

Profit Maximization in Mechanism Design

Jason D. Hartline and Anna R. Karlin

Abstract

We give an introduction to the design of mechanisms for profit maximization with a focus on single-parameter settings.

13.1 Introduction

In previous chapters, we have studied the design of truthful mechanisms that implement social choice functions, such as social welfare maximization. Another fundamental objective, and the focus of this chapter, is the design of mechanisms in which the goal of the mechanism designer is profit maximization. In economics, this topic is referred to as *optimal mechanism design*.

Our focus will be on the design of profit-maximizing auctions in settings in which an auctioneer is selling (respectively, buying) a set of goods/services. Formally, there are n agents, each of whom desires some particular service. We assume that agents are *single-parameter*; i.e., agent i 's valuation for receiving service is v_i and their valuation for no service is normalized to zero. A mechanism takes as input sealed bids from the agents, where agent i 's bid b_i represents his valuation v_i , and computes an outcome consisting of an allocation $\mathbf{x} = (x_1, \dots, x_n)$ and prices $\mathbf{p} = (p_1, \dots, p_n)$. Setting $x_i = 1$ represents agent i being allocated service whereas $x_i = 0$ is for no service, and p_i is the amount agent i is required to pay the auctioneer. We assume that agents have quasi-linear utility expressed by $u_i = v_i x_i - p_i$. Thus, an agent's goal in choosing his bid is to maximize the difference between his valuation and his payment.

To make this setting quite general, we assume that there is an inherent cost $c(\mathbf{x})$ in producing the outcome \mathbf{x} , which must be paid by the mechanism. Our goal is to design the mechanism, i.e., the mapping from bid vectors to price/allocation vectors so that the auctioneer's profit, defined as

$$\text{Profit} = \sum_i p_i - c(\mathbf{x}),$$

is maximized, and the mechanism is *truthful*.

Many interesting auction design problems are captured within this single-parameter framework. In what follows, we describe a number of these problems, and show that, for most of them, the VCG mechanism (Chapter 9), which maximizes social welfare, is a poor mechanism to use when the goal is profit maximization.

Example 13.1 (single-item auction) We can use the cost function $c(\mathbf{x})$ to capture the constraint that at most one item can be allocated, by setting $c(\mathbf{x}) = 0$ if $\sum_i x_i \leq 1$ and ∞ otherwise. The profit of the Vickrey auction (Chapter 9) is the second highest of the valuations in the vector \mathbf{v} . If prior information about agents' valuations is available, then there are auctions with higher profit than the Vickrey auction.

Example 13.2 (digital goods auctions) In a digital goods auction, an auctioneer is selling multiple units of an item, such as a downloadable audio file or a pay-per-view television broadcast, to consumers each interested in exactly one unit. Since the marginal cost of duplicating a digital good is negligible and digital goods are freely disposable, we can assume that the auctioneer has an *unlimited supply* of units for sale. Thus, for digital goods auctions $c(\mathbf{x}) = 0$ for all \mathbf{x} .

The profit of the VCG mechanism for digital goods auctions is zero. Indeed, since the items are available in unlimited supply, no bidder places any externality on any other bidder.

Example 13.3 (single-minded combinatorial auction, known bundles) In a combinatorial auction with single-minded agents, each agent has exactly one bundle of items that he is interested in obtaining. Agent i 's value for his desired bundle, S_i , is v_i . We use the cost function $c(\mathbf{x})$ to capture the constraint that each item can be allocated to at most one bidder. Thus, $c(\mathbf{x}) = 0$ if $\forall i, j, S_i \cap S_j \neq \emptyset \rightarrow x_i x_j = 0$, and $c(\mathbf{x}) = \infty$ otherwise.

Example 13.4 (multicast auctions) Consider a network with users residing at the nodes in the network, each with a valuation for receiving a broadcast that originates at a particular node, called the root. There are costs associated with transmitting data across each of the links in the network – the cost of transmitting across link e is $c(e)$. Our problem is then to design an auction that chooses a multicast tree, the set of users to receive the broadcast, and the prices to charge them. In this setting, $c(\mathbf{x})$ is the total cost of connecting all of the agents with $x_i = 1$ to the root (i.e., the minimum Steiner tree cost).

In most nondegenerate instances of this problem the VCG mechanism will run a deficit. One such example is the *public project* setting described in Chapter 9, Section 3.5 which can be mapped to a network with a single link of cost C , where one endpoint is the root and all the users are at the other endpoint.

All of the other examples detailed in Chapter 9, Section 3.5, i.e., *reverse auctions*, *bilateral trade*, *multiunit auctions*, and *buying a path in a network*, as well as many other problems can be modeled in this single-parameter agent framework.

13.1.1 Organization

Our discussion of optimal mechanism design will be divided up into three categories, depending on our assumptions about the agents' private values. On one hand, as is typical in economics, we can assume that agents' private values are drawn from a known prior distribution, the so-called Bayesian approach. Given knowledge of these prior distributions, the Bayesian optimal mechanism is the one that achieves the largest expected profit for these agents, where the expectation is taken over the randomness in the agents' valuations. In Section 13.2, we present the seminal result of Myerson, showing how to design the optimal, i.e., profit-maximizing, Bayesian auction given the prior distribution from which bidders' valuations are drawn.

On the other hand, in many cases, determining these prior distributions in advance may not be convenient, reasonable, or even possible. It is particularly difficult to collect priors in small markets, where the process of collecting information can seriously impact both the incentives of the agents and the performance of the mechanism. Thus, it is of great interest to understand to what extent we are able to design mechanisms for profit maximization even when we know very little about bidders' valuations. This approach leads us to the more traditional computer science approach of "worst-case analysis." While worst-case analysis could lead to results that are overly pessimistic, we shall see that in many cases we are able to obtain worst-case guarantees that are comparable to the optimal average-case guarantees for valuations from known distributions.

We begin our exploration of worst-case analysis in Section 13.3, where we survey techniques for approximating the optimal mechanism. We give natural mechanisms that approach optimality on large markets and a general formula for their performance as a function of the market size for small markets.

To obtain a theory of optimal mechanisms design without assumptions on the size of the market, we adopt a framework of *relative optimality*. This is motivated by two key observations. First, as we will explain later, there is no truthful mechanism that is best on every input. Second, in the worst case, all the agents' private values could be zero (or negligible) and thus no auction will be able to extract a high profit. In Section 13.4, we describe techniques for designing auctions that always (in worst case) return a profit that is within a small constant factor of some *profit benchmark* evaluated with respect to the agents' true private values.

Finally, in Section 13.5, we consider procurement settings where the auctioneer is looking to buy a set of goods or services that satisfy certain constraints, e.g., a path or a spanning tree in a graph. Specifically, we consider the problem of designing procurement auctions to minimize the total cost of the auctioneer (i.e., maximize their profit) relative to a natural benchmark.

We conclude the chapter with a discussion of directions for future research.

13.1.2 Preliminaries

In this section, we review basic properties of truthful mechanisms.

We will place two standard assumptions on our mechanisms. The first, that they are *individually rational*, means that no agent has negative expected utility for taking

part in the mechanism. The second condition we require is that of *no positive transfers* which restricts the mechanism to not pay the agents when they do not win, i.e., $x_i = 0 \rightarrow p_i = 0$.

In general, we will allow our mechanisms to be randomized. In a randomized mechanism, x_i is the probability that agent i is allocated the good, and p_i is agent i 's expected payment. Since x_i and p_i are outputs of the mechanism, it will be useful to view them as functions of the input bids as follows. We let $x_i(\mathbf{b})$, $p_i(\mathbf{b})$, and $u_i(\mathbf{b})$ represent agent i 's probability of allocation, expected price, and expected utility, respectively. Let $\mathbf{b}_{-i} = (b_1, \dots, b_{i-1}, ?, b_{i+1}, \dots, b_n)$ represent the vector of bids excluding bid i . Then with \mathbf{b}_{-i} fixed, we let $x_i(b_i)$, $p_i(b_i)$, and $u_i(b_i)$ represent agent i 's probability of allocation, expected price, and expected utility, respectively, as a function of their own bid. We further define the convenient notation $x_i(b_i, \mathbf{b}_{-i}) = x_i(\mathbf{b})$, $p_i(b_i, \mathbf{b}_{-i}) = p_i(\mathbf{b})$, and $u_i(b_i, \mathbf{b}_{-i}) = u_i(\mathbf{b})$.

Definition 13.5 A mechanism is *truthful in expectation* if and only if for all i , v_i , b_i , and \mathbf{b}_{-i} , agent i 's expected utility for bidding their valuation, v_i , is at least their expected utility for bidding any other value. In other words,

$$u_i(v_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i}).$$

For single-parameter agents, we restate the characterization of truthful mechanisms which was proven in Chapter 9, Section 5.6.

Theorem 13.6 A mechanism is *truthful in expectation* if and only if, for any agent i and any fixed choice of bids by the other agents \mathbf{b}_{-i} ,

- (i) $x_i(b_i)$ is monotone nondecreasing.
- (ii) $p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z) dz$.

Given this theorem, we see that once an allocation rule $\mathbf{x}(\cdot)$ is fixed, the payment rule $\mathbf{p}(\cdot)$ is also fixed. Thus, in specifying a mechanism we need specify only a monotone allocation rule and from it the truth-inducing payment rule can be derived.

It is useful to specialize Theorem 13.6 to the case where the mechanism is deterministic. In this case, the monotonicity of $x_i(b_i)$ implies that, for \mathbf{b}_{-i} fixed, there is some threshold bid t_i such that $x_i(b_i) = 1$ for all $b_i > t_i$ and 0 for all $t_i < b_i$. Moreover the second part of the theorem then implies that for any $b_i > t_i$, $p_i(b_i) = b_i - \int_{t_i}^{b_i} dz = t_i$. We conclude the following.

Observation 13.1.1 Any deterministic truthful auction is specified by a set of functions $t_i(\mathbf{b}_{-i})$ which determine, for each bidder i and each set of bids \mathbf{b}_{-i} , an offer price to bidder i such that bidder i wins and pays price t_i if $b_i > t_i$, or loses and pays nothing if $b_i < t_i$. (Ties can be broken arbitrarily.)

13.2 Bayesian Optimal Mechanism Design

In this section we describe the conventional economics approach of *Bayesian optimal mechanism design* where it is assumed that the valuations of the agents are drawn from a known distribution. The mechanism we describe is known as the Myerson mechanism: it is the truthful mechanism that maximizes the auctioneer’s expected profit, where the expectation is taken over the randomness in the agents’ valuations.

Consider, for example, a single-item auction with two bidders whose valuations are known to be drawn independently at random from the uniform distribution on $[0, 1]$. In Chapter 9, Section 6.3, it was shown that in this setting the expected revenue of both the Vickrey (second-price) auction and of the first-price auction is $1/3$. In fact, it was observed that any auction that always allocates the item to the bidder with the higher valuation achieves the same expected revenue.

Does this mean that $1/3$ is the best we can do, in expectation, with bidders of this type? The answer is no. Consider the following auction.

Definition 13.7 (Vickrey auction with reservation price r) The *Vickrey auction with reservation price r* , VA_r , sells the item to the highest bidder bidding at least r . The price the winning bidder pays is the maximum of the second highest bid and r .

It is a straightforward probabilistic calculation to show that the expected profit of the Vickrey auction with reservation price $r = 1/2$ is $5/12$. Thus, it is possible to get higher expected profit than the Vickrey auction by sometimes not allocating the item! This raises the problem of identifying, among the class of all truthful auctions, the auction that gives the optimal profit in expectation. The derivation in the next section answers this question and shows that in fact for this scenario $VA_{1/2}$ is the optimal auction.

13.2.1 Virtual Valuations, Virtual Surplus, and Expected Profit

We assume that the valuations of the agents, v_1, \dots, v_n , are drawn independently at random from known (but not necessarily identical) continuous probability distributions. For simplicity, we assume that $v_i \in [0, h]$ for all i . We denote by F_i the distribution function from which bidder i ’s valuation, v_i , is drawn (i.e., $F_i(z) = \Pr[v_i \leq z]$) and by f_i its density function (i.e., $f_i(z) = \frac{d}{dz}F_i(z)$). Since the agents’ valuations are independent, the joint distribution from which \mathbf{v} is drawn is just the product distribution $\mathbf{F} = F_1 \times \dots \times F_n$.

We now define two key notions: *virtual valuations* and *virtual surplus*.

Definition 13.8 The *virtual valuation* of agent i with valuation v_i is

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}.$$

Definition 13.9 Given valuations, v_i , and corresponding virtual valuations, $\phi_i(v_i)$, the *virtual surplus* of allocation \mathbf{x} is $\sum_i \phi_i(v_i)x_i - c(\mathbf{x})$.

As the *surplus* of an allocation is $\sum_i v_i x_i - c(\mathbf{x})$, the virtual surplus of an allocation is the surplus of the allocation with respect to agents whose valuations are replaced by their virtual valuations, $\phi_i(v_i)$.

We now show that any truthful mechanism has expected profit equal to its expected virtual surplus. Thus, to maximize expected profit, the mechanism should choose an allocation which maximizes virtual surplus. In so far as this allocation rule is monotone, this gives the optimal truthful mechanism!

Theorem 13.10 *The expected profit of any truthful mechanism, \mathcal{M} , is equal to its expected virtual surplus, i.e., $\mathbf{E}_v[\mathcal{M}(v)] = \mathbf{E}_v[\sum_i \phi_i(v_i)x_i(v) - c(\mathbf{x}(v))]$.*

Thus, if the mechanism, on each bid vector \mathbf{b} , chooses an allocation, \mathbf{x} , which maximizes $\sum_i \phi_i(b_i)x_i - c(\mathbf{x})$, the auctioneer's profit will be maximized. Notice that if we employ a deterministic tie-breaking rule then the resulting mechanism will be deterministic. Theorem 13.10 follows from Lemma 13.11 below, and the independence of the agents' valuations.

Lemma 13.11 *Consider any truthful mechanism and fix the bids \mathbf{b}_{-i} of all bidders except for bidder i . The expected payment of a bidder i satisfies:*

$$\mathbf{E}_{b_i}[p_i(b_i)] = \mathbf{E}_{b_i}[\phi_i(b_i)x_i(b_i)].$$

PROOF To simplify notation, we drop the subscript i and refer simply to the bid b being randomly chosen from distribution F with density function f .

By Theorem 13.6, we have

$$\mathbf{E}_b[p(b)] = \int_{b=0}^h p(b)f(b) db = \int_{b=0}^h bx(b)f(b) db - \int_{b=0}^h \int_{z=0}^b x(z)f(b) dz db.$$

Focusing on the second term and switching the order of integration, we have

$$\begin{aligned} \mathbf{E}_b[p(b)] &= \int_{b=0}^h bx(b)f(b) db - \int_{z=0}^h x(z) \int_{b=z}^h f(b) db dz. \\ &= \int_{b=0}^h bx(b)f(b) db - \int_{z=0}^h x(z) [1 - F(z)] dz. \end{aligned}$$

Now, we rename z to b and factor out $x(b)f(b)$ to get

$$\begin{aligned} \mathbf{E}_b[p(b)] &= \int_{b=0}^h bx(b)f(b) db - \int_{b=0}^h x(b) [1 - F(b)] db. \\ &= \int_{b=0}^h \left[b - \frac{1 - F(b)}{f(b)} \right] x(b)f(b) db. \\ &= \mathbf{E}_b[\phi(b)x(b)]. \quad \square \end{aligned}$$

13.2.2 Truthfulness of Virtual Surplus Maximization

Of course, it is not immediately clear that maximizing virtual surplus results in a truthful mechanism. By Theorem 13.6, this depends on whether or not virtual surplus maximization results in a monotone allocation rule. Recall that the VCG mechanism, which maximizes the actual surplus, i.e., $\sum_i v_i x_i - c(\mathbf{x})$, is truthful precisely because surplus maximization results in a monotone allocation rule. Clearly then, virtual surplus maximization gives an allocation that is monotone in agent valuations precisely when virtual valuation functions are monotone in agent valuations. Indeed, it is easy to find examples of the converse which show that nonmonotone virtual valuations result in a nonmonotone allocation rule. Thus, we conclude the following lemma.

Lemma 13.12 *Virtual surplus maximization is truthful if and only if, for all i , $\phi_i(v_i)$ is monotone nondecreasing in v_i .*

A sufficient condition for monotone virtual valuations is implied by the *monotone hazard rate* assumption. The hazard rate of a distribution is defined as $f(z)/(1 - F(z))$. Clearly, if the hazard rate is monotone nondecreasing, then the virtual valuations are monotone nondecreasing as well. There is a technical construction that extends these results to the nonmonotone case, but we do not cover it here.

Definition 13.13 Let \mathbf{F} be the prior distribution of agents' valuations satisfying the monotone hazard rate assumption. We denote by $\text{Mye}_{\mathbf{F}}(\mathbf{b})$ the *Myerson mechanism*: on input \mathbf{b} , output \mathbf{x} to maximize the virtual surplus (defined with respect to the distribution \mathbf{F}).

Thus, for single parameter problems, profit maximization in a Bayesian setting reduces to virtual surplus maximization. This allows us to describe Myerson's optimal mechanism, $\text{Mye}_{\mathbf{F}}(\mathbf{b})$, as follows:

- (i) Given the bids \mathbf{b} and \mathbf{F} , compute "virtual bids": $b'_i = \phi_i(b_i)$.
- (ii) Run VCG on the virtual bids \mathbf{b}' to get \mathbf{x}' and \mathbf{p}'
- (iii) Output $\mathbf{x} = \mathbf{x}'$ and \mathbf{p} with $p_i = \phi_i^{-1}(p'_i)$.

13.2.3 Applications of Myerson's Optimal Mechanism

The formulation of virtual valuations and the statement that the optimal mechanism is the one that maximizes virtual surplus is not the end of the story. In many relevant cases this formulation allows one to derive very simple descriptions of the optimal mechanism. We now consider a couple of examples to obtain a more precise understanding of $\text{Mye}_{\mathbf{F}}(\mathbf{b})$ and illustrate this point.

Example 13.14 (single-item auction) In a single-item auction, the surplus maximizing allocation gives the item to the bidder with the highest valuation, unless the highest valuation is less than 0 in which case the auctioneer keeps the item. Usually, we assume that all bidders' valuations are at least zero, or they would not want to participate in the auction, so the auctioneer never keeps the item.

However, when we maximize virtual surplus, it may be the case that a bidder has positive valuation but negative virtual valuation. Thus, for allocating a single item, the optimal mechanism finds the bidder with the largest nonnegative virtual valuation if there is one, and allocates to that bidder.

What about the payments? Suppose that there are only two bidders and we break ties in favor of bidder 1. Then bidder 1 wins precisely when $\phi_1(b_1) \geq \max\{\phi_2(b_2), 0\}$. This is a deterministic allocation rule, and thus the payment that a winning bidder 1 must make is the $p_1 = \inf\{b : \phi_1(b) \geq \phi_2(b_2) \wedge \phi_1(b) \geq 0\}$. Suppose that $F_1 = F_2 = F$, which implies that $\phi_1(z) = \phi_2(z) = \phi(z)$. Then this simplifies to $p_1 = \min(b_2, \phi^{-1}(0))$. Similarly, bidder 2's payment upon winning is $p_2 = \min(b_1, \phi^{-1}(0))$, thus we arrive at one of Myerson's main observations.

Theorem 13.15 *The optimal single-item auction for bidders with valuations drawn i.i.d. from distribution F is the Vickrey auction with reservation price $\phi^{-1}(0)$, i.e., $VA_{\phi^{-1}(0)}$.*

For example, when F is uniform on $[0, 1]$, we can plug the equations $F(z) = z$ and $f(z) = 1$ into the formula for the virtual valuation function (Definition 13.8) to conclude that $\phi(z) = 2z - 1$. Thus, the virtual valuations are uniformly distributed on $[-1, 1]$. We can easily solve for $\phi^{-1}(0) = 1/2$. We conclude that the optimal auction for two bidders with valuations uniform on $[0, 1]$ is the Vickrey auction with reservation price $1/2$, $VA_{1/2}$.

Example 13.16 (Digital goods auction) Recall that in a digital goods auction, we have $c(\mathbf{x}) = 0$ for all \mathbf{x} . Thus, to maximize virtual surplus, we allocate to each bidder such that $\phi_i(b_i) \geq 0$. As in the previous example, the payment a winning bidder must make is his minimum winning bid, i.e., $\inf\{b : \phi_i(b) \geq 0\}$, which is identically $\phi_i^{-1}(0)$.

Notice that with n bidders whose valuations are drawn independently from the same distribution function F , the reserve price for each bidder is $\phi^{-1}(0)$, the solution to $b - \frac{1-F(b)}{f(b)} = 0$. It is easy to check that this is precisely the *optimal sale price* for the distribution F : the take-it-or-leave-it price we would offer each bidder to maximize our expected profit.

Definition 13.17 (optimal sale price) *The optimal sale price for distribution F is $\text{opt}(F) = \arg\max_z z(1 - F(z))$.*

Summarizing, we obtain:

Theorem 13.18 *The optimal digital goods auction for n bidders with valuations drawn i.i.d. from distribution F is to offer each bidder the price $\text{opt}(F) = \phi^{-1}(0)$.*

13.3 Prior-Free Approximations to the Optimal Mechanism

In the previous section, we saw how to design the optimal mechanism when agents' valuations were drawn from known distributions. The assumption that the valuations are drawn from a known prior distribution makes sense in very large markets. In fact, as we shall see shortly, in large enough markets, a good approximation to the prior distribution can be learned on-the-fly and thus there are prior-free mechanisms that obtain nearly the optimal profit. We discuss these results in the first part of this section.

In small markets, on the other hand, incentives issues in learning an approximation of the prior distribution result in loss of performance and fundamental mechanism design challenges. Thus, new techniques are required in these settings. We develop an approach based on random sampling and analyze its performance in a way that makes explicit the connection between the size of the market and a mechanism's performance.

13.3.1 Empirical Distributions

The central observation that enables effective profit maximization without priors is Observation 13.1.1, which says that a truthful mechanism can use the reported bids of all the other agents in order to make a pricing decision for a particular agent.

Definition 13.19 (empirical distribution) For a vector of bids $\mathbf{b} = (b_1, \dots, b_n)$, the *empirical distribution* for these bids is $F_{\mathbf{b}}$ satisfying for $X \sim F_{\mathbf{b}}$, $\Pr[X > z] = n_z/n$, where n_z is the number of bids in \mathbf{b} above value z .

We now present a variant on Myerson's mechanism that can be used without any prior knowledge. As we shall see below, this mechanism has interesting interpretations in several contexts.

Definition 13.20 (empirical Myerson mechanism) The *empirical Myerson mechanism*, EM on input \mathbf{b} , for each i , simulates $\text{Mye}_{F_{\mathbf{b}_{-i}}}(\mathbf{b})$ to obtain outcome $\mathbf{x}^{(i)}$ and payments $\mathbf{p}^{(i)}$. It then produces outcome \mathbf{x} and \mathbf{p} with $x_i = x_i^{(i)}$ and $p_i = p_i^{(i)}$.

The outcome and payment for agent i in the empirical Myerson mechanism is based on the simulation of $\text{Mye}_{F_{\mathbf{b}_{-i}}}(\mathbf{b})$, and since agent i cannot manipulate $F_{\mathbf{b}_{-i}}$, this mechanism is truthful.

There are two issues that we need to address in order to understand the performance of the EM mechanism. First, we need to see if the outcomes it produces are feasible. The issue is that the allocation to different agents, say i and j , is determined from different information (\mathbf{b}_{-i} versus \mathbf{b}_{-j}). As we shall see, this inconsistency will sometimes produce allocations, \mathbf{x} , that are not feasible (i.e., $c(\mathbf{x}) = \infty$). Second, in those situations where it does produce feasible allocations, we need to understand how effective the mechanism is at profit maximization. The hope is that, in large markets, $F_{\mathbf{b}_{-i}}$ should

be close to $F_{\mathbf{b}}$ and hence the law of large numbers should imply good performance. Again, we will see that this does not hold in general.

We begin by considering the application of EM to digital goods auctions, where there is no feasibility issue.

Definition 13.21 (deterministic optimal price auction) We define the *deterministic optimal price auction* (DOP) as EM applied to the digital goods auction problem.

In the previous section, we saw that if each agent's valuation is drawn from the same distribution F , Myerson's mechanism offers price $\phi^{-1}(0) = \text{opt}(F)$ to each bidder. The deterministic optimal price auction, on the other hand, offers agent i price $\text{opt}(F_{\mathbf{b}_{-i}})$. Using the short-hand notation $\text{opt}(\mathbf{b}) = \text{opt}(F_{\mathbf{b}})$, Observation 13.1.1 allows us to express DOP quite simply as the auction defined by $t_i(\mathbf{b}_{-i}) = \text{opt}(\mathbf{b}_{-i})$. Since \mathbf{b}_{-i} is different for each agent, in general the prices offered the agents are different. Nonetheless, the law of large numbers implies the following result (which is a corollary of Theorem 13.30 proved in the next section).

Theorem 13.22 *For the digital goods setting and n bids \mathbf{b} distributed i.i.d. from distribution F with bounded support, the profit of DOP approaches the profit of Myerson in the limit as n increases.*

Unfortunately, the assumption that the input comes from an unknown, but i.i.d. distribution is crucially important to this result as the following example shows.

Example 13.23 With 10 bids at \$10, and 90 bids at \$1, consider the prices $t_i(\mathbf{b}_{-i})$ and $t_i(\mathbf{b}_{-10})$ that DOP offers bidders bidding \$1 and \$10 respectively:

- \mathbf{b}_{-1} is 89 bids at \$1 and 10 bids at \$10, so $\text{opt}(\mathbf{b}_{-1}) = \10 , and
- \mathbf{b}_{-10} is 90 bids at \$1 and 9 bids at \$10, so $\text{opt}(\mathbf{b}_{-10}) = \1 .

Thus, bids at \$10 are accepted, but offered price \$1, while bids at \$1 are rejected. The total profit is \$10 whereas the optimal is \$100. This example can be made arbitrarily bad.

What happened in this example is the result of the inconsistency between the distribution $F_{\mathbf{b}_{-i}}$ assumed when choosing a price for agent i , and the distribution $F_{\mathbf{b}_{-j}}$ assumed when choosing a price for agent j . Had we just run $\text{Mye}_{F_{\mathbf{b}_{-i}}}$ or $\text{Mye}_{F_{\mathbf{b}_{-j}}}$ on all bids, all would have been well. Indeed, in this example, we would have chosen either price \$1 for everyone or price \$10 for everyone. Both prices would have been fine.

This problem is not just one with DOP, but with any symmetric deterministic digital goods auction.¹ Indeed, the problem inherent in this example can be generalized to prove the following theorem.

¹ An auction is *symmetric* if the outcome and prices are not a function of the order of the input bids, but rather just the set of bids.

Theorem 13.24 *There do not exist constants β and γ and a symmetric deterministic truthful auction, \mathcal{A} , with profit at least $\text{OPT} / \beta - h\gamma$ on all bid vectors \mathbf{b} with $b_i \in [1, h]$.*

The inconsistency of EM can be more serious than just low profit on some, perhaps unlikely, inputs; if some outcomes are infeasible (i.e., $c(\mathbf{x}) = \infty$ for some \mathbf{x}) then EM may result in infeasible outcomes! In the next section we see how these consistency issues can be partially addressed through the use of random sampling.

13.3.2 Random Sampling

Random sampling plays an important role in the design of economic mechanisms. For example, during elections, polls that predict each candidate’s ranking affect the results of the elections; and in many settings, market analysis and user studies using a (small) random sample of the population can lead to good decisions in product development and pricing. In this section, we consider a natural extension of the empirical Myerson mechanism that uses random sampling to address the consistency issues raised in the preceding section.

Definition 13.25 (Random sampling empirical myerson) *The random sampling empirical Myerson mechanism (RSEM) works as follows:*

- (i) Solicit bids $\mathbf{b} = (b_1, \dots, b_n)$.
- (ii) Partition the bids into two sets \mathbf{b}' and \mathbf{b}'' uniformly at random.
- (iii) Compute empirical distributions for each set $F' = F_{\mathbf{b}'}$ and $F'' = F_{\mathbf{b}''}$.
- (iv) Run $\text{Mye}_{F''}(\mathbf{b}')$ and $\text{Mye}_{F'}(\mathbf{b}'')$.

For digital goods auctions, we can replace Steps iii and iv by their more natural interpretations (facilitated by the short-hand notation $\text{opt}(\mathbf{b}) = \text{opt}(F_{\mathbf{b}})$):

- (iii)’ Compute the optimal sale prices $p' = \text{opt}(\mathbf{b}')$ and $p'' = \text{opt}(\mathbf{b}'')$.
- (iv)’ Offer price p' to bidders in \mathbf{b}'' and price p'' to bidders in \mathbf{b}' .

We refer to the digital goods variant of the random sampling empirical Myerson mechanism as the *random sampling optimal price auction (RSOP)*. The randomization in RSOP allows it to bypass the deterministic impossibility for worst case settings leading to the following theorem. (Again, this is as a corollary of Theorem 13.30 which is proven in the next section.)

Theorem 13.26 *For \mathbf{b} with $b_i \in [1, h]$, the expected revenue of RSOP approaches that of the optimal single price sale as the number of bidders grows.*

Similar results do not necessarily hold for more general settings. It is easy to imagine situations where RSEM also gives infeasible outcomes as the following example illustrates.

Example 13.27 Consider the setting where we are selling a digital good in one of two markets, for convenience, bidders $1, \dots, i$ are in market A and bidders $i + 1, \dots, n$ are in market B . Due, perhaps, to government regulations, it is not legal to sell the good to bidders in both markets simultaneously. Thus, feasible solutions will have winners either only from market A or only from market B . It is easy to construct settings where RSEM will sell to one market in \mathbf{b}' and the other in \mathbf{b}'' . The combined outcome, however, is not feasible.

The biggest open problem in prior-free mechanism design is to understand how to approximate the optimal mechanism in more general settings.

13.3.3 Convergence Rates

As we have discussed above, the law of large numbers implies that the profit of the random sampling auction, say in the case of digital goods, is asymptotically optimal as the number of bidders grows. In this section, we study the rate at which the auction approaches optimal performance. The theorem we prove will enable us to obtain a precise relationship between the complexity of the class of outcomes considered by RSOP and its convergence rate. The results in this section will also give us a framework for evaluating the performance of random sampling-based mechanisms in very general contexts.

We make our discussion concrete, using the example of the digital goods auction problem. Recall that RSOP uses a subroutine that computes the optimal sale price for the bids in each partition of the bidders. Suppose that we allowed the auctioneer the ability to artificially restrict prices to be in some set \mathcal{Q} . For example, the auctioneer might only sell at integer prices, in which case \mathcal{Q} would be the set of integers. The auctioneer could further limit the set of possible prices, for example, by having \mathcal{Q} be powers of 2. We will see that different choices of \mathcal{Q} will give us different bounds on the convergence rate.

Given \mathcal{Q} , we define $\text{RSOP}_{\mathcal{Q}}$ as the random sampling auction that computes the optimal price from \mathcal{Q} on each partition and offers it to bidders in the opposite partition. We make use of the following notation. Let $q(b_i)$ be the payment made by bidder i when offered $q \in \mathcal{Q}$. That is, $q(b_i) = q$ if $b_i \geq q$ and $q(b_i) = 0$ otherwise. Let $q(\mathbf{b}) = \sum_i q(b_i)$. Finally, define $\text{opt}_{\mathcal{Q}}(\mathbf{b}) = \text{argmax}_{q \in \mathcal{Q}} q(\mathbf{b})$ as the q that gives the optimal profit for \mathbf{b} , and $\text{OPT}_{\mathcal{Q}}(\mathbf{b})$ to be this optimal profit, i.e., $\text{OPT}_{\mathcal{Q}}(\mathbf{b}) = \max_{q \in \mathcal{Q}} q(\mathbf{b})$.

The bounds we give in this section show the rate at which the profit of $\text{RSOP}_{\mathcal{Q}}(\mathbf{b})$ approaches $\text{OPT}_{\mathcal{Q}}(\mathbf{b})$ with some measure of the size of the market. The measure we use is $\text{OPT}_{\mathcal{Q}}$ itself, as this gives us the most general and precise result. Thus, these results show the degree to which $\text{RSOP}_{\mathcal{Q}}$ approximates $\text{OPT}_{\mathcal{Q}}$ as $\text{OPT}_{\mathcal{Q}}$ grows large in comparison to h , an upper bound on the payment of any agent, and the complexity of \mathcal{Q} .

Definition 13.28 Given partitions \mathbf{b}' and \mathbf{b}'' , price q in \mathcal{Q} is ϵ -good if

$$|q(\mathbf{b}') - q(\mathbf{b}'')| \leq \epsilon \text{OPT}_{\mathcal{Q}}(\mathbf{b})$$

Lemma 13.29 For \mathbf{b} and h satisfying $q(b_i) \leq h$, for all i , if bids \mathbf{b} are randomly partitioned into \mathbf{b}' and \mathbf{b}'' then q is not ϵ -good with probability at most $2e^{-\epsilon^2 \text{OPT}_{\mathcal{Q}}(\mathbf{b})/2h}$.

The proof of this lemma follows from McDiarmid’s inequality, see Exercise 13.5. The following is the main theorem of this section. A natural interpretation of h is an upper bound on the highest valuation, i.e., $h = \max_i v_i$.

Theorem 13.30 For \mathcal{Q} , \mathbf{b} , and h satisfying $q(b_i) \leq h$, for all q and i , and $\text{OPT}_{\mathcal{Q}}(\mathbf{b}) \geq \frac{8h}{\epsilon^2} \ln(\frac{2|\mathcal{Q}|}{\delta})$, with probability $1 - \delta$ the profit of $\text{RSOP}_{\mathcal{Q}}$ is at least $(1 - \epsilon) \text{OPT}_{\mathcal{Q}}(\mathbf{b})$.

PROOF Assume that $\text{OPT}_{\mathcal{Q}}(\mathbf{b}) \geq \frac{8h}{\epsilon^2} \ln(\frac{2|\mathcal{Q}|}{\delta})$. For random partitioning of \mathbf{b} into \mathbf{b}' and \mathbf{b}'' , Lemma 13.29 implies that the probability $q \in \mathcal{Q}$ is not $\frac{\epsilon}{2}$ -good is at most $\delta/|\mathcal{Q}|$. Using a union bound over all $q \in \mathcal{Q}$, we have that all $q \in \mathcal{Q}$ are $\frac{\epsilon}{2}$ -good with probability $1 - \delta$.

Let $q' = \text{opt}_{\mathcal{Q}}(\mathbf{b}')$, $q'' = \text{opt}_{\mathcal{Q}}(\mathbf{b}'')$, and $q^* = \text{opt}_{\mathcal{Q}}(\mathbf{b})$. By definition, $q'(\mathbf{b}') \geq q^*(\mathbf{b}')$ and likewise $q''(\mathbf{b}'') \geq q^*(\mathbf{b}'')$. Thus, $q'(\mathbf{b}') + q''(\mathbf{b}'') \geq q^*(\mathbf{b}) = \text{OPT}_{\mathcal{Q}}(\mathbf{b})$. If all q are $\frac{\epsilon}{2}$ -good, certainly q' and q'' are; therefore, $q'(\mathbf{b}'') \geq q'(\mathbf{b}') - \frac{\epsilon}{2} \text{OPT}_{\mathcal{Q}}(\mathbf{b})$ and $q''(\mathbf{b}') \geq q''(\mathbf{b}'') - \frac{\epsilon}{2} \text{OPT}_{\mathcal{Q}}(\mathbf{b})$. Thus, we conclude that our auction profit, which is $q'(\mathbf{b}'') + q''(\mathbf{b}')$ is at least $(1 - \epsilon) \text{OPT}_{\mathcal{Q}}(\mathbf{b})$ with probability $1 - \delta$ which gives the theorem. \square

Notice that this theorem holds for all ϵ and δ . In particular, it shows how big the optimal profit must be before we can guarantee a certain approximation factor. Of course, in the limit as the optimal profit goes to infinity, our approximation factor approaches one. We refer to the lower bound required of optimal profit, $\text{OPT}_{\mathcal{Q}}$, in the statement of the theorem as the *convergence rate*. Indeed, if the agents’ valuations are between 1 and h , the lower bound on the optimal profit can be translated into a lower bound on the size of the market needed to guarantee the approximation.

Let us now consider a few applications of the theorem: Suppose that $\mathcal{Q} = \{1, \dots, h\}$. Then $|\mathcal{Q}| = h$ and the convergence rate to a $(1 - \epsilon)$ -approximation with probability $1 - \delta$ is $O(h\epsilon^{-2} \log(2h/\delta))$. If instead \mathcal{Q} is powers of 2 on the interval $[1, h]$, then $|\mathcal{Q}| = \log h$ and the convergence rate for constant ϵ and δ is $O(h \log \log h)$.

It is worth noting that the particular bids \mathbf{b} that are input to any particular run of $\text{RSOP}_{\mathcal{Q}}$ may further restrict the set of possible prices in \mathcal{Q} that can be selected, say to some subset \mathcal{Q}' . We can apply Theorem 13.30 retrospectively to input \mathbf{b} to bound the performance of $\text{RSOP}_{\mathcal{Q}}$ in terms of $|\mathcal{Q}'|$. For example, in the original RSOP auction we consider all real numbers as prices; yet, $\text{opt}(\mathbf{b})$ is always one of the bids. Thus, using $\mathcal{Q}' = \{b_1, \dots, b_n\}$ and noting that $|\mathcal{Q}'| = n$, tells us that the convergence rate of our original RSOP digital good auction is $O(h\epsilon^{-2} \ln(2n/\delta))$. Even better bounds are possible using a notion called γ -covers (Exercise 13.6).

Corollary 13.31 For \mathcal{Q} , \mathcal{Q}' , \mathbf{b} , and h satisfying $q(b_i) \leq h$ for all q and i , $\text{opt}(\mathbf{b}') \in \mathcal{Q}'$ for all subsets \mathbf{b}' of \mathbf{b} , and $\text{OPT}_{\mathcal{Q}}(\mathbf{b}) \geq \frac{8h}{\epsilon^2} \ln(\frac{2|\mathcal{Q}'|}{\delta})$; with probability $1 - \delta$ the profit of $\text{RSOP}_{\mathcal{Q}}$ is at least $(1 - \epsilon) \text{OPT}_{\mathcal{Q}}(\mathbf{b})$.

Lemma 13.29 and Theorem 13.30 are quite general and can be applied, as written, to a wide variety of unlimited supply auction problems with rich structure on the class of allowable offers, \mathcal{Q} . Two examples are *attribute auctions* and *combinatorial auctions*.

13.4 Prior-Free Optimal Mechanism Design

In the previous sections, a number of results on approximating the optimal mechanism in worst-case settings were presented. Unfortunately, these results remain limited in their applicability. For example, what if $\text{OPT}_{\mathcal{Q}}(\mathbf{b})$ is too small, as might happen if the size of the market (i.e., the number of bidders) is too small? In such cases, Theorem 13.30 may give us no guarantee. Thus, a natural question to ask is: what is the best truthful mechanism? Can we design a truthful mechanism for which we can prove nontrivial performance guarantees under any market conditions?

The first observation that must be made is that there is no such thing as an absolute “best” truthful auction. To gain some intuition for this statement, recall that in any truthful auction, the offer price t_i to bidder i is a function of all other bids \mathbf{b}_{-i} , but not of b_i . Thus, given any particular auction, which is forced to fix the offer price t_i independently of b_i , (and hence always performs suboptimally for most values of b_i), there is always some input on which a different truthful auction performs better (see Exercise 13.8).

Given that there is no absolute best truthful mechanism on all inputs, we are left with the question of how we can arrive at a rigorous theoretical framework in which we can compare auctions and determine one to be better. The key to resolving this issue is in moving from absolute optimality to relative optimality. Indeed, whenever there is an information theoretic obstacle or computational intractability preventing an absolute optimal solution to a problem we can try to approximate. For example, in the design of online algorithms the objective is to find an online algorithm that performs comparably to an optimal offline algorithm. The notable analogy here is between the game theoretic constraint that a mechanism does not know the true bid values in advance and must solicit them in a truth-inducing manner, and the online constraint that an online algorithm does not have knowledge of the future.

13.4.1 Competitive Framework

The general approach will be to try to design an auction with profit that is always (in worst case) within a small constant factor of some *profit benchmark*.

Definition 13.32 A *profit benchmark* is a function $\mathcal{G} : \mathbb{R}^n \rightarrow \mathbb{R}$ which maps a vector of valuations to a target profit.

The following definition captures the intuition that an auction is good if it is always close to a reasonable profit benchmark.

Definition 13.33 The *competitive ratio* of auction \mathcal{A} (defined with respect to an implicit profit benchmark \mathcal{G}) is $\beta = \sup_{\mathbf{v}} \frac{\mathcal{G}(\mathbf{v})}{\mathcal{A}(\mathbf{v})}$.

Given a profit benchmark \mathcal{G} the task of an auction designer is to design an auction that achieves the minimum possible competitive ratio. This auction is the *optimal competitive auction* for \mathcal{G} .

13.4.2 A Competitive Digital Goods Auctions

In this section, we will see that the RSOP auction that was defined in Section 13.3.2 is in fact a competitive digital goods auction. To make this statement precise, we first need to define the profit benchmark we will be attempting to compete with. In the analysis of online algorithms it is not always best to gauge the performance of an online algorithm by comparing it to an unconstrained optimal offline algorithm. Similarly, in the analysis of truthful auctions, it sometimes makes sense to compare an auction's profit to a profit benchmark that is not necessarily the profit of the optimal algorithm that is given the bidders' true valuations in advance.

For digital goods auctions, natural profit benchmarks, such as (a) the maximum profit achievable with fully discriminating prices (where each bidder pays their valuation) or (b) the maximum profit achievable with a single price, are provably too strong in the sense that no truthful auction can be constant competitive with these benchmarks.

Thus, the profit benchmark we will use is the following.

Definition 13.34 ($\mathcal{F}^{(2)}$) The *optimal single priced profit with at least two winners* is

$$\mathcal{F}^{(2)}(\mathbf{v}) = \max_{i \geq 2} i v_{(i)},$$

where $v_{(i)}$ is the i th largest valuation.

Theorem 13.24 in Section 13.3.1 can be extended to this setting to show:

Corollary 13.35 *No symmetric deterministic truthful auction has constant competitive ratio relative to the profit benchmark $\mathcal{F}^{(2)}$.*

Thus, we turn to randomized auctions where we find the following theorem.

Theorem 13.36 *RSOP is 15-competitive with $\mathcal{F}^{(2)}$.*

We will not prove Theorem 13.36 here as it is primarily a technical probabilistic analysis. We do note, however, that 15 is likely to be a loose upper bound. On the other hand, it is easy to see that RSOP cannot have a competitive ratio better than 4, by considering the bid vector $\mathbf{b} = (\$1, \$2)$. With probability 1/2 both bids end up in the

same part and the RSOP profit is 0. Otherwise, with probability $1/2$ one bid is in each part. Without loss of generality, $\mathbf{b}' = \{\$1\}$ and $\mathbf{b}'' = \{\$2\}$, then $p' = \$1$ and $p'' = \$2$. Thus, the \$1-bid is rejected (because she cannot pay \$2) and the \$2-bid is offered a price of \$1 which she accepts. The RSOP profit in this case is \$1. The expected profit of RSOP is therefore \$0.50 while $\mathcal{F}^{(2)}(\mathbf{b}) = \2 , which shows that RSOP is at best 4-competitive. It is conjectured that this two bid input is in fact the worst case and that RSOP has a competitive ratio of 4.

13.4.3 Lower Bounds

Now that we have seen that there exists an auction that has constant competitive ratio to $\mathcal{F}^{(2)}$, it is interesting to ask: what is the *optimal auction* in terms of worst case competitive ratio to $\mathcal{F}^{(2)}$? What is the competitive ratio of this optimal auction? In this section, we approach this question from the other side, by looking for lower bounds on the competitive ratio. Specifically, we discuss a proof that shows that no auction is better than 2.42-competitive.

Theorem 13.37 *No auction has competitive ratio less than 2.42.*

The proof of this theorem involves a fairly complicated analysis of the expected value of $\mathcal{F}^{(2)}(\mathbf{b})$ when \mathbf{b} is generated from a particular probability distribution. We will instead prove a simpler result which highlights the main ideas of the theorem.

Lemma 13.38 *No 2-bidder auction has competitive ratio less than 2.*

PROOF The proof follows a simple structure that is useful for proving lower bounds for this type of problem. First, we consider bids drawn from a particular distribution. Second, we argue that for any auction \mathcal{A} , $\mathbf{E}_{\mathbf{b}}[\mathcal{A}(\mathbf{b})] \leq \mathbf{E}_{\mathbf{b}}[\mathcal{F}^{(2)}(\mathbf{b})]/2$. This implies that there exists a bid vector \mathbf{b}^* such that $\mathcal{A}(\mathbf{b}^*) \leq \mathcal{F}^{(2)}(\mathbf{b}^*)/2$.

We choose a distribution to make the analysis of $\mathbf{E}_{\mathbf{b}}[\mathcal{A}(\mathbf{b})]$ simple. This is important because we have to analyze it for all auctions \mathcal{A} . The idea is to choose the distribution for \mathbf{b} so that all auctions obtain the same expected profit. Consider \mathbf{b} with b_i satisfying $\Pr[b_i > z] = 1/z$. Note that whatever the price t_i is that \mathcal{A} offers bidder i , the expected payment made by bidder i is $t_i \times \Pr[b_i \geq t_i] = 1$. Thus, for $n = 2$ bidders the expected profit of any truthful auction is $\mathbf{E}_{\mathbf{b}}[\mathcal{A}(\mathbf{b})] = n = 2$.

We must now calculate $\mathbf{E}_{\mathbf{b}}[\mathcal{F}^{(2)}(\mathbf{b})]$. $\mathcal{F}^{(2)}(\mathbf{b}) = \max_{i \geq 2} i b_{(i)}$ where $b_{(i)}$ is the i th highest bid value. In the case that $n = 2$, this simplifies to $\mathcal{F}^{(2)}(\mathbf{b}) = 2b_{(2)} = 2 \min(b_1, b_2)$. We recall that a nonnegative random variable X has $\mathbf{E}[X] = \int_0^\infty \Pr[X \geq z] dz$ and calculate $\Pr[\mathcal{F}^{(2)}(\mathbf{b}) > z]$.

$$\begin{aligned} \Pr_{\mathbf{b}}[\mathcal{F}^{(2)}(\mathbf{b}) > z] &= \Pr_{\mathbf{b}}[b_1 \geq z/2 \wedge b_2 \geq z/2] \\ &= \Pr_{\mathbf{b}}[b_1 \geq z/2] \Pr_{\mathbf{b}}[b_2 \geq z/2] \\ &= 4/z^2. \end{aligned}$$

Note that this equation is valid only for $z \geq 2$. Of course for $z < 2$, $\Pr[\mathcal{F}^{(2)}(\mathbf{b}) \geq z] = 1$. Thus,

$$\mathbf{E}_{\mathbf{b}}[\mathcal{F}^{(2)}(\mathbf{b})] = \int_0^\infty \Pr[\mathcal{F}^{(2)}\mathbf{b} \geq z] dz = 2 + \int_2^\infty \frac{4}{z^2} dz = 4.$$

For this distribution and any auction \mathcal{A} , $\mathbf{E}_{\mathbf{b}}[\mathcal{A}(\mathbf{b})] = 2$ and $\mathbf{E}_{\mathbf{b}}[\mathcal{F}^{(2)}(\mathbf{b})] = 4$. Thus, the inequality $\mathbf{E}_{\mathbf{b}}[\mathcal{A}(\mathbf{b})] \leq \mathbf{E}_{\mathbf{b}}[\mathcal{F}^{(2)}(\mathbf{b})]/2$ holds and there must exist some input \mathbf{b}^* such that $\mathcal{A}(\mathbf{b}^*) \leq \mathcal{F}^{(2)}(\mathbf{b}^*)/2$. \square

For two bidders, this lower bound is tight. Indeed, it is trivial to check that for two bidders, the Vickrey auction has competitive ratio 2.

The lower bound proof given above can be generalized by a more complicated analysis to larger n . Such an analysis leads to bounds of $13/6$ for $n = 3$ and eventually to a bound of 2.42 for general n . It is conjectured that these bounds are tight. Indeed they are tight for $n \leq 3$.

13.4.4 The Digital Goods Auction Decision Problem

In the next sections, we derive an auction with a competitive ratio of 4. We do this by defining the notion of a decision problem for mechanism design and reducing the problem of designing a good competitive auction to it.

Definition 13.39 The *digital goods auction decision problem* is: given n bidders, n units of an item, and a target profit R , design a *truthful* mechanism that obtains profit R if possible, i.e., if $R \leq \mathcal{F}(\mathbf{v})$. Here, $\mathcal{F}(\mathbf{v}) = \max_{i \geq 1} i v_{(i)}$, where $v_{(i)}$ is the i th largest valuation.

This digital goods auction decision problem is also known as the *profit extraction* problem as its goal is to extract a profit R from a set of bidders. It turns out that this problem is solved by a special case of a general *cost-sharing* mechanism.

Definition 13.40 (ProfitExtract_R) The *digital goods auction profit extractor* with target profit R sells to the largest group of k bidders that can equally share R and charges each R/k .

It is straightforward to show that ProfitExtract_R is truthful and obtains a profit of R when $\mathcal{F}(\mathbf{b}) \geq R$ (see Exercise 13.10).

13.4.5 Reduction to the Decision Problem

A classical optimization problem can typically be phrased as follows: “find a feasible solution that maximizes some objective function.” The decision problem version of this is: “is there a feasible solution for which the objective function has value at least V ?” A standard reduction between the two involves solving the decision problem many times, using binary search over values V . Unfortunately, such an approach will not work for

mechanism design as it is not truthful to run several truthful mechanisms and then only take the output of the one that is the most desirable.

The following truthful auction avoids this problem.

Definition 13.41 (RSPE) The *Random Sampling Profit Extraction* auction (RSPE) works as follows:

- (i) Randomly partition the bids \mathbf{b} into two by flipping a fair coin for each bidder and assigning her to \mathbf{b}' or \mathbf{b}'' .
- (ii) Compute $R' = \mathcal{F}(\mathbf{b}')$ and $R'' = \mathcal{F}(\mathbf{b}'')$, the optimal profits for each part.
- (iii) Run ProfitExtract $_{R'}$ on \mathbf{b}'' and ProfitExtract $_{R''}$ on \mathbf{b}' .

The intuition for this auction is that ProfitExtract $_R$ allows us treat a set of bidders, \mathbf{b} , as one bidder with bid value $\mathcal{F}(\mathbf{b})$. Recall that a truthful auction must just offer a price t_i to bidder i who accepts if her value is at least t_i . This is analogous to trying to extract a profit R from bidders \mathbf{b} and actually getting R in profit when $\mathcal{F}(\mathbf{b}) \geq R$. The RSPE auction can then be viewed as randomly partitioning the bidders into two parts, treating one partition of the bids \mathbf{b}' as a single bid with value $R' = \mathcal{F}^{(2)}(\mathbf{b}')$, the other partition \mathbf{b}'' as a single bid with value $R'' = \mathcal{F}^{(2)}(\mathbf{b}'')$, and then running the Vickrey auction on these two “bids.” This intuition is crucial for the proof that follows as it implies that the profit of RSPE is the minimum of R' and R'' .

Theorem 13.42 *The competitive ratio of RSPE is 4.*

PROOF As we discussed above, the profit of RSPE is $\min(R', R'')$. Thus, we just need to analyze $\mathbf{E}[\min(R', R'')]$.

Assume that $\mathcal{F}^{(2)}(\mathbf{b}) = kp$ has with $k \geq 2$ winners at price p . Of the k winners in $\mathcal{F}^{(2)}$, let k' be the number of them that are in \mathbf{b}' and k'' the number that are in \mathbf{b}'' . Since there are k' bidders in \mathbf{b}' at price p , $R' \geq k'p$. Likewise, $R'' \geq k''p$. Thus,

$$\frac{\mathbf{E}[\text{RSPE}(\mathbf{b})]}{\mathcal{F}^{(2)}(\mathbf{b})} = \frac{\mathbf{E}[\min(R', R'')]}{kp} \geq \frac{\mathbf{E}[\min(k'p, k''p)]}{kp} = \frac{\mathbf{E}[\min(k', k'')]}{k} \geq \frac{1}{4}.$$

The last inequality follows from the fact that if $k \geq 2$ fair coins (corresponding to placing the winning bidders into either \mathbf{b}' or \mathbf{b}'') are flipped then $\mathbf{E}[\min\{\#\text{heads}, \#\text{tails}\}] \geq k/4$.

It is evident that RSPE is no better than 4-competitive via an identical proof to that of the analogous result for RSOP. \square

The currently best known competitive auction, which has a competitive ratio of 3.25, is based on generalizing the idea of RSPE: First, the bids are randomly partitioned into three parts, instead of two, with each part being treated as a single bid with value equal to its optimal single price revenue. Then the optimal 3-bidder auction is run on these three “bids.”

The random partitioning and profit extraction approach is fairly general. For it to work successfully, it needs to be shown that a profit extractor for the benchmark exists,

and that up to constant factors, the benchmark is preserved on a random sample of the agents. Notice that the consistency issue discussed in earlier sections is not relevant if only the agents in one partition win. This approach has been applied successfully to several other settings.

13.4.6 Consensus Estimation and Truthfulness with High Probability

We now look at an alternative reduction to the decision problem and an approach to competitive auctions that does not use random sampling. This approach leads to a truthful digital goods auction that is 3.39-competitive with $\mathcal{F}^{(2)}$. However, rather than presenting that result, we present a more general version of the approach with wider applicability. To achieve this greater level of generality, we will need to relax our solution concept and talk about *truthfulness with high probability*.

Definition 13.43 A randomized mechanism is *truthful with high probability*, say $1 - \epsilon$, if and only if for all i , v_i , b_i , and \mathbf{b}_{-i} , the probability that agent i benefits by bidding nontruthfully is at most ϵ , where the probability is taken over the coin flips of the mechanism. In other words, for all i , v_i , b_i , and \mathbf{b}_{-i} , $\Pr[u_i(v_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i})] \geq 1 - \epsilon$.

The techniques presented in this section, when applied to the digital goods auction, result in a mechanism that is truthful with probability $1 - O(1/m)$ where m is the number of winners in $\mathcal{F}^{(2)}$. Thus, as the input instance grows and there are more winners, the probability that nontruthful reporting by the agents is beneficial approaches zero.

Let us first describe the general idea. Consider attempting to design an auction to compete with profit benchmark \mathcal{G} . Suppose that there exists a profit extractor for \mathcal{G} , $\text{ProfitExtract}_{\mathcal{G},R}$, which obtains profit R from \mathbf{b} if $R \leq \mathcal{G}(\mathbf{b})$. Then the mechanism we would like to run is the following:

- (i) Compute $R = \mathcal{G}(\mathbf{b})$.
- (ii) Run $\text{ProfitExtract}_{\mathcal{G},R}$ on \mathbf{b} .

This fails of course because, generally, the R computed in Step (i) is a function of an agent's bid and therefore the agent could misreport their bid to obtain an R value that results in a more favorable outcome for them in Step (ii).

On the other hand, it is often the case that a single agent only contributes a small fraction to the profit $\mathcal{G}(\mathbf{b})$. In particular, suppose that there is some ρ such that for all i , $\mathcal{G}(\mathbf{b}_{-i}) \in [\mathcal{G}(\mathbf{b})/\rho, \mathcal{G}(\mathbf{b})]$. In this case $\mathcal{G}(\mathbf{b}_{-i})$ is a pretty good estimate of $\mathcal{G}(\mathbf{b})$. The idea then is to replace Step (i) above with

- (i)' Compute $R = r(\mathcal{G}(\mathbf{b}))$.

where the probabilistic function $r(\cdot)$ is a ρ -consensus β -estimate:

Definition 13.44 A (randomized) function $r(\cdot)$ is a ρ -consensus if for all $V > 0$ with high probability all $V' \in [V/\rho, V]$ satisfy $r(V') = r(V)$.

Intuitively, if $r(\cdot)$ is a ρ -consensus then with high probability $r(\mathcal{G}(\mathbf{b})) = r(\mathcal{G}(\mathbf{b}_{-i}))$ for all i . This will imply that bidder i has very small probability of being able to influence the value of $r(\mathcal{G}(\mathbf{b}))$ and thus we will be able to guarantee truthfulness with high probability.

Definition 13.45 A (randomized) function $r(\cdot)$ is a β -estimate if for all $V > 0$ it satisfies $r(V) \leq V$ and $\mathbf{E}[r(V)] \geq V/\beta$.

Intuitively, if $r(\cdot)$ is a β -estimate, then $r(\mathcal{G}(\mathbf{b}))$ is close to, but less than, $\mathcal{G}(\mathbf{b})$. If this is the case, then running ProfitExtract $_{\mathcal{G},R}$ on \mathbf{b} , with $R = r(\mathcal{G}(\mathbf{b}))$, will extract a revenue R which is close to $\mathcal{G}(\mathbf{b})$.

Of course, even in Step (i)', R is a function of all the bids, so the resulting auction is not truthful. However, under some mild assumptions² it is possible to show that in the case that $r(\mathcal{G}(\mathbf{b}))$ is a consensus no bidder has an incentive to deviate and misreport their valuation. The resulting mechanism is truthful with high probability.

We now show how to construct the function $r(\cdot)$.

Definition 13.46 (r_α) Given $\alpha > 1$, the randomized function $r_\alpha(\cdot)$ picks U uniformly from $[0, 1]$ and is

$$r_\alpha(V) = "V \text{ rounded down to the nearest } \alpha^{i+U} \text{ for integer } i."$$

Straightforward probabilistic analysis can be used to prove the following lemmas.

Lemma 13.47 r_α is a ρ -consensus with probability $1 - \log_\alpha \rho$.

Lemma 13.48 r_α is a β -estimate with $\beta = \frac{\alpha \ln \alpha}{\alpha - 1}$.

In the most general setting of single parameter agents, given the existence of a profit extractor for a benchmark \mathcal{G} , these lemmas can be combined with the *consensus estimate profit extraction* auction (CEPE) described above, to give the following theorem (see Exercise 13.11).

Theorem 13.49 Given a monotone profit benchmark \mathcal{G} for a single-parameter agent problem specified by cost function $c(\cdot)$ and a monotone profit extractor ProfitExtract $_{\mathcal{G},R}$, CEPE is $\frac{\alpha \ln \alpha}{\alpha - 1}$ -competitive and truthful with probability $1 - \log_\alpha \rho$ on inputs \mathbf{b} satisfying $\mathcal{G}(\mathbf{b}_{-i}) \in [\mathcal{G}(\mathbf{b})/\rho, \mathcal{G}(\mathbf{b})]$.

13.5 Frugality

We now turn to a radically different class of problems, in which the auctioneer is a *buyer* intent on hiring a team of agents to perform a complex task. In this model, each agent i

² What we need here is that the price offered to bidder i by ProfitExtract $_{\mathcal{G},R}$ is monotone in R , that $\mathcal{G}(\mathbf{b})$ is monotone in \mathbf{b} , and that $r(V)$ is monotone in V . See Exercise 13.11.

can perform a simple task at some cost $-v_i$ known only to himself. Based on the agents' bids b_i , the auctioneer must select a *feasible set* – a set of agents whose combined skills are sufficient to perform the complex task ($x_i = 1$ if agent i is selected) – and pay each selected agent some amount $-p_i$ (this is negative because we previously defined p_i as a transfer from the agent to the auctioneer). The setting is thus defined by the *set system* of feasible sets (E, \mathcal{S}) , where E represents the set of agents and \mathcal{S} represents the collection of feasible subsets of E . In terms of our single parameter framework, we have $c(\mathbf{x}) = 0$ if $\{i \mid x_i = 1\} \in \mathcal{S}$, and ∞ otherwise. Several special cases have received a great deal of attention.

Example 13.50 (path auctions) Here the agents own edges of a known directed graph (i.e., E is the set of edges) and the auctioneer wishes to purchase a path between two given nodes s and t (i.e., \mathcal{S} is the set of all s - t paths).

Example 13.51 (spanning tree auctions) Here the agents own edges of a known connected, undirected graph, so again E is the set of edges, and the auctioneer wishes to purchase a spanning tree.

Whereas when the auctioneer was a seller, our goal was to design a mechanism to maximize his profit, here our goal is to design a mechanism to minimize the payments the auctioneer makes, i.e., to hire the team of agents as cheaply as possible. Hence, analyzing the *frugality* of a mechanism – the amount by which it overpays – becomes an important aspect of mechanism design, analogous to profit maximization. We study frugality here using worst-case competitive analysis, as in Section 13.4.

A first observation is that here, unlike the digital goods auctions we focused on in the previous sections, the auctioneer is interested only in a single “object,” a feasible set. Thus, at a very high level, these problems are closest in spirit to the single item auction that we discussed in the context of profit maximization. For single-item auctions, in the absence of any prior information about agent's valuations, it is possible to show that the Vickrey auction is optimal, and, of course, achieves a profit equal to the value of the second highest bidder. Thus, a natural first mechanism to consider for hiring-a-team auctions is the VCG mechanism.

Consider a path auction where the graph consists of n parallel edges from s to t . This corresponds exactly to the case where the auctioneer is buying a single item, and the Vickrey mechanism will result in a payment equal to the cost of the second cheapest edge. Compare this to what happens in a graph consisting of two vertex disjoint s - t paths P and P' , each with n edges. Suppose that each edge on path P has cost zero, and each edge on path P' has cost one, so that the total cost of path P is zero and of path P' is n . Then the VCG mechanism will purchase path P , and each edge on that path will be paid n , for a total auctioneer payment of n^2 . Thus, here the VCG mechanism pays much more than the cost of the second cheapest path. Can we do better? How, in general, does the optimal truthful mechanism (in terms of competitive ratio) depend on the combinatorial structure of the set system?

13.5.1 Competitive Framework

As with our worst-case bounds from the previous section, the first issue that must be addressed to study frugality is the competitive framework and in particular the benchmark for comparison, which in this case is a *cost benchmark*.

We would like the frugality ratio to capture the overpayment of a mechanism with respect to a “natural” lower bound. One natural choice for this lower bound is the minimum payment by a nontruthful mechanism, in which case, the frugality ratio would characterize the cost of insisting on truthfulness.

Consider the mechanism \mathcal{N} which, given the bids \mathbf{b} , selects the cheapest feasible set with respect to these bids, and pays each winning agent his bid (ties are broken in favor of the efficient allocation). This mechanism is a *pay-your-bid auction* and is not truthful. However, it does have at least one (full information) pure *Nash equilibrium*, i.e., a bid vector \mathbf{b} such that, for each agent i , given the bids \mathbf{b}_{-i} by all other agents, i maximizes his profit by bidding b_i . A Nash equilibrium can be considered a natural outcome of the mechanism \mathcal{N} , and the resulting net payments are thus a good cost benchmark. As we are interested in a lower bound, we define the *cheapest Nash value* $\mathcal{N}(\mathbf{v})$ to be the minimum payments by \mathcal{N} over all of its Nash equilibria.³

To illustrate this definition, consider the case of an *s-t* path auction in which there are k parallel paths, as in our $k = 2$ path example above. Then, $\mathcal{N}(\mathbf{v})$ is precisely the cost of the second-cheapest path – the agents on the cheapest path will raise their bids until the sum of their bids equals the cost of the second-cheapest path, at which point they can no longer raise their bids. None of the other edges have incentive to raise their bids (as they are losing either way), nor to lower their bids, as they would incur a negative profit. Thus, the metric in this case makes perfect sense – it is the cost of the second cheapest solution disjoint from the actual cheapest.

With a cost benchmark in hand, we can now formalize a competitive framework for these problems.

Definition 13.52 The *frugality ratio* of truthful mechanism \mathcal{M} for buying a feasible set in set system (E, \mathcal{F}) is

$$\sup_{\mathbf{v}} \frac{\mathcal{M}(\mathbf{v})}{\mathcal{N}(\mathbf{v})},$$

where $\mathcal{M}(\mathbf{v})$ denotes the total payments of \mathcal{M} when the actual private values are \mathbf{v} , and $\mathcal{N}(\mathbf{v})$ is the cost benchmark, the cheapest Nash value with respect to the true values \mathbf{v} .

13.5.1.1 Bounds on the Frugality Ratio

The example we saw earlier shows that the VCG mechanism does not, in general, have small frugality ratio. There is, however, one class of set systems for which VCG is

³ Here we consider only Nash equilibria where nonwinners bid their true value, and ties are broken according to efficiency. We refer the reader to the relevant references for a justification of this restriction.

known to have optimal frugality ratio equal to 1, and is given in the following theorem (see Exercise 13.12).

Theorem 13.53 *VCG has frugality ratio one if and only if the feasible sets of the set system are the bases of a matroid.*

On the other hand, for path auctions, say when there are two parallel paths, each consisting of many agents, VCG can have frugality ratio $\Omega(n)$. The following lower bound shows that this bad case is not unique to the VCG mechanism.

Theorem 13.54 *Consider the path auction problem on a graph G consisting of two vertex disjoint s - t paths, P and P' , where $|P| = n$, ($|P|$ is the number of edges on the path P), and $|P'| = n'$. Then any truthful mechanism for buying a path in this graph has frugality ratio at least $\Omega(\sqrt{nn'})$.*

PROOF Define $\mathbf{v}^{(P,i)}$ to be the vector of private values for agents in P , in which edge i on P has cost $1/\sqrt{n}$ (so its value is $v_i = -1/\sqrt{n}$), and all the rest of the edges in P have cost zero. Similarly, let $\mathbf{v}^{(P',j)}$ be the vector of private values for agents in P' in which edge j on P' has cost $1/\sqrt{n'}$ and all the rest of the edges have cost zero. Let \mathcal{M} be an arbitrary deterministic truthful path auction applied to this graph. Define a bipartite graph G' with a node for each edge in G and directed edges defined as follows: there is an edge from node i (corresponding to edge i in P) to node j (corresponding to edge j in P') (respectively an edge from j to i), if when running \mathcal{M} on bid vector $(\mathbf{v}^{(P,i)}, \mathbf{v}^{(P',j)})$ path P' wins (resp. P wins).

Since there are nn' directed edges in this graph, there must be either a node i in P with at least $n'/2$ outgoing edges or a node j in P' with at least $n/2$ outgoing edges. In the former case, observe that, by the monotonicity of any truthful mechanism, P' must still win even if all edges in P' bid 0, and the payments to each of the relevant edges equal their threshold bid which is at least $1/\sqrt{n'}$. Thus the total payments are at least $\sqrt{n'}/2$. Since in this case the cheapest Nash equilibrium is $1/\sqrt{n}$, we obtain the desired lower bound. The analysis for the second case proceeds mutatis mutandis. \square

The previous lower bound can be generalized to randomized mechanisms. An immediate corollary of this lower bound is that any truthful mechanism has frugality ratio n on a graph consisting of two vertex disjoint paths of length n . Thus, for this graph, VCG achieves the optimal frugality ratio.

On the other hand, if $n' = 1$, the above lower bound on the frugality ratio of any mechanism is \sqrt{n} . However, for the case of two parallel paths, one of length 1 and one of length n , VCG has a frugality ratio of n – the worst case is when the long path wins. This raises the question of whether or not there is a better truthful mechanism for this graph.

The answer to this question is “yes.” The principle is fairly simple: if a large set is chosen as the winner, each of its elements will have to be paid a certain amount (depending on the other agent’s bids). Hence to avoid overpayment, a mechanism

should – within reason – give preference to smaller sets. Thus, rather than choosing the cheapest feasible set (i.e., the social welfare maximizing allocation), one could consider weighting the cost of feasible sets by weights that capture the relative sizes of those sets compared to other sets. To obtain a near-optimal mechanism for path auctions, the precise magnitude of these weights should be chosen to balance the worst-case frugality ratio over all potential combinations of winning sets.

To illustrate this, let us return to the graph consisting of two vertex disjoint paths. We can balance the worst-case frugality ratio by choosing the path that minimizes $\sqrt{|P|}c(P)$, where $c(P)$ is the cost of the path P , i.e., $c(P) = -\sum_{i \in P} v_i$. Notice that this mechanism uses a monotone allocation rule and hence is truthful. In this case, if the two paths are P and P' , and, say P is chosen, the payments to each edge on P will be upper bounded by $\frac{\sqrt{|P'|}c(P')}{\sqrt{|P|}}$. This is because the threshold bid, and hence the payment, to an edge e on P is the largest value they could bid and still win. Thus, the total payments are

$$|P| \frac{\sqrt{|P'|}c(P')}{\sqrt{|P|}} \leq \sqrt{|P||P'|}c(P').$$

Since $c(P')$ is a lower bound on the cheapest Nash of \mathcal{N} , the ratio of payments to cheapest Nash is upper bounded by $\sqrt{|P||P'|}$. The same bound holds when P' is the selected path, resulting in a frugality ratio matching the lower bound to within a constant factor.

These ideas can be generalized to get a mechanism whose frugality ratio is within a constant factor of optimal, for any path auction problem, as well as some other classes of “hiring-a-team” problems. For most set systems, however, the design of a truthful mechanism with optimal or near-optimal frugality ratio is open.

13.6 Conclusions and Other Research Directions

In this chapter, we have surveyed the primary techniques currently available for designing profit-maximizing (or cost-minimizing) auctions in single-parameter settings. Even in the single-parameter setting, finding mechanisms with optimal competitive ratio (for selling problems) or optimal frugality ratio (for buying problems) is challenging and largely open. The situation is much worse once we get to multiparameter problems such as various types of combinatorial auctions. In these settings, numerous new challenges arise. For example, we do not have a nice, clean, simple characterization of truthfulness. Another issue is that it is completely unclear what profit benchmarks are appropriate.

In the rest of this section, we briefly survey a number of other interesting research directions.

Profit Benchmarks. In our discussions of competitive mechanisms, we saw that the profit benchmark of a mechanism was a crucial component of the competitive approach to optimal mechanism design. This raises a fundamental issue (that has yet to be adequately resolved even in simple settings): what makes a profit benchmark the “right” one?

Pricing. In this chapter, we have discussed private-value mechanism design for profit maximization. However, even the public value versions of some of these problems, which are essentially algorithmic pricing problems, are open.

Consider, for example, the problem of pricing links in a network. We are given a graph, and a set of consumer valuations. Each valuation is given as a triple (s_i, t_i, v_i) , indicating that consumer i wishes to traverse a path from s_i to t_i and his value for traversing this path (i.e., the maximum price he is willing to pay) is v_i . With no restriction on pricing, the profit-maximizing solution to the public value problem is trivial: charge each consumer his value. However, such a pricing scheme is unreasonable for many reasons, the foremost of which is that this pricing scheme is highly unfair – different customers can get exactly the same product at different prices. An alternative pricing question is the following: define a set of prices for the edges in the graph (think of them as tolls) so as to maximize the total revenue collected. The model is that, for each consumer i , the network will collect the cost of the cheapest path from s_i to t_i with respect to the edge prices set, if that cost is at most v_i . This is just one example of an interesting algorithmic pricing problem that has recently received some attention. The vast majority of interesting combinatorial pricing problems are not well understood.

Derandomization. As we have seen, randomization is a very important tool in the design of competitive auctions. For example, randomization was used in digital goods auctions to skirt around impossibility results for deterministic symmetric auctions. Recently, however, deterministic constant competitive asymmetric digital goods auctions have been discovered. It is an interesting direction for future research to understand the general conditions under which one can derandomize competitive auctions, or design deterministic auctions from scratch. Unfortunately, standard algorithmic derandomization techniques do not work in truthful mechanism design because running the mechanism with the many possible outcomes of a randomized decision making procedure is no longer truthful. Thus, significant new ideas are required.

Fairness. We have focused our attention here on a single goal: profit maximization. In some situations, we desire that the mechanisms we design have other properties. For example, the randomized digital goods auctions that we have seen are not terribly fair – when we run, say, RSOP, some bidders pay a higher price than other bidders, and some bidders lose even though their value is higher than the price paid by other winning bidders. We say that outcomes of this type are not *envy-free*. (An auction is envy-free if after the auction is run, no bidder would be happier with someone else's outcome.)

It turns out that it is not possible to design a truthful, constant-competitive digital goods auction that is envy-free. Thus, alternative approaches have been explored for getting around this impossibility, including relaxing the solution concept to truthfulness with high probability, or allowing the mechanism to have a very small probability of producing a non-envy-free outcome.

More generally, designing auctions that both achieve high profit and are, in some sense, fair is a wide open direction for future research.

Collusion. All of the results presented in this chapter assume no collusion between the agents and indeed do not work properly in the presence of collusion. What can be done in the presence of collusion? For example, for digital goods auctions, it has been shown that it is not possible to design a truthful mechanism that is both profit-maximizing and collusion-resistant. However, using the approach of consensus estimates, it is possible to get around this impossibility with a mechanism that is truthful with high probability.

Bounded communication. How do we design revenue maximizing mechanisms when the amount of communication between the agents and the auctioneer is severely restricted? Bounded communication is particularly relevant in settings such as allocation of low-level resources in computer systems, where the overhead of implementing an auction will by necessity be severely restricted. Most of the work on this topic so far has focused on the trade-off between communication and efficiency. These results, of course, have implications for revenue maximization in a Bayesian setting due to the reduction from revenue maximization to surplus maximization via virtual valuations.

Bundling. Another interesting direction is bundling. It has been proved that in several settings, bundling items together may increase the revenue of the mechanism. However, the limits of this approach are not understood.

Repeated and online Games. Profit maximization (or cost minimization) in mechanism design arises in many settings, including resource allocation, routing and congestion control, and electronic commerce. In virtually every important practical application of mechanism design, the participants are dynamic. They arrive and depart over time, with decisions being made on an ongoing basis. Moreover, in many important applications, the same “game” is played over and over again. Our understanding of online, repeated games from the perspective of profit maximization is limited. For example, sponsored search auctions, discussed in Chapter 28, lead to many interesting open questions of this type.

Alternative solution concepts. Although truthfulness is not a goal in and of itself when the goal is profit maximization, it is a strong and appealing concept: First, truthful mechanisms obviate the need for agents to perform complex strategic calculations or gather data about their competitors. Second, in some cases, especially single-parameter problems, they simplify the design and analysis of protocols. Third, there is no loss of generality in restricting ourselves to truthful mechanisms if our plan is to implement a mechanism with dominant strategies (by the *revelation principle*). Fourth, in a number of settings, the revenue extracted by the natural truthful mechanism is the same as that extracted by natural nontruthful mechanisms (by the *revenue equivalence theorem*). A related point is that there are often natural and appealing variants of truthful mechanisms that achieve the same outcome (e.g., an English auction instead of a second-price auction). Finally, and this is important, *if we do not understand the incentive structure of a problem in a truthful setting, we are going to be very hard-pressed to understand it in any other setting.*

Having said all that, truthful mechanism design also has a number of significant drawbacks. For one thing, people often do not feel that it is safe to reveal their

information to an auctioneer. An interesting alternative is to use an *ascending auction*, where published prices can only rise over time, or an *iterative auction*, where the auction protocol repeatedly queries the different bidders, aiming to adaptively elicit enough information about the bidders' preferences to be able to find an optimal or near-optimal outcome. What is the power of ascending and iterative auctions when the auctioneer's goal is profit maximization?

Truthfulness may also needlessly limit our ability to achieve our goals. This is manifested in terms of extreme limitations on the mechanism, exceedingly high competitive ratios, or simply impossibility. In the repeated game setting, these issues are much more severe. Thus, one of the most important directions for future research is to consider alternative solution concepts.

It has been shown that taking a small step away from truthfulness, e.g., to truthfulness with high probability, can enable us to overcome some impossibility results. Other solution concepts that have received consideration in the literature include Nash equilibria, correlated equilibria, and extensions of these. However, very little work has been done concerning the design of profit-maximizing mechanisms using these solution concepts.

In summary, major directions for future research are to figure out the correct solution concepts for use in profit-maximizing auction design, and to develop techniques for designing profit-maximizing mechanisms with respect to these concepts, especially in online and repeated settings. The key desiderata of an equilibrium or solution concept are that (a) there exist mechanisms that in this equilibrium achieve or at least approximate our profit maximization goals (and whatever other goals we may have) and (b) there are simple, rational, i.e., utility-maximizing, strategies for the players that lead to outcomes in this equilibrium.⁴

13.7 Notes

Profit maximization in mechanism design has an extensive history beginning, primarily, with the seminal paper of Myerson (1981) and similar results by Riley and Samuelson (1981). These papers study Bayesian optimal mechanism design in the less restrictive setting of Bayes-Nash equilibrium. However, Myerson's optimal mechanism is precisely the optimal truthful mechanism we present here. This material is by now standard and can be found in basic texts on auction theory (Krishna, 2002; Klemperer, 1999).

The material on approximately optimal mechanism design, including the empirical Myerson mechanism and the random sampling optimal price auction comes from Baliga and Vohra (2003), Segal (2003), and Goldberg et al. (2006). Precise analysis of convergence rates for unlimited supply auction settings is given in Balcan et al. (2005).

The worst-case competitive approach to profit maximization, the proof that no symmetric, deterministic auction is competitive and the RSOP auction were first introduced in Goldberg et al. (1999), Goldberg et al. (2001), and Goldberg et al. (2006). The proof

⁴ Alternatively, we can ask that there are simple and reasonable behaviors that the players can follow that lead to outcomes in equilibrium and that the complexity of figuring out how to deviate advantageously is excessive.

of Theorem 13.36 can be found in Feige et al. (2005). The lower bound on the competitive ratio for digital goods auctions is taken from Goldberg et al. (2004). The notion of profit extraction, truthful mechanisms for reducing profit maximization to profit extraction, and the RSPE auction come from Fiat et al. (2002), Deshmukh et al. (2002), and Goldberg and Hartline (2003). The material on cost sharing that is the basis for many of the known profit extractors can be found in Moulin and Shenker (2001). The idea of consensus estimation and truthfulness with high probability come from Goldberg and Hartline (2003), Goldberg and Hartline (2003). Refinements and extensions of these results can be found in Goldberg and Hartline (2005) and Deshmukh et al. (2002). The material on frugality and path auctions is drawn from Archer and Tardos (2002), Elkind et al. (2004), and Karlin et al. (2005).

This survey focused primarily on auctions for digital goods. Further results on profit maximization (and cost minimization) in these and other settings can be found in Goldberg and Hartline (2001), Deshmukh et al. (2002), Fiat et al. (2002), Talwar (2003), Garg et al. (2002), Czumaj and Ronen (2004), Ronen and Talmann (2005), Balcan et al. (2005), Borgs et al. (2005), Hartline and McGrew (2005), Immorlica et al. (2005), Aggarwal and Hartline (2006), and Abrams (2006).

The research issues surveyed in the conclusions of this chapter are explored in a number of papers. Profit benchmarks are discussed in Goldberg et al. (2006), Deshmukh et al. (2002), Hartline and McGrew (2005), and Karlin et al. (2005); algorithmic pricing problems in Guruswami et al. (2005), Hartline and Koltrun (2005), Demaine et al. (2006), Briest and Krysta (2006), Balcan and Blum (2006), and Glynn et al. (2006); derandomization of digital goods auctions via asymmetry in Aggarwal et al. (2005); fairness in Goldberg and Hartline (2003a); collusion in Schummer (1980) and Goldberg and Hartline (2005); bounded communication in Blumrosen and Nisan (2002) and Blumrosen et al. (in press); and bundling in Palfrey (1983) and Jehiel et al. (in press). Studies of profit maximization in online auctions can be found in Bar-Yossef et al. (2002), Lavi and Nisan (2000), Blum et al. (2004), Kleinberg and Leighton (2003), Hajiaghayi et al. (2004), and Blum and Hartline (2005). Truthfulness with high probability was studied in Archer et al. (2003) and Goldberg and Hartline (2003a, 2005). Alternative solution concepts are explored in Osborne and Rubinstein (1994), Lavi and Nisan (2005), and Immorlica et al. (2005), among others.

Bibliography

- Z. Abrams. Revenue maximization when bidders have budgets. In *Proc. 17th ACM Symp. on Discrete Algorithms*, 2006.
- G. Aggarwal, A. Fiat, A. Goldberg, J. Hartline, N. Immorlica, and M. Sudan. Derandomization of Auctions. In *Proc. 37th ACM Symp. Theor. Comp.* ACM Press, New York, 2005.
- G. Aggarwal and J. Hartline. Knapsack auctions. In *Proc. 17th Annual ACM-SIAM Symp. Discrete Algorithms*, 2006.
- A. Archer, C. Papadimitriou, K. Talwar, and E. Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. In *Proc. 14th ACM Symp. on Discrete Algorithms. ACM/SIAM*, 2003.
- A. Archer and E. Tardos. Frugal path mechanisms. In *Proc. 13th ACM Symp. on Discrete Algorithms*, pp. 991–999. ACM/SIAM, 2002.

- N. Balcan and A. Blum. Approximation algorithms and online mechanisms for item pricing. In *Proc. 8th ACM Conf. on Electronic Commerce*, 2006.
- M. Balcan, A. Blum, J. Hartline, and Y. Mansour. Mechanism design via machine learning. In *Proc. 46th IEEE Symp. on Fdns. of Comp. Sci.*, 2005.
- S. Baliga and R. Vohra. Market research and market design. *Adv. Theor. Econ.*, 3, 2003.
- Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *Proc. 13th ACM Symp. on Discrete Algorithms*. ACM/SIAM, 2002.
- A. Blum and J. Hartline. Near-optimal online auctions. In *Proc. 16th ACM Symp. on Discrete Algorithms*. ACM/SIAM, 2005.
- A. Blum, V. Kumar, A. Rudra, and F. Wu. Online learning in online auctions. *Theoretical Computer Science*, 324:137–146, 2004.
- L. Blumrosen and N. Nisan. Auctions with severely bounded communication. In *Proc. 43rd IEEE Symp. on Fdns. of Computer Science*, 2002.
- L. Blumrosen, N. Nisan, and I. Segal. Auctions with severely bounded communication. *J. Artificial Intelligence Research*, in press.
- C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. In *Proc. 7th ACM Conf. on Electronic Commerce*, 2005.
- P. Briest and P. Krysta. Single-minded unlimited-supply pricing on sparse instances. In *Proc. 17th ACM Symp. on Discrete Algorithms*, 2006.
- A. Czumaj and A. Ronen. Towards generic low payment mechanisms for task allocation. In *23rd ACM SIGACT-SIGOPS Symp. on Princ. of Distributed Computing*, 2004.
- E. Demaine, U. Feige, M. Hajiaghayi, and M. Salavatipour. Combination can be hard: Approximability of the unique coverage problem. In *Proc. 17th ACM Symp. on Discrete Algorithms*, 2006.
- K. Deshmukh, A.V. Goldberg, J.D. Hartline, and A.R. Karlin. Truthful and competitive double auctions. In *Proc. 10th Europ. Symp. on Algorithms*. Springer, 2002.
- E. Elkind, A. Sahai, and K. Steiglitz. Frugality in path auctions. In *Proc. 15th ACM Symp. on Discrete Algorithms*. ACM/SIAM, 2004.
- U. Feige, A. Flaxman, J. Hartline, and R. Kleinberg. On the Competitive Ratio of the Random Sampling Auction. In *Proc. 1st Workshop on Internet and Network Economics*, 2005.
- A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Generalized competitive auctions. In *Proc. 34th ACM Symp. on Theory of Computing*. ACM Press, 2002.
- R. Garg, V. Kumar, A. Rudra, and A. Verma. Coalitional games on graphs: core structures, substitutes and frugality. Technical Report TR-02-60, UTCS, 2002.
- P.W. Glynn, P. Rusmevichientong, and B. Van Roy. A non-parametric approach to multi-product pricing. *Oper. Res.*, 54(1):82–98, 2006.
- A.V. Goldberg and J.D. Hartline. Competitive auctions for multiple digital goods. In *Proc. 9th Euro. Symp. on Algorithms*. Springer, 2001.
- A. Goldberg and J. Hartline. Envy-free auctions for digital goods. In *Proc. 5th ACM Conf. on Electronic Commerce*. ACM Press, 2003.
- A.V. Goldberg and J.D. Hartline. Competitiveness via consensus. In *Proc. 14th ACM Symp. on Discrete Algorithms*. ACM/SIAM, 2003.
- A. Goldberg and J. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *Proc. 16th ACM Symp. on Discrete Algorithms*, 2005.
- A. Goldberg, J. Hartline, A. Karlin, and M. Saks. A lower bound on the competitive ratio of truthful auctions. In *Proc. 21st Symp. on Theoretical Aspects of Computer Science*. Springer, 2004.
- A.V. Goldberg, J.D. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55:242–269, 2006.
- A.V. Goldberg, J.D. Hartline, and A. Wright. Competitive auctions and digital goods. Technical Report STAR-TR-99.09.01, STAR Laboratory, InterTrust Tech. Corp., Santa Clara, CA, 1999.

- A.V. Goldberg, J.D. Hartline, and A. Wright. Competitive auctions and digital goods. In *Proc. 12th ACM Symp. on Discrete Algorithms*, pp. 735–744. ACM/SIAM, 2001.
- V. Guruswami, J. Hartline, A. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proc. 16th ACM Symp. on Discrete Algorithms*, 2005.
- M. Hajiaghayi, D. Parkes, and R. Kleinberg. Adaptive limited-supply online auctions. In *Proc. 6th ACM Conf. on Electronic Commerce*, 2004.
- J. Hartline and V. Koltrun. Near-Optimal Pricing in Near-Linear Time. In *Proc. 9th Workshop on Algorithms and Data Structures*. Springer-Verlag, 2005.
- J. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. In *Proc. 7th ACM Conf. Electronic Commerce*, pp. 175–182, 2005.
- N. Immorlica, D. Karger, E. Nikolova, and R. Sami. First-price path auctions. In *Proc. 7th ACM Conf. on Electronic Commerce*, 2005.
- P. Jehiel, M. Meyer ter Vehn, and B. Moldovanu. Mixed bundling auctions. *J. Econ. Theory*, in press.
- A. Karlin, D. Kempe, and T. Tamir. Beyond vcg: Frugality in truthful mechanisms. In *Proc. 46th IEEE Symp. on Fdns. of Comp. Sci.*, 2005.
- R. Kleinberg and T. Leighton. The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. In *Proc. 44th IEEE Symp. on Fdns. of Comp. Sci.*, 2003.
- P. Klemperer. Auction theory: A guide to the literature. *J. Econ. Surveys*, 13:227–286, 1999.
- V. Krishna. *Auction Theory*. Academic Press, 2002.
- R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *The 2nd ACM Conf. Electronic Commerce*, pp. 233–241, 2000.
- R. Lavi and N. Nisan. Online ascending auctions for gradually expiring goods. In *SODA 05*, 2005.
- H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *J. Econ. Theory*, 18:511–533, 2001.
- R. Myerson. Optimal auction design. *Math. Operat. Res.*, 6:58–73, 1981.
- M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, MA, 1994.
- T.R. Palfrey. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica*, 51:463–484, 1983.
- J. Riley and W. Samuelson. Optimal auctions. *Amer. Econ. Rev.*, 71:381–92, 1981.
- A. Ronen and R. Tullisman. Towards generic low payment mechanisms for decentralized task allocation – A learning based approach. In *Proc. 7th Intl. IEEE Conf. E-Commerce Technology*, 2005.
- J. Schummer. Almost dominant strategy implementation. *Games Econ. Behav.*, 48:154–170, 1980.
- I. Segal. Optimal pricing mechanisms with unknown demand. *Amer. Econ. Rev.*, 93, 2003.
- K. Talwar. The price of truth: Frugality in truthful mechanisms. In *Proc. 20th Annual Symp. on Theoretical Aspects of Computer Science*, 2003.

Exercises

- 13.1** What is the optimal Bayesian single-item auction when the seller values the item at $v_0 > 0$ and bidder valuations are i.i.d?
- 13.2** What is the optimal Bayesian auction for a seller with k identical items and $n > k$ bidders with i.i.d. valuations drawn uniformly from $[0, 1]$?
- 13.3** Consider a discrete setting where bidder i 's probability of having valuation v_{ij} is f_{ij} . Derive the virtual valuations in this setting.
- 13.4** Show that the empirical Myerson mechanism, EM, applied to a single-item auction problem is identically the Vickrey auction.

- 13.5** The McDiarmid inequality is the following. Let Y_1, \dots, Y_n be independent random variables taking on values from a set A and $t : A^n \rightarrow \mathbb{R}$ a function satisfying

$$\sup_{\mathbf{y} \in A^n, y'_i \in A} |t(\mathbf{y}) - t(y'_i, \mathbf{y}_{-i})| \leq c_i$$

for all i . Then for all $\gamma \geq 0$ we have:

$$\Pr[|t(Y_1, \dots, Y_n) - \mathbf{E}[t(Y_1, \dots, Y_n)]| \geq \gamma] \leq 2e^{2\gamma^2 / \sum_{i=1}^n c_i^2}.$$

Prove Lemma 13.29 using the McDiarmid inequality.

- 13.6** Given a set of prices \mathcal{Q} and bids \mathbf{b} we say $\mathcal{Q}' \subset \mathcal{Q}$ is a γ -cover of \mathcal{Q} on \mathbf{b} if for all $q \in \mathcal{Q}$ there exists q' in \mathcal{Q}' such that

$$\sum_i |q(b_i) - q'(b_i)| \leq \gamma \text{OPT}_{\mathcal{Q}}(\mathbf{b}).$$

- (a) Prove that if \mathcal{Q}' is a γ -cover of \mathcal{Q} and all $q' \in \mathcal{Q}'$ are ϵ -good then all $q \in \mathcal{Q}$ are $(\epsilon + \gamma)$ -good.
 - (b) Show that RSOP $_{\mathcal{Q}}$ on input \mathbf{b} such that \mathcal{Q}' is a δ -cover of \mathcal{Q} is a $(1 - \epsilon - \gamma)$ -approximation with probability $(1 - \delta)$ when $\text{OPT}_{\mathcal{Q}}(\mathbf{b}) \geq \frac{8h}{\epsilon^2} \ln(\frac{2|\mathcal{Q}'|}{h})$.
 - (c) For any \mathbf{b} with $b_i \in [1, h]$, find a γ -cover of $\mathcal{Q} = \mathbb{R}$ of size $O(\frac{1}{\gamma} \log \log hn)$.
- 13.7** Give a deterministic asymmetric auction that is a 2-approximation to the optimal single price sale, $\text{OPT}_{\{1, h\}}(\mathbf{b})$, when \mathbf{b} satisfies $b_i \in \{1, h\}$ for all i and at least two bids have value h .
- 13.8** Prove that no truthful digital goods auction with 2 bidders is best. In other words, show that for any truthful auction \mathcal{A} , there is another auction \mathcal{A}' and an input \mathbf{v} such that the profit of \mathcal{A}' on input \mathbf{v} is higher than that of \mathcal{A} .
- 13.9** Show how to use a β -competitive digital goods auction (against benchmark $\mathcal{F}^{(2)}(\mathbf{v})$) to obtain a β -competitive auction for the limited supply setting where only k identical units are available for sale (use benchmark $\mathcal{F}^{(2, k)}(\mathbf{v}) = \max_{2 \leq i \leq k} i v_{(i)}$).
- 13.10** Prove the correctness of ProfitExtract $_R$ (Definition 13.40): prove that it is truthful and that it always obtains a profit of R when $\mathcal{F}(\mathbf{b}) \geq R$.
- 13.11** Given a monotone profit benchmark, \mathcal{G} ; a profit extractor ProfitExtract $_{\mathcal{G}, R}$ for \mathcal{G} that is monotone in R ; and a monotone function $r(V)$; consider the mechanism that (a) computes $R = r(\mathcal{G}(\mathbf{b}))$, and (b) runs ProfitExtract $_{\mathcal{G}, R}(\mathbf{b})$.
- (a) Prove that if $r(\mathcal{G}(\mathbf{v}_{-i})) = r(\mathcal{G}(\mathbf{v}))$ for particular bidder valuations \mathbf{v} that bidding $b_i = v_i$ is an ex-post-equilibrium, i.e., if $\mathbf{b}_{-i} = \mathbf{v}_{-i}$, then an optimal response for bidder i is to bid $b_i = v_i$.
 - (b) Prove Theorem 13.49.
- 13.12** Prove that the VCG mechanism has frugality ratio one for spanning tree auctions.