Beyond the Worst-Case Analysis of Algorithms

Edited by
Tim Roughgarden

Contents

1
Prior-Independent Auctions

Inbal Talgam-Cohen

Abstract

This chapter discusses prior-independent auctions. The goal is to design a single auction which, simultaneously for every distribution in a given class, approximates the expected revenue of the optimal auction designed specifically for that distribution. We consider two main approaches to designing such auctions: The first is *sample-based*, where the problem boils down to learning enough about the distribution (by performing "market analysis on the fly") in order to do almost as well as if the distribution were fully known. The second approach is *competition-based*, where the idea is to increase the competition in the market enough so as to drive up prices, while remaining blissfully ignorant about the distribution.

1.1 Introduction

Auctions as a meeting place for worst- and average-case analysis. Auctions are algorithms for resource allocation, with the extra complication that the input ("who values each resource by how much") comes from strategic agents who might misreport. Algorithmic research has only been studying auctions in the past two decades, spurred by their huge importance in the internet age as a main source of revenue for companies like Google. The classic theory of auctions was developed mainly within microeconomics (earning several Nobel prizes along the way).

Unlike the theory of algorithms, which has the worst-case approach at its core, in microeconomics the mainstream approach is average-case analysis. In the context of auctions, this means that buyers' values for different resources are assumed to come from known prior distributions. These priors are hard-coded into the auction in order to maximize the expected revenue from selling the resources. For example, if the known distribution tells us that a potential buyer's value for an item is likely to be high, the auction we design will charge a high price – how high exactly depends on the details of the distribution.

The assumption of distributional knowledge has all the usual downsides of the average-case approach, like being overly-brittle to noise. Of course, the worst-case approach has its own downsides like being overly-pessimistic. The premise of this chapter is that auctions offer a fascinating meeting place for the average-case economics approach and the worst-case computer science approach. In particular, auctions and other economic mechanisms are a natural testbed of what the algorithmic worst-case approach can contribute to other disciplines, e.g., in making various designs more robust or explaining the prevalence of simple designs in practice; and also of its limitations and how it can be made more suitable for practical applications by drawing closer to average-case analysis.

This chapter is organized as follows: Section 1.2 is a "crash course" in Nobel laureate Roger Myerson's theory of revenue-optimal auctions (which the familiar reader can safely skip). Section 1.3 defines prior-independence – a par excellence example of the semi-random models of Part III. Section 1.4 applies prior-independence to Myerson's theory. Section 1.5 tackles the much more challenging goal of maximizing revenue from selling several different items, to which Myerson's theory no longer applies – luckily resource augmentation (Chapter 4) comes to the rescue. Section 1.6 summarizes.

1.2 A Crash Course in Revenue-Maximizing Auctions

The basic problem. There is a single item for sale and a set of n bidders participating in an auction for the item. Every bidder i has a value $v_i \in \Re_{\geq 0}$ for winning the item, which is privately-known only to the bidder herself. This value is distributed according to a distribution F_i with positive density f_i,[1] and is reported (not necessarily truthfully) to the auctioneer as a *bid* for the item. There are usually two primary objectives when designing an auction: The first is *social welfare*, that is, the total value bidders gain from the auction (in our simple setting, if the item goes to bidder i then the welfare is v_i). The second is *revenue*, that is, the total payments the bidders transfer to the auctioneer (in our setting, what the winner of the item pays for it). An auction is carried out as follows: the auctioneer receives bids $\vec{b} = (b_1, \ldots, b_n)$ from the n bidders, applies an *allocation rule* x to the bids to decide how to allocate (in our setting, $x_i(\vec{b}) \in \{0, 1\}$ indicates whether or not bidder i is allocated the item),[2] and a *payment rule* p to decide how much to charge (where $p_i(\vec{b})$ is the payment of bidder i). An auction is thus an algorithm that gets value-bids as input, and returns an allocation and payments as output, with the goal of maximizing welfare or revenue.

Truthfulness. The "twist" in designing auctions compared to other algorithms is that the values are reported by *strategic* bidders, who will not report truthfully

[1] Extensions to discrete distributions are known.

[2] An allocation can also be randomized, in which case $x_i(\vec{b}) \in [0, 1]$ represents the probability with which bidder i is allocated the item.

unless it's in their best interest. This poses a challenge: E.g., for the goal of maximizing welfare, how you would bid in an auction where the item was given for free to the highest bidder? Indeed, you have a strong incentive to bid much higher than your true value for the item! How can the auctioneer find the bidder with the highest value if auction participants have an incentive to overbid to increase their chances of winning? An auction is (dominant-strategy[3]) *truthful* if for every i, no matter what the others bid, bidder i is (weakly) better off bidding her true value than over- or underbidding. A convenient notation for the vector of bids of all other bidders but i is b_{-i}. Using this notation, we can write bidder i's utility from bidding v_i as $v_i \cdot x_i(v_i, b_{-i}) - p_i(v_i, b_{-i})$. Truthfulness means this utility is at least as high as bidder i's utility $v_i \cdot x_i(b'_i, b_{-i}) - p_i(b'_i, b_{-i})$ if she were to bid $b'_i \neq v_i$. The goal of much of the theory of auction design is to get *truthful* auctions with good welfare or revenue guarantees. From now on we shall focus on truthful auctions and assume that bidders report their values (i.e., $b_i = v_i$ for every i); a discussion of this assumption appears in Section 1.3.1.

1.2.1 Welfare maximization: The second price and VCG auctions

The problem of designing a truthful auction that maximizes social welfare was solved by Vickrey in 1961. The high-level idea is to use the payments to align the interests of the bidders with those of society. The resulting auction for our single-item setting is very simple: Upon receiving the reported values $\vec{v}$, the allocation rule gives the item to a bidder who values it the most (i.e., bidder $i^* = \arg\max_i\{v_i\}$), and charges the second-highest bid as the item's price. This is called the *second price* auction. For example, if three bidders bid $\vec{v} = (5, 8, 3)$ for an item, then the second bidder wins and pays 5. Intuitively, neither the winner nor losers can gain by bidding higher or lower than their true value, and indeed the second price auction is truthful. Better yet, Vickrey's auction can be generalized beyond a single item to multiple items. The generalization is called the *VCG auction*, after Vickrey, Clarke and Groves. In the generalization, items are partitioned among bidders in a way that maximizes welfare.[4] Each bidder is charged her "externality" on the others (i.e., their lost welfare) caused by her participation in the auction. Notice that for a single item, the winner's externality on the second-highest bidder is precisely this bidder's loss from not being the winner herself, i.e., the second-highest value.

1.2.2 Worst-case revenue maximization

So far, the modeling assumption that v_i is drawn from a distribution F_i played no part in our account. Vickrey's second price auction allocates based only on

[3] A weaker requirement than dominant-strategy truthfulness is Bayesian truthfulness, which we return to briefly in Section 1.5.

[4] One caveat is that this allocation task is in general computationally intractable.

the reported value profile of the bidders, ignoring the distributions and always choosing the bidder with the highest value as winner. The second price auction thus maximizes welfare *pointwise*, i.e., for every realization of the random values. This ensures worst-case optimality: the auction is guaranteed to maximize welfare for every instance (value profile) of the problem.

Unfortunately, trying to apply the same approach to maximizing revenue rather than welfare is doomed to failure – there is no truthful auction that is worst-case optimal, or even approximately so, for revenue.[5] To see this consider the simplest possible setting with a single bidder interested in the item. Intuitively, all an auction can do to raise revenue is set a take-it-or-leave-it price for the item, which can't depend on the bidder's reported value (to maintain truthfulness). Setting the price to zero is clearly not worst-case optimal, and for any auction that sets a price $p > 0$, there is a worst-case instance on which this auction gets zero revenue (e.g. let the bidder's value be $v = p - \epsilon$). It thus seems to make sense to switch from the worst-case approach to the average-case one, in which the goal is to maximize *expected* revenue over the values' randomness. The average-case approach is indeed the standard one in the economics literature on revenue-maximizing auction design.

1.2.3 Average-case revenue-maximization and Myerson's theory

For a single item, Myerson solved the problem of designing a truthful auction that maximizes expected revenue in 1981. As discussed above the optimal auction *must* depend on the value distributions $F_1, \ldots, F_n$. The design thus relies on an extra assumption of full distributional knowledge.

In what follows we ignore for simplicity distributions that are "irregular" (e.g., distributions that are "too long-tailed" or bimodal; a formal definition of regularity appears below). Essentially what Myerson shows is that the dependence of the optimal auction on the distributions is very specific – they are used to transform the values into new ones called *virtual values*, by subtracting from each value a distribution-dependent "penalty" called the *information rent.* Once we have the virtual values, the allocation rule proceeds by simply maximizing welfare over these new values. In the remainder of the section we give the details of Myerson's theory.

Myerson's Lemma. Myerson's first contribution, which we refer to as Myerson's Lemma, is a *characterization* of truthful auctions for the single-item case. It turns out that truthful auctions are precisely those which allocate "monotonically in values", and charge according to a unique payment rule *whose formula depends only on the allocation rule* (provided that bidders who bid zero are not charged). By "allocating monotonically" we mean that for every bidder i and values v_{-i} of the other bidders, bidder i's allocation $x_i(b_i, v_{-i})$ is nondecreasing in her bid b_i. Intuitively, if bidding higher lowers your allocation, you will want to bid lower than

[5] One alternative suggested in the literature in order to stay in the worst-case regime is competitive analysis of online auctions.

your true value, violating truthfulness; Myerson shows that monotonicity is not only necessary but also sufficient, and that once a monotone allocation rule has been fixed, the payments are in effect fixed as well (so we never have to worry about designing payments!).

Application to a single bidder. Let's now use Myerson's characterization to find the optimal auction (pricing mechanism) for the single-bidder case. We focus for simplicity of exposition on deterministic allocation rules.[6] So a monotone allocation rule must assign 0 ("lose") to all values up to a certain threshold p, and 1 ("win") to all values above the threshold. This is equivalent to presenting the bidder with a price p for the item, and letting her decide whether or not to purchase at this price. We wish to optimize p for expected revenue given the bidder's value distribution F. The expected revenue given price p is $p(1-F(p))$, since $1-F(p)$ is precisely the probability that the bidder's value is at least p i.e. she purchases the item.

Recall we assumed F is regular. This assumption allows us to maximize the expression $p(1-F(p))$ – which we refer to as the *revenue curve in value space* – by taking the derivative $(1-F(p)) - pf(p)$ and setting it equal to zero. We get that the optimal price p^*, known as the *monopoly price* of F, is the solution to $p - \frac{1-F(p)}{f(p)} = 0$. This essentially concludes the single-bidder case. We now give another interpretation of this solution that will be helpful in the multi-bidder case.

Regularity and virtual values. Call $v - \frac{1-F(v)}{f(v)}$ the *virtual value* corresponding to value v drawn from distribution F. The information rent $\frac{1-F(v)}{f(v)}$ is subtracted from the value and so the virtual value may be negative.[7] We can now formally define *regularity* of a distribution F as the assumption that the virtual value function corresponding to F is (weakly) increasing in v. The uniform, Gaussian and exponential distributions are all examples of regular distributions. Long-tailed distributions like $F(v) = 1 - \frac{1}{\sqrt{v}}$, as well as bimodal distributions, are irregular.

Using the language of virtual values, the auction we ended up with in the single-bidder case maximizes the "virtual welfare" – i.e. the welfare with respect to the virtual value – since it allocates to the bidder iff her virtual value is ≥ 0. In the single-bidder case, we conclude that maximizing the virtual welfare is precisely what's needed in order to maximize the expected revenue. In fact, the revenue from the bidder is equal in expectation to her virtual value.

Multiple bidders and the i.i.d. assumption. Myerson shows that in the multi-bidder case, the same principle holds, and maximizing the virtual welfare (coupled with the unique payment rule coming from Myerson's Lemma) yields the optimal auction. In fact, the expected revenue of any auction is equal to the expected virtual welfare it induces by its allocation rule.

Perhaps the neatest conclusion from Myerson's theory is what auction to run when the bidder values are i.i.d., that is, all drawn independently from a single

[6] One of the conclusions from Myerson's theory is that randomization cannot help extract more expected revenue than deterministic auctions when selling a single item.

[7] This subtracted "penalty" is the inverse of what's known as the *hazard rate* of distribution F.

(regular) distribution F: In this case, we can skip the transformation to virtual values, since all values would be transformed using the same monotone transformation. We need only set a threshold such that no bidder with value below the threshold (read: virtual value below 0) can win. The optimal auction thus boils down to simply the second price auction with the monopoly price of F set as *reserve price*.[8] This is in fact a well-known auction format in practice, used for example on eBay. The i.i.d. case will play a big role in our account on prior-independence.

Take-away. An important take-away from the discussion of Myerson's theory in this section is that *the revenue-optimal auction is highly dependent on the value distributions and their knowledge*. In general, the distributional information is used to figure out the precise penalty to subtract from every value in order to get the virtual values. Even in the simple single-bidder or i.i.d. cases, the optimal reserve price closely depends on the distribution.

Beyond a single item. Unfortunately, the elegant theory of optimal auctions developed by Myerson does not extend (at least not in its simple and clean form) to settings in which buyers have different values for different items. Such settings are formally called *multi-parameter* settings, since their complexity stems from the preferences of the bidders being multi-dimensional rather than from the multiplicity of items per se.[9] In this chapter for simplicity we shall refer to such settings as *multiple item* (*multi-item*) settings, to differentiate them from the single item setting discussed so far. We address the complication introduced by multiple items (multi-dimensional values) in Section 1.5.

1.3 Defining Prior-Independence

A long-time goal of auction design has been to weaken the strong informational assumptions on which revenue-optimal auctions rely. *Robustness* with respect to distributional knowledge has been advocated across disciplines: In economics, Robert Wilson famously called for the weakening of auctions' dependence on the details of the economic environment, a position that has become known in the field as "Wilson's doctrine". In operations research, Herbert Scarf wrote in 1958 that we "have reason to suspect that the future demand will come from a distribution which differs from that governing past history in an unpredictable way." And in computer science, the discipline's general mistrust of average-case solutions (see Chapter 8) immediately extended to auction design.

But what exactly is robustness? Informally, what designers seek is auctions that have performance guarantees that are *in*sensitive to the environment details, i.e.,

[8] A reserve price is the lowest price at which the auctioneer is willing to sell.
[9] Myerson's theory does extend to all single-parameter settings, such as settings with multiple identical units of the same item for sale.

"perform well" for a "large range" of economic environments. We now formally define the robustness notion of *prior-independence*.

Definition. We focus for simplicity on the single item case with i.i.d. values. Consider first a particular distribution F from which the bidders' values are independently drawn. Let OPT_F be the optimal expected revenue that a truthful auction, which has full knowledge of F, can achieve in this environment, and let $\alpha \in (0, 1]$ be an approximation factor. An auction is *α-optimal* with respect to F if its expected revenue is at least $\alpha\mathrm{OPT}_F$ (so far this is the usual notion of approximation used in algorithms). Now let $\mathcal{F}$ be a family of value distributions, called *priors*. In particular, we shall focus on the family of distributions which satisfy the regularity property of virtual value monotonicity.

Definition 1.1 (Prior-independence) An auction is *robustly α-optimal* with respect to the family of priors $\mathcal{F}$ if for every prior $F \in \mathcal{F}$, the auction is α-optimal with respect to F.

Definition 1.1 fills with content the informal terms used above to describe robustness: for a prior-independent auction to "perform well", it must achieve expected revenue that approximates the optimal expected revenue simultaneously for every distribution in class $\mathcal{F}$; the "large range" of distributions is usually the class of all regular distributions. The definition extends naturally to multi-item settings.

1.3.1 Discussion of the definition

The definition of prior-independent robustness is an interesting mixture of average- and worst-case guarantees. On one hand, performance is measured in expectation over the random input (value profile); on the other, it is measured in the worst case over all distributions that belong to a class $\mathcal{F}$.

What are the rationales behind the definition of prior-independent robustness? In particular, why measure whether a robust auction has good performance by comparing it to the optimal auction with the "unfair" advantage of knowing the right distribution from $\mathcal{F}$? And given this measure, why take $\mathcal{F}$ to be the set of regular distributions? This can be split into two subquestions: First, why not take $\mathcal{F}$ to be even bigger? On the other hand, why allow $\mathcal{F}$ to encompass such a large range of distributions?

We begin by addressing the question of why prior-independence is a good idea as a design goal. The first reason is that existence of good prior-independence auctions according to the above definition is a powerful and useful result. An auction that performs well for all "reasonable" distributions is perfect for situations in which the seller has little to no information about the actual value distribution, such as a newcomer seller entering the market, a new item on the market, or an item whose distribution is constantly shifting. In other cases, the seller may be able to obtain some information regarding the value distribution, but at a prohibitively high cost or subject to substantial noise. Even assuming the seller somehow has reliable, affordable and up-to-date information on the value distribution, hard-wiring it into the auction can harm flexibility, since once an auction becomes the market standard it is not easy to make changes. A second important reason to set prior-independence as a design goal is that as it turns out, it often leads to auctions with a simple and natural format. Prior-independence thus gives a theoretical foundation for well-known auction formats and introduces promising new ones.

A regularity-type assumption on $\mathcal{F}$ is necessary and sufficient. The following example taken from (Dhangwatnotai et al., 2015) demonstrates the necessity of some kind of tail-restricting assumption on the class $\mathcal{F}$; regularity is a canonical example of such an assumption. Fix a number n of i.i.d. bidders. Consider a non-regular value distribution F_z, with probability $1/n^2$ for value z and probability $1 - 1/n^2$ for value zero. An auction with access to the prior distribution F_z can extract expected revenue of $\Omega(z/n)$ from the n bidders whose values are drawn from F_z. In contrast, any prior-independent truthful auction essentially has to guess the value of z, since using the winning bidder's bid violates truthfulness and all other bids are likely to be 0 (therefore providing no information about z). The conclusion is that absent a tail assumption, a prior-independent auction's expected revenue cannot be within a constant factor of z/n simultaneously for every F_z.[10]

On the other hand, the class $\mathcal{F}$ of regular distributions is not "too big": The results presented in this chapter show that the regularity assumption is sufficient to achieve a constant approximation to the ambitious benchmark of OPT_F, even in challenging environments like multi-item settings, for which optimal auctions remain elusive.

Alternative BWCA models. A natural alternative approach to robustness would be, rather than to approximate the optimal auction for every distribution in $\mathcal{F}$, to design an auction that *maximizes the minimum* expected revenue, where the minimum is taken over all distributions in $\mathcal{F}$. This approach is called *robust optimization* in the operations research literature (Bandi and Bertsimas, 2014), and has also been pursued in the context of auctions and mechanism design by economist Gabriel Carroll and others.

The prior-independence and max-min approaches are incomparable and both have already led to interesting insights. For prior-independence, notable successes in identifying natural and interesting mechanisms have been through either *approximation* or *resource augmentation.* Note that the smaller class $\mathcal{F}$ is, the easier it is to achieve a prior-independence result. For max-min, however, some of the most meaningful and interpretable results to date have been through characterizing the mechanism that achieves the *exact* maximum (where the minimum is over a judicious choice of distributions $\mathcal{F}$). Note that if $\mathcal{F}$ is too small, achieving an exact max-min result can become very challenging. For example, with multiple

[10] There are alternative tail assumptions to regularity that work too; e.g., Sivan and Syrgkanis (2013) show prior-independence results for convex combinations of a small number of regular distributions.

items, if $\mathcal{F}$ contains only a single distribution then the max-min mechanism is the revenue-optimal one for multiple items, which is known *not* to have a useful characterization. Interestingly, in either of the approaches, typically to single-out natural and interesting mechanisms, class $\mathcal{F}$ should exhibit sufficient richness.

Another alternative BWCA model to prior-independence is that of *prior-free* auctions. The prior-free approach makes no assumption that values come from an underlying (albeit unknown) distribution, and is thus fundamentally different from prior-independence. Yet there are interesting connections among the two approaches. In particular, prior-free analysis that evaluates mechanisms with respect to an *economically meaningful* benchmark will yield a prior-independent result as a corollary (see Hartline (2019a) for the definition of "economically meaningful"). Also, some techniques are relevant to both approaches, like that of randomly dividing the bidders into two and using one group as a "training set" to learn auction parameters applicable to the other (Balcan et al., 2008). This is similar in spirit to the single-sample methods in Sections 1.4.2 and 1.4.3, which use one bidder as a "training sample" to learn a reserve price applicable to the others. Of course, prior-free guarantees are more demanding than prior-independent ones (since less is assumed), and accordingly have less reach at the moment for settings such as multiple items.

A note on truthfulness. In our discussion in Section 1.2 of revenue maximization, we focused on truthful auctions. Truthfulness is an important property of auctions in itself – it simplifies participation, thus drawing more competition and leveling the playing field among sophisticated and naïve bidders. For optimal auction design it is remarkably without loss of generality due to a fundamental observation known as the *revelation principle*, by which a truthful mechanism can simulate the bidders' equilibrium strategies in a non-truthful mechanism to obtain the same outcome. Feng and Hartline (2018) note that in Bayesian settings where the agents' equilibrium strategies are a function of the prior, the Bayesian truthful mechanism (a weaker truthfulness guarantee than dominant-strategy) which is constructed via the revelation principle is *not* prior-independent. They show a gap of 1.013 between the robust approximate optimality of nontruthful mechanisms versus that of truthful mechanisms in welfare-maximization settings with budgeted bidders. Existence of a similar gap in revenue-maximization settings is left as an open question.

1.4 Sample-Based Approach: Single Item

In the era of big data, not having knowledge of a bidder's value distribution F may sound like a solvable problem – simply obtain sufficiently many samples, e.g. by interacting with other bidders of the same population (see also Chapter 29). This requires assuming the existence of a population of bidders with the same value distribution so that we can learn about the distribution of one bidder's value from the values of others. Indeed, such an assumption of i.i.d. values seems necessary to get positive prior-independence results. However, relying on large amounts of data is far from ideal in the context of auctions, as we discuss in Section 1.4.1. The main goal of this section is to minimize the required number of samples as much as possible, starting from as little as a single sample from the unknown distribution F (Sections 1.4.2 and 1.4.3). Clearly with so few samples, the empirical distribution will not in general resemble the ground truth one. This differentiates the computer science approach from previous efforts in economics, which typically rely on asymptotically large markets (e.g. Segal, 2003). We discuss a general measure of *sample complexity* in Section 1.4.4, and finish with lower bounds in Section 1.4.5.

1.4.1 How to get samples

Limitations on the number and nature of samples. As pointed out by Hartline (2019b), optimal auction design is probably most important in "thin" markets, i.e. markets in which there aren't many competitors for the item on sale, and thus insufficient past data. This could result from the nature of the item (say a unique modern art painting), or from deliberate targeting of a small set of bidders for whom the item is particularly well-suited (as done in the online ad market). In thick markets we could easily obtain many samples, but the folklore wisdom is that the auction format matters less in such markets (for example, bidders may not be able to effectively strategize regardless of the auction's truthfulness, since every bidder's action has very little influence on the outcome).

Another issue with relying on past data is that once repeated players realize that the seller is learning how to extract revenue from their bids, they have incentive to report untruthfully in order to gain in the long run. There is a growing literature on learning in the presence of strategic behavior, usually requiring behavioral assumptions on the bidders that are outside the scope of this chapter. Tang and Zeng (2018) abstract away from such assumptions, considering instead the equilibria of *distribution-reporting games* in which the distributions are endogenously reported by the bidders. They show that prior-dependent auctions are inferior to prior-independent ones when bidders are strategic about the distributions their bids reflect, a finding they dub as *price of prior-dependence*. If the setting is *not* a repeated one (imagine, e.g., tickets to a one-time event), long-run strategic behavior is less of an issue. But now there is no past data whatsoever.

Using extra bidders as samples. In light of the above discussion, we shall often assume that the samples come from the bidders themselves. That is, we randomly choose one (or more) of the bidders to excuse from the auction, in which case reporting truthfully becomes a dominant strategy for them. We then use these truthful reports as samples from F.

While providing invaluable information, throwing away bidders also loses a certain fraction of expected revenue. The following lemma shows a bound on how much expected revenue is lost:

Lemma 1.2 *In a single-item setting where the bidders have i.i.d. values, consider the optimal expected revenue OPT as a function of the number of bidders. Then* $OPT(k) \geq \frac{k}{k+\ell} OPT(k+\ell)$ *for every pair of integers* $k, \ell > 0$.

Lemma 1.2 can be interpreted as saying that the optimal expected revenue is subadditive in the number of bidders. We leave its proof as an exercise (see Exercise 1). Applications of Lemma 1.2 include showing that if we start out with a certain number of bidders, say $n = 2$, and throw out one of the bidders at random, we lose no more than 1/2 of the optimal expected revenue. It can also be applied to analyze the revenue effects of investing marketing efforts in recruiting extra bidders to the auction, then using the extras as data samples.

1.4.2 Single bidder, single sample

Recall from Section 1.2 that the optimal auction for a single bidder and known distribution F maximizes the expected revenue $p(1 - F(p))$ by offering the item at monopoly price p^*. For regular F, p^* is the value at which the corresponding virtual value becomes zero. Now assume we no longer know F but have access to a single sample $p \sim F$. A natural thing to try is to simply set the item's price to be this sample p. It turns out that this method achieves in expectation at least half of the optimal expected revenue. In this section we establish this result using a geometric proof by Dhangwatnotai et al. (2015).

Proposition 1.3 *Consider a single bidder with value drawn from F. Let F be a regular distribution and let p^* be its monopoly price. Then using a random price achieves half of the optimum in expectation:* $\mathbf{E}_{p\sim F}[p(1 - F(p))] \geq \frac{1}{2} p^*(1 - F(p^*))$.

Equivalent definition of regularity. For simplicity we assume throughout this section that F is continuous and has bounded support (the proof of Proposition 1.3 can be extended beyond these assumptions). In the proof below it will be convenient to use an alternative definition of regularity as follows.

Recall that regularity means that the virtual value function $v - \frac{1-F(v)}{f(v)}$ is (weakly) increasing in v. Recall also that in Section 1.2.3 we referred to $v(1 - F(v))$ as the "revenue curve in value space". We now switch to an equivalent, more convenient formulation for our purpose, which takes place in *quantile space*. The idea is that while the expected revenue can be represented as a function of the price p, where p ranges over all possible values, it can also be represented as a function of the *quantile* of the price, i.e., $q = 1 - F(p)$. The quantile q of p ranges between 0 and 1 and tells us *what fraction of the population would purchase at price p*. We denote by $R(q)$ the expected revenue as a function of the quantile q (also known as the *revenue curve in quantile space*). For example, $R(0.5)$ is the expected revenue if the

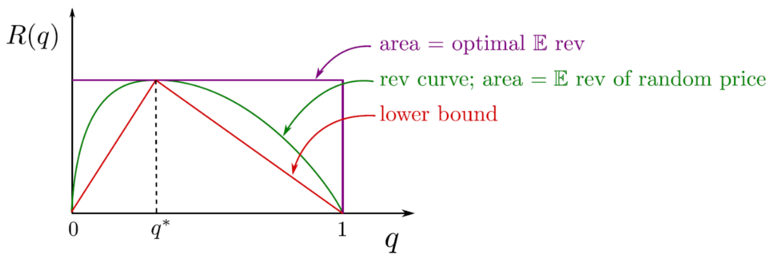

Figure 1.1 Geometric proof of Proposition 1.3.

price is set to be the value at quantile 0.5 – i.e., the median – so that a random bidder would purchase with probability 0.5.

We can now state the alternative definition of regularity: A distribution is regular if and only if its revenue curve in quantile space $R(q)$ is a concave function of q. To verify this one can check that the slope $R'(q)$ of the revenue curve at q is precisely the virtual value corresponding to the value $v = F^{-1}(1-q)$. We shall use this characterization in the following proof.

Proof of Proposition 1.3 Consider the revenue curve of F in quantile space. We can assume without loss that the lowest value in F's support is zero (this is the hardest case). Thus at the extreme quantiles 0 and 1, the expected revenue is zero (by definition nobody buys at quantile 0, and by assumption everybody buys but pays nothing at quantile 1). We plot the revenue curve in Figure 1.1. Due to regularity of F, the curve is concave. We now use this figure to visualize the two quantities we need to relate in order to prove the proposition.

The optimal expected revenue benchmark $p^*(1-F(p^*))$ can be written as $R(q^*)$, where q^* is the quantile of p^*. Geometrically, it is the area of the rectangle in the figure (since this rectangle's width is 1, and its height is $R(q^*)$). As for the expected revenue $\mathbf{E}_{p\sim F}[p(1-F(p))]$ from setting a randomly-drawn price, we can rewrite this as $\mathbf{E}_{q\sim U[0,1]}[R(q)]$. This is because randomly choosing a price according to the value distribution F is equivalent to choosing a quantile q uniformly at random and then taking the price $F^{-1}(1-q)$ that corresponds to it. Geometrically, $\mathbf{E}_{q\sim U[0,1]}[R(q)]$ is the area under the revenue curve in Figure 1.1.

To relate these two areas as required, we simply use the concavity of the revenue curve: By concavity, the triangle depicted in Figure 1.1 has area not greater than that under the curve. As its area is exactly half that of the rectangle, this completes the proof. □

The guarantee in Proposition 1.3 is tight: Consider the regular distribution F on the range $[0, H]$ where $F(v) = 1 - \frac{1}{v+1}$ for every $v \in [0, H)$ and $F(H) = 1$. As $H \to \infty$, the revenue curve in quantile space for this distribution is essentially a triangle, and so the analysis in the proof of Proposition 1.3 is tight.

1.4.3 Multiple bidders, single sample

In this section we build upon Proposition 1.3 from the previous section, in order to design a prior-independent auction for *multiple* bidders that is robustly approximately optimal. A slightly more general version of Proposition 1.3 establishes that under the same conditions and for every threshold $t \geq 0$, setting the maximum of t and a sample $p \sim F$ as the bidder's price achieves a $\frac{1}{2}$-approximation to setting the price to the maximum of t and the monopoly price p^* (Exercise 2). This can be used to establish that the following prior-independent auction, called the *single sample* auction, is robustly $\frac{n-1}{2n}$-optimal where n is the number of bidders.

Algorithm 1.4 The *single sample* auction has the following allocation rule: (1) Pick a "reserve" bidder i uniformly at random. (2) Run the second-price auction among the non-reserve bidders; let i^* be the tentative winner and let t be the second-highest bid. (3) Allocate the item to i^* if and only if $v_{i^*} \geq v_i$. The payment rule from Myerson's Lemma says that if i^* is allocated she pays $\max\{t, v_i\}$.

The single sample auction is clearly prior-independent as its allocation rule is defined with no mention of the value distributions. It is also not hard to verify its truthfulness using Myerson's Lemma. Its performance guarantee is as follows:

Theorem 1.5 *For every single-item setting with $n \geq 2$ i.i.d. bidders, the single sample auction is robustly $\frac{n-1}{2n}$-optimal with respect to regular distributions.*

Theorem 1.5 follows by first applying Lemma 1.2 to bound the loss from step (1) of the single sample auction, guaranteeing that a factor of $\frac{n-1}{n}$ of the optimal expected revenue is maintained, then applying the generalized version of Proposition 1.3 with random price $p = v_i$ and threshold t to bound the loss from steps (2) and (3) by another $\frac{1}{2}$-factor. The tightness of the guarantee in Theorem 1.5 is discussed shortly (Section 1.4.5).

We end this section by remarking that the single sample auction has also been extended to *multiple samples* – the generalized version is called the "empirical revenue maximization (ERM)" auction[11] – as well as to *multiple items* by Devanur et al. (2011) and Goldner and Karlin (2016).

1.4.4 Multiple samples: Sample complexity

The term sample complexity is borrowed from machine learning; in the context of auctions its study was initiated by Cole and Roughgarden (2014) (see also early works like (Elkind, 2007; Balcan et al., 2008), from which certain sample complexity results can be derived). The sample complexity of a family of settings measures how many samples from the prior distribution are needed in order to achieve, with

[11] ERM maximizes revenue given the empirical distribution, i.e., the uniform distribution over the samples. E.g., for a single bidder, the ERM auction sets the price to be the monopoly price of the empirical distribution, which for a single sample is simply the sample itself.

high probability, expected revenue that is close up to a given multiplicative factor (often, $1 - \epsilon$) to the target optimal expected revenue (when the distribution is known). Of course, the number of samples needed will grow with the inverse of the precision parameter ϵ; it is interesting to understand what else it depends on.[12] The sample complexity is an information theoretic measure, related but separate from the question of tractably learning from samples (this is another way in which the study of sample complexity diverges from the concrete prior-independent auctions we saw in previous sections).

There have been two main approaches in the literature to bounding the sample complexity, summarized nicely by Guo et al. (2019) (see references within). The first is to consider a class ("ϵ-net") of auctions, such that for every setting in the relevant family of settings, there always exists an approximately-optimal auction in the class. VC-like learning dimensions from statistical learning theory can then be applied to measure the complexity (or simplicity) of the auction class, telling us how many samples are needed to find the best auction in it.

An alternative approach, closer in spirit to the single sample method described in Sections 1.4.2 and 1.4.3, is to learn enough about the distributions to obtain an auction with approximately optimal expected revenue. The difference from the single sample approach is that the approximation factor is required to be very close to 1. Often the goal is to learn a small set of statistics, for instance sufficiently-many quantiles, which are relatively easy to estimate accurately via standard concentration inequalities, and sufficiently robust such that estimation errors don't harm revenue too much. For example, in the single item, single bidder case, it turns out that $\Theta(\log H)$ empirical quantiles suffice for a close-to-1 approximation when the support of the distribution is bounded by H, even without regularity.

We remark that sample complexity has also been studied for the challenging case of multiple items, largely focusing on information-theoretic rather than constructive results (for a single item, many of the most recent results are in fact constructive).

1.4.5 Lower bounds and tightness

We discuss lower bounds for the single item case, which has been the main focus of this section. There are three kinds of lower bounds, corresponding to the three types of results we have seen: (1) Robust approximation guarantees with a given number of samples (Proposition 1.3); (2) Robust approximation guarantees with a given number of i.i.d. bidders (Theorem 1.5); (3) Guarantees on the number of samples that suffice to obtain robust $(1 - \epsilon)$-optimal auctions (Section 1.4.4).

For (1), Fu et al. (2015) show that with a single sample, while no deterministic auction can do better than $\frac{1}{2}$, a randomized one can break the $\frac{1}{2}$ barrier. For (2), given $n \geq 2$ i.i.d. bidders, simply running the second-price auction with no reserve

[12] For example, with multiple non-identically distributed bidders, the number of samples required from every distributions turns out to depend polynomially on the total number of bidders.

achieves an approximation factor of $\frac{n-1}{n}$ (we further discuss this approach of running the welfare-maximizing auction with sufficient competition in Section 1.5). However the $\frac{n-1}{n}$ factor is not tight: in the special case of $n = 2$ bidders, Fu et al. (2015) show a prior-independent randomized auction with a guarantee of 0.512; Allouah and Besbes (2018) show that no auction can guarantee more than 0.556. Finally, for (3), the recent cutting-edge sample complexity results for a single item are all asymptotically tight up to polylogarithmic factors. For example, in (Guo et al., 2019) there is a lower bound of $\Omega(n\epsilon^{-3})$ samples if the prior distribution is regular, and they show a matching upper bound (up to polylogarithmic factors).

1.5 Competition-Based Approach: Multiple Items

We now shift our attention to settings with multiple different items. For such settings Myerson's theory no longer holds, and designing optimal auctions *even given full distributional knowledge* becomes a challenging task.[13]

A way to approach the multi-item challenge – one which we shall not pursue in this section – is to focus on simple, *approximately*-optimal auctions. These have been developed for many families of multiple-item markets in recent years. To make these auctions prior-independent, the "market analysis on the fly" methods from the previous section can be applied. The constant-fraction losses in expected revenue by the prior-dependent auctions are of course inherited by their prior-independent counterparts.

To avoid such losses, in this section we explore an alternative approach along the lines of *resource augmentation* (Chapter 4). Note that the model in this section is still a semi-random one, because the bidders' values are still assumed to be drawn from unknown prior distributions. What are the advantages of combining these two models?

Combining the semi-random and resource augmentation models. Recall that a main source for samples came from "throwing away" part of the demand-side of the market. Given that we are throwing away bidders who, besides being able to provide us with samples from the distributions we're interested in, also have purchasing power and willingness to pay for the items, is this the best way to use these extra bidders? Another idea would be to view extra bidders as increasing the competition for the items on the market. Intuitively, competition drives up revenue naturally, reducing the need to design careful pricing or lotteries. In fact, we can use the extra bidders to get revenue that is competitive with the optimal expected revenue benchmark – without extra bidders and with known distributions (this is the same benchmark approximated by the sample-based approach). This is achieved even while running extremely simple auctions.

[13] Indeed, even for a single bidder and two items, revenue-optimal auctions involve various complexities such as intricate menus of lotteries.

VCG. The simple auction we mostly focus on is the welfare-maximizing VCG auction. Recall from Section 1.2.1 that the VCG auction is a generalization of the second price auction to multiple items. As in the second price auction, the allocation rule of VCG partitions the items among the bidders in a way that maximizes the social welfare, in this case the bidders' aggregate value. The payment rule is not quite as simple as charging the second highest bid, but it is the natural analog for multiple items: each bidder is charged her externality (for a single item, the winner's externality on the second-highest bidder is precisely this bidder's value). The resulting VCG auction is deterministic, dominant-strategy truthful, and inherently prior-independent. And while our first-order goal is revenue, as an extra bonus by using VCG we get maximum welfare "for free".

1.5.1 Competition complexity

Warm-up: A seminal result from microeconomics. Bulow and Klemperer (1996) were the first to establish positive results using the competition-based approach sketched above. The main result in their paper is for a single item setting (they use Myerson's theory in their analysis):

Theorem 1.6 *For every single item setting with n i.i.d. bidders whose values are drawn from a regular distribution F, the optimal expected revenue is at most the expected revenue from running the second price auction with $n+1$ such bidders.*

In other words, the second price auction with a single additional bidder is robustly optimal. In the case of $n = 1$, Theorem 1.6 says that if we use the distribution-dependent monopoly price to sell the item to a single bidder, we get weakly less expected revenue than if we were to recruit a second bidder and run the second price auction. The proof of Theorem 1.6 for this case of $n = 1$ follows from Proposition 1.3, by observing that each bidder effectively faces a random price p drawn from F.

Multiple items. How far can we push the competition-based approach of Bulow and Klemperer (1996)? Consider complex settings with multiple items; is it always enough to add to the original n bidders a constant, or at least finite, number of extra bidders in order for the VCG auction to surpass the benchmark of optimal expected revenue for n bidders and known distributions? In other words, we seek Bulow-Klemperer-style results of the form: "The revenue of the welfare-maximizing auction with $n+C$ i.i.d. bidders whose values are drawn from a regular distribution F is at least as high in expectation as the optimal revenue with n such bidders." If such a statement holds for a family of auction settings, we call the smallest C for which it holds its *competition complexity* (Eden et al., 2017). Competition complexity falls under the resource augmentation umbrella, since we are comparing the performance of a simple auction with extra resources (bidders) to that of a complex auction with no extra resources.

Our model. We study the competition complexity of multi-item settings in the following model, which generalizes the i.i.d. assumption to multiple items: Consider m items, where each item $j \in [m]$ is associated with a regular distribution F_j. We assume that for every $i \in [n]$, bidder i's value $v_{i,j}$ for item j is drawn independently from F_j.

It is left to specify how to extend item values to values for sets of items. Let set function $v_i : 2^{[m]} \to \Re_{\geq 0}$ be bidder i's valuation function. We consider two cases, where v_i is either *unit-demand*, or *additive*, for every bidder i. In the former case, $v_i(S) = \max_{j \in S}\{v_{i,j}\}$, and in the latter case, $v_i(S) = \sum_{j \in S} v_{i,j}$. As an illustrative example, if the items are different desserts then a bidder on a low-sugar diet can be modeled as having a unit-demand valuation – she will enjoy no more than one dessert. A bidder who is not restricted by a diet can be modeled as additive, since she can enjoy any number of desserts. Notice that in the former case, the welfare-maximizing allocation is simply a maximum-value matching, and in the latter case the welfare-maximizing allocation simply gives every item to the bidder who values it most.

1.5.2 Unit-demand bidders

In this section we analyze the competition complexity of settings with n unit-demand bidders and m items. The first competition complexity results for multiple items were obtained by Roughgarden et al. (2019), demonstrating that the classic approach of Bulow and Klemperer (1996) is useful in more complex settings than previously realized. The benchmark used by Roughgarden et al. (2019) is the optimal *deterministic* auction, the expected revenue of which – in the unit-demand context – is up to a small constant fraction away from the optimal randomized auction. The restriction to a deterministic benchmark will be relaxed in Section 1.5.4.

Theorem 1.7 *For every setting with n unit-demand bidders and m items whose distributions $F_1, \ldots, F_m$ are regular, the optimal expected revenue by a deterministic truthful*[14] *auction is at most the expected revenue from running the VCG auction with $n + m$ such bidders.*

Theorem 1.7 extends Theorem 1.6 beyond $m = 1$, and shows that the VCG auction with m additional bidders is robustly optimal. In the remainder of this section we sketch the proof of Theorem 1.7 in three steps: upper-bounding the optimal expected revenue; lower-bounding the expected revenue of the VCG auction; and relating the two bounds.

Proof sketch. The first part of the proof relies on achieving a sufficient understanding of optimal auctions for multiple items to obtain a reasonable upper bound on their revenue. While no simple closed-form descriptions of these auctions are

[14] In this section it is important that we are focusing on the optimal *dominant-strategy* truthful auction.

known, there has been great progress on the approximation front in recent years. One useful bound was obtained for unit-demand bidders by Chawla et al. (2010):

Lemma 1.8 *The optimal expected revenue of a deterministic truthful auction with n unit-demand bidders and regular value distributions is upper-bounded by the expected revenue from selling every item j by the second price auction to $n+1$ bidders with i.i.d. values from F_j.*

The second part of the proof uses a simple bound on VCG's expected revenue, which follows from charging the winner of an item her externality on the others:

Lemma 1.9 *Let U be the set of* unallocated *bidders allocated no items by VCG. The expected revenue of the VCG auction is lower-bounded by $\sum_j \max_{i \in U}\{v_{i,j}\}$, i.e., by the sum over all items j of the highest value of a bidder in U for j.*

The third part of the proof relates the two bounds by utilizing the fact that the upper and lower bounds share a similar form. Fix an item j; the upper bound is the expected second-highest among $n+1$ values drawn independently from F_j; the lower bound is the highest among n values of unallocated bidders for item j, where bidders' values for j are drawn independently from F_j. This is where the proof uses the augmentation of the market with more bidders – in the augmented setting, only m out of $n+m$ bidders are allocated and so there are n unallocated bidders. However, a dependency issue arises: conditioned on the event that a bidder is unallocated by VCG, her value for item j is no longer distributed like a random sample from F_j. In other words, the losers in the VCG auction are likely to have lower values. Luckily, in unit-demand settings VCG allocates according to the maximum matching, and due to combinatorial properties of such matchings, the only thing that can be deduced about a losing bidder's value for item j is that it is lower than the value of item j's winner. Thus, appropriate coupling arguments can relate item j's expected contributions to the upper and lower bounds, completing the proof.

1.5.3 Lower bounds and tightness

Before we turn to competition complexity for additive bidders, let's briefly discuss lower bounds.

Unit-demand. Theorem 1.7 is tight, as adding fewer than m extra unit-demand bidders to the VCG auction may fail to guarantee the optimal expected revenue in the original environment. Consider the special case of a single unit-demand bidder ($n=1$), and m items whose values are all drawn from a point-mass distribution (and thus are identical). If no more than $m-1$ additional bidders are added to the original single bidder, each of the unit-demand bidders can get one of the m identical items, and so there is no competition to drive up the revenue. In fact, in this particular setting, the expected revenue achieved by the VCG auction with at most m (in total) unit-demand bidders is zero.

Additive. Consider again the $n = 1$ case. We now argue that the best lower bound on the extra number of additive bidders needed for the VCG auction to guarantee the optimal expected revenue is $\Omega(\log m)$ (compared to m for unit-demand).

Let the item distributions all be the regular "equal revenue" distribution F on range $[1, m^2]$, where $F(v) = 1 - \frac{1}{v}$ for every $v \in [1, m^2)$ and $F(m^2) = 1$. For sufficiently large m, the optimal expected revenue from selling the items to the original single bidder is $\Omega(m \log m)$.[15] Consider now the expected revenue of VCG with k extra bidders. For additive bidders, VCG is simply a collection of m second price auctions, one per item. Thus the expected revenue is m times the expected second price in a $(k+1)$-bidder auction with values drawn from F, which is k by direct calculation (omitted). For mk to match the benchmark $\Omega(m \log m)$, the number of extra bidders k must be $\Omega(\log m)$. This lower bound holds even if one is willing to lose an ϵ-fraction of the optimal expected revenue (Feldman et al., 2018).

1.5.4 Additive bidders

In this section we complete the picture of what's known on competition complexity by discussing settings with additive bidders and m items. For a single additive bidder, Beyhaghi and Weinberg (2019) show that the lower bound from the previous section is tight up to constant factors:

Theorem 1.10 *For every setting with a single additive bidder and m items with regular distributions $F_1, \ldots, F_m$, the optimal expected revenue by a truthful[16] auction is at most the expected revenue from running the VCG auction with $O(\log m)$ additional bidders.*

The proof follows the same general structure as that of Theorem 1.7, but the bounds and the arguments relating them are significantly more intricate.

For n additive bidders, we get an interesting dependence of the competition complexity on n: it is the minimum among $O(n \log \frac{m}{n})$ and $O(\sqrt{nm})$ (where the former is tight when $n \leq m$). Like in sample complexity, understanding which parameters factor into the competition complexity is one of the interesting insights that arise from the study of such complexity measures for auction settings.

1.6 Summary

In this chapter we surveyed prior-independent auctions, a par excellence example of a semi-random model: the objective is to maximize expected revenue, but the

[15] This is achieved by bundling together all m items, whose expected value separately is $\Omega(\log m)$ and so together is concentrated around $\Omega(m \log m)$.

[16] In this section the competition complexity results hold even for the benchmark of the optimal *Bayesian* truthful auction.

distribution over which the expectation is taken is adversarial. We started out with a single item, for which revenue-optimal auctions are well-understood. We saw what can be achieved with very little information about the distribution in the form of either a single sample or, better yet, a single additional bidder with the same value distribution. With so little information about the unknown distribution, we naturally aim for approximation results. However losing a constant approximation factor has its downsides, especially in economic context – e.g., companies usually would not be willing to settle for half of the optimal revenue.

One solution is to draw closer to the economics approach of a known distribution, by allowing the auction to learn a lot about the distribution rather than just a single sample. This can be achieved by accessing multiple samples or multiple extra bidders, and these two alternatives each have pros and cons. In particular, samples can be obtained from past data even if additional bidders cannot be recruited. But with extra bidders, the auction format becomes extremely simple and less open to strategic behavior. Perhaps the biggest advantage of the resource augmentation approach of extra bidders is that it does not lose an approximation factor compared to the benchmark; this is the only approach so far that enables such an achievement.

Open questions. We conclude with three directions for future research. The entire range between a single sample (or extra bidder) on the one hand, and an unlimited number of samples (or bidders) on the other, is interesting and relevant for prior-independence and has hardly been explored (see Babaioff et al. (2018) for a starting point). Developing competition complexity bounds for multi-item settings beyond unit-demand and additive presents new challenges (see Eden et al. (2017) for a starting point). In addition to completely eliminating distribution dependence as in prior-dependence, reducing such dependence is an important alternative that is largely open. This comes in two flavors: assuming very limited knowledge of the prior, e.g. just its mean; and allowing the auction to have a limited number of distribution-dependent parameters (see Azar et al. (2013) and Morgenstern and Roughgarden (2016), respectively, for starting points).

Chapter notes. For an excellent book chapter with an in-depth exposition of prior-independent auctions for mechanism designers, see (Hartline, 2019b). The interested reader is also referred to the PhD theses of Yan (2012) and Sivan (2013) on different aspects of prior-independence. Figure 1.1 is from Roughgarden (2017).

There are many extensions in the prior-independence literature to the results appearing in the chapter, we give here several examples: Utilizing limited, parametric knowledge on the prior distribution (Azar et al., 2013; Azar and Micali, 2013); risk-averse bidders (Fu et al., 2013); interdependent bidders (Roughgarden and Talgam-Cohen, 2016); bidders with non-identically distributed values (Fu et al., 2019); multiple items beyond additive or unit-demand bidders (Eden et al., 2017); dynamic auctions (Liu and Psomas, 2018); other objectives like makespan minimization in machine scheduling (Chawla et al., 2013); prior-independence for budgeted agents and welfare (Feng and Hartline, 2018); non-regular combinations of regular distributions (Sivan and Syrgkanis, 2013); limiting supply instead of adding bidders (Roughgarden et al., 2019).

Acknowledgments. This chapter has greatly benefited from the helpful and insightful comments of Maria-Florina Balcan, Jason Hartline, Balasubramanian Sivan, Qiqi Yan and Konstantin Zabarnyi.

References

Allouah, Amine, and Besbes, Omar. 2018. Prior-Independent Optimal Auctions. Page 503 of: *Proceedings of the 19th EC.*

Azar, Pablo Daniel, and Micali, Silvio. 2013. Parametric digital auctions. Pages 231–232 of: *Proceedings of the 4th ITCS.*

Azar, Pablo Daniel, Daskalakis, Constantinos, Micali, Silvio, and Weinberg, S. Matthew. 2013. Optimal and Efficient Parametric Auctions. Pages 596–604 of: *Proceedings of 24th SODA*.

Babaioff, Moshe, Gonczarowski, Yannai A., Mansour, Yishay, and Moran, Shay. 2018. Are Two (Samples) Really Better Than One? Page 175 of: *Proceedings of the 19th EC.*

Balcan, Maria-Florina, Blum, Avrim, Hartline, Jason D., and Mansour, Yishay. 2008. Reducing mechanism design to algorithm design via machine learning. *J. Comput. Syst. Sci.*, **74**(8), 1245–1270.

Bandi, Chaithanya, and Bertsimas, Dimitris. 2014. Optimal Design for Multi-Item Auctions: A Robust Optimization Approach. *Math. Oper. Res.*, **39**(4), 1012–1038.

Beyhaghi, Hedyeh, and Weinberg, S. Matthew. 2019. Optimal (and benchmark-optimal) competition complexity for additive buyers over independent items. Pages 686–696 of: *Proceedings of the 51st STOC.*

Bulow, Jeremy, and Klemperer, Paul. 1996. Auctions Versus Negotiations. *American Economic Review*, **86**(1), 180–194.

Chawla, Shuchi, Hartline, Jason D., Malec, David L., and Sivan, Balasubramanian. 2010. Multi-parameter mechanism design and sequential posted pricing. Pages 311–320 of: *Proceedings of the 42nd STOC.*

Chawla, Shuchi, Hartline, Jason D., Malec, David L., and Sivan, Balasubramanian. 2013. Prior-independent mechanisms for scheduling. Pages 51–60 of: *Proceedings of the 45th STOC.*

Cole, Richard, and Roughgarden, Tim. 2014. The sample complexity of revenue maximization. Pages 243–252 of: *Proceedings of the 46th STOC.*

Devanur, Nikhil R., Hartline, Jason D., Karlin, Anna R., and Nguyen, C. Thach. 2011. Prior-Independent Multi-parameter Mechanism Design. Pages 122–133 of: *Proceedings of the 7th WINE.*

Dhangwatnotai, Peerapong, Roughgarden, Tim, and Yan, Qiqi. 2015. Revenue maximization with a single sample. *Games and Economic Behavior*, **91**, 318–333.

Eden, Alon, Feldman, Michal, Friedler, Ophir, Talgam-Cohen, Inbal, and Weinberg, S. Matthew. 2017. The Competition Complexity of Auctions: A Bulow-Klemperer Result for Multi-Dimensional Bidders. Page 343 of: *Proceedings of the 18th EC.*

Elkind, Edith. 2007. Designing and learning optimal finite support auctions. Pages 736–745 of: *Proceedings of the 18th SODA*.

Feldman, Michal, Friedler, Ophir, and Rubinstein, Aviad. 2018. 99% Revenue via Enhanced Competition. Pages 443–460 of: *Proceedings of the 19th EC*.

Feng, Yiding, and Hartline, Jason D. 2018. An End-to-End Argument in Mechanism Design. Pages 404–415 of: *59th FOCS*.

Fu, Hu, Hartline, Jason D., and Hoy, Darrell. 2013. Prior-independent auctions for risk-averse agents. Pages 471–488 of: *Proceedings of the 14th EC*.

Fu, Hu, Immorlica, Nicole, Lucier, Brendan, and Strack, Philipp. 2015. Randomization Beats Second Price as a Prior-Independent Auction. Page 323 of: *Proceedings of the 16th EC*.

Fu, Hu, Liaw, Christopher, and Randhawa, Sikander. 2019. The Vickrey Auction with a Single Duplicate Bidder Approximates the Optimal Revenue. Pages 419–420 of: *Proceedings of the 20th EC*.

Goldner, Kira, and Karlin, Anna R. 2016. A Prior-Independent Revenue-Maximizing Auction for Multiple Additive Bidders. Pages 160–173 of: *Proceedings of the 12th WINE*.

Guo, Chenghao, Huang, Zhiyi, and Zhang, Xinzhi. 2019. Settling the sample complexity of single-parameter revenue maximization. Pages 662–673 of: *Proceedings of the 51st STOC*.

Hartline, Jason D. 2019a. Prior-free Mechanisms. Chap. 6 of: *Mechanism Design and Approximation*. Book draft available at http://jasonhartline.com/MDnA/.

Hartline, Jason D. 2019b. Prior-independent Approximation. Chap. 5 of: *Mechanism Design and Approximation*.

Liu, Siqi, and Psomas, Christos-Alexandros. 2018. On the Competition Complexity of Dynamic Mechanism Design. Pages 2008–2025 of: *Proceedings of the 29th SODA*.

Morgenstern, Jamie, and Roughgarden, Tim. 2016. Learning Simple Auctions. Pages 1298–1318 of: *Proceedings of the 29th COLT*.

Roughgarden, Tim. 2017. *Beyond Worst-Case Analysis*. Lecture notes available at http://timroughgarden.org/w17/w17.html.

Roughgarden, Tim, and Talgam-Cohen, Inbal. 2016. Optimal and Robust Mechanism Design with Interdependent Values. *ACM Trans. Economics and Comput.*, **4**(3), 18:1–18:34.

Roughgarden, Tim, Talgam-Cohen, Inbal, and Yan, Qiqi. 2019. *Robust Auctions for Revenue via Enhanced Competition*. To appear in Operations Research.

Segal, Ilya. 2003. Optimal pricing mechanisms with unknown demand. *American Economic Review*, **93**(3), 509–529.

Sivan, Balasubramanian. 2013. *Prior Robust Optimization*. PhD Thesis.

Sivan, Balasubramanian, and Syrgkanis, Vasilis. 2013. Vickrey Auctions for Irregular Distributions. Pages 422–435 of: *Proceedings of the 9th WINE*.

Tang, Pingzhong, and Zeng, Yulong. 2018. The Price of Prior Dependence in Auctions. Pages 485–502 of: *Proceedings of the 19th EC*.

Yan, Qiqi. 2012. *Prior Independence: A New Lens for Mechanism Design*. PhD Thesis.

Exercises

1. Prove Lemma 1.2.: In a single-item setting, let $\mathrm{OPT}(\kappa)$ be the optimal expected revenue from κ bidders with i.i.d. values. Show that for every pair of integers $k, \ell > 0$,
$$\mathrm{OPT}(k) \geq \frac{k}{k+\ell}\mathrm{OPT}(k+\ell).$$
2. Let F be a continuous, regular distribution with bounded support, and let p^* be its monopoly price. Fix a threshold $t \geq 0$. Show that
$$\mathbf{E}_{p\sim F}[\max\{t,p\}(1-F(\max\{t,p\}))] \geq \frac{1}{2}\max\{t,p^*\}(1-F(\max\{t,p^*\})).$$