design approach.

one natural technique, using random sampling to perform on-the-fly market analysis as the auction is being run, to give a solution to the multi-item auction problem [22]. This is a natural approach and it is likely be useful in other profit optimization problems.

As a building block for the design of profit maximizing mechanisms, we develop the notion of a *cancellable* mechanism. A cancellable mechanism is one that remains truthful even if its outcome can be cancelled if its profit does not meet some criterion. This gives a reduction from the private value optimization problems to public value optimization problems [20]. As an example, consider the multicast pricing problem. If a cancellable auction is run locally on the bidders in each market, this effectively allows us to replace the bidders' private values with a single public value representing the revenue possible from each market. The mechanism can then optimize which markets receive the good based on the cost function and these public revenue values. In any markets that turn out to be unprofitable, the local auction can be cancelled.

Finally, we give a reduction from the private value optimization problem to the private value decision problem [23]. In classical optimization, an optimization problem attempts to answer the question "find a feasible solution with the maximum value." The decision problem answers the question "find a feasible solution with value at least V if one exists." Likewise, we can define the decision problem version of a private value optimization problem. We call a mechanism that solves the private value decision problem a *profit extractor*. Such a mechanism given a target profit V should produce an outcome that achieves profit at least V if such an outcome is possible. We show how to construct an approximately optimal mechanism given a profit extractor for the same profit maximization problem. This technique gives the best known competitive ratio for the basic auction problem. We further demonstrate its usefulness by using it to solve the double auction problem as well [15].

Chapter 2

DEFINITIONS

In this chapter we formalize the definition of the basic auction problem. In doing so, we review the necessary game theoretic concepts including our bidder model and the solution concept of truthful mechanism design. We present the prototypical truthful auction, the Vickrey auction, as well as the truthful fixed price mechanism. We then describe the basic auction decision problem and the mechanism from the literature that solves it.

2.1 *The Basic Auction Problem*

Motivated by the problem of selling digital goods, such as music, video, and software; we consider the problem of a monopolistic seller attempting to maximize their profit in the sale a single good available in *unlimited supply*. In the digital good setting, it is natural to assume that the seller has enough copies of the good for sale to potentially sell one to each customer. Furthermore, in the digital good setting it is safe to assume that each consumer only desires one copy of the good, a.k.a, the *unit demand* case. We will assume that the seller knows nothing about the consumers preferences or willingness to pay in advance. We will be looking for solutions to this problem in the form of dynamic pricing mechanisms that adjust the sale prices of the good based on information received from the consumers in the form of bids, i.e., auctions.

As we will show later in this thesis, this auction problem is a special case of many interesting optimization problems. In fact in Chapter 5 we show that the *limited supply* case, when there are only a fixed number of identical items for sale, actually reduces to the unlimited supply case. We refer to the unlimited supply auction problem as the *basic auction problem*.

Definition 2.1 (The Basic Auction Problem) *Given:*

- *n identical items for sale and*

- *n bidders,*

design an auction to maximize profit of the sale.

2.2 The Game Theoretic Model and Truthful Mechanism Design

An auction outcome consists of a vector of prices, $\mathbf{p}$, with p_i the price for bidder i, and an allocation vector, $\mathbf{x}$, with $x_1 = 1$ if bidder i is allocated the good and $x_i = 0$ otherwise.

Definition 2.2 (Bidder Model) *We assume the following* private value model *for the bidders:*

- *Bidder i has a private value, u_i, representing the monetary value that the bidder associates with their possession of the item.*

- *Bidder i bids to maximize their profit defined as: $u_i x_i - p_i$. In particular this implies that there are* no externalities*, i.e., no bidder cares whether other bidders win or lose nor do they care what price other bidders might pay for the good.*

- *Bidders are rational.*

- *Bidders do not collude.*

A huge issue in the design and analysis of auctions is reasoning about how bidders will bid in the auction. As the bids dictate the profit that the auctioneer obtains it is impossible to analyze the auction profit without a model of how the bidders will bid. Since we have assumed that bidders will bid to try to maximize their own profit and their own profit is determined by the auction mechanism and the values of the other bidders' bids, a bidder's bid will be a function of the auction mechanism and what the bidder believes the other bidders are bidding. This is a daunting problem.

One solution from Economics is to look for a *Bayesian Nash equilibrium* of the mechanism. This approach assumes that the bidders' valuations are drawn from some probability distribution and that all bidders know what this distribution is. It then looks for a *Nash equilibrium*, bids for each bidder such that no bidder can obtain a higher expected profit by changing their bid. To be useful in practice, this solution relies on two things. First, it requires that the bidders know the probability distribution from which their valuations are drawn. Second, it assumes that they can actually find the Nash equilibrium. While these assumptions may be valid in some situations, we believe that there are many interesting situations in which they are not.

We adopt a different solution solution concept from Economics and game theory and alleviate the problem of having to have a model for the bidders' speculation by restricting our attention to mechanisms with *dominant strategies*.

Definition 2.3 (Dominant Strategy Mechanism) *A bidder's strategy is* dominant *if the strategy maximizes the bidder's profit for any possible strategies the other bidders may follow. A mechanism is a* dominant strategy mechanism *if all bidders have dominant strategies.*

Definition 2.4 *A* single-round, sealed-bid auction*, $\mathcal{A}$, is one where:*

1. *Each bidder submits a bid, representing the maximum amount they are willing to pay for an item. We denote by* $\mathbf{b}$ *the vector of all submitted bids, i.e., the* input*. The i-th component of* $\mathbf{b}$ *is b_i, the bid submitted by bidder i. We denote by n the number of bidders.*

2. *Given the bid vector* $\mathbf{b}$*, the auctioneer computes an* output *consisting of an* allocation*, $\mathbf{x} = (x_1, \ldots, x_n)$, and* prices*, $\mathbf{p} = (p_1, \ldots, p_n)$. The allocation x_i is an indicator for bidder i's receipt of the item (1 if bidder i receives the item and 0 otherwise). If $x_i = 1$, we say that* bidder i *wins.* Otherwise, bidder i *loses*, or is rejected*. The price, p_i, is what bidder i pays the auctioneer. We assume that $0 \leq p_i \leq b_i$ for all winning bidders and that $p_i = 0$ for all losing bidders (these are the standard assumptions of* no positive transfers *and* voluntary participation*. See, e.g., [37]).*

3. *The* profit *of the auction (or auctioneer) is $\mathcal{A}(\mathbf{b}) = \sum_i p_i$.*

An auction is *deterministic* if the allocation and prices are a deterministic function of the bid vector. An auction is *randomized* if the procedure by which the auctioneer computes the allocation and prices is randomized. Note that if the auction is randomized, the profit of the auction, the output prices, and the allocation are random variables.

The *revelation principle* says that any dominant strategy single or multi-round mechanism can be converted into a single round *truthful mechanism*, one where each bidders dominant strategy is to bid identically their valuation [38]. The proof of this principle is straightforward, given a dominant strategy mechanism $\mathcal{M}$, construct a truthful mechanism $\mathcal{M}'$ that simulates each bidders dominant strategy in $\mathcal{M}$ assuming that the bidder actually bid their true valuation.

Definition 2.5 (Truthful Mechanism) *We say that a deterministic auction is* truthful *if bidder i's dominant strategy is to bid their valuation, setting $b_i = u_i$. A randomized auction is truthful if it is a probability distribution over truthful deterministic auctions.*

We adopt the solution concept of *truthful mechanism design* and consider single-round, sealed-bid truthful auctions for the basic auction problem.

2.3 The Vickrey Auction

A 1-item auction is one that sells at most one item. The classic truthful auction from the Economics literature is the *1-item Vickrey auction.* The Vickrey auction is also known as the *second price auction* because it sells the item to the highest bidder at a price equal to the second highest bid value. It is interesting to note that the Vickrey auction can be derived by considering applying the revelation principle to the standard English auction.

The English Auction for one item is a standard multi-round auction mechanism. It is well known for being used in estate and art sales. It is an *open outcry* auction and typically bidders yell out their bids effectively raising the price of the item for sale until a price is reached at that nobody is willing to out bid. At the point the bidder that placed the highest bid obtains the item at their bid value.

In the English auction, a rational bidder's strategy would be to bid by raising the high bid placed by other bidders by the minimum increment until the high bid is above their valuation at which point they would cease bidding in the auction. The result of bidders in an English auction playing by this rational strategy is that the bidder with the highest utility will in fact win the item and the price they will pay will be approximately the second highest utility value. Applying the revelation principle we arrive at the Vickrey auction as the single round, sealed bid truthful equivalent of the multi-round English Auction.

It will be useful for later discussion to have a more direct proof that the Vickrey auction is truthful.

Lemma 2.1 *The Vickrey Auction is truthful.*

Proof: For a particular bidder i, fix the bids of all other bidders. Let $p = \max_{j \neq i} b_j$. Note that if bidder i bids $b_i > p$ then bidder i wins the auction and pays price p. This is because bidder i would be the highest bidder and p would be the second highest bid value. In this case, the bidder's profit is $u_i - p$. If bidder i bids below p then bidder i loses and pays nothing. Their profit is zero. Given this, we can consider the profit of bidder i with utility value u_i for any possible bid they might make. There are two cases of interest.

Case 1 ($u_i < p$): If the bidder bids above p, their profit is $u_i - p$ which is negative. Thus bidding below p and obtaining zero profit is preferable. Thus, any losing bid is an optimal strategy for bidder i, including bidding $b_i = u_i$.

Case 2 ($u_i > p$): If the bidder bids above p, their profit $u_i - p$ is positive. The is preferable to losing by biding below p. Thus, any bid greater than p is an optimal strategy for bidder i, including bidding $b_i = u_i$.

The third case is when $u_i = p$. Note that in this case the bidder cares not whether they win at price p or lose as both outcomes give zero profit. □

2.4 Bid-Independent Auctions

Throughout the remainder of this thesis, we will be developing truthful auctions. To get a better understanding of what makes a mechanism truthful we present an algorithmic characterization of truthful auctions. First observe that the fundamental property of the

Vickrey auction that made it truthful was that the price p that bidder i is compared to is not a function of that bidders bid value. In the Vickrey auction p is the maximum of all the other bids. However, the auction would still be truthful if p were the value of any function of all of the bids except for bidder i's bid.

Bid-independent auctions formalize this observation and give a characterization of truthful auctions. Related formulations to the one we give here have appeared in numerous places in recent literature (e.g., [3, 48, 25, 32]). To the best of our knowledge, the earliest dates back to the 1970s [35].

Definition 2.6 *Let $\mathbf{b}_{-i}$ denote the vector of bids $\mathbf{b}$ with b_i removed, i.e.,*

$$\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, ?, b_{i+1}, \ldots, b_n).$$

We call such a vector masked.

Definition 2.7 (Bid-independent Auction, BI_f) *Let f be a function from masked vectors to prices (non-negative real numbers). The* deterministic bid-independent auction defined by f, BI_f, *works as follows. For each bidder i:*

1. $t_i \leftarrow f(\mathbf{b}_{-i})$.

2. *If $t_i < b_i$, set $x_i \leftarrow 1$ and $p_i \leftarrow t_i$ (Bidder i wins).*

3. *If $t_i > b_i$ set $x_i = p_i = 0$ (Bidder i loses).*

4. *Otherwise, if $t_i = b_i$ the auction can either accept the bid at price t_i or reject it.*

For example, by setting $f = \max$ for all i and breaking ties arbitrarily, we obtain the 1-item Vickrey auction, i.e., the highest bidder wins at the second highest price.

Theorem 2.2 *A deterministic auction is truthful if and only if it is equivalent to a deterministic bid-independent auction.*

The theorem follows from the following two lemmas.

Lemma 2.3 *Any deterministic bid-independent auction is truthful.*

We omit this proof because it is identical to the proof that Vickrey is truthful. The following result completes the proof of equivalence of truthfulness and bid-independence for deterministic auctions.

Lemma 2.4 *A truthful deterministic auction is truthful if and only if it is equivalent to a deterministic bid-independent auction.*

Proof: Given any truthful deterministic $\mathcal{A}$ we can determine an f such that the bid-independent implementation, BI_f, is identical to $\mathcal{A}$. Let $\mathbf{b}_i^x = (b_1, \ldots, b_{i-1}, x, b_{i+1}, \ldots, b_n)$, the bid vector obtained by replacing b_i with x. If there is some value x^* such that in $\mathcal{A}(\mathbf{b}_i^{x^*})$ bidder i wins and pays p (note this requires $p \leq x^*$) then define $f(\mathbf{b}_{-i})$ to be p. To break ties, i.e., if $b_i = p$, consider whether bidder i wins in $\mathcal{A}$ on input $\mathbf{b}_i^p$.

Give this value of p, we now show for any x the outcome of $\mathcal{A}$ on $\mathbf{b}_i^x$ is such that:

1. If bidder i wins, they pay p.

2. Bidder i wins by bidding any value $x > p$ (and possibly by bidding $x = p$).

To see 1, assume to the contrary that there is some other bid value y such that running $\mathcal{A}(\mathbf{b}_i^y)$ results in bidder i winning and paying $q \neq p$. Without loss of generality $q > p$ so a bidder with utility y would have a higher profit by bidding x^*. This contradicts $\mathcal{A}$'s truthfulness.

To see 2, assume to the contrary that there is some bid value $y \in (p, \infty)$ such that bidder i does not win by bidding y. Notice that a bidder with utility y would have a higher profit by bidding x^*. Again this contradicts the $\mathcal{A}$'s truthfulness and gives the lemma. □

Definition 2.8 *A* randomized bid-independent auction *is a probability distribution over bid-independent auctions. For these auctions, $f(\mathbf{b}_{-i})$ is a non-negative real-valued random variable.*

Note that the random variables $f(\mathbf{b}_{-i})$ and $f(\mathbf{b}_{-j})$ need not be independent. It follows immediately from Definition 2.8 and Theorem 2.2 that:

Corollary 2.5 *A randomized auction is truthful if and only if it is equivalent to a randomized bid-independent auction.*

We note briefly that this definition of randomized truthful auctions is identical to requiring the probability distribution of a bidders profit when they bid their true value to *dominate*[1] the probability distribution of their profit bidding any other value.

2.5 More Truthful Auctions

In this section we review several other classical truthful auctions.

The Fixed Pricing Mechanism

While the *fixed pricing mechanism* is not technically an auction because it does not require any bids be submitted by the bidders, it will be useful for us to consider it as a truthful mechanism. The *fixed pricing mechanism with sale price* r sells to all bidders that bid at least r at price r. Clearly, this is truthful as it is implemented by the constant bid-independent function $f(\cdot) = r$.

The k-item Vickrey Auction

We define the notation $b_{(i)}$ to denote the ith largest bid value. If necessary, we break ties arbitrarily.

The k-item Vickrey auction is the natural extension of the Vickrey auction to sell k items instead of only one.

Definition 2.9 (k-item Vickrey Auction, $\mathcal{V}_k$) *The* k-item Vickrey Auction, $\mathcal{V}_k$, *sells to the highest* k *bidders at the* $k+1$*st bid value, i.e.,* $b_{(k+1)}$. *All other bidders are rejected.*

It is easy to see that the k-item Vickrey auction is truthful. It is implemented by the bid independent function that returns the kth largest bid value (and breaks ties arbitrarily).

[1] The random variable X dominates Y if for all v, $\mathbf{Pr}[X > v] \geq \mathbf{Pr}[Y > v]$.

The k-item Vickrey Auction with Reservation Price

A standard variant of k-Vickrey is parameterized by an a reservation price, r. This variant sells at most k items total to the highest k bidders that bid at least r. It uses a sale price of the larger of r and the $k+1$st highest bid, i.e., $\max(r, b_{(k+1)})$. This is easily implemented by the bid independent function that returns the maximum of the kth largest bid value and r. This is a natural combination of the fixed price mechanism and k-Vickrey.

Mechanisms with Prior Knowledge

As a brief aside, the standard approach to profit maximization in Economics is via the *Bayesian optimal auction.* The significant difference between this and the problem we solve in this thesis is that the Bayesian optimal auction is endowed with prior knowledge of the a probability distribution from which the bidders valuations are drawn. The Bayesian optimal auction problem is, given the prior distribution from which the bidders valuations are drawn, to come up with the auction that maximizes the expected profit. This expectation is taken over randomness in the bidders' valuations. For the case where the bidders' values are independent and identically distributed, the Bayesian optimal auction is just the k-item Vickrey with a reservation price judiciously chosen based on the prior distribution [38]. Note, that in the unlimited supply case this k, the number of items available, is n and this Bayesian optimal auction simplifies to the fixed price mechanism with the optimally chosen price, r.

An an interesting exercise we can consider using the fixed pricing mechanism by a seller with perfect information, i.e., knowledge of all bidders' exact valuations. Such a seller can simply pick the value of r that maximizes their profit. We denote this *optimal omniscient mechanism* by $\mathcal{F}$ and discuss it in more detail in Chapter 3.

2.6 Profit Extraction and the Basic Auction Decision Problem

In classical optimization theory, an optimization problem is to find a feasible solution with the maximum (or minimum) value. The corresponding decision problem is, given some prespecified value V, to find a feasible solution with value at least V if such a solution exists (or at most V for minimization problems). In classical optimization, the decision problem solution is useful in finding a solution to the optimization problem because the optimal value can be searched for by repeatedly solving the decision problem for different values of V, for example, by binary search.

Our basic auction problem is a private value optimization problem. One of the main result of this thesis is the development of a technique for using the solution to the decision problem version of the basic auction to solve the profit maximizing auctions.

Definition 2.10 (Basic Auction Decision Problem) *Given:*

- *n identical items for sale.*
- *n bidders, bidder i willing to pay at most u_i for an item.*
- *Target profit R.*

Design an auction mechanism that obtains profit R if R is less than the profit of $\mathcal{F}$, the fixed pricing mechanism with optimal price.[2]

We call a mechanism that solves the basic auction decision problem a *profit extractor* because it extracts the specified amount of profit whenever it is possible. It turns out that a special case of the general cost sharing mechanism of Moulin and Shenker [37] is a profit extractor for the basic auction problem.

Definition 2.11 (ProfitExtract$_R$) *Given bids* $\mathbf{b}$ *find the largest k such that the highest k bidders can equally split R, i.e., for $i \leq k$, $b_{(k)} \geq R/k$. Charge each of these bidders R/k.*

Observation 1 *ProfitExtract$_R$ obtains profit R if and only if $\mathcal{F}(\mathbf{b}) \geq R$.*

To see this, let k^* be the number of winners in $\mathcal{F}$. If $\mathcal{F}(\mathbf{b}) = k^* b_{(k^*)} \geq R$ then these k^* winners can equally split the cost R/k^*. On the other hand if $\mathcal{F}(\mathbf{b}) < R$ then no k bidders can equally split R.

For completeness, we now give a proof of the fact that ProfitExtract$_R$ is truthful, but first a lemma. A more general result is given by Moulin and Shenker in [37].

[2] $\mathcal{F}$ is described in detail in Chapter 3.

Lemma 2.6 *The number of winners in ProfitExtract$_R$ is a non-increasing function of any one bidder's bid.*

Proof: ProfitExtract$_R$ finds the largest set of bidders that can split R. If the highest k' bidders cannot equally split R then they certainly still cannot if some bidder lowers their bid. □

Theorem 2.7 *[37] ProfitExtract$_R$ is truthful.*

Proof: Pick any bidder i. We will show that ProfitExtract$_R$ is bid-independent. That is, there is a value p that is a function of $\mathbf{b}_{-i}$ and if bidder i bids at least p, i wins at price p and otherwise i loses. For the remainder of this argument fix $\mathbf{b}_{-i}$.

Assume bidder i bids ∞. Of course, i must win at price $p \leq R$, because at the very least, no other bidders bid above R and i can pay R. Let k be the number of winners in ProfitExtract$_R$ when i bids ∞. If bidder i were to lower their bid to any value at least p, bidder i would still win at price p: the same k winning bidders can still split the R, and by Lemma 2.6, no larger set of bidders can win.

To see that bidder i would lose if they were to lower their bid to any value strictly less than p, note that the top k bidders can no longer split R evenly as bidder i cannot pay $p = R/k$. Furthermore, by Lemma 2.6, no larger set of bidders can split R either. □

There are some other nice properties of ProfitExtract$_R$ such as the fact that it always produces a fair outcome in the sense that there is a single sale price which all winners pay and all losers would not prefer to win the item at this price.[3] Also ProfitExtract$_R$ is *group strategyproof* meaning that it is not possible for any bidders to collude for positive benefit without a loss being incurred by some member of the coalition (See [37] for proof).

---

[3] In Chapter 6 we formalize this fairness criterion as *envy-free*.

Chapter 3

ANALYSIS FRAMEWORK

The goal of the analysis of auctions, like the goal of analysis of algorithms, is to be able to consider a particular auction and determine whether it is good or not. We also might like to compare two auctions to determine which was is better, or even in what sense one auction is better than the other. The analysis framework we present here, that which we will be using for analyzing auctions, will give us a way to determine whether a particular auction is good and whether it is better or worse than another particular auction.

As mentioned in the introduction, the analysis of auctions is complicated by the fact that there is no auction that is the best on all inputs. For any truthful auction, $\mathcal{A}$, there exists an input, $\mathbf{b}$, and another truthful auction, $\mathcal{A}'$, such that the profit of $\mathcal{A}'$ on $\mathbf{b}$ is higher than that of $\mathcal{A}$. The fundamental obstacle that an auctioneer faces that prevents the auctioneer from obtaining a high profit is their lack of prior knowledge of the bidders' valuations.

The areas of *approximation algorithms* and *online algorithms* share similar obstacles. Approximation algorithms are useful for obtaining polynomial time algorithms for problems that are impossible to solve exactly in polynomial time. The performance of such an algorithm is gauged by measuring either the additive or multiplicative approximation factor that the algorithm obtains in comparison to that of the true optimal. In online algorithms the impossibility the algorithm faces is with knowing the future input. To gauge performance of an online algorithm, it it compared to the *optimal offline* algorithm's performance. The goal of this *competitive analysis* is to obtain the algorithm with the best *competitive ratio.*

We adopt this general approach as our means for evaluating profit maximizing mechanisms for private value problems. In this analysis we are interested in finding the algorithm that obtains a profit within the best multiplicative ratio of an "optimal" mechanism that is endowed with perfect prior knowledge of the private values of the bidders. We refer to such an optimal mechanism that knows the private values in advance as the *optimal omniscient mechanism* or the *optimal public value mechanism*, interchangeably.

In general, our goal is to get the strongest results we can, and thus we will try to compare our truthful auctions against the best possible metric that we can feasibly compete with. Once the optimal auction is defined, the best *competitive auction* is the one that minimizes the *competitive ratio*, the worst case ratio between the optimal auction profit and its profit. A fundamental goal of this research is to determine good metrics for comparison, the corresponding best competitive ratio, and the auction mechanism that achieves it.

3.1 Competitive Auction Framework

A key part of setting up a competitive framework for analyzing solutions to any problem is coming up with the right metric for comparison. As a starting point, we would like to take the strongest possible benchmark for comparison that we can: the profit of an auctioneer that is perfectly informed about the bidder's values. This leads us to consider as the two most natural metrics for comparison the optimal omniscient multi-price and single-price auctions, $\mathcal{T}$ and $\mathcal{F}$.

Recall that we use the notation $b_{(i)}$ to represent the ith largest bid value in $\mathbf{b}$. Where necessary we break ties arbitrarily.

Definition 3.1 *The* optimal single price omniscient auction*, $\mathcal{F}$, on input, $\mathbf{b}$, determines the value k such that $kb_{(k)}$ is maximized. All bidders with $b_i \geq b_{(k)}$ win at price $b_{(k)}$; all remaining bidders lose. The profit of $\mathcal{F}$ on input $\mathbf{b}$ is thus*

$$\mathcal{F}(\mathbf{b}) = \max_{1 \leq k \leq n} kb_{(k)}.$$

Definition 3.2 *The* optimal multiple price omniscient auction*, $\mathcal{T}$, is the auction that sells to each bidder at their bid value. Thus, the profit of $\mathcal{T}$ on input $\mathbf{b}$ is*

$$\mathcal{T}(\mathbf{b}) = \sum_{1 \leq i \leq n} b_i.$$

We first compare the performance of $\mathcal{F}$ relative to $\mathcal{T}$. Specifically, we observe that in the worst case, the maximum ratio of $\mathcal{T}$ to $\mathcal{F}$ is logarithmic in n, the number of bidders.

Lemma 3.1 *There exist bid vectors* $\mathbf{b}$ *for which*

$$\mathcal{F}(\mathbf{b}) = \Theta(\mathcal{T}(\mathbf{b})/\ln n).$$

Moreover, for all bid vectors $\mathbf{b}$

$$\mathcal{F}(\mathbf{b}) \geq \mathcal{T}(\mathbf{b})/\ln n.$$

Proof: For the first part, let $\mathbf{b}$ be the bid vector such that $b_i = n/i$. Then $\mathcal{F}(\mathbf{b}) = n$ and $\mathcal{T}(\mathbf{b}) = n(\ln(n) + \Theta(1))$.

For the second part, suppose that $\mathcal{F}(\mathbf{b}) = \max_i ib_{(i)} = kb_{(k)}$. Then for all i,

$$ib_{(i)} \leq kb_{(k)}.$$

Thus,

$$\mathcal{T}(\mathbf{b}) = \sum_{i=1}^{n} b_{(i)} \leq \sum_{i=1}^{n} \frac{kb_{(k)}}{i} \leq \mathcal{F}(\mathbf{b}) \sum_{j=1}^{n} \frac{1}{j} = \mathcal{F}(\mathbf{b})(\ln n + O(1)).$$

□

Now we show that no truthful auction can be competitive against $\mathcal{F}$ (and hence it can not be competitive against $\mathcal{T}$).

Lemma 3.2 *For any truthful auction* $\mathcal{A}$ *and any* $\beta \geq 1$*, there is a bid vector* $\mathbf{b}$ *such that the expected profit of* $\mathcal{A}$ *on* $\mathbf{b}$ *is less than* $\mathcal{F}(\mathbf{b})/\beta$*.*

Proof: Consider a bid-independent randomized auction on two bids, 1 and $x \geq 1$. Let f be the function that gives the bid-independent definition of $\mathcal{A}$.[1] Let h be the smallest value greater or equal to 1 such that $\mathbf{Pr}[f(1) \geq h] \leq \frac{1}{2\beta}$. Then the expected profit on input vector $\mathbf{b} = (1, H)$ with $H = 4\beta h$ is at most

$$\frac{H}{2\beta} + h(1 - \frac{1}{2\beta}) + 1 < 4h = \frac{H}{\beta} = \frac{\mathcal{F}(\mathbf{b})}{\beta}.$$

□

Lemma 3.2 shows that we cannot expect to come close to matching the performance of the optimal single price omniscient auction in the case where the optimal profit is generated

[1]We assume f is symmetric, thus, $f(b, ?) = f(?, b) = f(b)$. By Lemma 7.1 this assumption is without loss of generality.

from the single highest bid. Thus, we must set our sights slightly lower. Later in the paper we will present auctions that are *competitive* with $\mathcal{F}^{(2)}$, the optimal single price auction that sells at least two items. Such auctions perform comparably to $\mathcal{F}^{(2)}$ in that they achieve a constant fraction of the revenue of $\mathcal{F}^{(2)}$ on all inputs.

Definition 3.3 *The* optimal single price omniscient auction that sells at least two items, *$\mathcal{F}^{(2)}$, is defined as follows: Auction $\mathcal{F}^{(2)}$ on input* $\mathbf{b}$ *determines the value k such that $k \geq 2$ and $kb_{(k)}$ is maximized. All bidders with $b_i \geq b_{(k)}$ win at price $b_{(k)}$; all remaining bidders lose. The profit of $\mathcal{F}^{(2)}$ on input* $\mathbf{b}$ *is thus*

$$\mathcal{F}^{(2)}(\mathbf{b}) = \max_{2 \leq k \leq n} kb_{(k)}.$$

Note that for $\mathbf{b}$ where $\mathcal{F}$ elects to sell two or more items, $\mathcal{F}^{(2)}(\mathbf{b}) = \mathcal{F}(\mathbf{b})$. Thus, if we exclude bid vectors where only the highest bidder wins in the optimal auction, comparing auction performance to $\mathcal{F}^{(2)}$ is identical to comparing it to $\mathcal{F}$.

Next we generalize the definition of $\mathcal{F}^{(2)}$ to define $\mathcal{F}^{(m)}$, as the optimal single price omniscient auction that sells at least m items.

Definition 3.4 *The* optimal single price omniscient auction with at least m winners, *$\mathcal{F}^{(m)}$ is defined as follows: Auction $\mathcal{F}^{(m)}$ on input* $\mathbf{b}$ *determines the value k such that $k \geq m$ and $kb_{(k)}$ is maximized. All bidders with $b_i \geq b_{(k)}$ win at price $b_{(k)}$; all remaining bidders lose. The profit of $\mathcal{F}^{(m)}$ on input* $\mathbf{b}$ *is thus*

$$\mathcal{F}^{(m)}(\mathbf{b}) = \max_{m \leq k \leq n} kb_{(k)}.$$

Finally, we formalize the notion of a competitive auction.

Definition 3.5 *We say that auction $\mathcal{A}$ is β-*competitive against *$\mathcal{F}^{(m)}$ if for all bid vectors* $\mathbf{b}$*, the expected profit of $\mathcal{A}$ on* $\mathbf{b}$ *satisfies*

$$\mathbf{E}[\mathcal{A}(\mathbf{b})] \geq \frac{\mathcal{F}^{(m)}(\mathbf{b})}{\beta}.$$

We say that an auction is competitive against $\mathcal{F}^{(m)}$ if the auction is β-competitive against $\mathcal{F}^{(m)}$ for a constant β. We refer to β as the competitive ratio *of $\mathcal{A}$.*

The statement that an auction is competitive against $\mathcal{F}^{(m)}$ implies that, restricted to inputs $\mathbf{b}$ such that there are at least m items sold by the optimal auction, i.e., $\mathcal{F}^{(m)}(\mathbf{b}) = \mathcal{F}(\mathbf{b})$, our auctions are competitive against $\mathcal{F}$. For $m = 2$, this restriction precisely excludes the case where it is not possible to be competitive – when there is one bidder with very large utility.

Observe that since the profit of $\mathcal{F}^{(m)}$ decreases as m increases, as we compete against $\mathcal{F}^{(m)}$ for larger values of m, we are demanding less and less of the truthful auction. We will thus single out the case of competition against $m = 2$, as $\mathcal{F}^{(2)}$ is the strongest omniscient auction that we will be able to feasibly compete with.

Definition 3.6 *We say an auction is β-*competitive *if it is β-competitive against $\mathcal{F}^{(2)}$. In cases where we do not wish to specify the constant β, we simply say that the auction is* competitive.

By considering competitiveness against $\mathcal{F}^{(m)}$ for different values of m, we obtain results that are relevant to a wider range of applications. For example, in situations in which the auctioneer does not have prior distributions on bidders' bid values but does know that their profit will be maximized by selling at least m items, they can use an auction tailored to such a situation and obtain much stronger guarantees on the competitive ratio. One natural example of such a scenario is when there are a large number of bidders and the auctioneer is safe in assuming that all bid values come from a bounded range.

3.2 Mass Markets

We will demonstrate in Chapter 7 that no auction can achieve a worst case competitive ratio with $\mathcal{F}^{(2)}$ that is better than 2.42. In that we are attempting to design practical selling mechanisms, this is a rather discouraging result. It turns out that this competitive ratio comes from requiring the auction to perform well on inputs where only a small number of items are sold. Intuitively, if there are only small number of items sold, the contribution of the sale of each item to the total profit is large. Thus, small pricing mistakes made by the auction have a relatively large effect on the total profit. On the other hand, one of the original motivations of this work is on sales where the number of items sold is large, i.e., the *mass market* case.

Also in Chapter 7 provide some evidence supporting the conjecture that no auction that is constant competitive in worst case is almost optimal on mass markets (i.e., $(1+\epsilon)$-competitive as m, the number of winners, gets large). We show a weaker result that for this is true for a restricted class of auctions. For mass markets, it would be preferable to give up good worst case behavior when only a few items are sold to obtain better performance asymptotically. As such, we are motivated to look at a *promise* version of the competitive framework. In such a framework we look for auctions that perform well on certain classes of inputs, but allow them to perform poorly on other inputs. This motivates defining a version of the competitive ratio that is parameterized by m, the optimal number of items sold. We look for auctions that perform well on the restricted set of inputs in which the optimal number of items sold is at least m.

Definition 3.7 *We say that auction, $\mathcal{A}$, is $\beta(m)$-*competitive *for mass-markets if, for all bid vectors* $\mathbf{b}$ *such that $\mathcal{F}$ sells at least m items, we have*

$$\mathbf{E}[\mathcal{A}(\mathbf{b})] \geq \frac{\mathcal{F}^{(2)}(\mathbf{b})}{\beta(m)}.$$

This parameterized competitive framework is more expressive than the worst case framework discussed earlier. Here $\beta(1)$ is the worst case competitive ratio. Note that we have defined this competitiveness notion in terms of $\mathcal{F}^{(2)}$, it is possible to do this analysis with $\mathcal{F}^{(m')}$ for any m'. Note, however, that for bid vectors with $m \geq m'$ winners in $\mathcal{F}$, $\mathcal{F}^{(m')} = \mathcal{F}$.

Another variant of promise problem that we will consider is restricting the bids to be from a bounded range, i.e., $[1, h]$ and looking at auctions that perform well in the limit as the number of bidders increases. Note that for bids in a bounded range, as the number of bidders increases, the number of items sold also must increase.

3.3 Concentration

We have defined our notion of competitive analysis in terms of expected profit. However, often times it is necessary to have concentration results that show that with high probability the profit is close to its expectation.

These sorts of results are easy to obtain for all of the auctions that we design here. As we have defined them, a randomized auction is a randomization over deterministic auctions, i.e., first we flip some coins to pick a deterministic auction, second we run the deterministic auction on all the inputs. The total profit of the auctioneer is $\sum_i p_i$, the sum of the prices charged to the bidders. Note that by linearity of expectation, the auctions expected profit is $\sum_i \mathbf{E}[p_i]$. Given any auction, we can design an auction that achieves the same expected profit but also has simple to prove tail bounds on the profit we can just rerandomize the auction for each bidder. I.e., for each bidder i: flip some coins to pick a deterministic auction, simulate this auction on the entire input, and use the outcome p_i and x_i for this bidder. Note that the p_i are now independent random variables and therefore it is possible to apply a Chernoff bound to show that for the case the number of winners is large, that the auction profit is close to its expectation.

We note that this rerandomization may come at the expense of other auction properties, such as single pricedness as in Chapter 6. As well, this rerandomization is not possible when there is a cost function over the outcome allocation as in Chapter 10 or infeasible allocations as in the double auction problem in Chapter 11.

Chapter 4

COMPETITIVE AUCTIONS

In this chapter we develop a number of new techniques for designing profit maximizing auctions. First we discuss the use of both the fixed pricing mechanism with price r and the k-item Vickrey auction in obtaining auctions that perform well in worst case. We show that with deterministic choice of k or r neither mechanism gives an auction which is n-competitive in worst case. We then give a natural randomized technique for choosing k and p to give auctions that are $\Theta(\log n)$-competitive. To within constant factors this is the best possible using these methods. To obtain constant-competitive mechanisms we turn to more sophisticated techniques.

We first present a deterministic auction mechanism. While this mechanism can perform well in average case (e.g., when the bid values are chosen from independently from a bounded support probability distribution), it does not perform well in worst case. We prove this in general by showing that no symmetric[1] deterministic auction performs well in worst case [26].

We then discuss a number of techniques for designing auctions that perform well in worst case. Because of the deterministic impossibility, these techniques are all based on randomization procedures. The simplest of these randomization procedures is *random sampling.* Random sampling is a natural, well accepted technique in both algorithms and economics, in fact our use of random sampling in auctions is essentially performing market analysis dynamically as the auction is being run. A more sophisticated randomization technique is the *randomized consensus* procedure that we will develop later in this chapter.

We will show how these randomization procedures can be combined with existing mechanisms to get a profit maximizing auction. In particular we will utilize the simple selling mechanism of *fixed pricing* [26, 25] as well as the basic auction *profit extraction* mechanism

[1] A symmetric auction is one where the outcome is independent of the order in which the bids are input.

[20, 23].

4.1 Existing Methods

In this section we will discuss the use of existing methods to try to obtain profit maximizing mechanisms.

Profit Maximization via the Fixed Price Mechanism

Consider the fixed price mechanism with price r. The performance of the fixed price mechanism relies on a good choice of price. Given our assumption that the auctioneer knows nothing about the bidders' valuations in advance, there is no way that the auctioneer can intelligently set the price. For any price r, there is a very bad input $\mathbf{b}$ with all $b_i < r$ on which the fixed price mechanism with price r obtains no profit. Even if we were to assume that the auctioneer had some bound on the range of bids, say between 1 and h, it would still be impossible to choose r. For any choice of $r > 1$ the all ones input would result in the auction obtaining zero profit. Furthermore, for $r = 1$ the input with all bids at value h would result in the auction obtaining profit n, but the best possible is $\mathcal{F} = \mathcal{T} = hn$ and thus, the auction would only be h-competitive.

As one might expect it is possible to randomize this technique to obtain an auction that performs significantly better in expectation. Given that the bids are all in the interval $[1, h]$, we can chose p to be a random power of two between $[1, h]$, that is $r = 2^i$ for i uniformly distributed from $\{0, \ldots, \lfloor \log h \rfloor\}$. Note that the expected profit from bid b_i is $\Theta(b_i / \log h)$ thus, by linearity of expectation, the expected profit of the mechanism of $\Theta(\mathcal{T} / \log h)$, that is, the mechanism is obtains a profit that is a $\Theta(\log h)$ fraction of $\mathcal{T}$ on all inputs (which implies that it is $\Theta(\log h)$-competitive with $\mathcal{T}$, $\mathcal{F}$, and $\mathcal{F}^{(2)}$). These bounds are tight.

If the upper bound on the bid values, h, is not known, but the lower bound is of 1 is, Yossi Azar has observed that we can still get close to $\log h$-competitive [7]. This result is obtained by using the *classify and select* technique [33] (and a good description can be found in [6]).

Profit Maximization via k-Vickrey

Now we consider obtaining a profit maximizing auction via the k-Vickrey mechanism. Again, for any fixed k that is not a function of any of the bid values (though possibly a function of n), there is a bad input. Suppose $k \geq 2$. On the input $\mathbf{b}$ with exactly k bids at h and the remaining bids at value 1, the k-Vickrey uses the $k+1$st bid, i.e. value 1, as the sale price for the highest k bidders. Thus, its profit is k, yet the optimal fixed price profit $\mathcal{F} = \mathcal{F}^{(2)} = kh$, thus k-Vickrey is only h-competitive. As h is a free variable, the competitive ratio is unbounded for the $k \geq 2$ case. For the case that $k = 1$, we can consider the all ones input. On this input, 1-Vickrey obtains profit 1; however, $\mathcal{F}^{(2)} = \mathcal{F} = n$. Thus, 1-Vickrey is only n-competitive.

Again using standard techniques we can randomly choose k. One advantage of this technique over using fixed pricing with random r is that we do not need to know either upper or lower bounds on the range of bids.

Definition 4.1 (Randomized Vickrey) *The* Randomized Vickrey *auction picks i uniformly from $\{0, \dots, \lfloor \log n \rfloor\}$ and runs the 2^i-item Vickrey auction.*

Lemma 4.1 *Randomized Vickrey is $2 \log n$-competitive with $\mathcal{F}^{(2)}$.*

Proof: The worst-case competitive ratio for this auction occurs on $\mathbf{b}$ with $n - m$ bids at value 0 and m bids at value v for any positive v. If there are m winners in $\mathcal{F}^{(2)}$, the 2^i-Vickrey auction gets revenue $2^i v$ if $2^i < m$ and zero otherwise. Since we choose each auction with probability $1/\log n$, our expected revenue on $\mathbf{b}$ is at least:

$$\begin{aligned}
\mathbf{E}[\mathcal{R}] &= \tfrac{v}{\log n} \sum_{i=0}^{\lceil \log m \rceil - 1} 2^i = \tfrac{v}{\log n}\left(2^{\lceil \log m \rceil} - 1\right) \\
&\geq \tfrac{v}{\log n}(m-1) \geq \tfrac{m-1}{m \log n} \mathcal{F}(\mathbf{b}) \geq \tfrac{\mathcal{F}^{(2)}(\mathbf{b})}{2 \log n}.
\end{aligned}$$

□

Theorem 6.2 implies that up to a $O(\log \log n)$ factor, the randomized Vickrey auction is the best possible randomization over k-Vickrey auctions.

We note in passing that k-Vickrey requires tie breaking to occur if there the kth and $k+1$st highest bids are equal, i.e., $b_{(k)} = b_{(k+1)}$. This is necessary to enforce that exactly

k items are sold. Note that in the unlimited supply case it is not necessary to enforce that k-Vickrey sell only k items. It is possible to use a variant of k-Vickrey that does not break ties in place of k-Vickrey in the above Randomized Vickrey auction; however, this modification does not help the worst case competitive ratio.

None of these above techniques are all that satisfying as our goal is to obtain a mechanism that achieves close to the optimal revenue. As such, we must consider designing new mechanisms.

4.2 The Bid-independent Optimal Price Auction

The next auction we consider is motivated by the comparison between optimal fixed pricing and the bid-independent auction framework. Optimal fixed pricing uses the best single sale price for the input bids, $\mathrm{opt}(\mathbf{b})$, and sells to all bidders above that price.

Definition 4.2 *For bids,* $\mathbf{b}$*,* $\mathrm{opt}(\mathbf{b})$ *is the sale price that obtains the highest profit when used in the fixed pricing mechanism. I.e.,*

$$\mathrm{opt}(\mathbf{b}) = b_{(i^*)} \ : \ i^* = \mathrm{argmax}_i \, i b_{(i)}.$$

In the optimal fixed pricing mechanism, $\mathcal{F}$, bidder i is offered price $\mathrm{opt}(\mathbf{b})$. When we are restricted to designing bid-independent mechanisms, an approach that seems like it might perform similarly to optimal fixed pricing is to offer bidder i price $\mathrm{opt}(\mathbf{b}_{-i})$ instead of $\mathrm{opt}(\mathbf{b})$. The auction we have just described is $\mathrm{BI}_{\mathrm{opt}}$:

Definition 4.3 (Bid-independent Optimal Price Auction, $\mathrm{BI}_{\mathrm{opt}}$**)** *The* bid-independent optimal price auction *is the bid-independent auction defined by function* $\mathrm{opt}(\cdot)$*. I.e., for each bidder* i*:*

1. $t_i \leftarrow \mathrm{opt}(\mathbf{b}_{-i})$.

2. *If* $t_i \leq b_i$*, set* $x_i \leftarrow 1$ *and* $p_i \leftarrow t_i$ *(Bidder* i *wins).*

3. *If* $t_i > b_i$ *set* $x_i = p_i = 0$ *(Bidder* i *loses).*

The hope is that $\mathrm{opt}(\mathbf{b}_{-i})$ is similar to $\mathrm{opt}(\mathbf{b})$. Unfortunately, as the following lemma shows, this is not the case.

Lemma 4.2 *The Bid-independent Optimal Price auction is not competitive.*

Proof: Consider n bidders where n/h of them bid $h \gg 1$ and the remaining bidders bid 1. Running $\mathrm{BI}_{\mathrm{opt}}$ on this bid vector will result in the following: For a bid at value h, we remove it and compute $\mathrm{opt}(\cdot)$ of the remaining bids. Of the $n-1$ bids remaining $n/h-1$ of them are at h. Thus, $\mathrm{opt}(\cdot)$ outputs 1 as $n-1$ bids at price 1 gives a higher revenue than $n/h-1$ bids at price h. Similarly for a bid at value 1, we remove it and compute $\mathrm{opt}(\cdot)$ of the remaining bids. Of the $n-1$ of them, there are n/h of them at h, the rest at 1. Thus, $\mathrm{opt}(\cdot)$ chooses to output h because n/h bids at price h gives a higher revenue than $n-1$ bids at price 1. Thus, all bids at value 1 are rejected and all bids at value h win the auction and only have to pay 1. $\mathrm{BI}_{\mathrm{opt}}$'s profit is thus n/h (the number of bids at h) whereas the optimal single price profit is n. For $h = n/2$, $\mathrm{BI}_{\mathrm{opt}}$'s profit is 2, while $\mathcal{F}^{(2)} = 2h = n$. Thus, $n/2$ is a lower bound on the competitive ratio of $\mathrm{BI}_{\mathrm{opt}}$. □

We note that any input on which the single price omniscient auction finds two different bid values that give approximately the same revenue will lead to a similar worst case. We extend this result in below by showing that all symmetric deterministic auctions suffer from this type of problem.

Despite the fact that $\mathrm{BI}_{\mathrm{opt}}$ is not competitive in worst case, it nonetheless has some desirable features. It does performs well in expectation over a large class of natural randomized inputs. Another interesting and useful property of $\mathrm{BI}_{\mathrm{opt}}$ is that it produces outcomes with the following structure. There is a single sale price p and a threshold t with $p \leq t$ such that the outcome of $\mathrm{BI}_{\mathrm{opt}}$ sells the item at price p to all bidders bidding at least t. All other bidders are rejected. This is formally stated and proved in Section 8.3.1.

4.3 Deterministic Impossibility

In this section, we show that no symmetric deterministic auction can be competitive. An auction is *symmetric* if the outcome is independent of the order of the bids. More precisely, we say that $\mathcal{A}$ is symmetric if for all bid vectors $\mathbf{b}$ and permutations π of the bidders, the output of $\mathcal{A}$ on input $\pi(\mathbf{b})$ is price vector $\pi(\mathbf{p})$ and allocation $\pi(\mathbf{x})$ (given that the output of $\mathcal{A}$ on $\mathbf{b}$ is $\mathbf{p}$ and $\mathbf{x}$).

We now show that no symmetric deterministic auction is competitive. In contrast, later in this chapter, we will show that there are competitive (symmetric) randomized auctions.

Theorem 4.3 *Let* BI_f *be any symmetric deterministic auction defined by bid-independent function* f*. Then* BI_f *is not competitive: For any* $1 \leq m \leq n$ *there exists a bid vector* $\mathbf{b}$ *of length* n *such that the profit of* BI_f *on* $\mathbf{b}$ *is at most* $\mathcal{F}^{(m)}(\mathbf{b})\frac{m}{n}$*.*

Proof: Fix n and m and the symmetric bid-independent auction BI_f. Consider the set of bid vectors whose bids are either n or 1. Note that for a symmetric auction, the location of the '?' in the masked vector $\mathbf{b}_{-i}$ cannot affect the value of $f(\mathbf{b}_{-i})$, therefore, $f(\cdot)$ is a function only of how many of the remaining $n-1$ bids are at value h or 1. For $0 \leq j \leq n-1$, write $f(j)$ for the price the auction assigns to a masked vector with exactly j bids at n and $n-1-j$ bids at 1. Note that if $f(0) > 1$ then if all bids are 1, the auction has profit 0 and the conclusion of the theorem holds trivially. So assume $f(0) \leq 1$ and let k be the largest integer in $\{0, \ldots, n-1\}$ such that $f(k) \leq 1$. Thus, $f(k) \leq 1$ and $f(k+1) > 1$. Let $\mathbf{b}$ be the the bid vector with $k+1$ bids at n and $n-k-1$ bids at 1. The profit of BI_f on $\mathbf{b}$ is $(k+1)f(k) \leq k+1$ since the only winners are those who bid n. If $k \leq m-1$ then $\mathcal{F}^{(m)}$ has profit n. If $k \geq m$ then $\mathcal{F}^{(m)}$ has profit at least $(k+1)n$. In either case the profit is at most $\mathcal{F}^{(m)}(\mathbf{b})\frac{m}{n}$. □

The following observation is due to Amos Fiat [19]. On inputs with bids only at either of two values, i.e., 1 and h, there exists an asymmetric deterministic auction that is 2-competitive. To show this we will give a function f and show that BI_f is competitive. This auction will accept about half the h bids at h and half of all the bids at 1. To do this, we make sure that we get every other h. Fix a particular bidder i. Given $\mathbf{b}_{-i}$, define $d_{\triangleleft}$ to be the number of consecutive ones to the left of the ith position in $\mathbf{b}_{-i}$ and $d_{\triangleright}$ to be the number to the right. Let $\#h$ be the number of hs to the left of the ith position in bid vector

$\mathbf{b}_{-i}$.

$$f(b_{-i}) = \begin{cases} h & \#h \text{ even and } d_{\triangleright} \text{ odd} \\ 1 & \#h \text{ even and } d_{\triangleright} \text{ even} \\ 1 & \#h \text{ odd and } d_{\triangleleft} \text{ odd} \\ h & \#h \text{ odd and } d_{\triangleleft} \text{ even} \end{cases}$$

We omit a formal proof; however, it is simple to verify that from the left this pairs hs together such that exactly one of each pair will be charged h. Furthermore, of the 1s between two hs, half of them (rounded up if odd) will be accepted at price 1. We leave the problem of whether there exist competitive asymmetric deterministic auctions as an open question.

The deterministic impossibility result motivates the consideration of randomized mechanisms. As a randomized mechanism is just a randomization over deterministic auctions, it is still the case that a randomized auction can perform poorly on unlucky outcomes of our random coin flips. However, the auction can be designed such that these unlucky outcomes are very improbable. We note that this use of randomness in the mechanism is very different from assuming the bids come from a random distribution. The latter is making an external assumption on our inputs, while internal randomness in the mechanism is guaranteed and completely under the control of the mechanism. As we will show with our development of randomized auctions that are competitive on worst case inputs, there is no need to make the assumption that the bids are randomly generated when we can randomize the mechanism instead.

4.4 Random Sampling Auctions

We now discuss a number of ways to obtain competitive auctions via random sampling. Intuitively, the random sampling technique that we present here can be viewed as doing market analysis dynamically as the auction is being run.

4.4.1 Random Sampling Optimal Price Auction

The first auction mechanism we will discuss is the combination of the random sampling technique with the fixed price mechanism. This mechanism selects a random sample of the bids, computes the optimal sale price for this sample, and runs the fixed price mechanism with this optimal sale price on the bidders that were not sampled.

The Random Sampling Optimal Price (RSOP) auction is guaranteed in expectation to achieve a constant fraction of the profit of $\mathcal{F}^{(2)}$, the optimal single-price omniscient auction which sells at least two items, on all inputs. More importantly, on a large class of interesting and practical bid vectors, RSOP is guaranteed in expectation to get very close the profit of $\mathcal{F}$, the optimal single-price omniscient auction.

As a first attempt at designing the RSOP auction we consider a variant of it, RSOP$'$:

Definition 4.4 (RSOP$'$) *The* RSOP$'$ *auction works as follows:*

1. *Partition bids* $\mathbf{b}$ *uniformly at random into two sets: for each bid, with probability 1/2 put the bid in* $\mathbf{b}'$ *and otherwise* $\mathbf{b}''$*.*

2. *Let* $p' = \mathrm{opt}(\mathbf{b}')$ *the optimal fixed price for* $\mathbf{b}'$ *(See Definition 4.2).*

3. *Run the fixed pricing mechanism with price* p' *on bids* $\mathbf{b}''$*. I.e., all bids in* $\mathbf{b}''$ *of value below* p' *are rejected; all remaining bids win at price* p'*.*

4. *Reject all bidders in* $\mathbf{b}'$*.*

There are a number of observations that can be made about RSOP$'$. First, it is single-priced: all winners pay the same price. Second, it is not fair in the sense that some bidders willing to pay this sale price will be rejected (i.e., the bidders in $\mathbf{b}'$ that bid above p').[2] As well, Step 4 could be modified to give the item to each bidder in $\mathbf{b}'$ for free. This would not affect the auction revenue and can be interpreted as winning a lottery.

[2] In chapter 6 we generalize this notion of fairness as the *envy-free* property and discuss designing auctions that are simultaneously envy-free, truthful, and maximized the expected profit of the auctioneer.

We note that in this auction there is a tradeoff between the accuracy of the sample and the loss in profit of the auction due to rejecting all bidders in the sample. This tradeoff can be adjusted by choosing a larger or smaller sample size. Alternatively, this tradeoff can be eliminated by considering a version of the auction that treats $\mathbf{b}'$ and $\mathbf{b}''$ symmetrically. We do so and obtain RSOP.

Definition 4.5 (Random Sampling Optimal Price Auction, RSOP**)** *The* Random Sampling Optimal Price Auction *works as follows:*

1. *Partition bids* $\mathbf{b}$ *uniformly at random into two sets: for each bid, with probability 1/2 put the bid in* $\mathbf{b}'$ *and otherwise* $\mathbf{b}''$*.*

2. *Let* $p' = \mathrm{opt}(\mathbf{b}')$ *and* $p'' = \mathrm{opt}(\mathbf{b}'')$*, the optimal fixed price thresholds for* $\mathbf{b}'$ *and* $\mathbf{b}''$*, respectively.*

3. *Run the fixed pricing mechanism with price* p' *on bids* $\mathbf{b}''$*. I.e., all bids in* $\mathbf{b}''$ *of value below* p' *are rejected; all remaining bids win at price* p'*.*

4. *Run the fixed pricing mechanism with price* p'' *on bids* $\mathbf{b}'$*.*

It is readily apparent that RSOP can be implemented bid-independently and thus from Corollary 2.5:

Observation 2 *The Random Sampling Optimal Price auction is truthful.*

We now discuss the performance of RSOP. Our first result is that on every input, RSOP achieves a constant fraction of the profit of $\mathcal{F}^{(2)}$.

Theorem 4.4 RSOP *is constant competitive against* $\mathcal{F}^{(2)}$*.*

Proofs of this theorem and the rest of the theorems in this section are technical and thus are deferred to the end of this chapter (Section 4.7).

The constant bound we obtain in this theorem is quite weak. However, there are a number of interesting special cases in which RSOP's performance is significantly better. One such special case is presented in the following theorem.

Theorem 4.5 *Let* $\mathbf{b}$ *be any bounded-range bid vector, i.e., any bid vector of* n *bids with* $b_i \in [1, h]$ *for all* i*. Then*

$$\lim_{n\to\infty} \max_{\mathbf{b}} \frac{\mathcal{F}(\mathbf{b})}{\mathrm{RSOP}(\mathbf{b})} = 1.$$

To prove this theorem, and generalizations thereof, we consider the RSOP_ℓ auction, a parameterized version of the RSOP auction. To define RSOP_ℓ, we first generalize the definition of $\mathrm{opt}(\cdot)$ (Definition 4.2).

Definition 4.6 *Let* $\mathbf{b}$ *be a vector of bids. Denote by* $\mathrm{opt}_\ell(\mathbf{b})$ *the sale price for* $\mathbf{b}$ *that gives the optimal profit among those sale prices that result in the sale of at least* ℓ *items, i.e.,*

$$\mathrm{opt}_\ell(\mathbf{b}) = b_{(i^*)} \;\; : \;\; i^* = \mathrm{argmax}_{i \geq \ell}\, i b_{(i)}.$$

If $\ell > n$*, we arbitrarily define* $\mathrm{opt}_\ell(\mathbf{b}) = 0$.

Definition 4.7 (Parameterized Random Sampling Optimal Price Auction, RSOP_ℓ**)** *The* RSOP_ℓ *auction works as follows:*

1. *Partition bids* $\mathbf{b}$ *uniformly at random into two sets: for each bid, with probability 1/2 put the bid in* $\mathbf{b}'$ *and otherwise* $\mathbf{b}''$*.*

2. *Let* $p' = \mathrm{opt}_\ell(\mathbf{b}')$ *and* $p'' = \mathrm{opt}_\ell(\mathbf{b}'')$*, the optimal fixed price thresholds that sell at least* ℓ *items for* $\mathbf{b}'$ *and* $\mathbf{b}''$*, respectively.*

3. *Use* p' *as a threshold for all bids in* $\mathbf{b}''$ *(i.e., all bids in* $\mathbf{b}''$ *of value below* p' *are rejected; all remaining bids win at price* p'*).*

4. *Use* p'' *as a threshold for all bids in* $\mathbf{b}'$*.*

Theorem 4.6 *There is an absolute constant* C*, such that for any* $\epsilon > 0$*,* $\mathrm{RSOP}_{\frac{m}{2}-\epsilon m}$ *is* $(1+\epsilon)$*-competitive against* $\mathcal{F}^{(m)}$*, with probability at least* $1 - e^{-C\epsilon^2 m}$*.*

Thus, a parameterized version of RSOP asymptotically (as m gets large) matches the profit of $\mathcal{F}^{(m)}$. Recall that $\mathcal{F}^{(m)}$, the optimal single-price auction that is required to have at least m winners, is the same as $\mathcal{F}$ when $\mathcal{F}$ chooses to sell at least m items. Thus $\mathcal{F}^{(m)}$ is the best auction when the auctioneer is required to use a single price and sell at least m items. An immediate corollary of the above theorem is that on any bid vector $\mathbf{b}$ such that (a) $\mathcal{F}(\mathbf{b}) = \mathcal{F}^{(m)}(\mathbf{b})$ and (b) with high probability RSOP and $\text{RSOP}_{\frac{m}{2}-\epsilon m}$ exhibit the same behavior on $\mathbf{b}$ (i.e., for $\mathbf{b}$ partitioned into $\mathbf{b}'$ and $\mathbf{b}''$, $\text{opt}_\ell(\mathbf{b}') = \text{opt}(\mathbf{b}')$ and $\text{opt}_\ell(\mathbf{b}'') = \text{opt}(\mathbf{b}'')$, where $\ell = \frac{m}{2} - \epsilon m$), we can conclude that with high probability $\text{RSOP}(\mathbf{b}) \geq \frac{\mathcal{F}(\mathbf{b})}{(1+\epsilon)}$.

It is easy to check that conditions (a) and (b) hold for the case of bids of bounded support, as the number of bidders gets large, yielding Theorem 4.5 as an immediate corollary of Theorem 4.6.

Thus, for a large class of inputs, RSOP achieves essentially optimal performance. For worst-case inputs, however, the constant in the competitiveness of RSOP is weak. It is quite clear from studying the proof of Theorem 4.4 (in Section 4.7) that the analysis given there is not tight and the constant bound we obtain on RSOP's competitiveness is very weak. However, it is not hard to see that RSOP can not be better than 4-competitive. For example, when $\mathbf{b}$ consist of two very high bids h and $h + \epsilon$, and all other bids are negligible, the expected profit of RSOP $= \mathcal{F}^{(2)}/4$. In the next section, we present an auction which achieves this bound: the RSPE auction that is 4-competitive.

4.4.2 The Random Sampling Profit Extraction Auction

We now show how the basic auction profit extractor can be used in combination with random sampling to obtain a profit maximizing auction. The analysis of this auction is much simpler than the previous sampling based auction and the result we obtain is a tight bound of four on the worst case competitive ratio.

Definition 4.8 (Random Sampling Profit Extraction Auction, RSPE**)** *The* Random Sampling Profit Extraction Auction *works as follows:*

1. *Partition bids* $\mathbf{b}$ *uniformly at random into two sets, resulting in bid vectors* $\mathbf{b}'$ *and* $\mathbf{b}''$*.*

2. *Compute* $F' = \mathcal{F}(\mathbf{b}')$ *and* $F'' = \mathcal{F}(\mathbf{b}'')$*, the optimal fixed price profits for* $\mathbf{b}'$ *and* $\mathbf{b}''$*,*

respectively.

3. *Compute the auction results by running ProfitExtract*$_{F''}$ *on* $\mathbf{b}'$ *and ProfitExtract*$_{F'}$ *on* $\mathbf{b}''$.

Because ProfitExtract$_{F'}$ and ProfitExtract$_{F''}$ are truthful on their respective partitions, we have,

Theorem 4.7 RSPE *is truthful.*

Next we show that RSPE is competitive.

Theorem 4.8 RSPE *is 4-competitive, and this bound is tight.*

Proof: In the special case that $F' = F''$ the auction profit is $F' + F'' \geq \mathcal{F}(\mathbf{b})$ and we are done. Otherwise, the auction profit is $\mathcal{R} = \min(F', F'')$. To see this, suppose that $F' < F''$. Then ProfitExtract$_{F''}$ on $\mathbf{b}'$ will reject all bids in $\mathbf{b}'$. However, ProfitExtract$_{F'}$ on $\mathbf{b}''$ will be able to achieve profit F'.

By definition, $\mathcal{F}^{(2)}$ on $\mathbf{b}$ sells to $k \geq 2$ bidders at price p for a profit of $\mathcal{F}^{(2)} = kp$. These k bidders, all with bid value at least p, are divided uniformly at random between $\mathbf{b}'$ and $\mathbf{b}''$. Let k' be the number of them in $\mathbf{b}'$ and k'' the number in $\mathbf{b}''$. As such, $\mathcal{F}(\mathbf{b}') \geq pk'$ and $\mathcal{F}(\mathbf{b}'') \geq pk''$. Therefore,

$$\frac{\min(\mathcal{F}(\mathbf{b}'), \mathcal{F}(\mathbf{b}''))}{\mathcal{F}^{(2)}(\mathbf{b})} \leq \frac{\min(pk', pk'')}{pk} = \frac{\min(k', k'')}{k}.$$

Thus, the ratio

$$\frac{\mathbf{E}[\mathcal{R}]}{\mathcal{F}^{(2)}} = \tfrac{1}{k}\sum_{i=1}^{k-1} \min(i, k-i)\tbinom{k}{i}2^{-k} = \frac{1}{2} - \tbinom{k-1}{\lfloor \frac{k}{2} \rfloor}2^{-k}.$$

This ratio achieves its minimum of 1/4 for $k = 2$ and $k = 3$. Therefore RSPE is 4-competitive. As k increases, the ratio approaches 1/2.

To see that the bound presented on the competitive ratio is tight, consider the case where $\mathbf{b}$ consists of two very high bids h and $h + \epsilon$, and all other bids are negligibly small. In this case $\mathcal{F} = \mathcal{F}^{(2)} = 2h$, whereas the expected profit of the RSPE auction is $h \cdot \mathbf{Pr}[\text{two high bids are split between } \mathbf{b}' \text{ and } \mathbf{b}''] = h/2 = \mathcal{F}/4$. □

The RSPE auction, as described, gets no profit from one of the partitions, and therefore it will lose at least half of the potential profit. An alternative is to pick a parameter $\gamma < 1$

and run $\text{ProfitExtract}_{\gamma F'}$ and $\text{ProfitExtract}_{\gamma F''}$. The competitive ratio of the resulting auction is $4/\gamma > 4$. However, when there are many winners in $\mathcal{F}$ it would be possible to do better than $\mathcal{F}/2$ with an appropriate setting of γ.

4.4.3 Random Sampling and Truthfulness

It is natural to consider the generality of this random sampling technique for constructing auctions. From the above example of RSOP and RSPE the following framework is apparent.

Definition 4.9 (General Random Sampling Auction) *Given two auctions parameterized by bid vector* $\mathbf{v}$*,* $\mathcal{A}'_{\mathbf{v}}$ *and* $\mathcal{A}''_{\mathbf{v}}$*:*

1. *Partition the bids into* $\mathbf{b}'$ *and* $\mathbf{b}''$*.*

2. *Run* $\mathcal{A}'_{\mathbf{b}''}$ *on* $\mathbf{b}'$ *and* $\mathcal{A}''_{\mathbf{b}'}$ *on* $\mathbf{b}$*.*

Theorem 4.9 *If for any* $\mathbf{v}$*,* $\mathcal{A}''_{\mathbf{v}}$ *and* $\mathcal{A}'_{\mathbf{v}}$ *are truthful then the general random sampling auction based on* $\mathcal{A}''_{\mathbf{v}}$ *and* $\mathcal{A}'_{\mathbf{v}}$ *is truthful.*

In the auctions we considered above, RSOP and RSPE, both the parameterized auction were the same, i.e., $\mathcal{A}'_{\mathbf{v}} = \mathcal{A}''_{\mathbf{v}}$. Recall that in RSOP$'$ half of the bidders were rejects, i.e., $\mathcal{A}'_{\mathbf{v}}$ was the auction that always output the empty allocation. There are other interesting combinations of random sampling auctions we can consider. For example, recall that the Bid-independent Optimal Price (BI_{opt}) auction uses a single sale price (proof is given in Section 8.3.1). Consider the sampling auction in which $\mathcal{A}'_{\mathbf{v}}$ is BI_{opt} and $\mathcal{A}''_{\mathbf{v}}$ is the mechanism that simulates BI_{opt} on $\mathbf{v}$ to compute its single sale price, p, and then does fixed price sale with price p. The result is a single priced mechanism that does not lose half the revenue due to a rejected sample.[3] This mechanism is actually very similar to RSOP and the proof of competitiveness of RSOP implies that it is also competitive.

We have seen that it is possible to use the fixed price mechanism and the profit extraction mechanism to obtain auctions with provably good performance. It is also conceivable that

[3]Note that there may be bidders who bid above the sale price that lose as this is a possible outcome of BI_{opt}.

k-Vickrey can be used in a similar fashion. In this auction $\mathcal{A}'_{\mathbf{v}}$ would examine $\mathbf{v}$ to determine how many items to sell in a k-Vickrey auction. It may do so by computing the number of winners in $\mathcal{F}$, $\mathcal{F}^{(2)}$, or perhaps the k that gives the highest revenue in k-Vickrey on $\mathbf{v}$. A preliminary analysis reveals that these techniques do not yield an auction which is better than either of the aforementioned RSOP or RSPE auctions. As such, we will not prove any performance bounds on this auction here.

4.5 Randomized Consensus of Revenue Estimates

In this section we will consider another way of using the basic auction profit extractor to obtain a profit maximizing auction. We first give an auction that is competitive for mass markets and then show how to modify it to be competitive in worst case. To motivate the idea behind these mechanisms, we start by considering several failed attempts. Consider the mechanism that

1. Computes $R = \mathcal{F}(\mathbf{b})$, then
2. Runs ProfitExtract_R on $\mathbf{b}$.

Since $R = \mathcal{F}(\mathbf{b})$, this obtains the optimal revenue on $\mathbf{b}$. However, because R is a function of all of $\mathbf{b}$, running ProfitExtract_R is no longer truthful.

To remedy this we might consider the bid-independent version of this. Recall that the truthfulness of ProfitExtract_R implies that there is a bid-independent function, pe_R, such that ProfitExtract_R is identically BI_{pe_R}. Now consider BI_f with f defined as:

$$f(\mathbf{b}_{-i}) = \text{pe}_{\mathcal{F}(\mathbf{b}_{-i})}(\mathbf{b}_{-i}).$$

This is very subtly different from taking $f(\mathbf{b}_{-i}) = \text{pe}_{\mathcal{F}(\mathbf{b})}(\mathbf{b}_{-i})$, however this latter auction is actually identical to the non-truthful auction already dismissed above.

What BI_f is doing is, for each bidder i, obtaining a bid-independent revenue estimate, $\mathcal{F}(\mathbf{b}_{-i})$, of the true optimal revenue, $\mathcal{F}(\mathbf{b})$, and trying to profit-extract it. Unfortunately, since these revenue estimates are all different, the profit extraction procedure does not actually achieve the desired results. In fact, the proof that no deterministic and symmetric auction is competitive shows that BI_f is not competitive. If we could instead bid-independently agree (i.e., have a consensus) on a revenue estimate at some value R, i.e., via a function $r(\cdot)$ such that $r(\mathbf{b}_{-i}) = R$ for all i, then the auction outcome would be identical to that of ProfitExtract$_R$. Thus, if R is less than but close to $\mathcal{F}(\mathbf{b})$ then we would obtain the near optimal revenue R.

The reason we can hope to be able to obtain this bid-independent consensus estimate for the mass market case is because we can bound $\mathcal{F}(\mathbf{b}_{-i})$ close to $\mathcal{F}(\mathbf{b})$. I.e., there is some constant ρ, $\mathcal{F}(\mathbf{b}_{-i}) \in [\mathcal{F}^{(2)}(\mathbf{b})/\rho, \mathcal{F}^{(2)}(\mathbf{b})]$. Intuitively, this is because for mass markets, each bidder only contributes to a small fraction of the total optimal auction profit. This motivates the following consensus estimation problem.

4.5.1 The Consensus Estimate Problem

In this section we study the following consensus problem.

Definition 4.10 *For a given $\rho > 1$ and $v > 0$, we say that a function g is a ρ-*consensus estimate *of v if*

1. *g is a* consensus*: for any w such that $v/\rho \leq w \leq v$, we have $g(w) = g(v)$.*

2. *$g(v)$ is a nontrivial lower bound on v, i.e., $0 < g(v) \leq v$.*

We call $g(v)$ the consensus value*.*

Definition 4.11 *The* payoff*, γ_g, for a function g is $\gamma_g(v) = g(v)$ if g is a ρ-consensus estimate on v and $\gamma_g(v) = 0$ otherwise.*

Intuitively, we would like to find a good consensus estimate for v, where the higher the consensus value, the higher the quality of the estimate.

It is easy to see that no deterministic function g is a ρ-consensus estimate for all $v > 0$. First, a simple induction shows that for a deterministic $g(v)$ to be a consensus for all positive v it would have to be a constant function. The only constant values that are lower bounds on any positive v are trivial, i.e., non-positive.

Definition 4.12 *The* consensus estimate problem *is, for any ρ, to give a distribution $\mathcal{G}$ on functions g such that for any v the expected payoff is large relative to v. That is, over choices of v, the worst case value of $\mathbf{E}[\gamma_g(v)]/v$ is large.*

4.5.2 Consensus Estimate Algorithm.

We describe a distribution $\mathcal{G}_c$ that works well in the following sense: for any v, with g from $\mathcal{G}_c$ we have $\mathbf{E}[g(v)] = \Omega(v)$, where the constant hidden by the Ω notation is a function of ρ. Our solution uses an additional parameter $c > \rho$. The value of c is chosen as a function of ρ to maximize the quality of the estimate. Consider the following function g_u^c:

$$g_u^c(v) = v \text{ rounded down to nearest } c^{j+u} \text{ for integer } j.$$

Remark. The definition of g_u^c implies that for any v, $\frac{v}{c} \leq g_u^c(v) \leq v$. Thus if g_u^c is a consensus for v then it is a consensus estimate with value within a factor of c from v.

We define $\mathcal{G}_c$ as a distribution of functions of the form g_U^c with U chosen uniformly on $[0,1]$. We repeatedly make use of the following result.

Lemma 4.10 *For g from $\mathcal{G}_c$, $g(v)$ is distributed identically to $c^{U'}v/c$ for U' uniform on $[0,1]$.*

Proof: Consider a random variable $Y = \log_c g(v)$ and let $t = \log_c v - 1$. Then $\mathbf{Pr}[Y \leq t + x] = \mathbf{Pr}[U' \leq x]$ and therefore Y is uniformly distributed between t and $t+1$. Thus, $g(v)$ is identical to $c^{U'}v/c$. □

For U uniform $[0,1]$, the random variable c^U satisfies $\mathbf{Pr}\left[c^U \leq z\right] = \mathbf{Pr}[U \leq \log_c z] = \log_c z = \frac{\ln z}{\ln c}$. The probability density function for c^U is $\pi(x) = 1/(x \ln c)$ for $1 \leq x < c$. To see this, note that $\mathbf{Pr}\left[c^U \leq z\right]$ is $\int_1^z \frac{1}{(x \ln c)} dx = \frac{\ln z}{\ln c}$.

Next we bound the probability that g is a consensus.

Lemma 4.11 *For g from $\mathcal{G}_c$, the probability that g is a consensus estimate is $1 - \log_c \rho$.*

Proof: g is a consensus estimate for v if $g(v) \leq \frac{v}{\rho}$. Using Lemma 4.10, we get

$$\begin{aligned} \mathbf{Pr}\left[g(v) \leq \tfrac{v}{\rho}\right] &= \mathbf{Pr}\left[c^U v/c \leq \tfrac{v}{\rho}\right] = \mathbf{Pr}\left[c^U \leq \tfrac{c}{\rho}\right] \\ &= \log_c(c/\rho) = 1 - \log_c \rho. \end{aligned}$$

□

In the application to auctions, the value of ρ is fixed and we choose c to maximize the expectation of γ. Lemma 4.11 implies $\mathbf{E}[\gamma_g(v)] \geq \frac{v}{c}(1 - \log_c \rho)$. The following theorem gives a better bound.

Theorem 4.12 *For g from $\mathcal{G}_c$ defined above, for all v,* $\mathbf{E}[\gamma_g(v)] = \frac{v}{\ln c}\left(\frac{1}{\rho} - \frac{1}{c}\right)$.

Proof: By Lemma 4.10, $g(v)$ is distributed as $c^U v/c$ for U uniform on $[0, 1]$. Therefore,

$$\begin{aligned}\mathbf{E}[\gamma(v)] &= \frac{v}{c}\int_1^{c/\rho} x \cdot \pi(x)dx + \int_{c/\rho}^{c} 0 \cdot \pi(x)dx \\ &= \frac{v}{c}\int_1^{c/\rho} \frac{1}{\ln c}dx = \frac{v}{\ln c}\left(\frac{1}{\rho} - \frac{1}{c}\right).\end{aligned}$$

□

Note that for a fixed ρ, one can choose the value of c that maximizes $\mathbf{E}[\gamma_g(v)]/v$.

4.5.3 Consensus Estimates with One Random Bit.

Given the above consensus estimate solution that uses a random real number chosen uniformly from $[0, 1]$ and the fact that no purely deterministic consensus estimate exists, it is natural to ask how much randomness is necessary. We show how to give a consensus estimate with only one random bit. Choose a constant $c' > \rho$ and let $c = c'^2$. Pick the value of u from $\{0, 1/2\}$ with equal probability and use function g_u^c as defined above. Note that for these values of u, we round the revenue estimates to even and odd powers of c', respectively. Since $c' > \rho$, for any value of v, at most one of these values can be in the interval $[v/\rho, v]$ and therefore the revenue estimates agree with probability of at least 1/2 and the expected payoff is at least $v/(2c'\rho)$. This gives the following lemma:

Lemma 4.13 *The consensus estimate solution with g chosen from $\{g_0^c, g_{\frac{1}{2}}^c\}$ with equal probability and $c = \left(\rho + \frac{\epsilon}{2\rho}\right)^2$ gives $\mathbf{E}[\gamma_g(v)] = v/(2\rho^2 + \epsilon)$.*

4.5.4 CORE *Auctions*

In the mass market case, removing a bid does not change $\mathcal{F}^{(2)}$ much.

Lemma 4.14 *If $\mathcal{F}^{(2)}(\mathbf{b})$ has $k \geq 3$ winners, then for any i,*

$$\tfrac{k-1}{k}\mathcal{F}^{(2)}(\mathbf{b}) \leq \mathcal{F}^{(2)}(\mathbf{b}_{-i}) \leq \mathcal{F}^{(2)}(\mathbf{b}).$$

When the number of winners in $\mathcal{F}^{(2)}(\mathbf{b})$ is $k \geq 3$, Lemma 4.14 allows us to estimate $\mathcal{F}^{(2)}(\mathbf{b})$ from $\mathcal{F}^{(2)}(\mathbf{b}_{-i})$. In this case, we can use $\rho = \frac{3}{2} \geq \frac{k}{k-1}$ in combination with our ρ-consensus estimate solution from above as follows.

Definition 4.13 (CORE_c**)** *Let $g(\cdot)$ be a function picked from $\mathcal{G}_c$. Let $r(\cdot)$ be a function from masked bid vectors to reals defined by $r(\mathbf{b}_{-i}) = g(\mathcal{F}^{(2)}(\mathbf{b}_{-i}))$. The mass market auction,* CORE_c*, is the bid-independent auction given by function f_r:*

$$f_r(\mathbf{b}_{-i}) = \mathrm{pe}_{r(\mathbf{b}_{-i})}(\mathbf{b}_{-i}).$$

Since $g(\cdot)$, and therefore $r(\cdot)$, is chosen from a probability distribution, the auction is randomized. Intuitively, $r(\cdot)$ estimates the optimal revenue. The following results are immediate corollaries of Lemma 4.11, Theorem 4.12, Lemma 4.14, and the Taylor expansion of $\log(1 + 1/x) = \Theta(1/x)$.

Lemma 4.15 *Let $c > m/(m-1)$. For $\mathbf{b}$ with $m \geq 3$ winners in $\mathcal{F}^{(2)}$,* CORE_c *achieves consensus with probability*

$$1 - \log_c \tfrac{m}{m-1} = 1 - \Theta(1/m).$$

Theorem 4.16 *Let $c > m/(m-1)$. When $\mathcal{F}^{(2)}$ sells $m \geq 3$ items, the $CORE_c$ auction is $\beta(m)$-competitive with*

$$\beta(m) = \frac{\ln c}{1 - \frac{1}{c} - \frac{1}{m}}.$$

In the limit $CORE_c$ is $\lim_{m\to\infty} \beta(m) = c \ln c/(c-1)$ competitive.

The choice of c in the CORE_c auction is crucial. In order for consensus to work on bids $\mathbf{b}$ on which $\mathcal{F}$ sells at least m items, we need $c > m/(m-1)$. Otherwise there exists a set of bids such that no consensus is achievable.

4.5.5 Convex Combinations of Auctions

The mass market auction CORE_c is competitive on bids such that the optimal auction, $\mathcal{F}^{(2)}$, has at least three winners. We now show how to obtain an auction that is competitive with $\mathcal{F}^{(2)}$ on all inputs.

Let $\mathcal{A}'$ and $\mathcal{A}''$ be auctions such that $\mathcal{A}'$ is β'-competitive on $\mathbf{b} \in \mathcal{B}'$ and $\mathcal{A}''$ is β''-competitive on $\mathbf{b} \in \mathcal{B}''$. Consider the auction $\mathcal{A}$ that is a "convex combination" of $\mathcal{A}'$ and $\mathcal{A}''$: With probability p, $\mathcal{A}$ runs $\mathcal{A}'$ and otherwise $\mathcal{A}$ runs $\mathcal{A}''$. The following result is straight-forward.

Lemma 4.17 *$\mathcal{A}$ is $\max(\beta'/p, \beta''/(1-p))$-competitive on $\mathbf{b} \in \mathcal{B}' \cup \mathcal{B}''$. For the optimal choice of p, $\mathcal{A}$ is $(\beta' + \beta'')$-competitive on $\mathbf{b} \in \mathcal{B}' \cup \mathcal{B}''$.*

Note that the competitive ratio of $\mathcal{A}$ may be better than the lemma guarantees.

Recall that, by definition, $\mathcal{F}^{(2)}(\mathbf{b})$ has at least two winners. Let $\mathcal{B}_{3+}$ denote the set of all $\mathbf{b}$ such that $\mathcal{F}^{(2)}(\mathbf{b})$ has at least three winners and let $\mathcal{B}_2$ denote the set of all $\mathbf{b}$ such that $\mathcal{F}^{(2)}(\mathbf{b})$ has exactly two winners. As we have shown, the CORE_c auction is competitive on $\mathcal{B}_{3+}$. We obtain an auction competitive for all bids by combining the CORE_c auction with the 1-item Vickrey auction. Since the Vickrey auction sells to the highest bidder at the price equal to the second highest bid, we have the following result.

Lemma 4.18 *The Vickrey auction is 2-competitive on $\mathbf{b} \in \mathcal{B}_2$.*

4.5.6 CORE *in Worst Case*

Definition 4.14 *The* CORE *auction is a convex combination of Vickrey (with probability p) and the the mass market* CORE_c *auction (with probability $1-p$).*

To get a tighter analysis than that implied directly from Theorem 4.12 and Lemma 4.17 we utilize the observation that the Vickrey auction is k-competitive on $\mathcal{B}_{3+}$. Therefore if we run the Vickrey auction with probability p, this adds $p\mathcal{F}^{(2)}(\mathbf{b})/k$ to the expected revenue of the $\mathcal{B}_{3+}$ case. Judicious choices of p and c give the following result:

Theorem 4.19 *For an appropriate choice of c and p, the* CORE *auction is 3.39-competitive against $\mathcal{F}^{(2)}(\mathbf{b})$.*

Proof: Case 1 ($\mathbf{b} \in \mathcal{B}_{3+}$): Theorem 4.16 implies that if the CORE$_c$ auction is selected, the expected revenue is at least $\frac{\mathcal{F}^{(2)}(\mathbf{b})}{\ln c}\left(\frac{k-1}{k} - \frac{1}{c}\right)$. Furthermore, if the Vickrey auction is selected, the expected revenue is at least $\frac{\mathcal{F}^{(2)}(\mathbf{b})}{k}$. Thus the total expected revenue is at least

$$\mathcal{F}^{(2)}(\mathbf{b})\left(\frac{p}{k} + \frac{1-p}{\ln c}\left(1 - \frac{1}{k} - \frac{1}{c}\right)\right).$$

Case 2 ($\mathbf{b} \in \mathcal{B}_2$): The expected revenue for $\mathbf{b}$ is at least $\mathcal{F}^{(2)}(\mathbf{b})\frac{p}{2}$ due to the Vickrey auction.

We pick p and c to balance the Case 1 and Case 2 revenue. Numeric simulation shows that $c = 2.0$ and $p = 0.59$ is a near-optimal choice. This choice gives a competitive ratio of 3.39. □

This auction has several interesting properties. One property is that in the "normal" case, i.e., when the revenue estimates agree or when the Vickrey auction is used, the outcome of the auction is a single sale price with the property that every bidder that bid above this price wins. This is because the basic auction profit extractor is always single priced such that all bidders bidding above the sale price win. Unintuitively, in the case where consensus is not achieved CORE has at most two sale prices. These properties are described in greater detail and with proof in Chapter 6.

4.5.7 Random Reals vs. Random Bits.

The CORE auction of the previous section needs to select a real-valued random variable for selecting the consensus function. To combine this auction with the Vickrey auction, we use another random real. Consider a more realistic model of computation that does not allow infinite-precision reals. In this case our random variables must be rational. The CORE approach easily adapts to such a model.

Using Lemma 4.13, one can show the following result.

Theorem 4.20 *For any $\epsilon > 0$, and with appropriate parameter settings, a* CORE *auction that only uses two random bits is $(6+\epsilon)$-competitive against $\mathcal{F}^{(2)}$.*

Proof: If $\mathbf{b} \in \mathcal{B}_2$, the competitive ratio is at least four as we chose to run the 2-competitive Vickrey auction with probability 1/2. For the rest of the proof we assume $\mathbf{b} \in \mathcal{B}_{3+}$; this case determines the competitive ratio.

For $\mathbf{b} \in \mathcal{B}_{3+}$, the revenue in the case when the Vickrey auction is selected is $\mathcal{F}^{(2)}(\mathbf{b})/k$. In the other case, our one random bit consensus estimate algorithm with $c' \approx \frac{3}{2}$ and $\rho = k/(k-1)$ gets an expected consensus value of $\mathcal{F}^{(2)}(\mathbf{b})/(2c'\rho) = \mathcal{F}^{(2)}(\mathbf{b})\frac{k-1}{3k}$. Thus the expected profit is approximately

$$\frac{\mathcal{F}^{(2)}(\mathbf{b})}{2}\left(\frac{1}{k} + \frac{k-1}{3k}\right) \geq \frac{\mathcal{F}^{(2)}(\mathbf{b})}{6}.$$

□

By using more random bits, we get a better discrete approximation of the continuous distribution of U and of the optimal value of p of the previous section. With sufficiently many bits, we can get arbitrary close to the 3.39 competitive ratio.

4.6 Randomization via Weighted Pairing

All truthful auctions we have introduced so far have had only one or two sale prices. In this section we describe a multi-price auction. This auction is not competitive with $\mathcal{F}$ even in the case where $\mathcal{F}$ has many winners. However, an extension of this auction gave the first solution to the *online auction problem* with a sublograthmic competitive ratio [9]. In the online auction problem the auctioneer is presented the bids one at a time and must make a decision for each bidder before seeing any of the subsequent bids.

Definition 4.15 (Weighted Pairing) *The* weighted pairing *auction is the bid-independent auction,* BI_f*, with function* f *defined as follows:*

$$f(\mathbf{b}) = b_i \textit{ with probability } \frac{b_i}{\sum_j b_j}$$

Thus, to determine if bidder i wins the auction and at what price, pick a bid b_j from $\mathbf{b}_{-i}$ with probability proportional to the value of b_j, i.e., $b_j/(\mathcal{T} - b_i)$. This pairs b_i with b_j. If $b_j \leq b_i$, bidder i wins at cost b_j, otherwise bidder i loses. Note that the result of this selection for i does not affect the auction outcome for bidder j.

To understand the intuition behind this auction, consider a related *random pairing* auction. Assume that n is even and pair bidders at random, independent of their bid values. For each pair, conduct a 1-item Vickrey auction, that is, for a pair (b_i, b_j) with $b_i < b_j$, i loses and j wins at cost b_i.

Compared to the random pairing auction, a bidder in the weighted pairing auction is less likely to win, but when a bidder wins they are likely to pay more. It turns out that high bidders are still likely to win the auction, and the benefit of high bidders paying higher prices outweighs the benefit of more low bidders winning the auction. In particular, Theorems 4.21 and 4.22 imply that for the weighed pairing auction, in the worst case expected revenue is proportional to $\mathcal{T}/\log h$. For the (unweighted) pairing auction, we have shown that the worst case expected revenue is proportional to $\mathcal{T}/\sqrt{h}$. We do not include the latter result because it is dominated by the former.

Next we prove the following result.

Theorem 4.21 *If $4h \leq \mathcal{T}$, then for the weighted pairing auction* $\mathbf{E}[\mathcal{R}] = \Omega(\mathcal{T}/\log h)$.

Proof: Partition the n bids into $\log h$ bins as before such that bin j contains only bids in the interval $[2^{j-1}, 2^j]$. Recall that an important property of these bins is that bids in the same bin are within factor of two from each other. Let S_j be the sum of the bids in bin j. Note that the sum of the contents of bins that contain only one bid is upper bounded by $\sum_{j=1}^{\log h} 2^j = 2h - 1 < \mathcal{T}/2$. We will ignore such bins in our analysis below. Consider all bins with two or more bids, and let $\mathcal{T}'$ be the sum of bids in these bins. We have $\mathcal{T}' > \mathcal{T}/2$.

For each bin j, we look at pairings of bids that are both in j and we bound the expected revenue due to such pairings. First, the probability that a bid i in bin j is paired with another bid in bin j is $(S_j - b_i)/(\mathcal{T} - b_i) > S_j/(3\mathcal{T})$, since bin j contains at least two bids and $S_j - b_i \geq S_j/3$.

Let $b'_1, \ldots, b'_k$ be the values of bids in bin j in the increasing order. Given that a bid b'_i from the bin is paired with another bid in the same bin, the probability that the bid wins is at least

$$\frac{i-1}{2(k-i)+(i-1)} > \frac{i-1}{2k}.$$

This comes from assuming the worst case, namely that all bids below bid i are at value 2^{j-1} and all bids above bid i are at value 2^j. The expected number of bids in bin j that win when paired with other bids from the bin is at least

$$\sum_{i=1}^{k} \frac{i-1}{2k} = \frac{k-1}{4}.$$

The smallest bid in bin j has the value of at least $S_j/(2k)$. Let $\mathcal{R}_j$ be the revenue generated by bids in bin j being paired with other bids in bin j. We have:

$$\begin{aligned}\mathbf{E}[\mathcal{R}_j] &\geq \frac{S_j}{3\mathcal{T}} \cdot \frac{S_j}{2k} \cdot \frac{k-1}{4} \\ &\geq \frac{(k-1)S_j^2}{24k\mathcal{T}}\end{aligned}$$

Since $k \geq 2$, we have $(k-1)/k \geq 1/2$ so

$$\mathbf{E}[\mathcal{R}_j] \geq \frac{S_j^2}{48\mathcal{T}}$$

We are interested in $\mathbf{E}[\mathcal{R}] \geq \sum \mathbf{E}[\mathcal{R}_j] \geq \sum \frac{S_j^2}{48\mathcal{T}}$. Since S_j's sum up to $\mathcal{T}'$, the sum is minimized when all $\log h$ of the S_j's are equal to $\mathcal{T}'/\log h$. In this case,

$$\mathbf{E}[\mathcal{R}] \geq \frac{\mathcal{T}'^2(\log h)}{48\mathcal{T}\log^2 h}$$

Since $\mathcal{T}' > \mathcal{T}/2$, we have

$$\mathbf{E}[\mathcal{R}] \geq \frac{\mathcal{T}}{192\log h}$$

Thus, $\mathbf{E}[\mathcal{R}] = \Omega(\mathcal{T}/\log h)$. □

Note that the constant here is 1/192 which does not seem too good. However, the analysis was very loose in contributions to $\mathcal{R}$ that it considered.

Lower bounds on revenues of the weighted pairing and the random sampling optimal price auction, RSOP, and CORE, stated in terms of $\mathcal{T}$, are the same, $\Omega(\mathcal{T}/\log h)$. The next result shows that RSOP and CORE perform better than weighted pairing when compared with $\mathcal{F}$.

We show that the expected revenue of the weighted pairing auction is $\Omega(\mathcal{F}/\sqrt{\log h})$, and that this bound is tight in the worst case. This implies that the weighted pairing auction is not constant competitive.

We use the following notation. We denote the expected revenue of the weighted pairing auction by $\mathbf{E}[\mathcal{R}]$, let $k = \log h$ and $s = \sqrt{\log h}$. First we give bound the revenue from below.

Theorem 4.22 *If $\mathcal{F} \geq 2h$ then $\mathbf{E}[\mathcal{R}] = \Omega(\mathcal{F}/\sqrt{\log h})$ and this bound is tight.*

Proof: If $\mathcal{F} \leq \mathcal{T}/s$, then since $\mathbf{E}[\mathcal{R}] = \Omega(\mathcal{T}/k)$ we have $\mathbf{E}[\mathcal{R}] = \Omega(\mathcal{T}/s)$.

Assume $\mathcal{F} > \mathcal{T}/s$. Partition the bids over buckets, with bucket i containing bids in the range $[2^i, 2^{i+1})$. Let M be the set of bids in a bucket with the largest total bid value and assume that M is defined by the value range $[t, 2t)$. We show that the revenue due only to M is big. The assumption that $\mathcal{F} \geq 2h$ implies that $|M| \geq 2$. Recall that $|M|\, t = \Omega(\mathcal{F})$. Let M' be the highest $\lfloor |M|/2 \rfloor$ elements of M and let $M'' = M \setminus M'$.

The probability that an item in M' is paired up with an item in M'' is $\Omega(\mathcal{T}/\mathcal{F}) = \Omega(1/s)$, and the expected revenue is $\Omega(|M'|\, t/s) = \Omega(\mathcal{F}/s)$.

The above argument considers only contributions of items in M. Somewhat surprisingly, the resulting bound is tight up to a constant factor. We give an example where the expected revenue $\mathbf{E}[\mathcal{R}]$ of the weighted pairing auction is $O(\mathcal{T}/s)$. The bid values are 2^i for $i = 1, \ldots, k$. For each $i = 1, \ldots, k-1$ the number of bids of this value is 2^{k-i} and the total value of such bids is 2^k. For the last value, 2^k, there are s bids of this value for the total of $s2^k = \mathcal{F}$. Note that $\mathcal{T} = 2^k(s + k - 1)$.

We show that $\mathbf{E}[\mathcal{R}] = O(2^k)$ by first showing that the expected contribution of all bids of value less then 2^k to $\mathbf{E}[\mathcal{R}]$ is $O(2^k)$ and then showing that the contribution of all bids of value 2^k is $O(2^k)$.

For $i < k$, the probability that a bid of value 2^i wins is less than

$$\frac{i2^k - 2^i}{T - 2^i} \leq \frac{i2^k}{T} \leq \frac{i}{k}.$$

The expected contribution of a winning bid is less than

$$\frac{1}{i}\sum_{j=1}^{i} 2^j < \frac{1}{i}2^{i+1}.$$

This all bids of value 2^i contribute less than $2^{k+1}/k$. Summing over $i = 1, \ldots, k-1$, we conclude that the expected contribution of the corresponding bids is less than $2 \cdot 2^k$.

Next we consider the bids of value 2^k. The probability that such a bid wins is at least one and the expected revenue of a winning bid is at least

$$\frac{s}{k+s}2^k + \frac{1}{k+s}\sum_{i=1}^{k-1} 2^i < \frac{s+1}{k+s}2^k = \frac{2^k}{s}.$$

Thus the total contribution of such bids is less than 2^k. □

4.7 Analysis of RSOP

Below we prove that the RSOP auction presented in section 4.4.1 is competitive [25] (an early and weaker version of this proof appeared in [26]). Moreover, we show that, as m gets large, the expected profit of RSOP_ℓ tends to $\mathcal{F}^{(m)}$ for $\ell = \lceil \frac{m}{2}(1-\delta) \rceil$. This implies that for any $\epsilon > 0$, there is an m' such that for any $m \geq m'$, RSOP_ℓ is $(1+\epsilon)$-competitive against $\mathcal{F}^{(m)}$. The proof we give below is from [25] (a weaker analysis was originally given in [26]).

4.7.1 Preliminaries

We begin our analysis of RSOP_ℓ with some preliminary lemmas. We will use the following definitions:

- For any value v, let $F_v = vn_v$ (resp. $F'_v = vn'_v$ and $F''_v = vn''_v$) where n_v is the number of bids in $\mathbf{b}$ (resp. $\mathbf{b}'$ and $\mathbf{b}''$) greater than or equal to v. Thus, F_v is the profit from using v as a threshold for $\mathbf{b}$.

- Let R be the random variable representing the profit of the auction RSOP_ℓ (where the input $\mathbf{b}$ is implied by the context). Recall from the definition of the RSOP_ℓ auction that $p' = \mathrm{opt}_\ell(\mathbf{b}')$ and $p'' = \mathrm{opt}_\ell(\mathbf{b}'')$ are the thresholds used for $\mathbf{b}''$ and $\mathbf{b}'$ respectively. Thus, the profit of the auction is $R = F'_{p''} + F''_{p'}$.

- Let E_α be the event

$$E_\alpha : \quad R \geq (1-\alpha)\left(F_{p'} + F_{p''}\right). \tag{4.1}$$

  for $0 \leq \alpha \leq 1$. This event holding for small α indicates a high profit, a significant fraction of the total profit achievable using p' and p'' as thresholds.

- Let $\ell = \lceil \frac{m}{2}(1-\delta) \rceil$. Suppose that $\mathrm{opt}_m(\mathbf{b}) = b_{(k)}$, the k-th largest bid in $\mathbf{b}$. Define H to be the event that (a) there are at least ℓ bids in $\mathbf{b}'$ that are at least $b_{(k)}$ and (b) there are at least ℓ bids in $\mathbf{b}''$ that are at least $b_{(k)}$. Note that if H does not occur then RSOP_ℓ may be unable to pick a threshold of magnitude similar to $b_{(k)}$ from at least one of the partitions.

The key lemma we use is the following:

Lemma 4.23 *Let $0 < \delta < 1$ and $\ell = \lceil \frac{m}{2}(1-\delta) \rceil$. Then, for any* $\mathbf{b}$*,* RSOP_ℓ *satisfies*

$$\mathbf{Pr}\left[R \geq \frac{1-\alpha}{\alpha}\mathcal{F}^{(m)}(\mathbf{b})\right] \geq \mathbf{Pr}[E_\alpha \cap H]$$

and thus

$$\mathbf{E}[R] \geq \frac{1-\alpha}{\alpha}\mathcal{F}^{(m)}(\mathbf{b})\mathbf{Pr}[E_\alpha \cap H]\,.$$

Proof: Suppose that $\mathcal{F}^{(m)}(\mathbf{b}) = kb_{(k)}$. By definition, $\mathrm{opt}_\ell(\mathbf{b}')$ returns the p' that maximizes $F'_{p'}$ conditioned on there being ℓ bids that are at least p' (likewise for $\mathbf{b}''$ and p''). Notice that, conditioned on H, there are at least ℓ elements in each of $\mathbf{b}'$ and $\mathbf{b}''$ that are at least $b_{(k)}$. Thus, we can conclude that

$$F'_{b_{(k)}} \leq F'_{p'} \quad \text{and} \quad F''_{b_{(k)}} \leq F''_{p''}.$$

We also have

$$\mathcal{F}^{(m)}(\mathbf{b}) = F_{b_{(k)}} = F'_{b_{(k)}} + F''_{b_{(k)}},$$

and thus

$$\mathcal{F}^{(m)}(\mathbf{b}) \leq F'_{p'} + F''_{p''}. \tag{4.2}$$

Event E_α holding, equation (4.1) allows us to lower bound the profit R as

$$\begin{aligned} R &\geq (1-\alpha)[F_{p'} + F_{p''}] \\ &= (1-\alpha)[F'_{p'} + F''_{p'} + F'_{p''} + F''_{p''}] \\ &= (1-\alpha)[R + F'_{p'} + F''_{p''}] \end{aligned}$$

which by equation (4.2) gives

$$R \geq (1-\alpha)[R + \mathcal{F}^{(m)}].$$

Event H holding, we rearrange terms to obtain

$$R \geq \tfrac{1-\alpha}{\alpha}\mathcal{F}^{(m)}.$$

□

This lemma reduces our problem to studying the events E_α and H. To do so, for any j, we define B_j to be the j highest bids in $\mathbf{b}$ (i.e., $\{b_{(1)}, \ldots b_{(j)}\}$), and let $n'(B_j)$ be the number of these bids that are in $\mathbf{b}'$.

Definition 4.16 *Given $c: \ 0 < c < 1$, we say that B_j is c-*good *if*

$$\lceil cj \rceil \leq n'(B_j) \leq j - \lceil cj \rceil.$$

*Otherwise, B_j is c-*bad.

We prove that B_j is likely to be c-good using the following version of the Chernoff bound:

Theorem 4.24 (see e.g. [36], page 70) *Let X_i, $1 \leq i \leq n$ be independent Bernoulli trials such that for all i, $\mathbf{Pr}[X_i = 1] = 1/2$. Then for $X = \sum_{1 \leq i \leq n} X_i$, and $0 < \delta \leq 1$,*

$$\mathbf{Pr}\Big[X < (1-\delta)\frac{n}{2}\Big] < e^{-\frac{\delta^2 n}{4}}.$$

Since the partition of $\mathbf{b}$ into the two subvectors $\mathbf{b}'$ and $\mathbf{b}''$ is done by flipping a fair coin for each bid, we can conclude from the Chernoff bound that

$$\mathbf{Pr}[B_j \text{ is } c\text{-bad}] \leq 2e^{-\frac{(1-2c)^2 j}{4}}.$$

Thus, we can conclude by a simple union bound that

$$\mathbf{Pr}[B_j \text{ is not } c\text{-good for some } j > t\ ] \leq \sum_{j \geq t} 2e^{-\frac{(1-2c)^2 j}{4}}$$

and therefore we have

Lemma 4.25

$$\mathbf{Pr}[B_j \textit{ is not c-good for some } j > t] \leq \frac{2e^{-\frac{(1-2c)^2 t}{4}}}{1 - e^{-(1-2c)^2/4}}.$$

4.7.2 *Analysis of* RSOP

We are now ready to proceed with the analysis of RSOP_ℓ. First we show that as m increases, the competitive ratio of the auction approaches one.

Theorem 4.26 *Let $\delta:\ 0<\delta<1$ be a constant and let $\ell = \lceil \frac{m}{2}(1-\delta)\rceil$. Then in the limit as $m\to\infty$, RSOP_ℓ is $(1+\frac{\delta}{2})/(1-\frac{\delta}{2})$ competitive against $\mathcal{F}^{(m)}$.*

There is an absolute constant $C>0$ such that

$$\mathbf{Pr}\left[\text{RSOP}_\ell \geq \frac{(1-\frac{\delta}{2})}{(1+\frac{\delta}{2})}\mathcal{F}^{(m)}\right] \geq 1-e^{-C\delta^2 m}.$$

Proof: Fix $\epsilon = \delta/2$. An immediate corollary of Lemma 4.25 is that

$$\lim_{m\to\infty} \mathbf{Pr}\big[B_j \text{ is } \tfrac{1}{2}(1-\epsilon)\text{-good for all } j > \tfrac{m}{2}(1-2\epsilon)\big] = 1-o(1). \tag{4.3}$$

Thus, with probability $1-o(1)$, $F'_{p'} \leq \frac{1}{2}(1+\epsilon)F_{p'}$ and $F''_{p''} \leq \frac{1}{2}(1+\epsilon)F_{p''}$, and thus

$$\mathbf{Pr}\left[E_{\frac{1}{2}(1+\epsilon)}\right] = 1-o(1). \tag{4.4}$$

Finally, we show that $\mathbf{Pr}[H] = 1-o(1)$. As before, we assume that $\text{opt}_\ell(\mathbf{b})$ is the k-th largest bid $b_{(k)}$. We also assume, without loss of generality, that $b_{(k)}$ is in $\mathbf{b}'$, and that $b_{(k')}$ is the smallest bid larger than $b_{(k)}$ that is in $\mathbf{b}''$. Let G be the event that there is at least one element $b_{(i)}$ with $m(1-\epsilon)\leq i\leq m$ in each of $\mathbf{b}'$ and $\mathbf{b}''$. Then

$$\lim_{m\to\infty}\mathbf{Pr}[G] = \lim_{m\to\infty}\left(1-2^{1-\epsilon m}\right) = 1-o(1).$$

Thus, with probability $1-o(1)$,

$$m(1-\epsilon)\leq k' < k. \tag{4.5}$$

Moreover, from (4.3), we can conclude that with probability at least $1-o(1)$, there are at least $\frac{m}{2}(1-\epsilon) > \ell$ bids above $b_{(k)}$ in $\mathbf{b}'$. Also, from (4.3) and (4.5), we can conclude that there are at least $\frac{m}{2}(1-\epsilon)(1-\epsilon) > \ell$ bids above $b_{(\ell)}$ in $\mathbf{b}''$. Thus,

$$\mathbf{Pr}[H] = 1-o(1). \tag{4.6}$$

Finally, from Equations (4.4) and (4.6) and Lemma 4.23, we have

$$\mathbf{Pr}\left[R \geq \frac{(1-\epsilon)}{(1+\epsilon)}\mathcal{F}^{(m)}\right] = 1 - o(1),$$

where R is the profit of RSOP_ℓ. □

Next we show that RSOP is β-competitive for a (relatively large) constant β.

Theorem 4.27 *There is a constant β such that* RSOP *is β-competitive.*

Proof: Suppose that $\mathrm{opt}_2(\mathbf{b})$ is the k-th largest bid. We restrict our attention only to partitions of the bids in which the highest bid $b_{(1)}$ is in one subvector, without loss of generality, in $\mathbf{b}'$, and both $b_{(2)}$ and $b_{(k)}$ are in the other subvector. This event, which we shall denote G, occurs with probability 1/4.

We claim that

$$\mathbf{Pr}\left[E_{\frac{19}{20}}\right] \geq 0.05. \tag{4.7}$$

Indeed, an immediate corollary of Lemma 4.25 is that

$$\mathbf{Pr}\left[B_j \text{ is } \tfrac{1}{20}\text{-good for all } j > 20\right] \geq 0.8.$$

Moreover, if event G holds, all $j \leq 20$ are $\frac{1}{20}$-good, since $b_{(1)}$ is in $\mathbf{b}'$ and $b_{(2)}$ is in $\mathbf{b}''$. Thus, with probability at least $0.8 - \mathbf{Pr}[\neg G] = 0.8 - 0.75 = 0.05$, all B_k are $\frac{1}{20}$-good and event G holds, and thus, in particular, event $E_{\frac{19}{20}}$ holds.

For the case $m = 2$, event G implies event H. Thus, from (4.7) and Lemma 4.23,

$$\mathbf{Pr}\left[R \geq \frac{1}{19}\mathcal{F}^{(2)}\right] \geq 0.05,$$

and thus RSOP is 380-competitive.

□

Chapter 5

LIMITED SUPPLY

Up to this point, we have studied the unlimited supply case where the auction can potentially sell an item to every bidder. In this section, we consider the case where the number of items available for sale is bounded. This case is typical for physical goods markets. As in the unlimited supply case, we assume there is *free disposal*. That is, there is no cost to the auctioneer for unsold items. We denote the number of items available by k. As before, the seller wishes to maximize profit and is not required to sell all the items. The definitions of truthful and competitive auctions, stated for the unlimited supply case, also apply to the bounded supply case. We denote by $\mathcal{F}^{(m,k)}$ the profit for the optimal single price auction that sells at least m and at most k items. It is this quantity that we wish to be competitive with.

Obviously, the unlimited supply case is a special case of the limited supply case, i.e, take $k = n$. We now show the less obvious result that the limited supply auction problem reduces to the (unlimited supply) basic auction problem. To reduce the bounded supply case to the unlimited supply case, we can simply ignore, i.e., reject, all but the highest k bidders and run the unlimited supply auction on the remaining bids. We note that in order for this to be truthful we need to make sure that none of these k bidders win at a price lower than the highest ignored bid.

Definition 5.1 (Limited Supply Variant) *Given an unlimited supply auction $\mathcal{A}$, we define the limited supply variant of $\mathcal{A}$, $\mathcal{A}_k$, as follows.*

1. *Simulate k-Vickrey on $\mathbf{b}$ to get allocation $\mathbf{x}^{\mathcal{V}}$ and prices $\mathbf{p}^{\mathcal{V}}$.*

2. *Let $\mathbf{b}^{\mathcal{A}}$ be given by $b_i^{\mathcal{A}} = x_i^{\mathcal{V}} b_i$ (i.e., $\mathbf{b}$ with losers of k-Vickrey treated as zero).*

3. *Simulate $\mathcal{A}$ on $\mathbf{b}^{\mathcal{A}}$ to get allocation $\mathbf{x}^{\mathcal{A}}$ and prices $\mathbf{p}^{\mathcal{A}}$.*

4. Output $\mathbf{x}$ *with* $x_i = x_i^{\mathcal{V}} x_i^{\mathcal{A}}$ *and* $\mathbf{p}$ *with* $p_i = x_i \max(p_i^{\mathcal{V}}, p_i^{\mathcal{A}})$.

Lemma 5.1 *For any auction, $\mathcal{A}$, the limited supply variant, $\mathcal{A}_k$, is truthful.*

The proof of this lemma is straightforward and we omit a discussion of it except to note that it is crucial that no item is sold at a price lower than the k-Vickrey price.

Theorem 5.2 *Given $\mathcal{A}$ that is β-competitive with $\mathcal{F}^{(m)}$, the limited supply variant, $\mathcal{A}_k$, is β-competitive with $\mathcal{F}^{(m,k)}$.*

Proof: Let p be the sale price of k-Vickrey on $\mathbf{b}$. Note that by the definition of $\mathbf{b}_{\mathcal{A}}$, $\mathcal{F}^{(m,k)}(\mathbf{b}) = \mathcal{F}^{(m)}(\mathbf{b}_{\mathcal{A}})$ as there are only k positive bid values in $\mathbf{b}_{\mathcal{A}}$. By our assumption that $\mathcal{A}$ is β-competitive, running $\mathcal{A}$ on $\mathbf{b}_{\mathcal{A}}$ achieves a profit of $\mathcal{F}^{(m)}(\mathbf{b}_{\mathcal{A}})/\beta = \mathcal{F}^{(m,k)}(\mathbf{b})/\beta$. Further, since all non-zero bids in $\mathbf{b}_{\mathcal{A}}$ are at least p, raising the price payed by any winner of $\mathcal{A}$ on $\mathbf{b}_{\mathcal{A}}$ does not cause any bidder with possitive contribution to the auction profit to be rejected. Thus, the profit of $\mathcal{A}_k$ on $\mathbf{b}$ is at least $\mathcal{F}^{(m,k)}/\beta$. □

This reduction which originally appeared in [26] allows us to trivially extend all positive results for the unlimited supply case to the limited supply case.

It is interesting to point out that by using the limited supply variants of auctions like RSOP and RSPE in k-item auctions it is possible on many bid vectors to obtain profit significantly higher than that of the k-item Vickrey auction.

Chapter 6

ENVY-FREE

Consider the outcome obtained by running a particular basic auction. A natural fairness property of such an outcome is that it be defined by a single sale price such that all bidders with bids above this price win and pay the price and all bidders with bids below the price lose. We call such outcomes *envy-free* because, for bids equal to bidder utilities, no bidder would prefer another bidder's outcome. *Envy-free auctions*, auctions that always produce envy-free outcomes, are natural and desirable for consumer acceptance. While some auctions, such as the classical k-item Vickrey auction and some optimal Bayesian auctions [11, 38] are envy-free, other auctions, in particular profit-maximizing competitive auctions discussed in the previous chapter, are not. In this chapter we study compatibility of the envy-free property with the desirable auction properties of truthfulness and worst case profit maximization [24].

Unfortunately, all competitive auctions discussed thus far have the property that a bidder in the auction may get rejected while another bidder wins and pays a lower price than the first bidder's bid[1]. In this case, the first bidder would be envious of the second bidder's outcome. Thus, these auctions are not envy-free.

There is a good reason why none of the constant-competitive auctions considered to date have been envy-free. Our main result, presented in Section 6.2, is that it is not possible to design an auction that is constant competitive, truthful, and always has an envy-free outcome. We show that any auction that is always envy-free and truthful has a competitive ratio of $\Omega(\log n/ \log\log n)$, where n is the number of bids. This bound is close to being tight: the Randomized Vickrey auction discussed in Section 4.1 is $\Theta(\log n)$-competitive and always envy-free.

[1] In addition some of these auctions sell to different bidders at different prices; though, most of them can be adapted to have a single sale price for winners.

In order to obtain competitive auctions with envy-free-ness properties we will consider relaxing both the truthful requirement and the envy-free requirement. In relaxation of the envy-free requirement we will show that the CORE auction from Section 4.5, which is truthful and competitive, is envy-free with probability that approaches one as the number of items sold increases.

Alternatively we can consider relaxing the truthfulness requirement. When truthful mechanisms for a problem do not exist, it is natural to look at relaxations of the requirement that the mechanism be completely truthful. We adopt the notion from [2] that an auction is *truthful with probability* $1-\epsilon$ if the probability that any bidder can benefit from a untruthful bid is bounded by ϵ. An auction is *truthful with high probability* if ϵ tends to zero with some parameter in the input, e.g., the number of winners in an optimal auction. We show that a variant of the CORE auction that is truthful with high probability is always envy-free and also constant competitive assuming that bidders bid their true values.

Auctions that do not have the envy-free property are impractical selling mechanisms for markets where consumers object to differential pricing [5]. For these applications, auctions that are truthful and almost envy-free or envy-free and almost truthful may be acceptable.

6.1 Definitions

We now formalize our notion of an envy-free auction.

Definition 6.1 *The outcome of an auction is* envy-free *if there is a price, t, such that every winning bidder pays t, all bidders with bid values greater than t win, and all bidders with values lower than t lose. Bidders with bid values equal to t may either win or lose. Call such an outcome a t-*envy-free outcome.

From this definition it is clear that envy-free outcomes are completely specified by a number t and a description of whether tying bids, bids with $b_i = t$, win or lose. We note that our lower bounds will assume the more general case that tying bids can be treated differently. This is consistent with the envy-free definition as such bidders have zero profit if they win and zero profit if they lose and are therefore assumed to have no preference over the two possible outcomes. For our upper bounds we will give auctions that satisfy the stronger condition that tying bidders are treated the same way, either they all win or they all lose. In all of our analyses it will be largely irrelevant how tying bids are treated; therefore, the relevant characteristic of the outcome is the value t.

We will be gauging the performance of envy-free auctions using the mass market variant of our competitive analysis framework.

6.2 Lower Bound

In this section, we show that no envy-free auction can be $o(\log n/\log\log n)$ competitive. This result is with in a $\log\log n$ factor of being tight as the Randomized Vickrey auction (Definition 4.1) is envy-free and $\Theta(\log n)$-competitive.

For a particular input $\mathbf{b}$, an auction that always yields an envy-free outcome induces a probability distribution on values t such that the outcome is t-envy-free. Let $t(\mathbf{b})$ be the random variable for the price used by a auction with envy-free outcomes and define $p_{\mathbf{b}}(x)$ as

$$p_{\mathbf{b}}(x) = \mathbf{Pr}[t(\mathbf{b}) \leq x] \,.$$

We begin by showing that any auction with envy-free outcomes has the property that the distribution of $t(\mathbf{b})$ on values less than x is independent of all bids with values above x.

Lemma 6.1 *Let $\mathcal{A}$ be any truthful envy-free auction $\mathcal{A}$ with t and $p_{\mathbf{b}}$ defined as above. For all bid vectors $\mathbf{b}$, $\mathbf{b}'$, values x, and subsets of bidders S such that all $i \in S$ have $b_i > x$ and $b'_i > x$, and $i \notin S$ have $b_i = b'_i$, $t(\mathbf{b})$ and $t(\mathbf{b}')$ have the same distribution on values at least x. That is,*

$$\forall y \leq x,\ p_{\mathbf{b}}(y) = p_{\mathbf{b}'}(y).$$

Proof: Assume for a contradiction that there exists $\mathbf{b}$ and $\mathbf{b}'$ differing on some subset S (as defined above) such that $t(\mathbf{b})$ and $t(\mathbf{b}')$ do not have the same distribution on values at most x. Let $S = \{i_1, \ldots, i_k\}$ and $S_j = \{i_\ell \ :\ \ell \leq j\}$. Let $\mathbf{b}^{(j)}$ be equal to $\mathbf{b}$ except for $\mathbf{b}_{S_j} = \mathbf{b}'_{S_j}$. Note that $\mathbf{b}^{(0)} = \mathbf{b}$ and $\mathbf{b}^{(k)} = \mathbf{b}'$. It must be that for some j^* there exists $y \leq x$ such that $p_{\mathbf{b}^{(j^*-1)}}(y) \neq p_{\mathbf{b}^{(j^*)}}(y)$. However this violates truthfulness as the distribution of prices for bidder i_{j^*} when the bids are $\mathbf{b}^{(j^*)}$ is not bid-independent: if bidder i_{j^*} changes their bid between $b_{i_{j^*}}$ and $b'_{i_{j^*}}$ they receive a different distribution of prices. An auction that is not bid-independent is not truthful (Theorem 2.2). □

Theorem 6.2 *No truthful envy-free auction is $\beta(m)$-competitive with*

$$\beta(m) \in O(\log \tfrac{n}{m} / \log\log \tfrac{n}{m}).$$

In particular, if m is a constant, e.g., $m = 2$, no truthful envy-free auction is $\beta(m) \in O(\log n / \log\log n)$ competitive.

Proof: For any integer $d \geq 2$, we show that no envy-free auction, $\mathcal{A}$, is $d/2$-competitive on an input $\mathbf{b}$ of size $n = md^d$ that has at least m winners. Since $d > \log(n/m)/\log\log(n/m)$, this gives the theorem. Assume for a contradiction that some auction, $\mathcal{A}$, is $d/2$-competitive.

Let $\mathbf{b}^*$ be the bid vector with $n/d^k - n/d^{k+1}$ bids at value n^k for $0 \leq k < d$ and m bids at n^d. Let $\mathbf{b}^{(k)}$ be identical to $\mathbf{b}^*$ except for the largest n/d^{k+1} bids, which are decreased to $n^k + \epsilon$:

$$b_i^{(k)} = \begin{cases} n^k + \epsilon & \text{if } i \leq n/d^{k+1} \\ b_i^* & \text{otherwise.} \end{cases}$$

There are two observations about these bid vectors that will be useful.

- $\mathcal{F}(\mathbf{b}^{(k)}) = n^{k+1}/d^k$ as the highest n/d^k bidders win at price n^k.
- $\mathcal{A}$ on $\mathbf{b}^{(k+1)}$ may place probability mass on prices at n^k or below. An upper bound on the contribution of this mass to the expected revenue of $\mathcal{A}$ is $\mathcal{F}(\mathbf{b}^{(k)}) = n^{k+1}/d^k$.

Define $t(\mathbf{b})$ and $p_{\mathbf{b}}(x)$ for $\mathcal{A}$ as above. We will show by induction that $\mathcal{A}$ has

$$p_{\mathbf{b}^{(k)}}(n^k) \geq (k+1)/d.$$

This implies that for $\mathbf{b}^* = \mathbf{b}^{(d-1)}$ we have

$$p_{\mathbf{b}^*}(n^{d-1}) = 1.$$

This is a contradiction, because if this were the case, $\mathcal{A}$ would have expected revenue at most n^d/d^{d-1} and could not be $d/2$-competitive with $\mathcal{F}(\mathbf{b}^*) = mn^d$.

We will be using the fact that $\mathcal{A}$ is envy-free by invoking Lemma 6.1 in the inductive step to guarantee that $p_{\mathbf{b}^{(k+1)}}(n^k) = p_{\mathbf{b}^{(k)}}(n^k) \leq (k+1)/d$.

For the base case, we show that $p_{\mathbf{b}^{(0)}}(1) \geq 1/d$. Note that $\mathbf{b}^{(0)}$ is defined as:

$$b_i^{(0)} = \begin{cases} 1+\epsilon & \text{if } i \leq n/d \\ 1 & \text{otherwise.} \end{cases}$$

Here we have $\mathcal{F}(\mathbf{b}^{(0)}) = n$ and as a $d/2$-competitive auction, $\mathcal{A}$ must have expected revenue of at least $2n/d$. Suppose that $\mathcal{A}$ has $p_{\mathbf{b}^{(0)}}(1) = p$. Given this constraint, the best revenue is achieved by putting all probability mass p on 1 and the remaining probability mass $1-p$ on $1+\epsilon$. This gives the following bound on the expected revenue of $\mathcal{A}$:

$$\mathbf{E}\Big[\mathcal{A}(\mathbf{b}^{(0)})\Big] \leq np + (1-p)(1+\epsilon)(n/d).$$

This is an increasing function of $p \in [0,1]$. For $p = 1/d$ we have the following:

$$\mathbf{E}\Big[\mathcal{A}(\mathbf{b}^{(0)})\Big] \leq n/d + (1-1/d)(1+\epsilon)(n/d).$$

For ϵ chosen such that $(1+\epsilon)(1-1/d) = 1$ we have

$$\mathbf{E}\Big[\mathcal{A}(\mathbf{b}^{(0)})\Big] \leq 2n/d.$$

Therefore p must be at least $1/d$ for $\mathcal{A}$ to be $d/2$-competitive.

For the inductive step, assume that $p_{\mathbf{b}^{(k-1)}}(n^{k-1}) \geq k/d$ and consider running $\mathcal{A}$ on $\mathbf{b}^{(k)}$. By Lemma 6.1 and the definition of $\mathbf{b}^{(k)}$, $p_{\mathbf{b}^{(k-1)}}(n^{k-1}) = p_{\mathbf{b}^{(k)}}(n^k)$. Therefore, the probability mass remaining to be placed on values strictly larger than n^{k-1} is $P = 1 - p_{\mathbf{b}^{(k)}}(n^{k-1}) < 1 - k/d$. $\mathcal{A}$ must place p of this remaining mass on values at most n^k and the rest on higher values. Note that by definition, $\mathbf{b}^{(k)}$ has n/d^k bids at value strictly greater than n^{k-1}: it has n/d^{k+1} of them at $n^k + \epsilon$ and the remaining at n^k.

Thus, the most revenue can be obtained by placing all of p on n^k and all of the $P - p$ remaining mass on $n^k + \epsilon$. As discussed above, the expected revenue due to probability mass on values at most n^{k-1} is at most n^k/d^{k-1}. Thus, $\mathcal{A}$'s expected revenue is bounded as follows:

$$\mathbf{E}\left[\mathcal{A}(\mathbf{b}^{(k)})\right] \leq pn^k n/d^k + (P-p)(n^k+\epsilon)n/d^{k+1} + n^k/d^{k-1}$$

This is an increasing function of p. For $p = 1/d$ we have

$$\mathbf{E}\left[\mathcal{A}(\mathbf{b}^{(k)})\right] \leq n^k n/d^{k+1} + (P-p)(n^k+\epsilon)n/d^{k+1} + n^k/d^{k-1}.$$

Routine manipulation using the fact $d \geq 2$ shows that

$$\mathbf{E}\left[\mathcal{A}(\mathbf{b}^{(k)})\right] \leq 2n^{k+1}/d^{k+1}.$$

$\mathcal{F}$ of $\mathbf{b}^{(k)}$ is n^{k+1}/d^k, thus p is at least $1/d$. This gives the inductive claim that $p_{\mathbf{b}^{(k)}}(n^k) \geq k + 1/d$. □

6.3 Envy-free with High Probability

In this section, we relax the condition that the auction outcome is always envy-free and instead consider auctions that are truthful but only envy-free with high probability. One such auction is the Consensus Revenue Estimate (CORE) Auction. We review the mass-market version of the CORE auction (from Section 4.5) and discuss its properties.

Recall the definition of CORE_c:

- pick U uniformly at random from $[0, 1]$,
- let function $r(\cdot)$ be $\mathcal{F}(\cdot)$ rounded down to nearest c^{i+U} for integer i, and
- the CORE_c auction is defined by bid-independent function $f(\mathbf{b}_{-i}) = \text{cs}_{r(\mathbf{b}_{-i})}(\mathbf{b}_{-i})$.

We say that CORE_c achieves a consensus at value R if $r(\mathbf{b}_{-i}) = R$ for all i. In this case, by definition, CORE_c's outcome is identical to that of ProfitExtract_R. Since ProfitExtract_R always outputs envy-free outcomes, on consensus the outcome of CORE_c is envy-free. Thus, the following lemma follows directly from Lemma 4.15.

Lemma 6.3 *On* $\mathbf{b}$ *with* $\mathcal{F}$ *selling* $m \geq 2$ *items,* CORE_c *is envy-free (i.e., achieves a consensus) with probability at least*

$$1 - \log_c \tfrac{m}{m-1} = 1 - \Theta(1/m).$$

In the case that CORE does not achieve a consensus, there are two sale prices, and the auction is not envy-free.

Lemma 6.4 *The* CORE *auction has at most two sale prices.*

Proof: If we have a consensus estimate, the sale price is unique. Suppose the estimates disagree. Let k be the number of winners in $\mathcal{F}^{(2)}$ on $\mathbf{b}$ and consider two cases, $k \geq 3$ and $k = 2$.

In the first case, $r(\mathbf{b})$ is between $\frac{k-1}{k}\mathcal{F}^{(2)}(\mathbf{b})$ and $\mathcal{F}^{(2)}(\mathbf{b})$. Thus, some bids will use the sale price from $\mathrm{ProfitExtract}_{R(\mathbf{b})}$ and some will use the price from $\mathrm{ProfitExtract}_{R(\mathbf{b})/c}$. Lower bids will use the former (a higher price) and higher bids will use the latter (a lower price).

In the second case, $\mathcal{F}^{(2)}(\mathbf{b})$ is determined by the two highest bids. If b_i is not one of these bids, then $\mathcal{F}^{(2)}(\mathbf{b}_{-i}) = \mathcal{F}^{(2)}(\mathbf{b})$. The values of $\mathcal{F}^{(2)}$ when one or another of the two bids is removed are the same. It follows that there are only two possible values of $\mathcal{F}^{(2)}(\mathbf{b}_{-j})$, and therefore at most two distinct sale prices. $\square$

6.4 Truthful with high probability

Another approach to deal with the non-existence of truthful competitive auctions that always have envy-free outcomes is to relax the requirement that the auction always be truthful. To this end, we define a *probabilistically truthful* mechanism. This definition was first given by Archer et al. in [2]. Let m be the number of winners in the optimal auction on a given set of bids. We will be looking for a mechanism with good probabilistically truthful properties in terms of m. For other solution concepts related to approximate or probabilistic truthfulness, see for example [41, 47].

Definition 6.2 *[2] An auction is* truthful with probability $1-\epsilon$ *if the probability that any bidder can benefit from an untruthful bid is at most ϵ. An auction is* truthful with high probability *if $\epsilon \to 0$ as $m \to \infty$, where m is the number of winners in the optimal auction $\mathcal{F}$.*

Definition 6.3 (CORE'_c) *For constant c and input $\mathbf{b}$, CORE'_c is:*

1. *Pick U uniformly at random from $[0, 1]$.*
2. *Let function $r(\cdot)$ be $\mathcal{F}(\cdot)$ rounded down to nearest c^{i+U} for integer i.*
3. *Run ProfitExtract$_{r(\mathbf{b})}$ on $\mathbf{b}$.*

Lemma 6.5 *For bids $\mathbf{b}$ and choice of U fixed identically for both CORE_c and CORE'_c, if CORE_c is a consensus then CORE'_c is truthful.*

Proof: Let $R = r(\mathbf{b})$. To prove the lemma, we must show that for U such that CORE_c is a consensus, no bidder can benefit from bidding any value other than their true utility value. First we note that CORE'_c runs $\mathrm{ProfitExtract}_{r(\mathbf{b})}$ on $\mathbf{b}$. Consider the effect of bidder i changing their bid to b'_i resulting in bid vector $\mathbf{b}'$. Bidder i cannot benefit from bidding any b'_i such $r(\mathbf{b}') = r(\mathbf{b}) = R$ because $\mathrm{ProfitExtract}_R$ is truthful and therefore bidder i's best strategy in mechanism $\mathrm{ProfitExtract}_R$ is to bid b_i.

The fact that CORE_c is a consensus for this value of U means that $r(\mathbf{b}_{-i}) = r(\mathbf{b}) = R$ for all i. Again consider $\mathbf{b}'$ identical to $\mathbf{b}$ except for bidder i bidding b'_i. For $b'_i \in [0, b_i]$, since $r(\cdot)$ is a monotonically increasing function of the bids, $r(\mathbf{b}_{-i}) = r(\mathbf{b}') = r(\mathbf{b}) = R$. Thus, no bidder can lower the the value of $r(\mathbf{b})$.

We now consider bidder i raising their bid enough to make $r(\mathbf{b}') > r(\mathbf{b})$. Although we do not go into the details here, it is not difficult to show that $\mathrm{pe}_R(\mathbf{b}_{-i})$ (the bid-independent function implementing $\mathrm{ProfitExtract}_R$), for $\mathbf{b}_{-i}$ fixed, is a strictly monotone increasing function of R. Thus, if bidder i raises their bid to raise R to r', the price offered them by $\mathrm{pe}_{R'}$ will be higher than for pe_R and therefore, bidder i would be worse off. □

The following corollary follows from the above lemma and Theorem 4.15.

Corollary 6.6 *On* $\mathbf{b}$ *with* $\mathcal{F}$ *selling* $m \geq 2$ *items,* CORE'_c *is truthful with probability*

$$1 - \log_c \tfrac{m}{m-1} = 1 - \Theta(1/m).$$

We now consider the performance of CORE'_c. We can view the use of $r(\cdot)$ in CORE'_c as a consensus problem with $\rho = 1$ since the same value, $r(\mathbf{b})$, is used for all bidders. This allows us to make use of Theorem 4.12 directly to obtain the following corollary.

Corollary 6.7 *The expected revenue of* CORE'_c *is*

$$\frac{\mathcal{F}(\mathbf{b})}{\ln c}\left(1 - \frac{1}{c}\right).$$

Conclusions and Extensions

Recall that with high probability, the CORE auction is equivalent to the profit extraction mechanism that it is based on, which is a special case of the Moulin-Shenker [37] cost sharing mechanism. We exploited the fact that the latter mechanism is envy-free. The mechanism is also *group strategy-proof*: no coalition of bidders can collude by bidding untruthfully so as some members of the coalition strictly benefit without causing other members of the coalition to be strictly worse off. As such, both CORE_c and CORE'_c have collusion resistant properties.

We have considered truthfulness with high probability to get an auction that is both envy-free and competitive. In our analysis we make the assumption that bidders will still reveal their true utility values even though the mechanism is only truthful with high probability. In the case where each bidder is perfectly informed as to the bidding strategies and utility values of other bidders, any non-truthful auction would fail to obtain true bids, as each bidder could calculate their own optimum bid. However, in the presence of bidder uncertainty about the strategies or utility values of other bidders, this calculation is no longer straight-forward. A bidder in an auction that is truthful with high probability is faced with a choice. The bidder can bid truthfully, which with high probability is an optimal strategy, or alternatively the bidder could try to gain by manipulating their bid on the basis of available information. If information is incomplete, the latter may be impossible.

More generally, we have shown that one can get stronger results by relaxing the notion of truthfulness. Relaxed truthfulness is an interesting direction for future research. For example, is it possible to get better competitive ratios under a reasonable relaxation? Which relaxations are reasonable?

Chapter 7

ON THE COMPETITIVE RATIO

One goal of algorithm design (and in our case auction design) in a competitive analysis framework is to obtain the algorithm (auction) that achieves the best competitive ratio. In this chapter we consider this endeavor. We will look at several restrictions on auctions and consider the case when there are only two or three bidders. For the two bidder case we show that the Vickrey auction obtains the optimal competitive ratio. For the three bidder case, we show that auctions that only sell the good at prices that are equal to bid values are not as good as auctions that may sell the good for any price [21]. We conclude by giving a bound on the the best competitive ratio for an auction on n bidders and show that in the limit this competitive ratio is 2.42 [25].

7.1 Asymmetry

Recall the definition of a symmetric auction as one that gives an output that is not a function of the order of the bids in the input vector, $\mathbf{b}$. In Chapter 4 we saw that for a certain class of inputs asymmetric deterministic auctions perform better in worst case than symmetric deterministic auctions. Here we briefly show that asymmetry does not help randomized auctions. This result is trivial, but for completeness we give a proof.

Lemma 7.1 *For any asymmetric auction $\mathcal{A}$ with competitive ratio β, there is a symmetric randomized auction with competitive ratio at least β.*

Proof: Given asymmetric auction $\mathcal{A}$ with competitive ratio β, we construct a symmetric auction $\mathcal{A}'$ that first permutes the input bids $\mathbf{b}$ at random to get $\pi(\mathbf{b})$ and then runs $\mathcal{A}$ on $\pi(\mathbf{b})$. Note, $\mathcal{F}^{(2)}(\mathbf{b}) = \mathcal{F}^{(2)}(\pi(\mathbf{b}))$ and since $\mathcal{A}$ is β-competitive on $\pi(\mathbf{b})$ for any choice of

π, $\mathcal{A}'$ is β-competitive on $\mathbf{b}$. □

Throughout the remainder of this chapter when we are proving lower bounds on auctions we will assume that the auction is symmetric.

7.2 Auctions for two items

Given two bids, $\mathbf{b} = (b_1, b_2)$, we look for an auction that sells one or two items and attains a revenue competitive with $\mathcal{F}^{(2)}(\mathbf{b}) = 2\min(b_1, b_2)$.

Observation 3 *The Vickrey Auction is 2-competitive.*

One might wonder if it is possible to do better with some other auction, perhaps even one that is randomized. However, Theorem 7.10 shows us that this is not the case: no auction as a competitive ratio better than two in the two bidder case.

7.3 Auctions for three items

Observation 4 *The one item Vickrey auction, a deterministic auction, is 3-competitive.*

Proof: If the auction bid are $\{b_1, b_2, b_3\}$ with $b_1 \geq b_2 \geq b_3$ then $\mathcal{F}^{(2)} = \max(2b_2, 3b_1) \leq 3b_2$. Vickrey gets a revenue of b_2 which is at least a third of $\mathcal{F}^{(2)}$. □

Definition 7.1 *We say a bid-independent auction is* restricted *if the function f is required to return a bid value,* unrestricted *otherwise.*

Lemma 7.2 *Assuming f is continuous (I.e. $\lim_{\epsilon \to 0} f(a, b) = f(a, b + \epsilon)$), no restricted auction can achieve a competitive ratio better than* $5/2$*.*

Proof: For h fixed, any restricted auction must be of the form:

$$f(1,1) = 1$$

$$f(1,h) = \begin{cases} 1 & \text{with probability } p \\ h & \text{otherwise} \end{cases}$$

The expected revenue for the auction on $\{1, 1+\epsilon, h\}$ is $1+p$ (the 1 is from the h and the p is from the $1+\epsilon$). In this case, $\mathcal{F}^{(2)} = 3$ so the competitive ratio is $3/(1+p)$.

The expected revenue for the auction on $\{1, h, h+\epsilon\}$ is $2p + h(1-p)$ (the $p + (1-p)h$ is from the $h+\epsilon$ and the p is from the h). For large h, $\mathcal{F}^{(2)} = 2h$ so the competitive ratio is $2h/(2p + h(1-p)) \approx 2/(1-p)$.

We set these two ratios equal and solve for p obtaining a value of $p = 1/5$. Plugging this value in for p in either of the above equations, we get a competitive ratio of $5/2$. □

This upper bound is tight as the following theorem proves.

Lemma 7.3 *The bid-independent auction with f as*

$$f(1,h) = \begin{cases} 1 & \textit{with probability } 1/5 \\ h & \textit{otherwise} \end{cases}$$

achieves a competitive ratio of $5/2$.

Proof: Running this auction on bids $\{b_1, b_2, b_3\}$ with $b_1 \le b_2 \le b_3$ gives a revenue of

$$\begin{aligned} \mathbf{E}[\mathcal{R}] &= \underbrace{0}_{\text{for } b_1} + \underbrace{b_1/5}_{\text{for } b_2} + \underbrace{b_1/5 + 4b_2/5}_{\text{for } b_3} \\ &= 2b_1/5 + 4b_2/5 \end{aligned}$$

Thus, $\mathbf{E}[\mathcal{R}] \ge 6b_1/5$ (because $b_2 \ge b_1$) and $\mathbf{E}[\mathcal{R}] \ge 4b_2/5$ which means

$$\begin{aligned} \mathbf{E}[\mathcal{R}] &\ge \tfrac{2}{5}\max(3b_1, 2b_2) \\ &= \tfrac{2}{5}\mathcal{F}^{(2)} \end{aligned}$$

□

What is interesting to note is that this auction is essentially doing a 1-item Vickrey auction with probability $4/5$ and a 2-item Vickrey auction with probability $1/5$.

Lemma 7.4 *A* non-restricted *auction can achieve a better competitive ratio than 5/2.*

Proof: Define a bid-independent auction with f as

$$f(1,h) = \begin{cases} \alpha & \text{with probability } p \\ h & \text{otherwise} \end{cases}$$

for $\alpha \in [0,1]$.

Running this auction on bids $\{b_1, b_2, b_3\}$ with $b_1 \leq b_2 \leq b_3$ Gives a revenue of

$$\mathbf{E}[\mathcal{R}] = \underbrace{0}_{\text{for } b_1} + \underbrace{p\alpha b_1}_{\text{for } b_2} + \underbrace{p\alpha b_1 + (1-p)b_2}_{\text{for } b_3} + b_1\text{-term}$$

The competitive ratio of this auction is maximized at $\alpha = 6/7$ and $p = 36/43$. This gives a ratio of 2.39 (which is $2/p$). □

Other similar techniques to the above have yielded slightly better competitive ratios; however, none are much closer to the lower bound of 13/6 that we prove later in this chapter.

7.4 Worst Case Verse Mass Market Analysis

We have presented auction that perform well in worst-case and others that are competitive for mass-markets. We have come up with auctions that are constant competitive in worst case, as well as auctions that approach 1-competitive for mass markets. In this section we show that for *restricted* auctions (see Definition 7.1, it is not possible for an auction to both be constant competitive in worst case, and be 1-competitive for mass markets. The intuition behind this argument is that even when it may be optimal to sell many items, a bid-independent auction may be fooled into thinking that selling a small number of items is better. We conjecture that this result holds with out the assumption that the restricted auction assumption as well.

Lemma 7.5 *No restricted auction, $\mathcal{A}$, is better than ϵ-competitive in worst case and better than $(1-\epsilon)$ competitive on mass markets.*

Proof:

$\mathcal{A}$ is truthful, so $\mathcal{A}$ is equivalent to the bid-independent mechanism BI_f. We will assume without loss of generality that the auction, $\mathcal{A}$, is symmetric and that the $f(\mathbf{b}_{-i})$ is independent from $f(\mathbf{b}_{-j})$ for all $i \neq j$. This assumption can be made without loss of generality, because given any auction $\mathcal{A}'$ that is not symmetric and independent we can make it symmetric and independent without worsening its competitive ratio by first creating $\mathcal{A}''$ as the auction that randomly permutes the input and then runs $\mathcal{A}'$ (Lemma 7.1). Finally, we can create $\mathcal{A}$ as the auction that, for each bidder i, simulates $\mathcal{A}''$ with new random coins to obtain the outcome for bidder i (as discussed in Section 3.3).

We will now consider the symmetric restricted auction $\mathcal{A}$. We will fix n, the number of bidders, and consider inputs with only bids at value 1 and h. As such, $f(\mathbf{b}_{-i})$ is a function from the number of hs (given n and the number of hs the number of 1s is just the difference less one) to a probability that $f(\mathbf{b}_{-i})$ outputs 1 or h. So, let p_k be the probability that $f(\mathbf{b}_{-i})$ outputs h if there are k bids at h and $n-k-1$ bids at 1 in $\mathbf{b}_{-i}$. Of course, the auction could reject some bidder by offering price ∞, however, this is never good in terms of profit maximization, so we will assume that $\mathcal{A}$ does not. Let $\mathbf{b}^{(k)}$ be the input with k bids at h and the remaining $n-k$ bidders at value 1.

Note that because $\mathcal{A}$ is restricted, on the all ones input it must offer price 1. Therefore, $p_0 = 0$. We now show that if the auction is to be better than $(1-\epsilon)$-competitive, it must be that $p_1 < \epsilon$. Note that the revenue of this auction is thus,

$$\mathcal{R}_1 = 1 + (1-p_1)(n-1).$$

This approaches $1-p_1$-competitive in the limit so it must be that $p_1 < \epsilon$ for it to be better than $(1-\epsilon)$-competitive for mass markets. We now show that this auction with $p_1 < \epsilon$ is not ϵ competitive in worst case (against $\mathcal{F}^{(2)}$). Consider $\mathbf{b}^{(2)}$. When we remove an h, we are left with a masked bid vector with one h bid and $n-2$ bids at value 1. When we remove a 1 bid, we are left with 2 bids at value h and $n-3$ bids at value 1. Thus, the revenue is,

$$\mathcal{R}_2 = 2(hp_1 + 1 - p_1) + p_2(n-2).$$

In the limit as h gets large, this quantity approaches $2hp_1$, which by our previous argument is less than $2h\epsilon$. Since $\mathcal{F}^{(2)} = 2h$, our auction is not ϵ-competitive with $\mathcal{F}^{(2)}$ in worst case.

The contrapositive shows that any restricted auction that is better than ϵ-competitive with $\mathcal{F}^{(2)}$ in worst case is not $(1-\epsilon)$-competitive for mass markets. □

7.5 A Lower Bound on the Competitive Ratio

Now we prove a lower bound on the competitive ratio of any truthful auction, even a randomized one, in comparison to $\mathcal{F}^{(2)}$: We show that for any randomized truthful auction, $\mathcal{A}$, there exists an input bid vector $\mathbf{b}$ on which

$$\mathbf{E}[\mathcal{A}(\mathbf{b})] \leq \frac{\mathcal{F}^{(2)}(\mathbf{b})}{2.42}$$

.

To prove the lower bound, we analyze the behavior of $\mathcal{A}$ on a bid vector chosen from a probability distribution over bid vectors. The outcome of the auction is then a random variable depending on both the randomness in $\mathcal{A}$ and the randomness in $\mathbf{b}$. We show that $\mathbf{E}_{\mathbf{b}}[\mathbf{E}_{\mathcal{A}}[\mathcal{A}(\mathbf{b})]] \leq \frac{\mathbf{E}_{\mathbf{b}}[\mathcal{F}^{(2)}(\mathbf{b})]}{2.42}$. We then use the following fact to claim that there must exist a fixed choice of $\mathbf{b}$ (depending on $\mathcal{A}$) for which $\mathbf{E}[\mathcal{A}(\mathbf{b})] \leq \frac{\mathcal{F}^{(2)}(\mathbf{b})}{2.42}$.

Fact 1 *Given random variable X and to functions f and g, $\mathbf{E}[f(X)] \leq \mathbf{E}[g(X)]$ implies that there exists x such that $f(x) \leq g(x)$.*

As a quick proof, if for all x, $f(x) > g(x)$ then it would be the case that $\mathbf{E}[f(x)] > \mathbf{E}[g(x)]$ instead of the other way around.

Consider n i.i.d. bids $\mathbf{b}^{(n)}$ generated from the distribution with each bid b_i satisfying $\mathbf{Pr}[b_i > y] = 1/y$ for all $y \geq 1$. Consider a truthful auction $\mathcal{A}$. Let V_i be the price offered to b_i in the bid-independent implementation of $\mathcal{A}$. V_i is a random variable depending on $\mathcal{A}$ and all of the b_j other than b_i. Let P_i be the price paid by bidder i, i.e., 0 if $b_i < V_i$ and V_i otherwise. For $v \geq 0$, $\mathbf{E}[P_i \mid V_i = v] = v \cdot \mathbf{Pr}[b_i > v | V_i = v] = v \cdot \mathbf{Pr}[b_i > v] \leq 1$, since b_i is independent of V_i. Therefore $\mathbf{E}[P_i] \leq 1$. Thus we have:

Lemma 7.6 *For $\mathbf{b}^{(n)}$ defined above, the expected revenue of any truthful auction, $\mathcal{A}$, is at most n.*

Note that for any deterministic bid-independent auction that offers prices of at least one, the expected revenue is exactly n.

Lemma 7.7 *For n bids from the above distribution, the expected value of $\mathcal{F}^{(2)}$ is*

$$\mathbf{E}\Big[\mathcal{F}^{(2)}(\mathbf{b}^{(n)})\Big] = n - n\sum_{i=2}^{n}\left(\frac{-1}{n}\right)^{i-1}\frac{i}{i-1}\binom{n-1}{i-1}.$$

Proof: In this proof we will get a closed form expression for $\mathbf{Pr}\big[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\big]$ and then integrate to obtain the expected value. Note that all bids are at least one and therefore, we will assume that $z \geq n$. Clearly for $z < n$, $\mathbf{Pr}\big[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\big] = 1$. To get a formula for $\mathbf{Pr}\big[\mathcal{F}^{(2)}(\mathbf{b}^{(n)})\big]$, we define a recurrence based on $F_{n,k}$ defined as

$$F_{n,k} = \max_i (k+i)b_i$$

for bids $\mathbf{b}^{(n)}$ sorted from highest to lowest (i.e., $b_i \geq b_{i+1}$). Intuitively, $F_{n,k}$ represents the optimal single price revenue from $\mathbf{b}^{(n)}$ and an additional k consumers each of which has a value equal to the highest bid, b_1. To define the recurrence, fix n, k, and z and define the events $\mathcal{H}_i$ for $1 \leq i \leq n$. Intuitively, the event $\mathcal{H}_i$ represents the fact that i bidders in $\mathbf{b}^{(n)}$ and the k additional consumers have bid high enough to equally share z, while no larger set of $j > i$ bidders of $\mathbf{b}^{(n)}$ can do the same.

$$\begin{aligned}\mathcal{H}_i &= b_i \geq z/(k+i) \wedge \bigvee_{j=i+1}^{n} b_j < z/(k+j)\\ &= \binom{n}{i}\left(\frac{k+i}{z}\right)^i \mathbf{Pr}[F_{n-i,k+i} < z]\,.\end{aligned}$$

Note that events $\mathcal{H}_i$ are disjoint and that $F_{n,k}$ is at least z if and only if one of the $\mathcal{H}_i$ occurs. Thus,

$$\begin{aligned}\mathbf{Pr}[F_{n,k} > z] &= \mathbf{Pr}\left[\bigwedge_{i=1}^{n}\mathcal{H}_i\right] = \sum_{i=1}^{n}\mathbf{Pr}[\mathcal{H}_i]\\ &= \sum_{i=1}^{n}\binom{n}{i}\left(\frac{k+i}{z}\right)^i \mathbf{Pr}[F_{n-i,k+i} < z]\,. \end{aligned} \tag{7.1}$$

Also, note that $F_{0,k} = 0$. For n bids $\mathbf{b}^{(n)}$, $\mathcal{F}(\mathbf{b}^{(n)}) = F_{n,0}$. We are interested in $\mathcal{F}^{(2)}(\mathbf{b}^{(n)})$

which is the same as $\mathcal{F}(\mathbf{b}^{(n)}) = F_{n,0}$ except that we ignore the $\mathcal{H}_1$ case. This gives

$$\begin{aligned}\mathbf{Pr}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\right] &= \mathbf{Pr}[F_{n,0} > z] - \mathbf{Pr}[\mathcal{H}_1] \\ &= \mathbf{Pr}[F_{n,0} > z] - \tfrac{n}{z}\mathbf{Pr}[F_{n-1,1} < z]\,. \end{aligned} \tag{7.2}$$

So in order to obtain $\mathbf{Pr}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)})\right]$ we need to solve the recurrence $F_{n,k}$, i.e., Equation (7.1). We will show that the solution is:

$$\mathbf{Pr}[F_{n,k} > z] = 1 - \left(\frac{z-k}{z}\right)^n \left(\frac{z-k-n}{z-k}\right). \tag{7.3}$$

Note that our solution for the recurrence is correct for $n = 0$. We show that it is true in general inductively.

$$\mathbf{Pr}[F_{n,k} > z] = \sum_{i=1}^{n} \binom{n}{i} \left(\frac{k+i}{z}\right)^i \mathbf{Pr}[F_{n-i,k+i} < z]\,.$$

Substituting in our solution, we get

$$\begin{aligned}\mathbf{Pr}[F_{n,k} > z] &= \sum_{i=1}^{n} \binom{n}{i} \left(\frac{k+i}{z}\right)^i \left(\frac{z-k-i}{z}\right)^{n-i} \left(\frac{z-k-n}{z-k-i}\right) \\ &= \frac{z-k-n}{z^n} \sum_{i=1}^{n} \binom{n}{i} (k+i)^i (z-k-i)^{n-i-1}\,.\end{aligned}$$

We now apply the following version of Abel's Identity: [1]

$$\frac{(x+y)^n}{x} = \sum_{j=0}^{n} \binom{n}{j} (x+j)^{j-1} (y-j)^{n-j}.$$

Making the change of variables, $j = n - i$, $x = z - k - n$, and $y = k + n$ we get:

$$\frac{z^n}{z-k-n} = \sum_{i=0}^{n} \binom{n}{i} (k+i)^i (z-k-i)^{n-i-1}.$$

We plug this in above and subtract out the $i = 0$ term to get

$$\begin{aligned}\mathbf{Pr}[F_{n,k} > z] &= \frac{z-k-n}{z^n} \left(\frac{z^n}{z-k-n} - (z-k)^{n-1}\right) \\ &= 1 - \left(\frac{z-k}{z}\right)^n \frac{(z-k-n)}{(z-k)}.\end{aligned}$$

Thus, our closed form expression for the recurrence is correct.

Recall our goal is to compute $\mathbf{Pr}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\right]$. Equation (7.3) shows that $\mathbf{Pr}[F_{n,0} > z] = n/z$. This combined with Equation (7.2) and Equation (7.3) gives the following for $z \geq n$:

$$\begin{aligned}
\mathbf{Pr}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\right] &= \tfrac{n}{z} - \tfrac{n}{z}\mathbf{Pr}[F_{n-1,1} < z] \\
&= \tfrac{n}{z}\mathbf{Pr}[F_{n-1,1} > z] \\
&= \frac{n}{z}\left(1 - \left(\frac{z-1}{z}\right)^{n-1}\left(\frac{z-n}{z-1}\right)\right).
\end{aligned}$$

Recall that for $z \leq n$, $\mathbf{Pr}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\right] = 1$. To complete this proof, we use the formula $\mathbf{E}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)})\right] = \int_0^\infty \mathbf{Pr}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\right] dz = n + \int_n^\infty \mathbf{Pr}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\right] dz$. In the form above, this is not easily integrable; however, we can transform it back into a binomial sum which we can integrate:

$$\begin{aligned}
\mathbf{Pr}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\right] &= n\sum_{i=2}^{n}\left(\frac{-1}{z}\right)^{i} i\binom{n-1}{i-1}. \\
\mathbf{E}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)}) > z\right] &= n + n\int_n^\infty \sum_{i=2}^{n}\left(\frac{-1}{z}\right)^{i} i\binom{n-1}{i-1} dz. \\
&= n - n\sum_{i=2}^{n}\left(\frac{-1}{n}\right)^{i-1}\frac{i}{i-1}\binom{n-1}{i-1}.
\end{aligned}$$

□

Combining this with the previous lemma we get:

Lemma 7.8 *For bids as defined above, we have*

$$\frac{\mathbf{E}\left[\mathcal{F}^{(2)}(\mathbf{b}^{(n)})\right]}{\mathbf{E}\left[\mathcal{A}(\mathbf{b}^{(n)})\right]} = 1 - \sum_{i=2}^{n}\left(\frac{-1}{n}\right)^{i-1}\frac{i}{i-1}\binom{n-1}{i-1}.$$

Interesting special cases are $n = 2$ where this gives a lower bound of 2 which matches the best competitive auction for two bids, the 1-item Vickrey auction. For $n = 3$ this gives a lower bound of 13/6. A lower bound for the ratio of the best competitive auction on general n is obtained by taking the limit.

Lemma 7.9

$$\lim_{n\longrightarrow\infty}\left(1 - \sum_{i=2}^{n}\left(\frac{-1}{n}\right)^{i-1}\frac{i}{i-1}\binom{n-1}{i-1}\right) = 1 + \sum_{i=2}^{\infty}(-1)^i\frac{i}{(i-1)(i-1)!}$$

Proof: It is sufficient to show that

$$\left|\left(1+\sum_{i=2}^{n}(-1)^i\frac{i}{(i-1)(i-1)!}\right)-\left(1-\sum_{i=2}^{n}\left(\frac{-1}{n}\right)^{i-1}\frac{i}{i-1}\binom{n-1}{i-1}\right)\right|=O\left(\frac{1}{n}\right).$$

We use the following fact below: If, for $1 \leq k \leq K$, $0 < a_k < 1$, then

$$\prod_{k=1}^{K}(1-a_k) \geq 1-\sum_{k=1}^{K}a_k.$$

$$\begin{aligned}
&\left|\left(1+\sum_{i=2}^{n}(-1)^i\frac{i}{(i-1)(i-1)!}\right)-\left(1-\sum_{i=2}^{n}\left(\frac{-1}{n}\right)^{i-1}\frac{i}{i-1}\binom{n-1}{i-1}\right)\right| \\
&\leq \sum_{i=2}^{n}\left|\frac{i}{(i-1)(i-1)!}-\left(\frac{1}{n}\right)^{i-1}\frac{i}{i-1}\binom{n-1}{i-1}\right| \\
&= \sum_{i=2}^{n}\left|\frac{i}{(i-1)(i-1)!}\left(1-\frac{n(n-1)\dots(n-i+2)}{n^{i-1}}\right)\right| \\
&= \sum_{i=2}^{n}\left|\frac{i}{(i-1)(i-1)!}\left(1-\left(1-\frac{1}{n}\right)\left(1-\frac{2}{n}\right)\cdots\left(1-\frac{i-2}{n}\right)\right)\right| \\
&\leq \sum_{i=2}^{n}\left|\frac{i}{(i-1)(i-1)!}\left(1-\left(1-\sum_{j=1}^{i-2}\frac{j}{n}\right)\right)\right| \\
&\leq \sum_{i=2}^{n}\left|\frac{i}{(i-1)(i-1)!}\left(\frac{i^2}{n}\right)\right| = \frac{1}{n}\sum_{i=2}^{n}\frac{i^3}{(i-1)(i-1)!} \leq \frac{1}{n}\sum_{i=2}^{\infty}\frac{i^3}{(i-1)(i-1)!}
\end{aligned}$$

As $(i-1)!$ grows exponentially, $\sum_{i=2}^{\infty}\frac{i^3}{(i-1)(i-1)!}$ is bounded by a constant and we have the desired result. □

Routine calculations then show that the limit value in Lemma 7.9 is at least 2.42. Combining this with Fact 1 we obtain:

Theorem 7.10 *Let $\mathcal{A}$ be any truthful randomized auction. The competitive ratio of $\mathcal{A}$ is at least* 2.42.

Chapter 8

MONOTONICITY

In this section we give further justification for the use of $\mathcal{F}$ and $\mathcal{F}^{(2)}$ as the benchmark in our competitive analysis [25].

Thus far we have demonstrated several auction mechanisms that perform comparably to $\mathcal{F}^{(2)}$. In that $\mathcal{F}^{(2)}$ is comparable to $\mathcal{F}$, these auctions perform comparably to $\mathcal{F}$. As $\mathcal{F}$ is the optimal single-price auction it is clear that no auction that always uses a single price can achieve a higher revenue than $\mathcal{F}$. This gives good justification for comparing the revenue of single-price auctions to $\mathcal{F}$ (or $\mathcal{F}^{(2)}$).

The question remains of whether or not we have chosen the "best" metric possible for comparison when considering multi-priced auctions. There are two primary goals in choosing the metric. First, we wish to achieve the best possible performance and thus we would like to compare ourselves against the strongest possible benchmark. On the other hand, we would like to find a natural metric that comes as close as possible to capturing the performance of the best truthful auction across a wide range of inputs.

Let us reconsider for a moment the strongest possible benchmark, the optimal omniscient multi-price auction, $\mathcal{T}$, that achieves as revenue the sum of the bids. Lemma 3.1 showed that the profit of $\mathcal{T}$ can be a factor of $\Theta(\log n)$ more than $\mathcal{F}$. There remains the question: Does there exist an alternative metric on bids $\mathbf{b}$ that captures the benefit of using multiple prices (and hence is sometimes larger than the $\mathcal{F}(\mathbf{b})$) that we could compare our truthful multi-priced auctions against?

As evidence to the contrary, we will show that no *monotone* auction can achieve an expected profit higher than $\mathcal{F}$ *on any input*. Monotone auctions are a large class of natural multi-priced truthful auctions; all the competitive auctions presented in this thesis are monotone. This result supports the conjecture that there is no systematic way for an auction to achieve a higher profit than $\mathcal{F}$.

8.1 Hard-coded Auctions

We begin with some motivation for the notion of monotonicity by presenting some examples of non-monotone auctions. These auctions are not natural in the context of worst-case profit maximization in that while they can achieve significantly higher profit than that of $\mathcal{F}$ on certain bid sets (in fact they can achieve $\mathcal{T}(\mathbf{b})$), they do so at the cost of having horrendously low profit (and hence very bad competitive ratio) on other inputs.

We first observe that for every $\mathbf{b}$ there is a symmetric truthful auction that achieves a profit of $\mathcal{T}(\mathbf{b})$. For example, consider the n-tuple $\mathbf{b}$ with half of the bids at value one and half at value $h > 2$. Thus, $\mathcal{T}(\mathbf{b}) = (h+1)n/2$ and $\mathcal{F}(\mathbf{b}) = hn/2$. Consider the symmetric auction given by bid-independent function f:

$$f(\mathbf{b}_{-i}) = \begin{cases} 1 & \text{if more } h\text{s than 1s in } \mathbf{b}_{-i}. \\ h & \text{otherwise.} \end{cases}$$

This auction achieves profit $\mathcal{T}(\mathbf{b})$ on our particular input: If $b_i = h$ then $f(\mathbf{b}_{-i})$ outputs h; on $b_i = 1$, $f(\mathbf{b}_{-i})$ outputs 1. Note that on most other inputs, $\mathbf{b}'$, this auction performs much worse than $\mathcal{F}$.

Now we generalize this result and show that for any set of bids $\mathbf{b}^*$ there exists a truthful (symmetric) auction that achieves a revenue of $\mathcal{T}(\mathbf{b}) = \sum_i b_i^*$. This is exemplified in the bid-independent auction $\mathcal{A}_{f_{\mathbf{b}^*}}$ parameterized by $\mathbf{b}^*$ and defined as follows:

$$f_{\mathbf{b}^*}(\mathbf{b}_{-i}) = \begin{cases} b_j & \text{if } \pi(\mathbf{b}_{-i}) = \mathbf{b}^*_{-j} \text{ for some permutation } \pi \\ \infty & \text{otherwise.} \end{cases}$$

The "otherwise" case is arbitrarily chosen. In fact any number of bid vectors that have a pairwise difference of at least two bid values can be hard-coded into an auction in this manner. The auction will perform very poorly on any input that differs on only one bid value from one of the hard-coded bid vectors.

For worst case profit maximization, the mechanism of both of these auctions is counter intuitive. For the case that half the bids are at one and half at h, the bid-independent function sees more h values and outputs one. When it sees less h values, it outputs h. A more intuitive output for a profit maximizing auction would be to output h when there are more bids at value h.

Note that we can combine a hard-coded auction and a competitive auction by flipping a fair coin and running the former or the latter depending on the outcome of the toss. The resulting auction is competitive, with the competitive ratio twice that of the underlying competitive auction. Furthermore, on the hard-coded input $\mathbf{b}^*$, the expected revenue of the auction is at least $\mathcal{T}(\mathbf{b}^*)/2$, which can be significantly bigger than $\mathcal{F}(\mathbf{b}^*)$.

Although a competitive auction can outperform $\mathcal{F}$ on some inputs, we conjecture that this happens at the expense of the competitive ratio and auctions designed to achieve high competitive ratios do not outperform $\mathcal{F}$ in this sense. In the next section we introduce the class of *monotone auctions* that includes $\mathrm{BI}_{\mathrm{opt}}$, RSOP, RSPE, CORE, and Vickrey. We show that no monotone auction can outperform $\mathcal{F}$.

8.2 Monotonicity

The intuition underlying our notion of monotonicity is that the bid-independent function defining the auction should output higher prices when it sees higher bid values.

Definition 8.1 *An auction is* monotone *if for any pair of bidders i and j with $b_i \leq b_j$, we have:*

$$\forall x \leq b_i,\ \mathbf{Pr}[\textit{bidder } i \textit{ wins at price} \leq x] \leq \mathbf{Pr}[\textit{bidder } j \textit{ wins at price} \leq x]\,.$$

The intuition is that if $b_i < b_j$, then $\mathbf{b}_{-i}$ looks like a higher set of bids than $\mathbf{b}_{-j}$. Therefore, the price bidder i pays if he wins will tend to be higher than the price bidder j pays if he wins.

The class of monotone auctions is very general. It is not difficult to verify that the Vickrey auction with a reservation price is monotone. Thus, the optimal Bayesian auctions for i.i.d. prior distributions are monotone. Analysis of $\mathrm{BI}_{\mathrm{opt}}$, RSOP, RSPE, and CORE shows that they are also monotone. Proofs are given in Section 8.3. The optimal single-price omniscient auction, $\mathcal{F}$, is also monotone.

$\mathcal{F}$ is clearly monotone as all bidders bidding above the sale price win and pay the same price. We now show that $\mathcal{F}$ is the optimal monotone auction. This further justifies our comparison to $\mathcal{F}$ and the related metric $\mathcal{F}^{(2)}$.

Theorem 8.1 *Let $\mathcal{A}$ be any monotone (truthful) randomized auction. For all bid vectors* $\mathbf{b}$*, the revenue $\mathcal{R} = \sum_i p_i$ of $\mathcal{A}$ on input* $\mathbf{b}$ *satisfies*

$$\mathbf{E}[\mathcal{R}] \leq \mathcal{F}(\mathbf{b}).$$

Proof: Let f be the bid-independent function defining $\mathcal{A}$. For each i, define $g_i(x) = \mathbf{Pr}[f(b_{-i}) \leq x]$.

Now consider the following thought experiment. Let U be a random variable that is uniform on $[0, 1]$. Imagine running the bid-independent auction that for each i uses $g_i^{-1}(U)$ to set the threshold for bidder i, with g_i^{-1} defined as $g_i^{-1}(y) = \inf\{x : g_i(x) = y\}$. We denote by $\mathcal{R}_U$ the resulting auction revenue. We observe that the threshold distribution for bidder i in this experiment is precisely the same as the original threshold distribution for bidder i:

$$\mathbf{Pr}\left[g_i^{-1}(U) \leq x\right] = \mathbf{Pr}[U \leq g_i(x)] = g_i(x).$$

Therefore, by summing the expectations for the bidders, we obtain

$$\mathbf{E}[\mathcal{R}_U] = \mathbf{E}[\mathcal{R}].$$

We complete the proof by showing that the expected revenue from our thought experiment $\mathbf{E}[\mathcal{R}_U]$ is at most $\mathcal{F}(\mathbf{b})$. Conditioned on $U = u$, let k be the index of the smallest winning bid. Thus, $g_k^{-1}(u) \leq b_k$. Since $\mathcal{A}$ is monotone, for $x \leq b_k$ and all j with $b_j \geq b_k$, we have $g_k(x) \leq g_j(x)$. Furthermore, $g_k(x)$ and $g_j(x)$ are monotone non-decreasing functions. Therefore, it must be that $g_j^{-1}(u) \leq g_k^{-1}(u) \leq b_k \leq b_j$ and therefore all bidders with bid values at least b_k win at a price at most b_k. Thus, the revenue, $\mathcal{R}_u$, is at most b_k times the number of bids with bid value least b_k which totals to at most $\mathcal{F}(\mathbf{b})$. This holds for all $u \in [0, 1]$, and thus $\mathbf{E}[\mathcal{R}_U] \leq \mathcal{F}(\mathbf{b})$. □

8.3 Analysis of Auctions

In this section we show that the auctions $\mathrm{BI}_{\mathrm{opt}}$, RSOP, RSPE, and CORE are monotone.

8.3.1 BI$_{\mathrm{opt}}$ *is monotone*

As BI$_{\mathrm{opt}}$ is a symmetric auction, the order of the bids in the input does not affect the outcome. For convenience of notation we will assume that they are indexed from highest, b_1, to lowest, b_n. We denote by $t = \mathrm{opt}(\mathbf{b})$ and $t_1 = \mathrm{opt}(\mathbf{b}_{-1})$ the optimal sale price for $\mathbf{b}$ and $\mathbf{b}_{-1}$, respectively.

Theorem 8.2 BI$_{\mathrm{opt}}$ *is monotone as*

- *It uses the single sale price t_1 for all winners.*

- *All bidders that bid above the lowest winning bid also win.*

Proof: Recall that the bid-independent function that implements BI$_{\mathrm{opt}}$ is $f(\mathbf{b}) = \mathrm{opt}(\mathbf{b}) = \mathrm{argmax}_{b_k}\, kb_k$. We will show that for any bidder i that wins the auction, bidder $i-1$ also wins the auction and at the same price. From this, a simple induction gives the theorem.

In the case that $b_i = b_{i-1}$ the fact that BI$_{\mathrm{opt}}$ is symmetric implies that they must both win at the same price. Now consider the case that $b_{i-1} > b_i$. Since bidder i wins, the computation of $\mathrm{opt}(\mathbf{b}_{-i})$ must find the maximum of kb_k for $k < i$ and $(k-1)b_k$ for $k > i$ to be $(k^*-1)b_{k^*}$ for $k^* > i$. The only difference between the computation of $\mathrm{opt}(\mathbf{b}_{-i})$ and of $\mathrm{opt}(\mathbf{b}_{-i+1})$ is that for $\mathrm{opt}(\mathbf{b}_{-i})$ we consider $(i-1)b_{i-1}$ and for $\mathrm{opt}(\mathbf{b}_{-i+1})$ we consider $(i-1)b_i$. Since $\mathrm{opt}(\mathbf{b}_{-i})$ finds $(k^*-1)b_{k^*}$ bigger than the other values it considers and since $(i-1)b_{i-1} > (i-1)b_i$, it must be that $(k^*-1)b_{k^*}$ is the biggest value that $\mathrm{opt}(\mathbf{b}_{-i+1})$ considers as well. □

8.3.2 RSOP *Auction*

We will show that RSOP is monotone with respect to any two bidders i and j. First fix i and j such that $b_i \leq b_j$. The randomness of RSOP is in the partitioning. We verify the monotonicity of RSOP by presenting a disjoint grouping of the partitionings such that each group is monotone with respect to i and j and the RSOP auction is just a randomization over monotone groups.

We put partitionings with b_i and b_j in the same partition in their own group. Note that these groups are monotone with respect to b_i and b_j because both b_i and b_j are offered the same price.

Any partitioning with b_i and b_j in opposite partitions we will pair in a group with the partitioning we get when we swap b_i with b_j. This gives us a group with two partitionings:

Partitioning P: $B_1 \cup \{b_i\}$ and $B_2 \cup \{b_j\}$.

Partitioning P': $B_1 \cup \{b_j\}$ and $B_2 \cup \{b_i\}$.

We show that this group $\{P, P'\}$ is monotone with respect to b_i and b_j if one of the two partitionings is chosen by flipping a fair coin. Consider the thought experiment where for b_i a "heads" coin flip means use P and "tails" means use P', but for b_j a "heads" coin flip means use P' and "tails" for P. In this experiment, "heads" results in b_i getting the optimal price for $B_2 \cup \{b_j\}$ and b_j getting the optimal price of $B_2 \cup \{b_i\}$. This is just the outcome of $\mathrm{BI}_{\mathrm{opt}}$ on the bid set $B_2 \cup \{b_i, b_j\}$. Likewise on "tails", the outcome is that of $\mathrm{BI}_{\mathrm{opt}}$ on the bid set $B_1 \cup \{b_i, b_j\}$. By Theorem 8.2, both of these outcomes are monotone. Note that the outcome for b_i and b_j as random variables in our thought experiment are the same as the random variables for their actual outcomes for this grouping. Thus, this grouping is monotone with respect to bidders i and j.

8.3.3 RSPE *Auction*

To show that RSPE is monotone, we will use the same general approach as for RSOP of finding disjoint groupings in the partitionings and showing that each is itself monotone with respect to two bids b_i and b_j with $b_i \leq b_j$.

As with RSOP, partitionings with both b_i and b_j on the same partition are themselves monotone as the cost sharing mechanism is monotone: it gives an outcome such that all winning bidders pay the same price and all bidders whose bid value is above this price win.

Likewise, we pair a partitioning with b_i and b_j in different partitions, with the partitioning with b_i and b_j swapped. We get partitionings P and P' as defined above.

$$F_j = \mathcal{F}(B_2 \cup \{b_j\}) \qquad\qquad F_i = \mathcal{F}(B_1 \cup \{b_i\})$$
$$F'_i = \mathcal{F}(B_2 \cup \{b_i\}) \qquad\qquad F'_j = \mathcal{F}(B_1 \cup \{b_j\})$$

Since $b_i \leq b_j$ we have:

$$F_j \geq F'_i \qquad\qquad F'_j \geq F_i$$

Note that if b_i does not win for either P or P' then this grouping is trivially monotone with respect to bidders i and j. Otherwise, suppose b_i wins in partitioning P and pays price p_i. We will show that b_j wins in P' and pays $p'_j \leq p_i$.

Let pe_R be the bid-independent function for $\mathcal{O}$-Extract$_R$. Recall that for partitioning P, the price for b_i is computed by running $\mathcal{O}$-Extract$_{F_j}$ on $B_1 \cup \{b_i\}$. For partitioning P' the price for b_j is computed by running $\mathcal{O}$-Extract$_{F'_i}$ on $B_1 \cup \{b_j\}$. Thus,

$$p_i = \mathrm{pe}_{F_j}(B_1) \geq \mathrm{pe}_{F'_i}(B_1) = p_j.$$

The intermediate step here follows because $F_j \geq F'_i$ and because pe_R is monotone in R.

8.3.4 CORE *Auction*

We now consider the Consensus Revenue Estimate auction (CORE). This auction performs the 1-item Vickrey auction with some probability, and otherwise does the bid-independent cost sharing with a bid-independent revenue estimate computed as $\mathcal{F}^{(2)}$ rounded down to the nearest c^{i+U} for integer i, $c > 1.5$, and U uniform $[0, 1]$. We show that this is just a randomization over deterministic monotone auctions. First, the 1-item Vickrey auction is monotone. Second, on success, the outcome of CORE is the same as the outcome of $\mathcal{O}$-Extract$_R$ with R the largest c^{i+U} less than $\mathcal{F}^{(2)}(\mathbf{b})$. $\mathcal{O}$-Extract$_R$ is monotone for any R, thus on success, the consensus revenue estimate technique is monotone. It remains to show that on failure, when there is no consensus, that the outcomes are monotone. Order the bids from larger to smaller, that is $b_j \geq b_i$ for $j < i$. Let R_i be the consensus value used for b_i. That is R_i is the largest value of the form c^{i+U} that is at most $\mathcal{F}^{(2)}(\mathbf{b}_{-i})$. $\mathcal{F}^{(2)}(\mathbf{b}_{-i})$ is non-decreasing in i so R_i is as well. Consider any bid that wins, b_i. Since b_i wins, it must be that $p_i = \mathrm{pe}_{R_i}(\mathbf{b}_{-i}) \leq b_i$. We now show that all bids higher than i win and they do so at price at most p_i. To do so, we note that for all $j < i$ (i.e., $b_j \geq b_i$) $\mathrm{pe}_{R_i}(\mathbf{b}_{-j}) = \mathrm{pe}_{R_i}(\mathbf{b}_{-i})$, and that $\mathrm{pe}_{R_j}(\mathbf{b}_{-j}) \leq \mathrm{pe}_{R_i}(\mathbf{b}_{-j})$. The first is true from the definition $\mathcal{O}$-Extract$_{R_i}$. The second is true because $R_i \geq R_j$ and by the definition of $\mathcal{O}$-Extract$_{R_i}$ as the auction that sells to largest set of bidders that can afford to equally share the price R_i. Certainly the sale price can not go up when R_i is lowered to R_j. Thus CORE is monotone.

8.4 Non-monotone variant of $\mathrm{BI}_{\mathrm{opt}}$

As a proof of concept, we now exhibit an auction which is a variant of $\mathrm{BI}_{\mathrm{opt}}$ that is both non-monotone and performs achieves a revenue that is higher than $\mathcal{F}$ on some bid sets.

Recall that for each bid, $\mathrm{BI}_{\mathrm{opt}}$ computes the sale price as the optimal sale price for the remaining bids. Define $\mathrm{BI}_{\mathrm{opt}}{}'$ as the auction for each bid, uses a sale price that is 2/3 of the optimal sale price for the other bids. Consider the bids $\mathbf{b} = (9/4, 3/2, 1)$. The following table shows the result of $\mathrm{BI}_{\mathrm{opt}}{}'(\mathbf{b})$.

b_i	b_1	b_2	b_3	$2\,\mathrm{opt}\,/3$
$\mathbf{b}$	9/4	3/2	1	
$\mathbf{b}_{-1}$	?	3/2	1	2/3
$\mathbf{b}_{-2}$	9/4	?	1	3/2
$\mathbf{b}_{-3}$	9/4	3/2	?	1

Note that for all i, $b_i \geq 2\,\mathrm{opt}(\mathbf{b}_{-i})/3$ so all bidders win. Also note that the sale prices are not monotone. In particular, bidder 2 bid higher than bidder 3 but is required to pay a higher price. Note that $\mathcal{F}(\mathbf{b}) = 3$ and the revenue of $\mathrm{BI}_{\mathrm{opt}}{}'(\mathbf{b}) = 1 + 3/2 + 2/3 > 3$.

Part II

Extensions

Chapter 9

THE MULTI-ITEM AUCTION PROBLEM

Consider an airplane flight where passengers have individual movie screens and can choose to view one out of a dozen movies that are broadcast simultaneously. The flight is only long enough for one movie to be seen. The airline wants to price movies to maximize its revenue. Currently, airlines charge a flat fee for movies. Even if the fee is based on a careful marketing study, passenger demographics may vary from one flight to another, and individual utilities can vary with flight route, time of the year, etc. Therefore a non-adaptive pricing is unlikely to be optimal for the seller. We investigate adaptive pricing via auctions.

In the scarce supply case, multiple item auctions have been studied by Shapley and Shubik [49]. (See [44] for a survey of the area.) Results for the scarce case, however, do not directly apply to the unlimited supply case. Consider the case where each item for sale is unique – for example the real estate market considered in [49]. In this case consumers will bid heavily for highly desirable items, which will sell for a high price. In contrast, in the unlimited supply case the seller can in principle give every consumer a copy of the item the consumer desires most. However, in such an auction, the consumer has no incentive to bid high. Thus a good auction mechanism must in some cases limit the number of copies of each item.

In this chapter we extend some of the results for basic auctions to unlimited supply multi-item auctions. In particular, we show how the random sampling technique along with the fixed pricing mechanism can be used to obtain a solution to the multi-item auction problem.

One of our main results is on a sampling problem that may be of independent interest. A variant of the sampling problem is as follows. Suppose we have n applicants and m tests. Each applicant takes each test and gets a real-valued score. We have to select k applicants based on the results of these scores. Furthermore suppose that we choose a random subset of the applicants, call the applicants in the subset red, and call the remaining applicants blue. After the results of the tests are known and the subset is selected, an adversary selects the k winning applicants while obeying the following restriction: If an applicant x is accepted and for every test, applicant y get a score that is at least as good as the score of x, then y must be accepted as well. The adversary's goal is to bias the admission in favor of red applicants. Although we study a slightly different problem, our techniques can be used to show that if $k = o(m^2 \log n)$, then with high probability the ratio of the number of red applicants to the number of blue applicants is bounded by a constant.

This problem seems natural. One can view candidates as points in m-dimensional space, and view the adversary as selecting a shift of the positive quadrant so that the shifted quadrant contains k points total and as many red points as possible.

9.1 Definitions

The input to an auction is a number of bidders, n, a number of items, m and a set of bids $\{a_{ij}\}$. We study the case when each bidder wants only a single item, i.e., the *unit demand* case. We say that an auction is *single-priced* if the sale prices for copies of the same item are the same, and *multi-priced* otherwise.

Our analysis framework for auction mechanisms for the multi-item problem will be similar to that used for the basic auction except that we will be considering a promise version of the problem. We will give auctions that achieve a profit that is a constant fraction of $\mathcal{F}$, the profit of the optimal envy-free outcome under the assumption that the input is well behaved. Our promise will be based on $\mathcal{F}$ being large compared to h, the ratio between the highest and lowest bid.

For convenience, we assume that the input bids are *non-degenerate*, i.e., all input bids values a_{ij} are distinct or zero. This assumption can be made without loss of generality because we can always use lexicographic tie-breaking to achieve it.

Our competitive multi-item auctions are randomized. We use the following lemma, which is a variation of the Chernoff bound (see e.g. [12, 36]), as the main tool in our analysis.

Lemma 9.1 *Consider a set A and its subset $B \subset A$. Suppose we pick an integer k such* *that* $0 < k < |A|$ *and a random subset (sample)* $S \subset A$ *of size* k*. Then for* $0 < \delta \leq 1$ *we have*

$$\mathbf{Pr}[|S \cap B| < (1-\delta)|B| \cdot k/|A|] < \exp(-|B| \cdot k\delta^2/(2|A|)).$$

Proof: We refer to elements of A as points. Note that $|S \cap B|$ is the number of sample points in B, and its expected value is $|B| \cdot k/|A|$. Let $p = k/|A|$. If instead of selecting a sample of size exactly k we choose each point to be in the sample independently with probability p then the Chernoff bound would yield the lemma.

Let $A = \{a_1, \ldots, a_n\}$ and without loss of generality assume that $B = \{a_1, \ldots, a_k\}$. We can view the process of selecting S as follows. Consider the elements of A in the order induced by the indices. For each element a_i considered, select the element with probability p_i, where p_i depends on the selections made up to this point.

At the point when a_{i+1} is considered, let t be the number of currently selected points. Then $i-t$ is the number of points considered but not selected. Suppose that $t/i < p$. Then $p_{i+1} > p$.

We conclude that when we select the sample as a random subset of size k, the probability that the number of sample points in B is less than the expected value is smaller than in the case we select each point to be in the sample with probability p.

□

9.2 Fixed Price Auction and Optimal Prices

Consider the following *fixed price auction.* The bidders supply the bids and the seller supplies the *sale prices*, r_j, $1 \leq j \leq m$. Define $c_{ij} = a_{ij} - r_j$. The auction assigns each bidder i to the item j with the maximum c_{ij}, if the maximum is nonnegative, and to no item otherwise. Note that this is precisely the item that maximizes the bidder's profit. In case of a tie, we choose the item with the maximum j. If a bidder i is assigned item j, the corresponding sale price is r_j.

Lemma 9.2 *Suppose the sale prices are set independently of the input bids. Then the fixed price auction is truthful.*

Proof: If bidder i gets object j, the bidder's price is at least r_j and the bidder's profit is at most $a_{ij} - r_j$. The best possible profit for i is $\max_j(u_{ij} - r_j)$. If the bidder bids $a_{ij} = u_{ij}$, this is exactly the profit of the bidder.

□

Remark. Although we assume that the bidders do not see sale prices before making their bids, the lemma holds even if the bidders do see the prices.

Now consider the following *optimal pricing* problem: Given a set of bids, find the set of prices such that the fixed price auction brings the highest revenue. Suppose an auction solves this problem and uses the resulting prices. We call this auction the *optimal omniscient single-price auction* and denote its revenue by $\mathcal{F}$. We can interpret $\mathcal{F}$ as the revenue of fixed pricing using perfect market analysis.

9.3 The Random Sampling Auction

We use random sampling to make the optimal single-price auction truthful.

Definition 9.1 (Random Sampling Optimal Price Multi-item Auction) *The* Random Sampling Optimal Price Multi-item Auction *works as follows:*

1. *Partition bids uniformly at random into two sets, S and N.*
2. *For both sets, compute the optimal prices, $\mathbf{r}_S$ and $\mathbf{r}_N$.*
3. *Run the fixed pricing mechanism with prices $\mathbf{r}_S$ on bidders in N.*
4. *Run the fixed pricing mechanism with prices $\mathbf{r}_N$ on bidders in S.*

Lemma 9.3 *The random sampling auction is truthful.*

Remark. Another natural way of sampling is to sample bids instead of bidders. However, this does not lead to a truthful auction, because bidder's bids selected in the sample may influence the price used to satisfy the bidder's remaining bids.

Next we show that, under certain assumptions, in expectation the auction's revenue is within a constant factor of $\mathcal{F}$. Without loss of generality, for every $1 \leq i \leq n, 1 \leq j \leq m$, if a_{ij} is undefined (not in the input) we define a_{ij} to be zero. For every bidder i, we view $(a_{i1}, \ldots, a_{im})$ as a point in the m-dimensional space and denote this point by v_i. Thus v_i is in the quadrant Q of the m-dimensional space where all coordinates are nonnegative. We denote the set of all input points by B.

Note that by the symmetry in the problem and linearity of expectation, the expected revenue collected from S and N are equal and sum to the total expected revenue of the auction. For ease of discussion we will analyze the contribution of the revenue from S''. Let $\mathcal{R}$ be the random variable representing this revenue. As such, S and N will play different roles in our analysis. We will refer to S as the *sample* as it will be used to compute prices for N. We will refer to N as the *non-sample.*

For a fixed m and sale prices $(r_1, \ldots, r_m)$, let R_j be a region in the m-dimensional space such that if $v_i \in R_j$, then bidder i prefers j to any other item, i.e., for any $1 \leq k \leq m$, $c_{ij} \geq c_{ik}$ (recall that $c_{ij} = a_{ij} - r_j$). We would like $\{R_j : 1 \leq j \leq m\}$ to be a partitioning of Q. We achieve this by assigning every boundary point to the highest-index region containing the point. (This is consistent with our tie-breaking rule for the fixed price auction.) R_j is a convex (and therefore connected) region in Q. In fact, the region R_j is as follows:

$$R_j = \{x : x_j \geq r_j \; \& \; x_j - r_j \geq x_k - r_k \; \forall k \neq j\}. \tag{9.1}$$

Figure 9.1 shows a two item auction with prices r_1 and r_2 for items 1 and 2 respectively. These prices induce the regions $R_1 = R_1' \cup R_1''$ and $R_2 = R_2' \cup R_2''$. Arrows point to selling prices for the bidders in each region.

Thus sampling and computing r_j's partitions Q into the regions, and each bidder i in N gets the item corresponding to the region that i is in. Intuitively, our analysis says that if a region has many sample points, it must have a comparable number of non-sample points – even though the regions are defined based on the sample. The latter fact makes the analysis difficult by introducing conditioning. Intuitively, we deal with the conditioning by considering regions defined by the input independently of the sample.

For a given input, let $(q_1, \ldots, q_m)$ be the optimal prices for the input bids that yield revenue $\mathcal{F}$. These prices induce the regions discussed above. Bidders in region R_j pay q_j for the item j. If we sample half of the points, the expected number of sample points in a

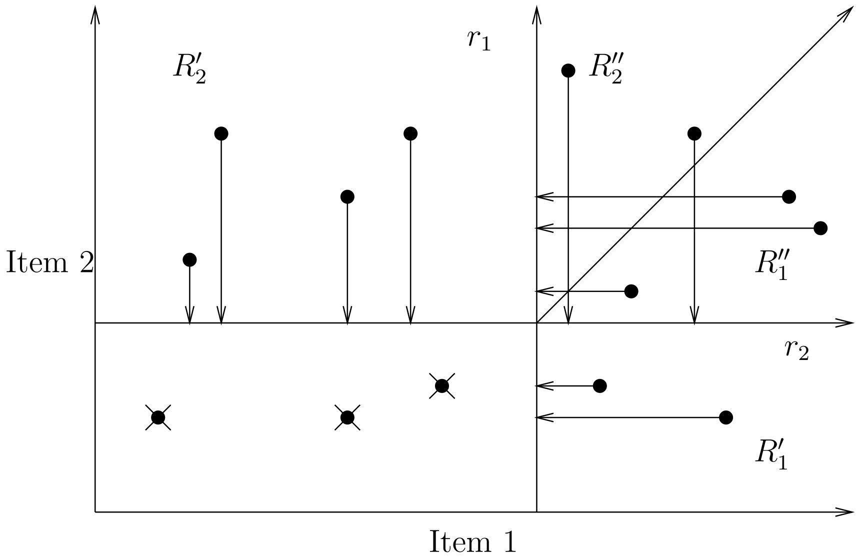

Figure 9.1: Two item auction with regions R_1 and R_2

region R_j is half of the total number of points in the region, and for the prices $q_1, \ldots, q_m$, the expected revenue is $\mathcal{F}/2$. The optimal fixed pricing on the sample does at least as well. Thus the expected revenue of optimal fixed pricing of the sample, $\mathbf{E}[\mathcal{F}_s]$, is at least $\mathcal{F}/2$. However, we need a high-probability result. Our goal is to show that with high probability $\mathbf{E}[\mathcal{F}_s]$ is close to $\mathcal{F}/2$ and that $\mathbf{E}[\mathcal{R}]$ is close to $\mathbf{E}[\mathcal{F}_s]$, where $\mathcal{R}$ is the revenue of the random sampling auction.

We say that a set $A \subseteq B$ is *t-feasible* if A is nonempty and for some set of sale prices, A is exactly the set of points in R_t. For each feasible set A, we will define its signature $S_A = (s_1, \ldots, s_m)$ such that s_i's are (not necessarily distinct) elements of A and, for a fixed t, different t-feasible sets have different signatures. In the following discussion, s_{ij} denotes the j-th coordinate of s_i.

We construct signatures as follows. Let R_t be a region defining A. R_t is determined by a set of prices $(r_1, \ldots, r_m)$. We first increase all r_j's by the same amount (moving R_t diagonally) until some point in A is on the boundary of R_t. Note that since we change all prices by the same amount, the limiting constraint from (9.1) is $x_t \geq r_t$. Thus the stopping is defined by $x_t = r_t$, and the point on the boundary has the smallest t-th coordinate among

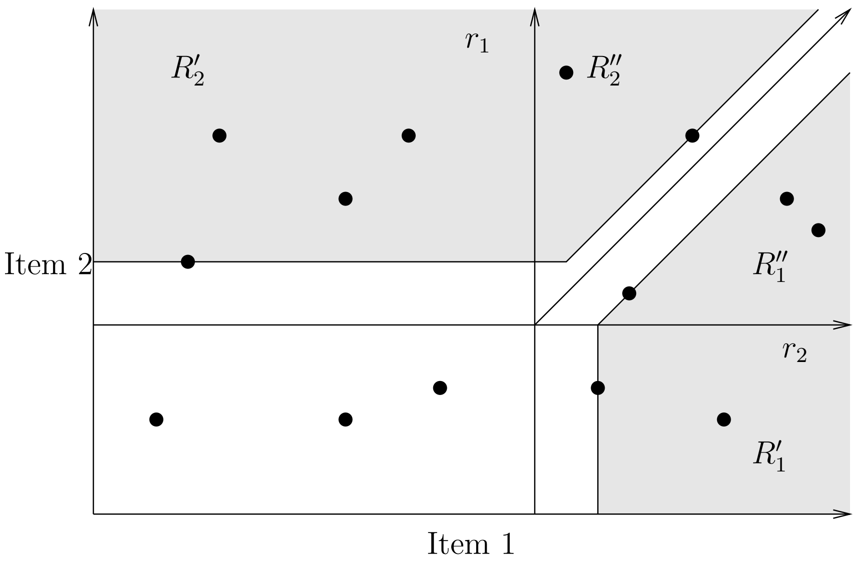

Figure 9.2: Signatures in a two item auction

the points in A. We set s_t to this point.

Then for $j \neq t$, we move the the region starting at its current position down the j-th coordinate direction by reducing r_j until the first point hits the boundary. The boundary we hit is defined by $x_t - r_t = x_j - r_j$, and the point that hits it first has the minimum $x_j - x_t + s_{tt}$ among the points in A. Observe that the point s_t remains on the boundary $x_t = r_t$, and therefore we stop before r_j becomes negative. When we stop, we take a point that hits the boundary and assign s_j to it.

Consider the set of points in the signature, $S_A = \{s_1, \ldots, s_m\}$. Define R to be the region we got at the end of the procedure that computed S_A. R is defined by

$$R = \{x : x_t \geq s_{tt} \ \& \ x_t - s_{tt} \geq x_j - s_{jj} \ \forall j \neq t\}.$$

It follows that R can be constructed directly from S_A.

Figure 9.2 shows the signatures we get from the prices r_1 and r_2. The points on the boundary of the shaded region are the signature of that region. Note, for example, that there are no points in R_1 that are not inside the boundary induced by the signature for R_1.

The next two lemmas are simple, so we omit the proofs.

Lemma 9.4 *For each t, $1 \leq t \leq m$, there are at most n^m t-feasible sets.*

Lemma 9.5 *For every t-feasible subset C of the sample S there is a t-feasible subset A of the input such that $C = A \cap S$.*

For $k \geq 1$ and $0 < \delta < 1$, we say that a sample S is (k, δ)-*balanced* if for every $1 \leq t \leq m$ and for every t-feasible subset of the input, A, such that $|A| \geq k$, we have

$$(1 - \delta) \leq (|A \cap S|)/(|A \cap N|) \leq 1/(1 - \delta).$$

Lemma 9.6 *The probability that a sample containing half of the input points is (k, δ)-balanced is at least $1 - 2mn^m \exp(-k\delta^2/8)$.*

Proof: Lemma 9.1 implies that the probability that for a set A with $|A| \geq k$,

$$\mathbf{Pr}[|A \cap S| < (1 - \delta)|A \cap N|] < \exp(-k\delta^2/8)$$

and

$$\mathbf{Pr}[|A \cap N| < (1 - \delta)|A \cap S|] < \exp(-k\delta^2/8).$$

Note that the fact that the number of sample points in one subset is close to its expectation makes it no less likely that the number of sample points in another subset is close to expectation. Thus the conditioning we get is favorable. By Lemma 9.5, there are at most n^m t-feasible subsets for every t, so the total number of feasible subsets is mn^m. These observations imply the lemma. $\square$

Theorem 9.7 *Assume $\alpha h m^2 \ln n \leq \mathcal{F}$ and $m \geq 2$. Then $\mathcal{R} \geq \mathcal{F}/24$ with probability at least $1 - \exp(-\alpha/1728)$ (for some constant $\alpha > 1$).*

Proof: Consider Lemma 9.6 with $\delta = 1/2$ and $k = \alpha m \log n/12$. The probability that the sample is (k, δ)-balanced is

$$1 - 2mn^m \exp(-k\delta^2/8) = 1 - 2mn^m \exp(-\alpha m \log n/864) \geq 1 - \exp(-\alpha/1728)$$

for $m \geq 2$. For the rest of the proof we assume that the sample is (k, δ)-balanced; we call this the balanced sample assumption.

Next we show that the revenue of the auction on the sample, $\mathcal{F}_s$, satisfies $\mathcal{F}_s \geq \mathcal{F}/6$. Let Q_i be the set of bidders who get item i when computing $\mathcal{F}$ on the entire bid set. Consider sets Q_i containing less than $(\alpha m \log n)/2$ bidders. The total contribution of such sets to $\mathcal{F}$ is less than $\mathcal{F}/2$. This is because there are at most m such sets and each bid is at most h giving a maximum possible revenue of $\alpha h m^2 \log n/2 = \mathcal{F}/2$. Thus the contribution of the sets with at least $(\alpha m \log n)/2$ bidders is more than $\mathcal{F}/2$, and we restrict our attention to such sets. By the balanced sample assumption, each such set contains at least 1/3 sample points, and thus $\mathcal{F}_s \geq (1/3)\mathcal{F}/2 = \mathcal{F}/6$.

Finally we show that $\mathcal{R} \geq \mathcal{F}/24$ using a similar argument. Let R_i be the regions defined by the prices computed by the auction on the sample. Consider the regions containing less than $(\alpha m \log n)/12$ sample points. The total contribution of such sets to the revenue is less then $\mathcal{F}/12$. The remaining regions contribute at least $\mathcal{F}/12$ (out of $\mathcal{F}/6$). Each remaining region contains at least $(\alpha m \log n)/12$ sample points. By the balanced sample assumption, each such region contains at least one non-sample point for every two sample point and thus $\mathcal{R} \geq \mathcal{F}/24$. □

Lemma 9.3 and Theorem 9.7 imply that if the assumptions of the theorem hold, the random sampling auction is competitive.

9.4 Concluding Remarks

Our analysis of the random sampling auction is somewhat brute-force, and a more careful analysis may lead to better results, both in terms of constants and in terms of asymptotic bounds. In particular, the assumption $\alpha h m^2 \ln n \leq \mathcal{F}$ in Theorem 9.7 may be stronger than necessary. One can prove that $\mathcal{F}_s = \Omega(\mathcal{F})$ assuming $\alpha h m \leq \mathcal{F}$. We wonder if the theorem holds under this weaker assumption.

Although our theoretical bounds require m to be small compared to n and the optimal fixed price solution to contain a large number of items, it is likely that in practice our auctions will work well for moderately large m and moderately small optimal fixed price solutions. This is because our analysis is for the worst-case. In many real-life applications, bidder utilities for the same item are closely correlated and our auctions perform better.

Another open problem is a generalization of our results. One possible generalization is to the case when some items are in fixed supply. Another generalization is to the case when consumer i wants up to k_i items.

Chapter 10

AUCTION PROBLEMS WITH PRODUCTION COSTS AND MARKET SEGMENTATIONS

In this section we discuss a natural generalization of the basic auction problem to the case where the bidders are divided into markets and there may be a cost to the auctioneer that is a function of which markets have winning bidders. In particular, this generalization includes the *multicast pricing problem* (defined below) of Feigenbaum et al. [18]. To consider this problem, we first look at a special case, the *fixed cost basic auction problem* which can represent the problem of selling a digital good when there is a fixed cost to producing an initial copy, but once produced, duplicating is free. The goal then would be to ascertain via an auction mechanism whether the good should be created and sold or whether there is not enough demand for the good to cover the costs.

10.1 Introduction

We briefly present a few motivating problems. These problems are all extensions of the basic auction problem.

Conditional Financing

A company is considering an initial public offering of its shares, or a venture capitalist is trying to sell fixed return junk bonds for some venture. The company would like to sell the shares only if the revenue raised is sufficiently high. How can this be done so that the company raises as much money as is possible, given the potential shareholders' honest valuations of the company, but so that, if the company is not valued sufficiently highly by the public, the IPO can be cancelled?

Offering a pay-per-view broadcast in a segmented market

A company offering a pay-per-view broadcast needs to formulate a pricing scheme for its potential viewers. Suppose that the potential viewers are partitioned into markets (e.g., by location or by some measure of the quality of the good they are receiving). Further, suppose that the cost of providing the broadcast to viewers in the ith market is C_i, a fixed cost which is paid once if and only if there are any viewers in the ith market. The goal of the company is to maximize its profit, the sum of the prices paid by each of the viewers minus the costs of providing the broadcast to those markets in which there are viewers.

Multicast pricing

Feigenbaum, Papadimitriou and Shenker [18] initiated the study of pricing algorithms for multicast transmission. The model is a network with users residing at nodes in the network. There are costs associated with transmitting data across each of the links in the network. Each user has a utility for receiving the broadcast. The problem is to choose the multicast tree and the prices to charge each of the recipients of the broadcast. Feigenbaum et al. focus on the network complexity of implementing solutions that are either *budget-balanced*, in which the broadcaster precisely recovers the cost of the transmission, or *efficient*, in which the receivers chosen are the set which maximizes the difference between the sum of the utilities of the receivers and the cost of multicasting to that set of receivers. For budget-balanced solutions there is no profit to the broadcaster, and for efficient solutions the broadcaster may run a deficit. Our interest is in designing truthful pricing mechanisms which maximize the broadcaster's profit.

10.1.1 Generalized Auction Problems

We define a framework for modeling dynamic pricing problems for revenue maximization as *generalized auctions.* This framework captures all of the problems just described, and many others.

A *generalized auction problem* $\mathcal{G} = (\mathcal{S}, c(\cdot))$ is described by the following parameters:

- $\mathcal{S} = \{S_i \ : \ 1 \leq i \leq m\}$. The sets S_i partition the n bidders into m *market segments.*

Importantly, within each market, bidders are indistinguishable from one another.[1]

- $c(\cdot)$, a cost function mapping vectors in $\{0,1\}^m$ to non-negative real numbers. The domain of $c(\cdot)$ describes each possible market allocation, so, for example, a vector $\mathbf{r} = (r_1, \ldots, r_m)$ in the domain indicates for each market i whether or not goods are allocated to that market ($r_i = 1$ or $r_i = 0$, respectively). The cost, $c(\mathbf{r})$, is the cost to the auctioneer (the service provider) of providing the goods assuming the market allocation, $\mathbf{r}$.

So, for example, for the basic unlimited supply auction problem, $m = 1$, and $c(\mathbf{r}) = 0$ for all $\mathbf{r}$ as duplicating and distributing the goods is assumed to be free for the auctioneer. For the multicast pricing problem, the number of markets is the number of nodes in the network, and the viewers at the ith node form S_i, the ith market. The cost of a market allocation $c(\mathbf{r})$ is exactly the cost of transmitting along the multicast tree defined by the nodes (markets) receiving the broadcast.

10.1.2 Cancellable Auctions

A natural approach to the design of generalized auction mechanisms that do not incur deficits is to run a basic auction on each market and then cancel the results if the revenue raised from the bidders does not exceed the costs that the auctioneer will have to pay to provide goods to the selected receivers. Consider, for example, the problem of selling a digital good via a truthful auction, where the cost of producing the good is C, and, once produced, the good can be duplicated at negligible cost (so that effectively the auctioneer has an unlimited supply of the good). In this case, the auctioneer (who will bear the burden of producing the good), does not even want to produce the good unless they can recover their cost of C.

We are thus motivated to introduce a stronger form of truthful auctions, cancellable auctions. Given a truthful auction $\mathcal{A}$ and a parameter C, we define an auction $\mathcal{A}_C$ as follows: On input $\mathbf{b}$, run $\mathcal{A}$ on $\mathbf{b}$. If the resulting revenue is at least C, return the outcome

[1] Bidders in different markets may also turn out to be indistinguishable from one another if the cost structure of providing services to them is identical.

of $\mathcal{A}(\mathbf{b})$. Otherwise, *cancel* $\mathcal{A}$ by returning the outcome with no winners. We say that $\mathcal{A}$ is *cancellable* if, for any value of C, $\mathcal{A}_C$ is truthful. As we shall see in Section 10.3.1, not every truthful auction is cancellable.

10.1.3 Contributions

We present a truthful mechanism that obtains a constant factor of the optimal truthful profit on *any* input set of bids for which (a) there is competition in each market, i.e., there is not one bidder whose bid dwarfs all other bids; and (b) there is a significant profit margin to be had, i.e., the optimal profit is at least a constant factor more than the cost incurred by the optimal allocation. Furthermore, on any set of bids not satisfying these conditions, the mechanism incurs no deficit (See Section 10.4).

Thus, we obtain mechanisms that are profit maximizing for a broad class of generalized auction problems, including all the problems mentioned earlier in this section. For example, we obtain the first results on profit maximization for the conditional financing problem and for the multicast pricing problem in a worst-case, competitive analysis framework.[2] Although it is not the primary focus of this chapter, the multicast pricing mechanism we present also has low network complexity, in the sense of Feigenbaum et al. [18].

Optimal fixed-pricing

We first recall the definition of optimal fixed pricing for basic auctions, $\mathcal{F}(\mathbf{b}) = \max_i i b_{(i)}$.

We now extend this to general auctions:

Definition 10.1 *The optimal profit for any selling mechanism for $\mathcal{G} = (\mathcal{S}, c(\cdot))$ that uses a single price for each market is*

$$\mathcal{F}_{\mathcal{G}}(\mathbf{b}) = \max_{\mathbf{r} \in \{0,1\}^m} \left(\sum_{1 \le j \le m} r_j \mathcal{F}(\mathbf{b}_{S_j}) - c(\mathbf{r}) \right).$$

where $\mathbf{b}_{S_j}$ is the restriction of $\mathbf{b}$ to the jth market S_j.

[2] Recent work by Mehta et al. studies profit maximization for the multicast pricing problem in a Bayesian setting [34]

10.2 Competitive Framework

Given a mechanism $\mathcal{M}$, for each input $\mathbf{b} \in \mathbb{R}^n$, the *profit to the auctioneer*, $\mathcal{M}_{\mathcal{G}}(\mathbf{b})$, is the revenue from the bidders minus the cost of providing the goods, i.e.,

$$\mathcal{M}_{\mathcal{G}}(\mathbf{b}) = \sum_i p_i - c(\mathbf{r}).$$

The fact that $\mathcal{F}_{\mathcal{G}}(\mathbf{b})$ is an upper bound on the profit of any monotone truthful auction on input set $\mathbf{b}$ motivates the evaluation of truthful mechanisms by considering their competitive ratio relative to $\mathcal{F}_{\mathcal{G}}$, i.e., the supremum over all bid vectors of $\mathcal{F}_{\mathcal{G}}(\mathbf{b})/\mathcal{M}_{\mathcal{G}}(\mathbf{b})$. Unfortunately, in many cases, there is no constant bound achievable on this competitive ratio. For example, for the $\mathcal{G}$ as the basic auction problem (with no costs), $\mathcal{F}_{\mathcal{G}}(\mathbf{b})$ is equal to the optimal fixed price revenue $\mathcal{F}(\mathbf{b})$. In Chapter 3 we showed that no randomized basic auction can obtain a constant fraction of $\mathcal{F}(\mathbf{b})$. Recall that intuition that if there is one very high bidder that completely dominates all other bidders, there is no way to truthfully extract a constant fraction of their bid value. Instead of $\mathcal{F}$ we considered $\mathcal{F}^{(2)}$ as our metric for comparison for the basic auction problem. Recall, that $\mathcal{F}^{(2)}$ is defined as the optimal fixed price revenue assuming at least two items are sold (again, $b_{(i)}$ is the ith largest bid).

$$\mathcal{F}^{(2)}(\mathbf{b}) = \max_{i \geq 2} i b_{(i)}.$$

We conjecture that for $\mathcal{G}$ the fixed cost unlimited supply auction where the cost of producing the good is C and $\mathcal{F}_{\mathcal{G}}(\mathbf{b}) = \max(0, \mathcal{F}(\mathbf{b}) - C)$ it is not possible to attain a constant fraction of $\max(0, \mathcal{F}^{(2)}(\mathbf{b}) - C)$ as this difference gets arbitrarily small. Intuitively, it is much easier to be competitive with gross profit than it is to be competitive with the net profit.

We are thus motivated to define the following weaker notion of competitiveness:

Definition 10.2 *For generalized auction problems, we say that a truthful mechanism $\mathcal{M}$ is* β-competitive *if, for all bid vectors* $\mathbf{b}$*,*

$$\mathbf{E}[\mathcal{M}_{\mathcal{G}}(\mathbf{b})] \geq \frac{\text{profit}_\beta(\mathbf{b})}{\beta},$$

where

$$\text{profit}_\beta(\mathbf{b}) = \max_{\mathbf{r}\in\{0,1\}^m}\left(\sum_{1\le i\le m} r_i\mathcal{F}^{(2)}(\mathbf{b}_{S_i}) - \beta c(\mathbf{r})\right),$$

As noted previously, we cannot be constant competitive against $\mathcal{F}_{\mathcal{G}}(\mathbf{b})$. Nonetheless, for a large class of inputs, achieving a constant fraction of $\mathcal{F}_{\mathcal{G}}(\mathbf{b})$ implies achieving a constant fraction of $\text{profit}_\beta(\mathbf{b})$ and vice versa. This is the class of bids for which

- For each market S_j, $\mathcal{F}(\mathbf{b}_{S_j}) = \mathcal{F}^{(2)}(\mathbf{b}_{S_j})$, meaning there is no single bidder whose bid completely dominates all others in that market, and

- The optimal profit margin is a constant factor of the cost of the optimal allocation.

Thus, for an interesting class of problems and inputs to these problems, we will present mechanisms that obtain profit that is within a constant factor of optimal. An interesting problem left open by our work is whether or not there is a stronger performance measure one can compare against and still obtain strong competitiveness guarantees.

10.3 The Fixed Cost Basic Auction

In this section, we consider the fixed cost basic auction problem:

Definition 10.3 (Fixed Cost Basic Auction Problem) *Given:*

- *n identical items for sale.*

- *n bidders, bidder i willing to pay at most u_i for an item.*

- *fixed cost C if any items are sold (zero cost otherwise).*

Design an auction mechanism that maximized the profit of sale (the sum of the revenue less the cost, C.

A natural attempt to solve this problem would be to run a competitive basic auction $\mathcal{A}$ on the input bids $\mathbf{b}$ to obtain profit $\mathcal{A}(\mathbf{b}) - C$. Unfortunately, this may come out negative and the auctioneer may take a loss. As we are interested in how auctioneer perform in the worst case, this is clearly unacceptable.

A simple modification we might think to make to the above proposed solutions is to run simulate the auction $\mathcal{A}$ on $\mathbf{b}$ and compute its revenue. If its revenue is less than C then we cancel the auction and content ourselves with not having a loss. Otherwise, we use effect the outcome of the auction and obtain the non-negative profit $\mathcal{A}(\mathbf{b}) - C$. The problem with this technique is that the resulting mechanism may not be truthful. This motivates the definition of a *cancellable* auction as one that remains truthful even if its outcome may be cancelled if its revenue does not meet some prespecified criterion.

10.3.1 Cancellability

Cancellable auctions are crucial building blocks in the design of mechanisms for more general auctions. As discussed in the introduction, we define cancellable auctions as follows.

Definition 10.4 *Given a truthful auction $\mathcal{A}$ and a parameter C, we define the auction $\mathcal{A}_C$ as follows: On input $\mathbf{b}$, run $\mathcal{A}$ on $\mathbf{b}$. If the resulting revenue is at least C, return the outcome of $\mathcal{A}(\mathbf{b})$. Otherwise,* cancel *$\mathcal{A}$ by returning the outcome with no winners. We say that $\mathcal{A}$ is* cancellable *if and only if, for any value of C, $\mathcal{A}_C$ is truthful.*

There are two natural ways to prove that an auction is cancellable: The first is to show that for each C, there is a bid-independent function f such that the outcome (as a random variable) of the bid-independent auction defined by f is the same as that of $\mathcal{A}_C$. A second approach is to show directly that in $\mathcal{A}_C$ any bidder's profit is maximized by bidding their utility value. Specifically it suffices to use the following:

Proposition 10.1 *The truthful auction $\mathcal{A}$ is cancellable if for any bid vector $\mathbf{b}_{-i}$ and any bidder i with utility u_i, if bidder i wins upon bidding u_i, then $R_{u_i} = R_b$ for any b such that bidder i wins when bidding b. For the fixed values of $\mathbf{b}_{-i}$, R_b denotes the revenue of $\mathcal{A}$ when bidder i bids b.*

To see this, observe that, by assumption, $\mathcal{A}$ is truthful, and thus, bidder i's profit before the possible cancellation is maximized by bidding their utility value. To ensure that bidder i's profit is still maximized after cancellation by bidding u_i, it must be that if the auction is cancelled when bidder i bids u_i, then bidder i loses or otherwise has non-positive profit when bidding any value that results in the auction not being cancelled. Another way to put this is that a truthful auction is cancellable if its revenue is not a function of the bid values of the winning bids.

We next observe that the notion of a cancellable auction is strictly stronger than that of a truthful auction.

Observation 5 *Not all truthful auction mechanisms are cancellable.*

In particular we show that the RSOP auction is not cancellable. RSOP is obviously truthful, as it is bid-independent. We observe now that the auction obtained by running RSOP and cancelling the outcome if the revenue is less than C is not truthful.

Consider $\mathbf{b} = \{1, 1, 1, \ldots, 1, h\}$. Clearly, for h sufficiently large, h is the optimal threshold for its partition, and bids in the other partition will be rejected and the revenue of the auction is close to $n/2$. However, if the high bidder bids 1 instead of h, $\mathbf{b}' = \{1, 1, 1, \ldots, 1, 1\}$ and all bidders win at price 1. Thus the auction revenue is n.

Suppose the auction is to be cancelled if its revenue is less than $C = \frac{3}{4}n$. Then,

- the high bidder's profit is 0 when bidding h, and
- the high bidder's profit is $h - 1$ when bidding 1.

Thus RSOP is not cancellable.

It turns out that none of the auctions presented thus far are cancellable except for the fixed price mechanism, the k-item Vickrey auction, and Randomized Vickrey. We now show that a variant of the RSPE auction is cancellable. This variant is identical to RSPE except for a subtle tie-breaking step.

10.3.2 The Random Sampling Profit Extraction Auction with Tie-Breaking

A formal description of the Random Sampling Profit Extraction Auction with Tie-Breaking (RSPE') is given in Definition 10.5. This auction achieves the same competitive ratio as RSPE, i.e., four. Recall that RSPE computes the optimal fixed price revenues for two random partitions of the bidders and then on both partitions it tries to profit extract the revenue from the opposite partition. Typically, these optimal fixed price revenues will be different and the partition with the lesser revenue will be completely rejected as it cannot afford to cost share the optimal revenue of partition with the higher optimal revenue. In the case where the optimal revenues are the same, we could either allow both partitions to have winners or impose some tie breaking criterion to make sure that only one of the partitions has winners. If we allow both partitions to have winners as in RSPE, the resulting mechanism, while being truthful and 4-competitive, is not cancellable.[3] Thus, if we want a cancellable auction, we must do some sort of tie breaking. The technique we choose, as described in the formal definition of RSPE', is to impose a total order, $\prec$, on fixed price revenues.

Definition 10.5 (Random Sampling Profit Extraction with Tie-Breaking (RSPE')**)**
The RSPE' *auction works as follows:*

1. *Partition bids* $\mathbf{b}$*, into two sets* S' *and* S''*, by flipping a fair coin for each bid. Let the resulting bid vectors be* $\mathbf{b}_{S'}$ *and* $\mathbf{b}_{S''}$*.*
2. *Compute* $F' = \mathcal{F}(\mathbf{b}_{S'})$ *and* $F'' = \mathcal{F}(\mathbf{b}_{S''})$*, the optimal fixed price revenues for* $\mathbf{b}_{S'}$ *and* $\mathbf{b}_{S''}$*, respectively.*
3. *Compute the auction results by running* $\mathrm{CostShare}_{F''}$ *on* $\mathbf{b}_{S'}$ *and* $\mathrm{CostShare}_{F'}$ *on* $\mathbf{b}_{S''}$*.*

---

[3] For example, consider two bidders with utilities $u_1 = 4$ and $u_2 = 2$ and a cost $C = 3$. Assume we sample with $S' = \{1\}$ and $S'' = \{2\}$. If both bidders bid their utility we have $\mathcal{F}' = 4$ and $\mathcal{F}'' = 2$ and the revenue is $\mathcal{R} = 2$. Since the revenue is less than the cost C, the auction is cancelled and the first bidder's profit is zero. If we allow ties, the first bidder could lower their bid to $b_1 = 2$ causing $\mathcal{F}' = \mathcal{F}'' = 2$. This results in both bidders winning and yields a revenue of $\mathcal{R} = 4$. In this case, the auction is not cancelled and the first bidder's profit is two.

We impose a total ordering, "$\prec$", on values of the form "kb_i" that respects their natural partial ordering given by "$<$". We define $\prec$ as:

$$kb_i \prec \ell b_j \iff kb_i < \ell b_j \vee kb_i = \ell b_j \wedge i < j.$$

Note that F' and F'' are of the form "kb_i" for some k and i, so using this total ordering in Steps 2 and 3 guarantees $F' \neq F''$.

We prove the following.

Lemma 10.2 RSPE$'$ *is a cancellable auction.*

Proof: Let RSPE'_C be the auction that runs RSPE$'$ and cancels the outcome if its revenue is not at least C. By definition, RSPE is cancellable if and only if RSPE'_C is truthful for all C. We prove this using Proposition 10.1.

Consider any fixed outcome of the coin flips of RSPE$'$, and any bid vector $\mathbf{b}_{-i}$. In addition, suppose that bidder i's utility value is u_i. For this particular execution, suppose that bidder i is in S'. If bidder i wins in this execution of RSPE$'$ (prior to the possible cancellation) then $F'' \prec F'$. Now suppose that bidder i changes his bid, resulting, possibly, in a new value of F'. If we still have $F'' \prec F'$, then the total revenue of the auction is unchanged, equal to F''. Otherwise, if we now have $F' \prec F''$, then bidder i loses. Thus, in terms of Proposition 10.1, $R_{u_i} = R_b$ for any b such that bidder i wins when bidding b. $\square$

10.4 Applications of Cancellable Auctions

In this section we show how cancellable auctions can be used as a building block for the construction of profit optimizing mechanisms. For generalized auction problems, we define the following *Local Random Sampling Profit Extraction* auction (Auction 10.6).

Definition 10.6 (Local Random Sampling Profit Extraction auction, LRSPE) *On input* $\mathbf{b}$*:*

1. *Run* RSPE$'$ *on each market,* $\mathbf{b}_{S_j}$*, to get revenue* $\mathcal{R}_j$*.*

2. *Compute the maximizing (or approximate) market allocation* $\mathbf{r}^* = \mathrm{argmax}_{\mathbf{r}} \sum_j r_j \mathcal{R}_j - c(\mathbf{r})$.

3. *For each* j*, if* $r^*_j = 0$*, cancel auction on* S_j*. Else collect* $\mathcal{R}_j$ *in revenue from market* S_j*.*

Lemma 10.3 *The Local Random Sampling Profit Extraction auction is truthful.*[4]

This proof is immediate from the cancellability of RSPE.

Theorem 10.4 LRSPE *is 4-competitive.*

Proof: Since RSPE is 4-competitive, for all j we can expect at least $\mathcal{F}^{(2)}(\mathbf{b}_{S_j})/4$ from the jth market. Consider the allocation $\mathbf{r}'$ used by $\mathrm{profit}_4(\mathbf{b})$, that is, $\mathbf{r}'$ maximizes $\max_{\mathbf{r}} \sum_j r_j \mathcal{F}(\mathbf{b}_{S_j}) - 4c(\mathbf{r})$.

Our revenue is

$$\begin{aligned} \mathrm{LRSPE}_{\mathcal{G}}(\mathbf{b}) &= \max_{\mathbf{r}} \sum_j \mathcal{R}_j r_i - c(\mathbf{r}) \\ &\geq \sum_j \mathcal{R}_j r'_i - c(\mathbf{r}') \end{aligned}$$

and

$$\begin{aligned} \mathbf{E}\left[\sum_j \mathcal{R}_j r'_i - c(\mathbf{r}')\right] &\geq \sum_j \mathcal{F}^{(2)}(\mathbf{b}_{S_j}) r'_i/4 - c(\mathbf{r}') \\ &= \mathrm{profit}_4(\mathbf{b})/4 \end{aligned}$$

and so

$$\mathbf{E}[\mathrm{LRSPE}_{\mathcal{G}}(\mathbf{b})] \geq \mathrm{profit}_4(\mathbf{b})/4.$$

□

[4] In fact, any cancellable auction could be used in place of RSPE and the resulting local mechanism would be truthful.

Corollary 10.5 LRSPE *is 4-competitive for the multicast problem.*

We note that LRSPE for the multicast problem of [18] can be implemented with low network complexity, 2 messages per link, using the natural prize-collection algorithm for trees (E.g. [18]).

Corollary 10.6 *Consider the algorithm* LRSPE′ *that differs from* LRSPE *in that rather than actually compute the optimal allocation, it makes use of an α-approximation to the optimal allocation.* LRSPE′ *is* 4α*-competitive.*

As we have seen, the LRSPE auction works very well for generalized auction problems when each market has size at least 2 and there is competition in each market.

These assumptions might be valid for multicast in the Internet today where the content provider is charged for usage of the backbone, while consumers are located at single ISPs and are already paying a flat rate for their service. In this case, the consumers at each ISP would form markets and there would possibly be a large number of them at each.

A key open question is how well a more global auction mechanism can do when these assumptions do not hold. For example, there is a relatively simple global mechanism for the basic market segmentation problem (where there are m markets but the cost, $c(\mathbf{r})$, is zero for all $\mathbf{r}$). It remains open whether this or any other more global mechanism works well in any quantifiable sense.

10.5 An Upper Bound on the Profit of Truthful Mechanisms

As for basic auctions, there is a natural definition of a *monotone* auction. For basic auctions, demonstrating that it was not possible for a monotone auction to perform better than the optimal single-price omniscient mechanism on any input reinforced our framework for competitive analysis. In this section we extend the definition of monotonicity and the bound on profit to the generalized auction setting.

10.5.1 Monotonicity

We are now extend the definition of monotonicity from basic auctions to generalized auctions.

Definition 10.7 *A generalized auction, $\mathcal{A}$, is monotone on markets $\mathcal{S}$ if it is defined by a bid-independent function f with the property that for any bid vector $\mathbf{b}$, and any pair of bidders i and j within the same market such that $b_i \leq b_j$ then*

$$\forall x \leq b_i,\ \mathbf{Pr}[\textit{bidder } i \textit{ wins at price} \leq x] \leq \mathbf{Pr}[\textit{bidder } j \textit{ wins at price} \leq x]\,.$$

10.5.2 Bound on Profit

Theorem 10.7 *Let $\mathcal{A}$ be any monotone auction for the basic auction problem on bids $\mathbf{b}$, and let $\mathcal{W}$ be the event that there is at least one winner, i.e., there exists i such that $x_i = 1$. The revenue $\mathcal{R} = \sum_i p_i$ of $\mathcal{A}$ on input $\mathbf{b}$ satisfies:*

$$\mathbf{E}[\mathcal{R} \mid \mathcal{W}] \leq \mathcal{F}(\mathbf{b}).$$

Proof: Let f be the bid-independent function defining $\mathcal{A}$. Let q be the probability that the event $\mathcal{W}$ occurs. We define $g_i(x)$ as follows:

$$\begin{aligned} g_i(x) &= \mathbf{Pr}[f(\mathbf{b}_{-i}) \leq x \mid \mathcal{W}] \\ &= \tfrac{1}{q}\mathbf{Pr}[f(\mathbf{b}_{-i}) \leq x \cap \mathcal{W}] \end{aligned}$$

For $x \leq b_i$ bidder i is a winner and thus $f(\mathbf{b}_{-i}) \leq x$ implies event $\mathcal{W}$. So we can conclude that for $x \leq b_i$,

$$g_i(x) = \tfrac{1}{q}\mathbf{Pr}[f(\mathbf{b}_{-i}) \leq x].$$

By the monotonicity of f, for all i and j, g_i satisfies the property that if $b_i \geq b_j$ then $g_i(x) \geq g_j(x)$ for $x \leq b_j$.

Now consider the following thought experiment. Let U be a random variable that is uniform on $[0, 1]$. Imagine running the bid-independent auction that for each i using $g_i^{-1}(U)$ to set the threshold for bidder i, with g_i^{-1} defined as $g_i^{-1}(y) = \inf\{x : g_i(x) = y\}$. We denote by $\mathcal{R}_U$ the resulting auction revenue. We observe that the threshold distribution for bidder i in this experiment is precisely the same as the original threshold distribution for bidder i conditioned on $\mathcal{W}$:

$$\begin{aligned}\mathbf{Pr}\left[g_i^{-1}(U) \leq x\right] &= \mathbf{Pr}[U \leq g_i(x)] \\ &= g_i(x).\end{aligned}$$

Therefore, by summing the expectations for the bidders, we obtain

$$\mathbf{E}[\mathcal{R}_U] = \mathbf{E}[\mathcal{R} \mid \mathcal{W}].$$

We complete the proof by showing that the expected revenue from our thought experiment $\mathbf{E}[\mathcal{R}_U]$ is at most $\mathcal{F}(\mathbf{b})$. Conditioned on $U = u$, let k be the index of the smallest winning bid. Thus, $g_k^{-1}(u) \leq b_k$. Since $g_k(x) \leq g_j(x)$ for all j such that $b_j \geq b_k$, and $g_k(x)$ and $g_j(x)$ are monotone non-decreasing functions, we must have $g_j^{-1}(u) \leq g_k^{-1}(u) \leq b_k \leq b_j$ and therefore all bidders with bid values at least b_k win at a price at most b_k. Thus, the revenue, $\mathcal{R}_u$, is at most b_k times the number of bids with bid value least b_k which totals to at most $\mathcal{F}(\mathbf{b})$. This holds for all $u \in [0, 1]$, and thus $\mathbf{E}[\mathcal{R}_U] \leq \mathcal{F}(\mathbf{b})$. □

We now give an upper bound on profit for generalized auctions.

Theorem 10.8 *Let $\mathcal{G} = (\mathcal{S}, c(\cdot))$ be a generalized auction problem. For any truthful mechanism $\mathcal{M}$ that is monotone on $\mathcal{S}$, and for any input bid vector $\mathbf{b}$, we give an upper bound $\mathcal{M}_{\mathcal{G}}(\mathbf{b})$, the profit of $\mathcal{M}$ on bid set $\mathbf{b}$ as:*

$$\mathbf{E}[\mathcal{M}_{\mathcal{G}}(\mathbf{b})] \leq \mathcal{F}_{\mathcal{G}}(\mathbf{b}).$$

Proof: Define the following:

- $q_{\mathbf{r}}$ with $\mathbf{r} = (r_1, \ldots, r_m) \in \{0, 1\}^m$ – the probability that the result of the auction is the allocation $\mathbf{r}$ to the markets.

For $1 \leq j \leq m$ define:

- q_j – the probability that there is at least one winner in market S_j. Note that $q_j = \sum_{\mathbf{r}\in\{0,1\}^m} r_j q_{\mathbf{r}}$.
- $\mathcal{F}_j = \mathcal{F}(\mathbf{b}_{S_j})$ – the optimal fixed price revenue for market j.
- $\mathcal{R}_j$ – the revenue collected from the jth market.

Note that from Theorem 10.7 we have

$$\begin{aligned}\mathbf{E}[\mathcal{R}_j] &= q_j \mathbf{E}[\mathcal{R}_j \mid r_j = 1] \\ &\leq q_j \mathcal{F}_j \qquad (10.1)\end{aligned}$$

Now we bound the expected profit:

$$\begin{aligned}\mathbf{E}[\mathcal{M}_{\mathcal{G}}(\mathbf{b})] &= \mathbf{E}\left[\sum\nolimits_j \mathcal{R}_j - c(\mathbf{r})\right] \\ &= \sum\nolimits_j \mathbf{E}[\mathcal{R}_j] - \mathbf{E}[c(\mathbf{r})]\,.\end{aligned}$$

By equation (10.1),

$$\begin{aligned}\mathbf{E}[\mathcal{M}_{\mathcal{G}}(\mathbf{b})] &\leq \sum\nolimits_j q_j \mathcal{F}_j - \mathbf{E}[c(\mathbf{r})] \\ &= \sum\nolimits_j \mathcal{F}_j \left(\sum\nolimits_{\mathbf{r}} r_j q_{\mathbf{r}}\right) - \sum\nolimits_{\mathbf{r}} q_{\mathbf{r}} c(\mathbf{r}) \\ &= \sum\nolimits_{\mathbf{r}} q_{\mathbf{r}} \left(\sum\nolimits_j r_j \mathcal{F}_j - c(\mathbf{r})\right).\end{aligned}$$

This is a convex combination, so we have

$$\begin{aligned}\mathbf{E}[\mathcal{M}_{\mathcal{G}}(\mathbf{b})] &\leq \max_{\mathbf{r}\in\{0,1\}^m} \left(\sum_{1\leq j\leq m} r_j \mathcal{F}_j - q_{\mathbf{r}} c(\mathbf{r})\right) \\ &= \mathcal{F}_{\mathcal{G}}(\mathbf{b}).\end{aligned}$$

□

The proof just presented may seem at first glance to be more complicated than necessary. However, we note that the most obvious approach of using the equation

$$\mathbf{E}[\mathcal{M}_{\mathcal{G}}(\mathbf{b})] = \sum_{\mathbf{r}} q_{\mathbf{r}} \left(\sum_j \mathbf{E}[\mathcal{R}_j \mid \mathbf{r}] - c(\mathbf{r})\right)$$

and then showing that $\mathbf{E}[\mathcal{R}_j \mid \mathbf{r}] \leq \mathcal{F}_j$ fails: $\mathbf{E}[\mathcal{R}_j \mid \mathbf{r}]$ can in fact significantly exceed $\mathcal{F}_j$.

10.6 Conclusions

For the fixed cost unlimited supply auction, i.e., one market with cost C if there are any winners, we have shown that it is possible to get a constant fraction of $\mathcal{F}^{(2)}(\mathbf{b}) - 4C$ in worst case. We conjecture that this scaling of C is necessary in the sense that it is not possible to get a constant fraction of $\mathcal{F}^{(2)}(\mathbf{b}) - C$ in worst case. This result seems likely because $\mathcal{F}^{(2)}(\mathbf{b}) - C$ can be arbitrarily small compared to $\mathcal{F}^{(2)}(\mathbf{b})$ and we know from Theorem 7.10 that it is not possible for a basic auction to be better than 2.42-competitive against $\mathcal{F}^{(2)}(\mathbf{b})$.

For general auctions we gave a local mechanism that is competitive if there is competition amongst the bidders in each market and if there is a significant profit margin. This solution can be applied directly to the multicast cost sharing problem of [18]. To design competitive mechanisms for generalized auction problems that do not require the assumption of competition among the bidders in each market (e.g., allowing for singleton markets), global mechanisms must be employed. Even for simple special cases of the multicast tree, this problem is difficult. Special cases, e.g., segmented markets with no cost (for all $\mathbf{r}$, $c(\mathbf{r}) = 0$), that do permit natural global mechanisms are difficult to generalize.

Our formulation for generalized auctions captures a wide variety of allocation mechanisms; however, more general models can be considered. In particular, our model could be extended to allow cost functions that can take into account how many items are allocated to each market instead of just whether or not any goods were allocated to the market.

Chapter 11

THE DOUBLE AUCTION PROBLEM

In this chapter we consider the problem of designing a mechanism for double auctions where bidders each bid to buy or sell one unit of a single commodity. The profit of the auctioneer is the difference between the total payments from buyers and the total payments to the sellers. We aim to maximize this profit. We extend the competitive analysis framework of basic auctions and give an upper bound on the profit of any truthful double auction. We then reduce the competitive double auction problem to basic auctions by showing that any competitive basic auction can be converted into a competitive double auction with a competitive ratio of twice that of the basic auction. In addition, we show that better competitive ratios can be obtained by directly adapting basic auction techniques to the double auction problem. In doing so, we show that the consensus and revenue estimate technique from Section 4.5 gives a general result for profit maximizing mechanism design: it reduces private value optimization problems to private value decision problems.

11.1 Introduction

For double auctions, the auctioneer, acting as a broker, is faced with the task of matching up a subset of the buyers with an equal-sized subset of the sellers. The auctioneer decides on a price to be paid to each seller and received from each buyer in exchange for the transfer of one item from each of the selected sellers to each of the selected buyers. The *profit of the auctioneer* is the difference between the prices paid by the buyers and the prices paid to the sellers. We assume that each buyer wishes to purchase exactly one item, each seller wishes to sell exactly one item, and that the items are indistinguishable, i.e., no buyer has reason to prefer one seller's item over that of another.

We note that the basic auction problem is the special case of the double auction problem where all sellers are known to have zero value for their items.

11.2 Preliminaries

We consider single-round, sealed-bid double auction mechanisms in which each bidders wants to either buy or sell one out of a set of identical items. Bidders submit one sealed bid each, and publicly declare themselves to be either a *buyer* or a *seller*. For buyers, the bid represents the maximum amount they are willing to pay for an item, whereas for sellers, the bid represents the minimum amount they are willing to sell the item for.

We denote by $\mathbf{b}$ the vector of bid values associated with buyers and by $\mathbf{s}$ the vector of all bid values associated with sellers. The ith component of $\mathbf{b}$ (resp. $\mathbf{s}$) is b_i (resp. s_i), the bid value submitted by the ith buyer (resp. seller). We assume that the number of buyers is equal to the number of sellers, and we use n to denote this number.

Definition 11.1 (Double Auction) *A single-round sealed-bid double-auction mechanism is one where:*

- *Given the two bid vectors $\mathbf{b} = (b_1, \ldots, b_n)$ and $\mathbf{s} = (s_1, \ldots, s_n)$, the mechanism computes a pair of* allocation *vectors, $\mathbf{x}$ and $\mathbf{y} \in \{0,1\}^n$, and* payment vectors $\mathbf{p}$ *and* $\mathbf{q} \in \mathbb{R}^n$*, subject to the constraints that:*
  - *The number of winning buyers is equal to the number of winning sellers, i.e.,* $\sum_i x_i = \sum_i y_i$.[1]
  - $0 \leq p_i \leq b_i$ *(resp. $s_i \leq q_i$) for all winning buyers (resp. sellers) and $p_i = 0$ (resp. $q_i = 0$) for all losing buyers (resp. sellers). These are the standard assumptions of* no positive transfers *and* voluntary participation*. See, e.g., [37].*
- *If $x_i = 1$ buyer i* wins *(i.e., receives the item) and pays price p_i, otherwise we say that buyer i* loses *or is* rejected*. If $y_i = 1$ seller i* wins *(i.e., sells the item) and gets paid q_i, otherwise we say that seller i* loses *or is* rejected.

[1] We assume that the auctioneer neither has any items for sale nor is willing to purchase any. For this reason, we can also assume that the number of buyer bids equals the number of seller bids. If there are any extra buyers or sellers, the auctioneer can earn the same amount of profit by ignoring the extra low bidding buyers or high bidding sellers.

- *The* profit of the mechanism *is* $\mathcal{M}(\mathbf{b},\mathbf{s}) = \sum_i p_i - \sum_i q_i$.

It is assumed that input order of the bids $\mathbf{b}$ and $\mathbf{s}$ is arbitrary. Throughout our discussion of auction problems we will use the following notation for the ith highest bidding buyer and the ith lowest bidding seller.

Definition 11.2 *The ith highest bidding buyer is $b_{(i)}$. The ith smallest seller bid is $s_{(i)}$.*

Definition 11.3 (Vickrey Double Auction) *The k-item* Vickrey double auction *on bids $\mathbf{b}$ and $\mathbf{s}$, $\mathcal{V}_k(\mathbf{b},\mathbf{s})$, sells to the highest k buyers at price $b_{(k+1)}$ and buys from the lowest k sellers at price $s_{(k+1)}$. Its revenue is*

$$\mathcal{V}_k(\mathbf{b},\mathbf{s}) = k(b_{(k+1)} - s_{(k+1)}).$$

In typical examples it is important that the k-Vickrey basic auction only sells k items. It is likewise important that the k-Vickrey double auction has the same number of winning buyers as winning sellers, e.g. k of them. As specified above, the k-Vickrey auctions are not well defined if there are several bidders with identical bid values. It is simple to fix this problem by assuming that the k-Vickrey auction breaks ties arbitrarily. This tie breaking is natural and necessary for double auction problem and we will be assuming throughout this chapter that bid values are distinct. This can be achieved by assuming an arbitrary (or random) total order on the bidders that respects the partial order given by their actual bid values. Thus, $b_{(i)} > b_{(i+1)}$ and $s_{(i)} < s_{(i+1)}$.

11.2.1 Bid Independence

We now extend the characterization of truthful mechanisms using the notion of *bid independence* to the double auction problem.

Definition 11.4 *Let f and g be a functions from bid vectors to prices (non-negative real numbers). The* deterministic bid-independent double auction defined by f and g *is* $\mathrm{BI}_{f,g}$. *For each buyer i,*

1. *Compute bid-independent threshold $t_i = f(\mathbf{b}_{-i}, \mathbf{s})$.*
   (where $\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, ?, b_{i+1}, \ldots, b_n)$)

2. If $t_i < b_i$, set $x_i \leftarrow 1$ and $p_i \leftarrow t_i$. (Buyer i wins.)

3. If $t_i > b_i$ set $x_i = p_i = 0$. (Buyer i is rejected.)

4. Otherwise, if $t_i = b_i$ the auction can either accept the bid at price t_i or reject it.

We treat the sellers symmetrically using threshold $v_i = g(\mathbf{b}, \mathbf{s}_{-i})$ and buying from seller i at price v_i if v_i is more than s_i.

A randomized bid-independent auction is a probability distribution over deterministic bid-independent auctions.

The following theorem, which is a trivial generalization of the equivalent result for basic auctions, relates bid independence to truthfulness.

Theorem 11.1 *A double auction is truthful if and only if it is bid-independent.*

11.2.2 Single Price Omniscient Mechanisms

As the basic auction problem is a special case of the double auction problem, we extend the competitive framework for the basic auction problem to the double auction problem via the following definitions of $\mathcal{F}$ and $\mathcal{F}^{(2)}$ for the double auction problem.

Definition 11.5 *The* optimal single price omniscient mechanism*, $\mathcal{F}$, is the mechanism that uses the optimal single buy price and single sell price. It achieves the* optimal single price profit *of*

$$\mathcal{F}(\mathbf{b}, \mathbf{s}) = \max_i i(b_{(i)} - s_{(i)}).$$

Definition 11.6 *The* optimal fixed price mechanism that transfers at least two items*, $\mathcal{F}^{(2)}$, has profit*

$$\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) = \max_{i \geq 2} i(b_{(i)} - s_{(i)}).$$

11.2.3 Competitive Mechanisms

We now formalize the notion of a competitive mechanism:

Definition 11.7 *We say that a truthful double auction $\mathcal{M}$ is β-*competitive against $\mathcal{F}^{(2)}$ *if, for all bid vectors* $\mathbf{b}$ *and* $\mathbf{s}$ *the expected profit of $\mathcal{M}$ satisfies*

$$\mathbf{E}[\mathcal{M}(\mathbf{b},\mathbf{s})] \geq \mathcal{F}^{(2)}(\mathbf{b},\mathbf{s})/\beta.$$

We say that $\mathcal{M}$ is competitive if $\mathcal{M}$ is β-competitive for some constant β.

11.3 Reducing Competitive Double Auctions to Competitive Basic Auctions

In this section we describe a general technique for converting any β-competitive basic auction into a 2β-competitive double auction.

Definition 11.8 ($\mathcal{M}_{\mathcal{A}}$) *Given,*

- *basic auction, $\mathcal{A}$,*
- *input* $\mathbf{b}$ *and* $\mathbf{s}$*, and*
- *ℓ, the largest value such that $b_{(\ell)} \geq s_{(\ell)}$;*

the double auction, $\mathcal{M}_{\mathcal{A}}$, does as follows:

Case 1 ($\ell = 1$): *Output the empty allocation.*

Case 2 ($\ell = 2$): *Simulate the 1-item Vickrey double auction, $\mathcal{V}_1(\mathbf{b},\mathbf{s})$, and output its outcome.*

Case 3 ($\ell \geq 3$): *Let $\mathbf{b}'$ and $\mathbf{s}'$ be n-dimensional vectors with components by $b'_i = b_i - s_{(\ell)}$ and $s_i = b_{(\ell)} - s_i$. Let $\mathbf{b}''$ (resp. $\mathbf{s}''$) be the $(\ell-1)$-dimensional vector consisting of the largest $\ell - 1$ bids in $\mathbf{b}'$ (resp. $\mathbf{s}'$).*

1. *With probability* $1/2$ *simulate* $\mathcal{A}(\mathbf{b}'')$*. If buyer* i *wins the simulation of* $\mathcal{A}$ *at price* p_i'' *then buyer* i *wins* $\mathcal{M}_\mathcal{A}$ *at price* $p_i = \max(b_{(\ell)}, p_i'' + s_{(\ell)})$*. All other buyers lose. Let* $k < \ell$ *be the number of winners in* $\mathcal{A}(\mathbf{b}'')$*. To determine the outcome for the sellers, run the* k*-Vickrey basic auction on* $\mathbf{s}$*.*
2. *Otherwise (with probability* $1/2$*) simulate* $\mathcal{A}(\mathbf{s}'')$*. If seller* i *wins the simulation of* $\mathcal{A}$ *at price* q_i'' *then seller* i *wins* $\mathcal{M}_\mathcal{A}$ *at price* $q_i = \min(s_{(\ell)}, b_{(\ell)} - q_i'')$*. As in Step 1, we run the* k*-Vickrey basic auction on the buyers to determine the outcome for buyers, where* $k < \ell$ *is the number of winners in* $\mathcal{A}(\mathbf{s}'')$*.*

Theorem 11.2 $\mathcal{M}_\mathcal{A}$ *is truthful.*

Proof: We show that $\mathcal{M}_\mathcal{A}$ is truthful for the buyers. The result for sellers is symmetric. First note all buyers that win the auction pay at least $b_{(\ell)}$ and that the buyer with the ℓth highest bid and all buyers with lower bids always lose the auction. In the case that $\ell \leq 2$ this is obvious. For $\ell \geq 3$ we have:

- In Step 1, since $b_{(\ell)}$ is excluded from $\mathbf{b}''$, $b_{(\ell)}$ loses. In this case by definition, all winners pay at least $b_{(\ell)}$.

- In Step 2, since $k < \ell$ the k-Vickrey auction on buyers rejects $b_{(\ell)}$ (and winners must pay at least $b_{(\ell)}$).

We now argue that $\mathcal{M}_\mathcal{A}$ is truthful for the buyer with the ℓth highest bid and all lower bidding buyers. Suppose buyer i is one if these lower bidding buyers. Hold all other bid values fixed. To win the auction, buyer i would have to bid higher than $b_{(\ell-1)}$. In this case, all winners would pay at least $b_{(\ell-1)}$ which is more than this buyer's utility value, b_i. Such a sale price would give buyer i a negative profit, hence buyer i prefers losing the auction.

Now we show that the mechanism is truthful for the remaining $\ell-1$ high bidding buyers. First, none of these bidders can change the value of $b_{(\ell)}$ or ℓ without lowering their bid value below $b_{(\ell)}$ which would cause them to lose the auction.

We now argue (the simpler case) that Step 2 is truthful for the $\ell-1$ high bidding buyers. In this step the truthful k-Vickrey auction is run on these buyers. Since k-Vickrey is truthful and because the k is determined from $\mathbf{s}$, ℓ, and $b_{(\ell)}$ which we have shown to be unchangeable by any winning buyers, Step 2 is truthful.

The truthfulness of Step 1 is similar. Suppose basic auction $\mathcal{A}$ is defined by bid-independent function f'. Then we can define the bid-independent function for this case as $f(\mathbf{b}_{-i}, \mathbf{s}) = \max(b_{(\ell)}, f'(\mathbf{b}_{-i} - s_{(\ell)}) + s_{(\ell)})$. Given that $b_{(\ell)}$, ℓ, and thus $s_{(\ell)}$ have been shown to be unchangeable by any winning buyers, this shows that Step 1 is truthful. □

Theorem 11.3 *If $\mathcal{A}$ is β-competitive, $\mathcal{M}_{\mathcal{A}}$ is 2β-competitive against $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$.*

Proof: If $\ell = 1$, $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \leq 0$ so the null allocation is competitive. If $\ell = 2$, $\mathcal{M}_{\mathcal{A}}$ runs the 1-item Vickrey double auction which is 2-competitive when $\ell = 2$. For the rest of the proof assume $\ell \geq 3$. Let $k \in [2, \ell]$ be the number of items sold by $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$. Thus,

$$\begin{aligned} \mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) &= k(b_{(k)} - s_{(k)}) \\ &= k(b_{(k)} - s_{(\ell)}) + k(b_{(\ell)} - s_{(k)}) - k(b_{(\ell)} - s_{(\ell)}). \end{aligned}$$

But by definition $\mathcal{F}^{(2)}(\mathbf{b}') \geq k(b_{(k)} - s_{(\ell)})$ and likewise for $\mathbf{s}'$, therefore

$$\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \leq \mathcal{F}^{(2)}(\mathbf{b}') + \mathcal{F}^{(2)}(\mathbf{s}') - k(b_{(\ell)} - s_{(\ell)}). \tag{11.1}$$

Note that for the buyers (and similarly for sellers):

$$\mathcal{F}^{(2)}(\mathbf{b}') \leq \mathcal{F}^{(2)}(\mathbf{b}'') + b_{(\ell)} - s_{(\ell)}. \tag{11.2}$$

Because $k \geq 2$, from Equations (11.1) and (11.2) we have

$$\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \leq \mathcal{F}^{(2)}(\mathbf{b}'') + \mathcal{F}^{(2)}(\mathbf{s}'').$$

Note that because $\mathcal{A}$ is β-competitive, the expected revenues from Step 3 and Step 4 are at least $\mathcal{F}^{(2)}(\mathbf{b}'')/2\beta$ and $\mathcal{F}^{(2)}(\mathbf{s}'')/2\beta$ respectively. Thus,

$$\mathbf{E}[\mathcal{M}_{\mathcal{A}}(\mathbf{b}, \mathbf{s})] \geq \tfrac{1}{2\beta}(\mathcal{F}^{(2)}(\mathbf{b}'') + \mathcal{F}^{(2)}(\mathbf{s}'')) \geq \tfrac{1}{2\beta}\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}).$$

□

Plugging in the 4-competitive Random Sampling Profit Extraction basic auction, RSPE, we get a double auction with a competitive ratio of 8. Plugging in the 3.39-competitive

Consensus Revenue Estimate basic auction, CORE, we get a competitive ratio of 6.78. We show below that we can do better than this by using the CORE technique customized to the double auction problem.

11.4 The General Reduction to the Private Value Decision Problem

We now discuss a reduction from private value decision problems to private value optimization problems. Using this reduction we obtain a result that is similar to that which applies to classical optimization: for private value optimization problems we can obtain a near optimal solution using a solution to the decision problem. The reduction for classical optimization involves some sort of search (binary or otherwise) over the value of the optimal solution and repeatedly invoking the decision problem solution until the optimal solution is found. This is not possible in truthful mechanism design, as running several truthful mechanisms and taking the output of the best one will not, in general, result in a truthful mechanism overall. Instead we use a technique based on truthfully (i.e., bid-independently) estimating the profit of the optimal solution and using this estimate with a truthful mechanism for the decision problem. The general technique developed here will follow the same approach as that of the CORE auction from Chapter 4.

11.4.1 Profit Extraction: The Private Value Decision Problem

We now formalize the notion of a private value decision problem by defining a *profit extractor* as its solution. Let $\mathbf{I}$ represent the private value input to a generic private value optimization problem. For example $\mathbf{I} = \mathbf{b}$ for the basic auction problem and $\mathbf{I} = (\mathbf{b}, \mathbf{s})$ for the double auction problem.

Definition 11.9 (Profit Metric) *A* profit metric*, $\mathcal{O}$, is a function from the input specification of a private value optimization problem, $\mathbf{I}$, to a real number (to be interpreted as an amount of profit).*

As an example, $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ is a profit metric for the double auction problem and $\mathcal{F}(\mathbf{b})$ is a profit metric for the basic auction problem.

Definition 11.10 (Profit Extractor) *A* profit extractor *for profit metric $\mathcal{O}$ is a truthful mechanism that given target profit R has profit at least R on all inputs, $\mathbf{I}$, such that $\mathcal{O}(\mathbf{I}) \geq R$.*

For the basic auction problem a profit extractor for metric $\mathcal{F}$ was given in Chapter 2. For some problems and metrics, exact profit extractors do not exist and approximate profit extractors may be useful instead.

Definition 11.11 (Approximate Profit Extractor) *An α-*approximate profit extractor *for profit metric $\mathcal{O}$ is a truthful mechanism that given target profit R has profit at least R/α on all inputs, $\mathbf{I}$, such that $\mathcal{O}(\mathbf{I}) \geq R$.*

11.4.2 The Revenue Estimation Technique

We now sketch the reduction of the private value optimization problem to the private value decision problem by showing how to truthfully estimate the optimal revenue for a given input. Assume for some private value decision problem we have a truthful profit extractor $\mathcal{O}$-Extract$_R$ for metric $\mathcal{O}$.

The characterization of truthful mechanisms as those that have bid-independent implementations (Theorem 11.1) and its truthfulness imply that $\mathcal{O}$-Extract$_R$ is implemented by some bid-independent function, pe$_R$. That is, the bid-independent auction defined by $f(\mathbf{I}_{-i}) = \text{pe}_R(\mathbf{I}_{-i})$ is exactly $\mathcal{O}$-Extract$_R$. Consider the auction, $\mathcal{M}$, parameterized by function $r(\cdot)$ that is defined by bid-independent function

$$f(\mathbf{I}_{-i}) = \text{pe}_{r(\mathbf{I}_{-i})}(\mathbf{I}_{-i}).$$

Note that if r is *a consensus*, i.e., $r(\mathbf{I}_{-i}) = R$ for all i, then $\mathcal{M}$ is identically $\mathcal{O}$-Extract$_R$. Furthermore, if $r(\mathbf{I})$ is a *revenue estimate*, i.e. R is a constant fraction of $\mathcal{O}(\mathbf{I})$, then the mechanism $\mathcal{O}$-Extract$_R$ (and thus, so does $\mathcal{M}$) achieves revenue a constant fraction of $\mathcal{O}(\mathbf{I})$. Below we show a probability distribution from which to pick $r(\cdot)$ such that r is a consensus with high probability and also a revenue estimate.

Our revenue estimation technique is based on the observation that the value of one of the bidders in the mechanism should not have much effect on the mechanism's overall revenue.

In many cases of interest, for example the mass market sale of a digital good, this does indeed hold true. In cases where this does not hold true, it is unlikely that any mechanism can obtain a near optimal revenue. This motivates the following definition:

Definition 11.12 (Smoothness) *For $\rho > 1$, the private value input, $\mathbf{I}$, is ρ-smooth for metric, $\mathcal{O}$, if for all bidders, i, we have:*

$$\mathcal{O}(\mathbf{I})/\rho \leq \mathcal{O}(\mathbf{I}_{-i}) \leq \mathcal{O}(\mathbf{I}).$$

In any ρ-smooth input $\mathbf{I}$ with metric, $\mathcal{O}$, the following consensus estimate function has the properties that we desire.

Definition 11.13 (Consensus Estimate Function Distribution) *Parameterized by $c > 1$ and the metric $\mathcal{O}$ we define the distribution of consensus estimate functions, $\mathcal{D}_{\mathcal{O},c}$, as follows:*

- *For U be a uniform random variable from $[0, 1]$,*
- *$r(\cdot) \sim \mathcal{D}_{\mathcal{O},c}$ is defined as $r(\mathbf{I}_{-i})$ is $\mathcal{O}(\mathbf{I}_{-i})$ rounded down to the nearest c^{j+U} for integer j.*

Lemma 11.4 *[23] For ρ-smooth input $\mathbf{I}$, let random variable X be $r(\mathbf{I})$ if $r(\mathbf{I}) = r(\mathbf{I}_{-i})$ for all i and zero otherwise. The consensus estimate function $r(\cdot)$ parameterized by $c > \rho$ satisfies:*

$$\mathbf{E}[X] = \frac{\mathcal{O}(\mathbf{I})}{\ln c}\left(\frac{1}{\rho} - \frac{1}{c}\right).$$

The following general definition of the Consensus Revenue Estimate (CORE) mechanism gives the reduction from the private value optimization problem to the private value decision problem.

Definition 11.14 (CORE$_c$) *For constant c, metric $\mathcal{O}$, and profit extractor $\mathcal{O}$-Extract$_R$ for $\mathcal{O}$ that is defined bid-independently by pe_R; the mechanism CORE$_c$ is defined bid-independently by*

$$f(\mathbf{I}_{-i}) = \mathrm{pe}_{r(\mathbf{I}_{-i})}(\mathbf{I}_{-i})$$

for $r(\cdot)$ sampled from $\mathcal{D}_{\mathcal{O},c}$.

Before we can claim that CORE_c performs well, we must take care of one additional potential pit-fall. Consider CORE_c on ρ-smooth inputs $\mathbf{I}$ with $c > \rho$. In the case where there is no consensus, there will be two values R and R' such that some bidders will have a price and outcome that given from $\mathcal{O}$-$\mathrm{Extract}_R$ and other from $\mathcal{O}$-$\mathrm{Extract}_{R'}$. A problem might arise if some outcomes are infeasible or if there is a cost function that associates large negative costs with certain outcomes. In this case, the CORE auction could output an infeasible outcome or acquire a large cost. We define a profit extractor to be *safe* if this can never happen.

Definition 11.15 (Safety) *A profit extractor $\mathcal{O}$-$Extract_R$ is* safe *if it is impossible for $CORE_c$ to output an infeasible outcome or an outcome with negative profit.*

Of course, CORE_c can never output infeasible outcomes or ones with negative profit when there is consensus. To prove that a profit extractor is safe, one must only consider the behavior of the profit extractor when used in CORE_c when there is no consensus.

Theorem 11.5 *Given*

- *metric $\mathcal{O}$,*
- *safe profit extractor $\mathcal{O}$-$Extract_R$ for $\mathcal{O}$, and*
- *and input $\mathbf{I}$ that is ρ-smooth for metric $\mathcal{O}$;*

For $c > \rho$, the CORE_c mechanism has an expected revenue and competitive ratio, respectively, of:

$$\mathbf{E}[\mathrm{CORE}_c(\mathbf{I})] = \frac{\mathcal{O}(\mathbf{I})}{\ln c}\left(\frac{1}{\rho} - \frac{1}{c}\right), \qquad \frac{\mathcal{O}(\mathbf{I})}{\mathbf{E}[\mathrm{CORE}_c(\mathbf{I})]} = \ln c \left(\frac{1}{\rho} - \frac{1}{c}\right)^{-1}.$$

Given ρ, the value of c can be chosen to give the optimal competitive ratio when restricted to ρ-smooth inputs.

11.4.3 Truthfulness with High Probability

As we saw in Chapter 6, there is a version of the CORE_c auction that is not completely truthful. Instead it is *truthful with high probability* (see Definition 6.2). A key benefit of going from the truthful version of CORE_c to CORE'_c, the version that is only truthful with high probability, is that the outcome is always identical to the outcome of the profit extractor, $\mathcal{O}$-Extract_R, for some value of R. For general optimization problems, this may be important if there is no profit extractor that is *safe* for the desired metric, $\mathcal{O}$. This gives us the following results.

Definition 11.16 (CORE'_c) *For constant c, metric $\mathcal{O}$, and profit extractor $\mathcal{O}$-$\mathit{Extract}_R$ for $\mathcal{O}$ that is defined bid-independently by* pe_R*; the mechanism* CORE'_c *is:*

1. *Choose $r(\cdot)$ from $\mathcal{D}_{\mathcal{O},c}$.*

2. *Run $\mathcal{O}$-$\mathit{Extract}_{r(\mathbf{I})}$ on* $\mathbf{I}$.

Theorem 11.6 *Given*

- *metric $\mathcal{O}$,*

- *profit extractor $\mathcal{O}$-$\mathit{Extract}_R$ (not necessarily safe) for $\mathcal{O}$, and*

- *and input* $\mathbf{I}$ *that is ρ-smooth for metric $\mathcal{O}$;*

for $c > \rho$, the CORE'_c *mechanism has an expected revenue and competitive ratio, respectively, of:*

$$\mathbf{E}\big[\text{CORE}'_c(\mathbf{I})\big] = \frac{\mathcal{O}(\mathbf{I})}{\ln c}\left(1 - \frac{1}{c}\right), \qquad\qquad \frac{\mathcal{O}(\mathbf{I})}{\mathbf{E}[\text{CORE}'_c(\mathbf{I})]} = \ln c\left(1 - \frac{1}{c}\right)^{-1}.$$

The probability that CORE'_c *is truthful is:* $1 - \log_c \rho$.

11.5 Consensus and Revenue Estimate for Double Auctions

We now apply the Consensus and Revenue Estimate technique to the double auction problem. First we discuss profit extraction for the double auction problem and then we discuss smoothness.

11.5.1 Profit Extraction for Double Auctions

As discussed in Section 2.6, for the basic auction problem, ProfitExtract$_R$ is a profit extractor for metric $\mathcal{F}$. In the following lemma we show that for the double auction problem, it is not possible to even approximately profit extract $\mathcal{F}$. We then extend this simple argument to show that it is not possible to exactly profit extract $\mathcal{F}^{(2)}$.

Lemma 11.7 *For any value R and $\alpha \geq 1$, there is no truthful mechanism for the double auction problem that achieves a profit of at least R/α on all inputs $\mathbf{b}$ and $\mathbf{s}$ with $\mathcal{F}(\mathbf{b}, \mathbf{s}) \geq R$.*

Proof: Suppose for a contradiction that such a mechanism $\mathcal{M}_R$ did exist. Consider the single buyer, single seller case with $b_1 = s_1 + R$. On this input, to achieve revenue R/α for any $\alpha \geq 1$, both the buyer and the seller must win the auction. Theorem 11.1 and $\mathcal{M}_R$'s truthfulness implies that the price for b_1 is given by a bid-independent function $f(s_1)$. Since b_1 wins the auction and the auctioneer has a positive profit, it must be that $f(s_1) \in [s_1, b_1] = [s_1, s_1 + R]$. Symmetrically, we must have $g(b_1) \in [b_1 - R, b_1]$. Now consider inputs $b'_1 = s'_1 + 2R$. Given f and g above, the sell price is in $p'_1 \in [s'_1, s'_1 + R]$ and the buy price is in $q'_1 \in [b'_1 - R, b'_1] = [s'_1 + R, s'_1 + 2R]$. The auctioneer's profit is $p'_1 - q'_1 \leq 0$ which gives a contradiction. □

Theorem 11.8 *For input $\mathbf{b}$ and $\mathbf{s}$, let $k \geq 2$ be the number of items sold by $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$. Let $\mathcal{M}_R$ be a truthful double auction mechanism that is guaranteed to extract a profit of at least R/α_k on any $\mathbf{b}$ and $\mathbf{s}$ with $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \geq R$. Then, $\alpha_k \geq k/(k-1)$.*

Proof: We prove the lemma assuming the mechanism is deterministic. Since randomized truthful mechanisms are just distributions over deterministic truthful mechanisms, the lemma follows in general from the consideration of this case.

We prove the $k = 2$ case that $\alpha_2 \geq 2$ below. The general theorem follows by a simple, but tedious inductive argument that follows the same lines. For the case $k = 2$, with a bid vector (b_1, b_2, s_1, s_2) consisting of two buyers and two sellers where $b_1 = b_2 = b$ and $s_1 = s_2 = s = b - R/2$. If $\alpha_2 < 2$, then both buyers and both sellers must win. Let $f(\cdot)$ and $g(\cdot)$ denote the bid-independent functions defined by the truthful mechanism $\mathcal{M}_R$. Since all bidders must win, we have that

$$b \geq f(?, b, s, s),\ f(b, ?, s, s) \geq g(b, b, ?, s),\ g(b, b, s, ?) \geq s.$$

Now consider instead the bid vector (b, b', s, s), where $b' > b$. Then since the price offered buyer 2 is the same as it was in the previous case, namely $f(b, ?, s, s)$, and both buyers and sellers must win in order to achieve a revenue which exceeds $R/2$, we still have

$$b \geq f(b, ?, s, s),\ f(?, b', s, s) \geq g(b, b', ?, s),\ g(b, b', ?, s) \geq s.$$

Similarly, if we consider the bid vector (b', b, s, s), we must have

$$b \geq f(?, b, s, s)\ ,\ f(b', ?, s, s) \geq g(b', b, ?, s),\ g(b', b, s, ?) \geq s.$$

Together the facts $b \geq f(b', ?, s, s)$ and $b \geq f(?, b', s, s)$ imply that for the bid vector (b', b', s, s), with $b' > b$, the price offered the buyers is at most b, which in turn implies that the prices paid to the sellers are both at most b.

An analogous argument, starting with a bid vector in which both buyers bid b' and both sellers bid $s' = b' - R/2$ and then moving the seller bids down shows that for any pair of buyers with bid values b' and sellers with bid values $s < b' - R/2$, the sellers both win at prices that are at least $b' - R/2$. This is true in particular for the bid vector (b', b', s, s).

For $b' > b + R/2$, this leads to a contradiction: in the first case, we argued that the prices offered to the sellers were at most b, whereas the second argument shows that the prices offered to the sellers are strictly greater than b. Such prices would give negative profit. □

Faced with these impossibilities, for the double auction problem we can either look for profit extractors for other metrics or approximate profit extractors for $\mathcal{F}^{(2)}$. We show that both of these techniques give constant competitive double auctions; however, the latter technique gives the better competitive ratio.

Definition 11.17 (Vickrey Optimal Double Auction) *The Vickrey optimal public value mechanism, $\mathcal{V}_{\text{opt}}$, chooses the outcome of the k-Vickrey double auction that gives the most revenue. Its profit is*

$$\mathcal{V}_{\text{opt}}(\mathbf{b}, \mathbf{s}) = \max_k \mathcal{V}_k(\mathbf{b}, \mathbf{s}) = \max_k k(b_{(k+1)} - s_{(k+1)})$$

It is easy to see from the definition of $\mathcal{F}^{(2)}$ and $\mathcal{V}_{\text{opt}}$ that for all $\mathbf{b}$ and $\mathbf{s}$, and $\mathcal{F}^{(2)}$ exchanging at least k items, $\mathcal{V}_{\text{opt}}(\mathbf{b},\mathbf{s}) \geq \frac{k-1}{k}\mathcal{F}^{(2)}(\mathbf{b},\mathbf{s})$. This allows us to compare the profit of a CORE mechanism using a profit extractor for $\mathcal{V}_{\text{opt}}$ to $\mathcal{F}^{(2)}$.

We now give the definition of a double auction profit extractor for $\mathcal{V}_{\text{opt}}$. From this definition it is obvious that the revenue of the mechanism is R is $R \leq \mathcal{V}_{\text{opt}}$ and zero otherwise.

Definition 11.18 *The double auction profit extractor, $\mathcal{V}_{\text{opt}}$-Extract$_R$, on input $\mathbf{b}$ and $\mathbf{s}$ gives the outcome of $\mathcal{V}_k$ for the largest k such that $\mathcal{V}_k(\mathbf{b},\mathbf{s}) \geq R$.*

Lemma 11.9 *$\mathcal{V}_{\text{opt}}$-Extract$_R$ is truthful.*

Proof: We show that the auction is bid-independent for buyers. The case of sellers is analogous. Define the bid-independent function $f(\mathbf{b}_{-i},\mathbf{s})$ for buyers as follows for input $\mathbf{b}_{-i}$ and $\mathbf{s}$:

- Compute $\mathbf{b}^{(i)}$ from $\mathbf{b}_{-i}$ as
$$b_j^{(i)} = \begin{cases} b_j & \text{if } i \neq j \\ \infty & \text{otherwise.} \end{cases}$$

- Simulate $\mathcal{V}_{\text{opt}}$-Extract$_R$ on $\mathbf{b}^{(i)}$ and $\mathbf{s}$ to compute a buyer allocation, $\mathbf{x}$, and prices, $\mathbf{p}$, (and seller allocation, $\mathbf{y}$, and prices, $\mathbf{q}$).

- If $x_i = 0$, i.e., buyer i is not allocated the item in the simulation, then output price ∞.

- Otherwise, $x_i = 1$. Output price p_i, the price buyer i paid in the simulation.

We now show that the bid-independent mechanism, BI$_f$, is identical to $\mathcal{V}_{\text{opt}}$-Extract$_R$.

Suppose buyer i bidding b_i is allocated an item at price p_i by $\mathcal{V}_{\text{opt}}$-Extract$_R(\mathbf{b},\mathbf{s})$. This occurs because there is some largest k with $b_{(k+1)} < b_i$ such that $\mathcal{V}_k(\mathbf{b},\mathbf{s}) \geq R$. It is easy to see that for $k' \geq k$, $\mathcal{V}_{k'}(\mathbf{b},\mathbf{s}) = \mathcal{V}_{k'}(\mathbf{b}^{(i)},\mathbf{s})$. This is because these Vickrey auction revenues are not a function of the values of the top k buyer bids and bidder i's bid is one of the top k

in both $\mathbf{b}$ and $\mathbf{b}^{(i)}$. Thus, the simulation of $\mathcal{V}_{\text{opt}}$-Extract$_R(\mathbf{b}^{(i)}, \mathbf{s})$ also finds k as the largest value with $\mathcal{V}_k(\mathbf{b}^{(i)}, \mathbf{s}) \geq R$, and the sale price offered buyer i in BI$_f$ is $b_{(k+1)}$ also.

Suppose now that buyer i bidding b_i is not allocated an item by $\mathcal{V}_{\text{opt}}$-Extract$_R(\mathbf{b}, \mathbf{s})$. This means that for all k such that $b_{(k)} < b_i$, $\mathcal{V}_k(\mathbf{b}, \mathbf{s}) < R$. Since for all k, $\mathcal{V}_k(\mathbf{b}^{(i)}, \mathbf{s}) \geq \mathcal{V}_k(\mathbf{b}, \mathbf{s})$, the simulation of $\mathcal{V}_{\text{opt}}$-Extract$_R(\mathbf{b}^{(i)}, \mathbf{s})$ either finds a sale price that that is more than b_i or it does not sell any items. In either case buyer i is rejected by BI$_f$. $\square$

Definition 11.19 *The approximate double auction profit extractor, $\mathcal{F}^{(2)}$-Extract$_R$, on input* $\mathbf{b}$ *and* $\mathbf{s}$ *computes the largest k such that $k(b_{(k)} - s_{(k)}) \geq R$, buys from the lowest $k-1$ bidding sellers at price $s_{(k)}$, and sells to the top $k-1$ bidding buyers at price $b_{(k)}$. All other buyers and sellers (including the kth) are rejected.*

$\mathcal{F}^{(2)}$-Extract$_R$ is closely related to $\mathcal{V}_{\text{opt}}$-Extract$_R$ and it is not difficult to adapt the proof of Lemma 11.9 to show that $\mathcal{F}^{(2)}$-Extract$_R$ is truthful. We omit the details.

The following theorem shows that $\mathcal{F}^{(2)}$-Extract$_R$ is the optimal approximate profit extractor for $\mathcal{F}^{(2)}$ since its approximation ratio matches the lower-bound of Theorem 11.8. Thus, Theorem 11.8 is tight.

Lemma 11.10 *On bids* $\mathbf{b}$ *and* $\mathbf{s}$ *such that $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ exchanges k items, $\mathcal{F}^{(2)}$-Extract$_R$ is a $\frac{k}{k-1}$-approximate profit extractor for $\mathcal{F}^{(2)}$. I.e., it achieves revenue at least $R(k-1)/k$ if $R \leq \mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$.*

Proof: Suppose that $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \geq R$. Let k be the number of items exchanged by $\mathcal{F}^{(2)}$ and let k' be the number of items exchanged by $\mathcal{F}^{(2)}$-Extract$_R$. Note that $\mathcal{F}^{(2)}$-Extract$_R$ finds the largest k' such that $k'(b_{(k')} - s_{(k')}) \geq R$. By the definition of k, $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) = k(b_{(k)} - s_{(k)}) \geq R$. Thus, $k' \geq k$. The revenue of $\mathcal{F}^{(2)}$-Extract$_R$ is

$$\begin{aligned} \mathcal{F}^{(2)}\text{-Extract}_R(\mathbf{b}, \mathbf{s}) &= (k'-1)(b_{(k')} - s_{(k')}) \\ &\geq \tfrac{k'-1}{k'} R \geq \tfrac{k-1}{k} R. \end{aligned}$$

$\square$

Lemma 11.11 *$\mathcal{F}^{(2)}$-Extract$_R$ is safe.*

Proof: We must show that CORE_c with $\mathcal{F}^{(2)}$-Extract_R does not effect infeasible outcomes, i.e., outcomes where the number of winning buyers is not equal to the number of winning sellers. Let k be the number of items sold by $\mathcal{F}^{(2)}$ on input $\mathbf{b}$ and $\mathbf{s}$ and let $F = \mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ be its profit.

1. For buyer i not allocated items by $\mathcal{F}^{(2)}$ we have $\mathcal{F}^{(2)}(\mathbf{b}_{-i}, \mathbf{s}) = F$. Likewise for sellers.

2. $\mathcal{F}^{(2)}$-Extract_R for $R \leq F$ has at least k items exchanged.

Note that the top k buyers and sellers always exchange items regardless of consensus. Further, by 1 above, the bottom $n-k$ buyers and sellers always have consensus. Thus the same number of these additional buyers and sellers exchange items. □

11.5.2 CORE for double auctions

We will show below how to use the CORE framework and approximate profit extractor $\mathcal{F}^{(2)}$-Extract_R to obtain a 3.75-competitive double auction. We can apply a similar technique to the exact profit extractor for $\mathcal{V}_{\text{opt}}$; however, this only gives a provable competitive ratio of 4.15 so we omit the details.

As discussed above, the CORE with profit extractor approach only works on $\mathbf{b}$ and $\mathbf{s}$ that are ρ-smooth for $\mathcal{F}^{(2)}$. Our goal, however, is an auction that is constant-competitive with $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$ for any $\mathbf{b}$ and $\mathbf{s}$. We make the following observations about $\mathcal{F}^{(2)}$. Let k be the number of winners in $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$.

- If $k \geq 3$ then $\mathbf{b}$ and $\mathbf{s}$ are $\frac{k}{k-1}$-smooth. I.e., for all i,

$$\frac{k-1}{k}\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \leq \mathcal{F}^{(2)}(\mathbf{b}_{-i}, \mathbf{s}), \mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}_{-i}) \leq \mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}).$$

  In this case, for $c > 3/2$, CORE_c with profit extractor $\mathcal{F}^{(2)}$-Extract_R is constant-competitive:

$$\mathbf{E}[\text{CORE}_c(\mathbf{b}, \mathbf{s})] = \frac{\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})}{\ln c}\left(\frac{k-1}{k}\right)\left(\frac{k-1}{k} - \frac{1}{c}\right). \tag{11.3}$$

- For $k = 2$, $\mathcal{V}_1$ is 2-competitive with $\mathcal{F}^{(2)}$.

Thus, it is possible to take a convex combination of the CORE_c auction with $\mathcal{V}_1$ to get a constant-competitive auction in worst-case.

Definition 11.20 ($\mathbf{CORE}_{c,p}$) *The double auction $CORE_{c,p}$ parameterized by $c > 3/2$ and $p \in (0,1)$ flips a coin and with probability $1-p$ runs $CORE_c$ with profit extractor $\mathcal{F}^{(2)}$-$Extract_R$ and with probability p runs $\mathcal{V}_1$.*

Theorem 11.12 *With a near optimal choice of $c = 2.62$ and $p = 0.54$, the $CORE_{c,p}$ double auction with profit extractor $\mathcal{F}^{(2)}$-$Extract_R$ is 3.75-competitive against $\mathcal{F}^{(2)}$.*

Proof: Let k be the number of items sold by $\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s})$. We consider the case $k = 2$ and $k \geq 3$ separately.

Case 1 ($k = 2$): Our expected profit is $p\mathcal{F}^{(2)}/2$.

Case 2 ($k \geq 3$): From Vickrey we get $p\mathcal{F}^{(2)}/k$ and from CORE_c we get $(1-p)$ times the quantity in Equation (11.3) for a combined expected profit of:

$$\mathcal{F}^{(2)}(\mathbf{b}, \mathbf{s}) \left(\frac{p}{k} + \frac{1-p}{\ln c} \left(\frac{k-1}{k} \right) \left(\frac{k-1}{k} - \frac{1}{c} \right) \right).$$

Our choice of p and c optimizes and balances the two cases. Numerical analysis gives $c = 2.62$ and $p = 0.54$ as a near-optimal choice. This choice gives a competitive ratio of 3.75. □

Note that the competitive ratio of the CORE basic auction is better than the competitive ratio of the CORE double auction (3.39 vs. 3.75). This difference is due to the fact that the former uses an exact revenue extractor and the latter uses an approximation.

11.6 Monotonicity

We now extend the notion of monotonicity to double auctions. The intuition underlying our notion of monotonicity is that if an auction is to achieve a large profit, the bid-independent function defining the auction should output higher prices for buyers when it sees higher buyer bid values.

Definition 11.21 *A double auction is* monotone *if:*

- *For any pair of buyers i and j with $b_i \leq b_j$, we have:*

$$\forall x \leq b_i,\ \mathbf{Pr}[\textit{buyer } i \textit{ wins at price} \leq x] \leq \mathbf{Pr}[\textit{buyer } j \textit{ wins at price} \leq x]\,.$$

- *For any pair of sellers i and j with $s_i \geq s_j$, we have:*

$$\forall x \geq s_i,\ \mathbf{Pr}[\textit{seller } i \textit{ wins at price} \geq x] \leq \mathbf{Pr}[\textit{seller } j \textit{ wins at price} \geq x]\,.$$

11.6.1 Upper Bound on the Profit of Truthful Mechanisms

In this section, we show that the profit for all monotone double auction mechanisms is bounded by $2\mathcal{F}(\mathbf{b}, \mathbf{s})$. Recall in Chapter 8 we showed that no monotone basic auction could achieve an expected revenue more than $\mathcal{F}(\mathbf{b})$ for any $\mathbf{b}$ (Theorem 8.1).

We conjecture that the same bound holds for double auctions as well, though what we prove below is a factor of two worse.

Lemma 11.13 *For any value v and buyer and seller bids $\mathbf{b}$ and $\mathbf{s}$, define $\mathbf{b}'$ and $\mathbf{s}'$ as $b'_i = b_i - v$ and $s'_i = v - s_i$ for $1 \leq i \leq n$. Then for any monotone double auction, $\mathcal{M}$:*

$$\mathbf{E}[\mathcal{M}(\mathbf{b}, \mathbf{s})] \leq \mathcal{F}(\mathbf{b}') + \mathcal{F}(\mathbf{s}').$$

Proof: Let $\mathbf{x}$, $\mathbf{y}$, $\mathbf{p}$, and $\mathbf{q}$ be the outcome and prices when $\mathcal{M}$ is run on $\mathbf{b}$ and $\mathbf{s}$. Let $X = \{i\ :\ x_i = 1\}$ and $Y = \{i\ :\ y_i = 1\}$. Note $|X| = |Y|$. Thus,

$$\begin{aligned}\mathcal{M}(\mathbf{b}, \mathbf{s}) &= \sum_i p_i - \sum_i q_i = \sum_{i \in X} p_i - \sum_{i \in Y} q_i \\ &= \sum_{i \in X} (p_i - v) + \sum_{i \in Y} (v - q_i).\end{aligned}$$

Let $\mathcal{A}_{v,\mathbf{s}}$ be the basic auction that on $\mathbf{b}'$ simulates $\mathcal{M}(\mathbf{b}, \mathbf{s})$ to compute prices p_i for each bidder b'_i and then offers them $p_i - v$. It is easy to see that this is truthful, monotone (since $\mathcal{M}$ is), and gives revenue

$$\mathcal{A}_{v,\mathbf{s}}(\mathbf{b}') = \sum_{i \in X} (p_i - v).$$

Using the bound on the revenue of any monotone basic auction (Theorem 8.1) we get:

$$\mathbf{E}\left[\sum_{i\in X}(p_i - v)\right] = \mathbf{E}\big[\mathcal{A}_{v,\mathbf{s}}(\mathbf{b}')\big] \le \mathcal{F}(\mathbf{b}').$$

Combining this with the analogous argument for $\mathbf{s}'$ we have:

$$\mathbf{E}[\mathcal{M}(\mathbf{b},\mathbf{s})] = \mathbf{E}\left[\sum_{i\in X}(p_i - v)\right] + \mathbf{E}\left[\sum_{i\in Y}(v - q_i)\right] \le \mathcal{F}(\mathbf{b}') + \mathcal{F}(\mathbf{s}').$$

□

Theorem 11.14 *For any bid vectors* $\mathbf{b}$ *and* $\mathbf{s}$*, any truthful monotone double auction,* $\mathcal{M}$*, has expected profit at least* $2\mathcal{F}(\mathbf{b},\mathbf{s})$*.*

Proof: Find the largest ℓ such that $b_{(\ell)} \ge s_{(\ell)}$ and choose $v \in [s_{(\ell)}, b_{(\ell)}]$. Now we let $\mathbf{b}'$ and $\mathbf{s}'$ be $b'_i = b_i - v$ and $s'_i = v - s_i$ for $1 \le i \le n$. For our choice of v, Lemma 11.13 gives $\mathbf{E}[\mathcal{M}(\mathbf{b},\mathbf{s})] \le \mathcal{F}(\mathbf{b}') + \mathcal{F}(\mathbf{s}')$.

Note that $\mathcal{F}(\mathbf{b},\mathbf{s}) = \max_i i(b_{(i)} - s_{(i)})$. Let k be the number of winners in $\mathcal{F}(\mathbf{b},\mathbf{s})$. Note that by our choice of v, we have $b_{(k)} \ge v$ and $s_{(k)} \le v$. This gives:

$$\mathcal{F}(\mathbf{b}') = \max_i i(b_{(i)} - v) \le \max_i i(b_{(i)} - s_{(i)}) = \mathcal{F}(\mathbf{b},\mathbf{s}), \text{ and}$$
$$\mathcal{F}(\mathbf{s}') = \max_i i(v - s_{(i)}) \le \max_i i(b_{(i)} - s_{(i)}) = \mathcal{F}(\mathbf{b},\mathbf{s}).$$

Thus, $\mathbf{E}[\mathcal{M}(\mathbf{b},\mathbf{s})] \le \mathcal{F}(\mathbf{b}') + \mathcal{F}(\mathbf{s}') \le 2\mathcal{F}(\mathbf{b},\mathbf{s})$. □

11.7 Conclusions

In this chapter we have given a game theoretic treatment of the off-line problem of matching up buyers and sellers of a single identical commodity so as to maximize the worst case profit of the arbitrating auctioneer. Open questions related to the double auction problem include considering similar questions in an on-line setting where the buyers, sellers, or both arrive one at a time and the auctioneer must decide whether to sell and at what price before the arrival of the next customer.

We have also shown again how the consensus and revenue estimate technique reduces private value optimization problems to private value decision problems. We applied this technique to the double auction problem using an approximate profit extractor (i.e., an approximate solution to the "decision problem"). These results motivate the very natural general questions about profit extractors. For what problems do exact or approximate profit extractors exist? What characterizes such problems? With the reduction, answering these questions gives an answer to the question of what private value optimization problems can be solved almost as well as their public value equivalents.

Chapter 12

CONCLUSIONS AND FUTURE RESEARCH DIRECTIONS

In this work, we have shown competitive analysis to be a successful technique for obtaining provable performance bounds on worst case and promise versions of private value profit maximization problems. The important distinction between our work and prior work is that the mechanisms we design are not assumed to have any prior knowledge about the input.

We develop a number of techniques for solving private value profit maximization problems including the use of randomized sampling, the private value decision problem (i.e., profit extractors), cancellable mechanisms, and the technique of obtaining a consensus on revenue estimates. Random sampling is a general technique for performing market analysis on the fly as the mechanism is running. The technique of cancellable mechanisms gave rise to a reduction from private value optimization to public value optimization for the case that the bidders are divided into markets – running a cancellable auction locally in each market replaces the private values in the market with a single public value that can be collected (if profitable) or cancelled (if not profitable). The technique of obtaining a consensus of revenue estimates leads to a general reduction from private value optimization problems to private value decision problems.

Through these techniques we have given solutions to the basic auction problem, the limited supply auction problem, the multi-item auction problem, the multicast pricing problem, and the double auction problem.

A number of interesting questions still remain open. Recall, that in Chapter 4 we proved that no symmetric basic auction is competitive. Further, we gave a symmetric auction that uses two random bits and is competitive in worst case. A variant of this auction uses only one random bit and is competitive on mass markets. There is still the question of whether there exists a competitive deterministic auction that instead of randomness uses asymmetry, i.e., makes use of the arbitrary ordering of the bids in determining its outcome, to obtain worst case or mass market competitiveness.

For randomized auctions, we have shown a lower bound on the competitive ratio of 2.42. We have an auction that obtains a competitive ratio of 3.39. It is an open question as to what the optimal worst case competitive ratio is – even for the case where there are only three bidders. Recall that for the three bidder case the lower bound on the competitive ratio is 13/6.

One key assumption in the relevance of our competitive framework of comparing the auction revenue to an optimal single priced auction is that the bidders are a priori indistinguishable. Indeed, for the case where the bidders are divided into markets where bidders within a single market are indistinguishable in advance but bidders in different markets can be distinguished, it makes more sense, as we have done in Chapter 10, to compete with an optimal auction that is allowed to set different prices for bidders that can be distinguished. We leave as an open question a solution to the *basic auction problem with attributes.* In this variant of the basic auction problem each bidder i is accompanied by an attribute a_i. Two bidders with the same attribute $a_i = a_j$ cannot be distinguished but bidders with distinct attributes can. This problem arises as a special case of the multicast pricing problem. We note that our solution to the multicast pricing problem does not handle this special case.

Our solution to the multi-item unlimited supply auction problem requires, as a subroutine, an algorithm for computing the optimal sale prices for each item. Obtaining a polynomial time algorithm for this public value problem that gives better than a $\log n$-approximation is an open question. Note that obtaining a log approximation can be trivially achieved by computing the optimal price vector restricted such that all items are sold at the same price, i.e, $r_j = r_{j'}$ for all j and j'. The computation problem of computing the optimal prices is APX-hard (i.e., there is no polynomial time approximation scheme) [8].

An interesting question which we have not considered at all in this thesis is the *multi-item limited supply auction problem.* In this private value problem there are m non-identical items for sale and n bidders. Each bidder has a utility value for each item, but the items are interchangeable in the sense that no bidder wishes to acquire more than one item, i.e., the unit demand case. The goal is to design an auction that obtains the maximum profit.

As this is a special case of the unlimited supply problem, an efficient approximation to the public value version of the problem is also an open question.

In Chapter 6 we showed that no truthful envy-free auction could be competitive. As we were still interested in auctions with envy-free outcomes, we considered auctions that are truthful with high probability instead of fully truthful. Indeed, under this relaxed notion of truthfulness, we obtained competitive and envy-free auctions for mass markets. This solution concept of truthfulness with high probability gives some of the power lost by the fully truthful requirement back to the auctioneer and allows them to both obtain a higher profit and have more flexibility over the nature of the solution such as restricting it to be envy-free. A justification for truthfulness with high probability as a valid solution concept is as follows. We have assumed that the auctioneer does not know a distribution from which the bidders valuations are drawn. It is natural in such situations to assume that the bidders also do not know the values of the other bidders. In this case, in a non-truthful mechanism there is no obvious strategy for a bidder to follow to maximize their own profit. In a mechanism that is truthful with high probability we give the bidders an obvious strategy that is almost always optimal, i.e., bidding their true value. They can they choose between this obvious strategy and trying to determine a non-obvious optimal strategy that is a complex function of the values of the other bidders. As this latter task is not possible unless each bidder knows the other bidders values, we can argue that in auctions that are truthful with high probability, rational bidders will bid their true values. An interesting area for future research is in developing alternative solution concepts to truthfulness that allow problems to be solved that are impossible to solve with truthful mechanisms.

One assumption we made throughout this thesis is that the bidders do not collude. A very interesting direction for future research involves relaxing this assumption and looking at mechanisms that are collusion resistant. A promising direction along these lines is to look for mechanisms that have the property that with high probability (in random coin flips made in the mechanism) any coalition of colluding bidders of size t has no incentive to collude. As an example, a weak notion of collusion resistance, that of *group strategyproofness* which requires that if any member of a coalition gains from the coalition's non-truthful bidding strategy then some member of the coalition is hurt by it. The Consensus and Revenue

Estimate Auction (CORE) is t-group strategyproof with high probability. We conjecture that similar result can be shown for stronger notions of collusion resistance than group strategyproofness.

With the multicast pricing problem and the double auction problem, we have begun to look at optimization problems with limited feasible solutions and structured cost functions. Of course, considering additional structured optimization problems in a private value profit maximization framework is of interest. Along these lines there is the question of finding profit extractors for other structured optimization problems. Such a profit extractor, if found, could be used to solve the profit maximization problem using the consensus and revenue estimation techniques developed in Chapter 11.

Research in the area of intersection between the fields of Economics, Game Theory, and Computer Science is still quite new and there are many open questions in all aspects of the area. Furthermore, we believe that the importance of understanding these problems is only going to increase in the future as the need for solutions to problems with resource sharing between parties with diverse and selfish interests becomes more crucial.

BIBLIOGRAPHY

[1] N. Abel. Beweis eines Ausdrucks von welchem die Binomial-Formel ein einzelner Fall ist. *Crelles Journal für die Reine und Angewandte Mathematik*, 1:159–160, 1826.

[2] A. Archer, C. Papadimitriou, K. Tawar, and É. Tardos. An Approximake Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents. In *Proc. 14th Symp. on Discrete Algorithms*. ACM/SIAM, 2003.

[3] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.

[4] A. Archer and É. Tardos. Frugal Path Mechanisms. In *To appear, SODA*, 2002.

[5] Associated Press. Amazon may spell end for 'dynamic' pricing. *USA Today*, 29 September 2000. Available at: `http://www.usatoday.com/life/cyber/tech/cti595.htm`.

[6] B. Awerbuch, Y. Azar, and A. Meyerson. Reducing truth-telling online mechanisms to online optimization. In *Proc. of 35th Symposium on Theory of Computing*, pages 503–510. ACM Press, New York, 2003.

[7] Y. Azar. Private communication. 2003.

[8] Y. Azar, A. Fiat, J. Hartline, A. Karlin, and C. Kenyon. Envy-free Pricings and Matchings. Unpublished Manuscript, 2003.

[9] Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *Proc. 13th Symp. on Discrete Algorithms*. ACM/SIAM, 2002.

[10] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online Learning in Online Auctions. In *Proc. 14th Symposium on Discrete Algorithms*. ACM/SIAM, 2003.

[11] J. Bulow and J. Roberts. The Simple Economics of Optimal Auctions. *The Journal of Political Economy*, 97:1060–90, 1989.

[12] H. Chernoff. A Measure of Asymptotic Efficiency for Test of a Hypothesis Based on the Sum of Observations. *Anals of Math. Stat.*, 23:493–509, 1952.

[13] E. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11:17–33, 1971.

[14] X. Deng, C. Papadimitriou, and M. Safra. On the Complexity of Equilibria. In *Proc. of 34th Symposium on Theory of Computing*. ACM Press, New York, 2002.

[15] K. Deshmukh, A. Goldberg, J. Hartline, and A. Karlin. Truthful and Competitive Double Auctions. In *Proc. 10th European Symposium on Algorithms*. Springer–Verlag, 2002.

[16] N. Devanur, C. Papadimitriou, V. Vazirani, and A. Saberi. Market Equilibria Via Primal-Dual-Type Algorithm. In *Proc. of the 43rd IEEE Symp. on Foundations of Computer Science*, 2002.

[17] E. Elkind, A. Sahai, and K. Steiglitz. Frugality in Path Auctions. Unpublished Manuscript, 2003.

[18] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the Cost of Multicast Transmissions. In *Proc. of 32nd Symposium Theory of Computing*, pages 218–226. ACM Press, New York, 2000.

[19] A. Fiat. Private communication. 2001.

[20] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. In *Proc. 34th ACM Symposium on the Theory of Computing*. ACM Press, New York, 2002.

[21] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Unpublished manuscript. 2003.

[22] A. Goldberg and J. Hartline. Competitive Auctions for Multiple Digital Goods. In *Proc. 9th European Symposium on Algorithms*, pages 416–427. Springer–Verlag, 2001.

[23] A. Goldberg and J. Hartline. Competitiveness via Concensus. In *Proc. 14th Symp. on Discrete Algorithms*. ACM/SIAM, 2003.

[24] A. Goldberg and J. Hartline. Envy-Free Auction for Digital Goods. In *Proc. of 4th ACM Conference on Electronic Commerce*. ACM Press, New York, 2003.

[25] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions and digital goods. *Games and Economic Behavior*, 2002. Submitted for publication. An earlier version available as InterTrust Technical Report STAR-TR-99.09.01.

[26] A. Goldberg, J. Hartline, and A. Wright. Competitive Auctions and Digital Goods. In *Proc. 12th Symp. on Discrete Algorithms*, pages 735–744. ACM/SIAM, 2001.

[27] T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.

[28] D. Johnson, M. Minkoff, and S. Phillips. The prize collecting Steiner tree problem: theory and practice. In *Symposium on Discrete Algorithms*, pages 760–769, 2000.

[29] R. Kleinberg and T. Leighton. The Value of Knowing a Demand Curve: Bounds on Regret for Online Posted-Price Auctions. In *Proc. of the 44rd IEEE Symp. on Foundations of Computer Science*, 2003.

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

[31] E. Koutsoupias and C. Papadimitriou. The Price of Truth: Frugality in Truthful Mechanisms. In *Proc. of 16th International Symposium on Theoretical Aspects of Computer Science*. Springer, 1999.

[32] D. Lehmann, L. O'Callaghan, and Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. In *Proc. of 1st ACM Conf. on E-Commerce*, pages 96–102. ACM Press, New York, 1999.

[33] R. Lipton and A. Tomkins. Online interval scheduling. In *Proc. 5th Symp. on Discrete Algorithms*. ACM/SIAM, 1994.

[34] A. Mehta, S. Shenker, and V. Vazirani. Profit-Maximizing Multicast Pricing by Approximating Fixed Points. In *Proc. of 4th ACM Conference on Electronic Commerce*. ACM Press, New York, 2003.

[35] J. Mirrlees. An Exploration into the Theory of Optimal Income Taxation. *Review of Economics Studies*, 38:175–208, April 1971.

[36] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

[37] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance vs. Efficiency. *Economic Theory*, 18:511–533, 2001.

[38] R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.

[39] N. Nisan and A. Ronen. Algorithmic Mechanism Design. In *Proc. of 31st Symposium on Theory of Computing*, pages 129–140. ACM Press, New York, 1999.

[40] N. Nisan and A. Ronen. Computationally feasible VCG Mechanisms. In *ACM Conference on Electronic Commerce*, pages 242–252, 2000.

[41] R. Radner. Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives. *J. Economic Theory*, 22:136–154, 1980.

[42] J. Riley and W. Samuelson. Optimal Auctions. *American Economic Review*, 71:381–92, 1981.

[43] A. Ronen. On Approximating Optimal Auctions. In *Proc. of Third ACM Conference on Electronic Commerce*. ACM Press, New York, 2001.

[44] A. Roth and M. Oliveira Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, 1990.

[45] T. Roughgarden. *Selfish Routing*. PhD thesis, Cornell University, May 2002.

[46] T. Roughgarden and É. Tardos. How Bad Is Selfish Routing? In *Proc. of the 41st IEEE Symp. on Foundations of Computer Science*, 2000.

[47] J. Schummer. Almost dominant strategy implementation. Technical report, MEDS, Kellogg Graduate School of Management, 2001. Available from `http://www.kellogg.nwu.edu/faculty/schummer/ftp/research/esp.ps`.

[48] I. Segal. Optimal Pricing Mechanism with Unknown Demand. *American Economic Review*, 93:509–29, 2003.

[49] L. Shapley and M. Shubik. The Assignment Game I: The Core. *Int. J. of Game Theory*, 1:111–130, 1972.

[50] K. Talwar. The Price of Truth: Frugality in Truthful Mechanisms. In *Proc. of 20th International Symposium on Theoretical Aspects of Computer Science*. Springer, 2003.

[51] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16:8–37, 1961.

VITA

Jason Hartline was born in Seattle in 1976. Jason graduated from Cornell University in 1997 with a B.S. in both Computer Science and Electrical Engineering. He is receiving his Ph.D. from The University of Washington where he has been worked on extending ideas from algorithmic optimization to economically motivated problems in game theoretic settings, e.g., auction problems. This thesis research was conducted in part at University of Washington with his advisor Anna Karlin and at InterTrust's STAR Lab and later Microsoft Research with Andrew Goldberg.