

Approximately Optimal Mechanism Design: Motivation, Examples, and Lessons Learned

TIM ROUGHGARDEN
Stanford University

This survey describes the approximately optimal mechanism design paradigm and uses it to investigate two basic questions in auction theory. First, when is complexity — in the sense of detailed distributional knowledge — an essential feature of revenue-maximizing single-item auctions? Second, do combinatorial auctions require high-dimensional bid spaces to achieve good social welfare?

Categories and Subject Descriptors: F.0 [**Theory of Computation**]: General

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Mechanism design, auctions, approximation

1. INTRODUCTION

1.1 Preamble

Optimal mechanism design enjoys a beautiful and well-developed theory, and also a number of killer applications. Rules of thumb produced by the field influence everything from how governments sell wireless spectrum licenses to how the major search engines auction off online advertising.

There are, however, some basic problems for which the traditional optimal mechanism design approach is ill-suited — either because it makes overly strong assumptions, or because it advocates overly complex designs. The thesis of this paper is that *approximately optimal* mechanisms allow us to reason about fundamental questions that seem out of reach of the traditional theory.

1.2 Organization

This survey has three main parts. The first part reviews a couple of the greatest hits of optimal mechanism design, the single-item auctions of Vickrey and Myerson. We'll see how taking baby steps beyond these canonical settings already highlights limitations of the traditional optimal mechanism design paradigm, and motivates a more relaxed approach. This part also describes the approximately optimal mechanism design paradigm — how it works, and what we aim to learn by applying it.

The second and third parts of the survey cover two case studies, where we instantiate the general design paradigm to investigate two basic questions. In the

This paper is based on a talk given by the author at the 15th ACM Conference on Economics and Computation (EC), June 2014.

This work is supported in part by NSF grants CCF-1016885 and CCF-1215965, an ONR PECASE Award.

Author's address: T. Roughgarden, Computer Science Department, Department of Computer Science, 462 Gates Building, Stanford, CA 94705 USA. Email: tim@cs.stanford.edu.

first example, we consider revenue maximization in a single-item auction with heterogeneous bidders. Our goal is to understand if complexity — in the sense of detailed distributional knowledge — is an essential feature of good auctions for this problem, or alternatively if there are simpler auctions that are near-optimal. The second example considers welfare maximization with multiple items. Our goal here is similar in spirit: when is complexity — in the form of high-dimensional bid spaces — an essential feature of every auction that guarantees reasonable welfare? Are there interesting cases where low-dimensional bid spaces suffice?

2. THE OPTIMAL AND APPROXIMATELY OPTIMAL MECHANISM DESIGN PARADIGMS: VICKREY, MYERSON, AND BEYOND

2.1 Example: The Vickrey Auction

Let's briefly recall the Vickrey or second-price single-item auction [Vickrey 1961]. Consider a single seller with a single item; assume for simplicity that the seller has no value for the item. There are n bidders, and each bidder i has a valuation v_i that is unknown to the seller. Vickrey's auction is designed to maximize the welfare, which in a single-item auction just means awarding the item to the bidder with the highest valuation. This sealed-bid auction collects a bid from each bidder, awards the item to the highest bidder, and charges the second-highest price. The point of the pricing rule is to ensure that truthful bidding is a dominant strategy for every bidder. Provided every bidder follows its dominant strategy, the auction maximizes welfare ex post (that is, for every valuation profile).

In addition to being theoretically optimal, the Vickrey auction has a simple and appealing format. Plenty of real-world examples resemble the Vickrey auction. In light of this confluence of theory and practice, what else could we ask for? To foreshadow what lies ahead, we mention that when selling multiple non-identical items, the generalization of the Vickrey auction is much more complex.

2.2 Example: Myerson's Auction

What if we want to maximize the seller's revenue rather than the social welfare? Since there is no single auction that maximizes revenue ex post, the standard approach here is to maximize the expected revenue with respect to a prior distribution over bidders' valuations. So, assume bidder i 's valuation is drawn independently from a distribution F_i that is known to the seller. For the moment, assume also that bidders are homogeneous, meaning that their valuations are drawn i.i.d. from a known distribution F .

Myerson [1981] identified the optimal auction in this context, and under mild conditions on F it is a simple twist on the Vickrey auction — a second-price auction with a reserve price r .¹ Moreover, the optimal reserve price is simple and intuitive — it is just the *monopoly price* $\operatorname{argmax}_p [p \cdot (1 - F(p))]$ for the distribution F , the optimal take-it-or-leave-it offer to a single bidder with valuation drawn from F . Thus, to implement the optimal auction, you don't need to know much about the valuation distribution F — just a single statistic, its monopoly price.

¹That is, the winner is the highest bidder with bid at least r , if any. If there is a winner, it pays either the reserve price or the second-highest bid, whichever is larger.

Once again, in addition to being theoretically optimal, Myerson’s auction is simple and appealing. It is more or less equivalent to an eBay auction, where the reserve price is implemented using an opening bid. Given this success, why do we need to enrich the traditional optimal mechanism design paradigm? As we’ll see, when bidders’ valuations are not i.i.d., the theoretically optimal auction is much more complex and no longer resembles the auction formats that are common in practice.

2.3 The Optimal Mechanism Design Paradigm

Having reviewed two famous examples, let’s zoom out and be more precise about the optimal mechanism design paradigm. The first step is to identify the design space of possible mechanisms, such as the set of all sealed-bid auctions. The second step is to specify some desired properties. In this talk, we focus only on cases where the goal is to optimize some objective function that has cardinal meaning, and for which relative approximation makes sense. We have in mind objectives such as the seller’s revenue (in expectation with respect to a prior) or social welfare (ex post) in a transferable utility setting. The goal of the analyst is then to identify one or all points in the design space that possess the desired properties — for example, to characterize the mechanism that maximizes the welfare or expected revenue.

What can we hope to learn by applying this framework? The traditional answer is that by solving for the optimal mechanism, we hope to receive some guidance about how to solve the problem. With the Vickrey and Myerson auctions, we can take the theory quite literally and simply implement the mechanism advocated by the theory. More broadly, one looks for features present in the theoretically optimal mechanism that seem broadly useful — for example, Myerson’s auction suggests that combining welfare maximization with suitable reserve prices is a potent approach to revenue-maximization.

There is a second, non-traditional answer that we exploit explicitly when we extend the paradigm to accommodate approximation. Even when the theoretically optimal mechanism is not directly useful to the practitioner, for example because it is too complex, it is directly useful to the analyst. The reason is that the performance of the optimal mechanism can serve as a benchmark, a yardstick against which we measure the performance of other designs that are more plausible to implement.

2.4 The Approximately Optimal Mechanism Design Paradigm

To study approximately optimal mechanisms, we again begin with a design space and an objective function. Often the design space is limited by side constraints such as a “simplicity” constraint. For example, we later consider mechanisms with limited distributional knowledge, and those with low-dimensional bid spaces.

The new ingredient of the paradigm is a *benchmark*. This is a target objective function value that we would be ecstatic to achieve. Generally, the working hypothesis will be that no mechanism in the design space realizes the full value of the benchmark, so the goal is to get as close to it as possible. In the two examples we discuss, where the design space is limited by a simplicity constraint, a simple and natural benchmark is the performance achieved by an unconstrained, arbitrarily complex mechanism. The goal of the analyst is to identify a mechanism in the

design space that approximates the benchmark as closely as possible. For example, it is clearly interesting to establish that there is a “simple” mechanism with performance almost as good as an arbitrarily complex one.

What is the point of applying this design paradigm? The first goal is exactly the same as with the traditional optimal mechanism design paradigm. Whenever you have a principled way of selecting out one mechanism from many, you can hope that the distinguished mechanism is literally useful or highlights features that are essential to good designs. The approximation paradigm provides a novel way to identify candidate mechanisms.

There is a second reason to use the approximately optimal mechanism design paradigm, which has no analog in the traditional approach. The approximation framework enables the analyst to quantify the cost of imposing side constraints on a mechanism design space. For example, if there is a simple mechanism with performance close to that of the best arbitrarily complex mechanism, then this fact suggests that simple solutions might be good enough. Conversely, if every point in the design space is far from the benchmark, then this provides a forceful argument that complexity is an essential feature of every reasonable solution to the problem.

2.5 Two Case Studies

Sections 3 and 4 instantiate the approximately optimal mechanism design paradigm to study two fundamental questions. We first study expected revenue-maximization in single-item auctions, with bidders that have independent but not necessarily identically distributed valuations. The theoretically optimal mechanism can be complex, in the sense that it requires detailed distributional knowledge. We use the approximation paradigm to identify when such complexity is an inevitable property of every near-optimal auction.

Our second case study concerns welfare maximization. Here, the complexity stems from selling multiple non-identical items. Again, the theoretically optimal mechanism is well known but suffers from several drawbacks that preclude direct use. We apply the approximation paradigm to identify when simpler mechanisms, meaning mechanisms with low-dimensional bid spaces, can perform well, versus when complex bid spaces are necessary for non-trivial welfare guarantees.

2.6 Other Applications of the Approximation Paradigm

An enormous amount of research over the past fifteen years, largely but not entirely in the computer science literature, can be viewed as instantiations of the approximately optimal mechanism design paradigm. This paper merely singles out two recent examples that are near and dear to the author’s heart.

For example, all of the following questions have been studied through the lens of approximately optimal mechanisms.

- (1) What is the cost of imposing bounded communication in settings with very large type spaces, such as combinatorial auctions? This line of research originated in Nisan and Segal [2006] and is surveyed by Segal [2006].
- (2) What is the cost of imposing bounded computation in settings that involve computationally difficult optimization problems, such as combinatorial auctions? Two early papers are Lehmann et al. [2002] and Nisan and Ronen [2001]

and a recent survey is Nisan [2014].

- (3) What is the cost of limiting the distributional knowledge of a mechanism? Several papers in the economics literature shed light on this question [Baliga and Vohra 2003; Bulow and Klemperer 1996; Neeman 2003; Segal 2003]. The approximation interpretation is explicit in Dhangwatnotai et al. [2010]; see also Hartline [2013] for a survey. The case study in Section 3 is another example of work in this vein.
- (4) Are there auctions that achieve good revenue in the worst case (i.e., ex post)? This question was formalized using the approximately optimal mechanism design framework in Goldberg et al. [2006]; see Hartline [2013] for a recent survey.
- (5) Are there mechanisms with “simple” allocation rules that perform almost as well as arbitrarily complex mechanisms? For example, see Chawla et al. [2007] and Hartline and Roughgarden [2009] for revenue guarantees for auctions that make use only of welfare-maximization supplemented by reserve prices.
- (6) Are there mechanisms with “simple” pricing rules that perform almost as well as arbitrarily complex mechanisms? See Lucier and Borodin [2010] and Caragiannis et al. [2012] for case studies in combinatorial and keyword auctions, respectively. The case study in Section 4 is another example of this and the preceding directions.

3. CASE STUDY: DO GOOD SINGLE-ITEM AUCTIONS REQUIRE DETAILED DISTRIBUTIONAL KNOWLEDGE?

This section applies the approximately optimal mechanism design paradigm to the problem of revenue-maximization in single-item auctions. The take-away from this exercise is that the amount of distributional knowledge required for near-optimal revenue is governed by the degree of bidder heterogeneity.

3.1 Optimal Single-Item Auctions

We now return to expected revenue-maximization in single-item auctions, but allow *heterogeneous* bidders, meaning that each bidder i 's private valuation v_i is drawn independently from a distribution F_i that is known to the seller. Myerson [1981] characterized the optimal auction, as a function of the distributions F_1, \dots, F_n .

The trickiest step of Myerson's optimal auction is the first one, where each bid b_i is transformed into a *virtual bid* $\varphi_i(b_i)$, defined by

$$\varphi_i(b_i) = b_i - \frac{1 - F_i(b_i)}{f_i(b_i)}. \quad (1)$$

The exact functional form in (1) is not important for this paper, except to notice that computing $\varphi_i(b_i)$ requires knowledge of the distribution, namely of $f_i(b_i)$ and $F_i(b_i)$.

Given this transformation, the rest of the auction is straightforward. The winner is the bidder with the highest positive virtual bid (if any). To make truthful bidding a dominant strategy, the winner is charged the minimum bid at which it would continue to be the winner.²

²We have only described the optimal auction in the special case where each distribution F_i is *reg-*

When all the distributions F_i are equal to a common F , and hence all virtual valuation functions φ_i are identical, the optimal auction simplifies and is merely a second-price auction with a reserve price of $\varphi^{-1}(0)$, which turns out to be the monopoly price for F . In this special case, the optimal auction requires only modest distributional knowledge — a single statistic, the monopoly price. In general, the optimal auction does not simplify further than the description above, and detailed distributional knowledge is required to compute and compare the virtual bids of bidders with different valuation distributions.

3.2 Motivating Question

This section uses the approximately optimal mechanism design paradigm to study the following question.

Does a near-optimal single-item auction require detailed distributional knowledge?

To study this question formally, we need to parameterize the “amount of knowledge” that the seller has about the valuation distributions. We look to computational learning theory, a well-developed branch of computer science [Valiant 1984], for inspiration. We consider a seller that does not know the valuation distributions F_1, \dots, F_n , except inasmuch as it knows s valuation profiles $\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(s)}$ that have been sampled i.i.d. from these distributions. In an auction context, an obvious interpretation of these samples is as the valuations of comparable bidders in past auctions for comparable items, as inferred from bid data. See Ostrovsky and Schwarz [2009] for a real-world example of this approach, in the context of setting reserve prices in Yahoo! keyword auctions.

Thus, our design space is the set of auctions that depend on the valuation distributions only through samples. Formally, for a parameter $s \geq 1$, a point in the design space is a function from s valuation profiles (the samples) to a single-item auction, which is then run tomorrow on a fresh valuation profile drawn from $F_1 \times \dots \times F_n$. Our objective function is the expected revenue, where the expectation is over both the samples (which determines the auction used) and the final valuation profile (which determines the revenue earned by the chosen auction).

Our benchmark — the highest expected revenue we could conceivably obtain — is simply the expected revenue earned by Myerson’s optimal auction for the distributions F_1, \dots, F_n . We call this the *Myerson benchmark*. Thus, we are comparing the optimal expected revenue obtainable by a seller with partial distributional knowledge to that by a seller with full distributional knowledge. The goal is to understand the amount of knowledge (i.e., the number of samples) needed to earn expected revenue at least $(1 - \epsilon)$ times the Myerson benchmark, where ϵ is a parameter such as 0.1 or 0.01.

3.3 Formalism: One Bidder

To make sure that the formalism is clear, let’s warm up with a simple example. In addition to only one seller with one item, suppose there is also only one bidder, with

ular, meaning that the virtual valuation functions φ_i are nondecreasing. The general case “monotonizes” the virtual valuation functions — monotonicity is essential for incentive-compatibility — but otherwise applies the same three steps [Myerson 1981].




Fig. 1. Single-bidder formalism. A take-it-or-leave-it offer p is chosen as a function of s i.i.d. samples v_1, \dots, v_s from an unknown distribution F , and applied to a fresh sample v_{s+1} from the same distribution. The goal is to understand, given a parameter $\epsilon > 0$, how many samples s are needed so that there exists a price function with expected revenue at least $1 - \epsilon$ times that of the monopoly price for F .

valuation drawn from a distribution F unknown to the seller. With only one bidder, auctions are merely take-it-or-leave-it offers.³ The goal is to design a function $p(v_1, \dots, v_s)$ from samples $v_1, \dots, v_s \sim F$ to prices that, for every F , achieves expected revenue $\mathbf{E}_{v_1, \dots, v_s} [p(v_1, \dots, v_s) \cdot (1 - F(p(v_1, \dots, v_s)))]$ close to that achieved by the monopoly price $\operatorname{argmax}_p p \cdot (1 - F(p))$ of F . In other words, given data from s past transactions, the goal is to set a near-optimal price for a new bidder encountered tomorrow. See also Figure 1.

3.4 Results for a Single Bidder

We next state a series of results for the single-bidder special case. These results are not the main point of this case study, and instead serve to calibrate our expectations for what might be possible for single-item auctions with multiple bidders.

The bad news is that, without any assumptions about the unknown distribution F , no finite number of samples yields a non-trivial expected revenue guarantee for every F . That is, for every finite s , there is a valuation distribution F such that you learn essentially nothing about F from s samples.⁴ This observation motivates restrictions on the unknown distribution.

The good news is that under a standard “regularity” condition, intuitively stating that the tail of F is no heavier than a power-law distribution, is sufficient for interesting positive results.⁵ Even just *one* sample can be used to obtain a non-trivial revenue guarantee for unknown regular distributions: for every such F , the function $p(v_1) = v_1$ — using yesterday’s bid as tomorrow’s price — yields expected revenue at least 50% times that of the monopoly price.⁶

³Probability distributions over take-it-or-leave-it-offers are also allowed. We discuss only deterministic auctions for simplicity of presentation, but the results of this section also apply to randomized auctions.

⁴For example, for a parameter $M \rightarrow \infty$, consider distributions F that put a point mass of $1/M$ at M and are otherwise zero.

⁵Formally, a distribution F is *regular* if its virtual valuation distribution (1) is nondecreasing.

⁶This is a consequence of the following special case of the Bulow-Klemperer theorem on auctions vs. negotiations [Bulow and Klemperer 1996]: the expected revenue of a Vickrey auction with two




Fig. 2. Multiple-bidder formalism. A single-item auction A is chosen as a function of s i.i.d. samples $\mathbf{v}_1, \dots, \mathbf{v}_s$ from an unknown distribution $F_1 \times \dots \times F_n$, and applied to a fresh sample \mathbf{v}_{s+1} from the same distribution. The benchmark is the expected revenue of the Myerson-optimal auction for F_1, \dots, F_n .

What if we want a better revenue guarantee, like 90% or 99% of this benchmark? To achieve a $(1 - \epsilon)$ -approximation guarantee, we expect the number of samples required to increase with $1/\epsilon$. Happily, the amount of data required is relatively modest, scaling as a polynomial function of $1/\epsilon$. For an unknown regular distribution, this function is roughly ϵ^{-3} [Dhangwatnotai et al. 2010; Huang et al. 2014]. The sample complexity improves if we impose stronger conditions on the tails of the valuation distribution. For example, if F satisfies the monotone hazard rate condition — meaning $f(x)/(1 - F(x))$ is nondecreasing in x — then roughly $\epsilon^{-3/2}$ samples are necessary and sufficient to achieve a $(1 - \epsilon)$ -approximation of the benchmark [Huang et al. 2014]. The upper bounds on sample complexity follow from natural pricing strategies, such as choosing the monopoly price for the empirical distribution of the samples.

3.5 Formalism: Multiple Bidders

Generalizing the formalism to single-item auctions with multiple bidders proceeds as one would expect. The seller is now given s samples $\mathbf{v}_1, \dots, \mathbf{v}_s$, where each sample \mathbf{v}_j is a valuation profile, comprising one valuation (drawn from F_i) for each bidder i . The seller picks an auction $A(\mathbf{v}_1, \dots, \mathbf{v}_s)$ that is a function of these samples only. Recall that the Myerson benchmark is the expected revenue of the optimal auction for F_1, \dots, F_n . The goal is to design a function $A(v_1, \dots, v_s)$ from samples $\mathbf{v}_1, \dots, \mathbf{v}_s \sim F_1 \times \dots \times F_n$ to single-item auctions that, for every F_1, \dots, F_n , achieves expected revenue close to this benchmark. As in the single-bidder case, the expectation is over both the past bid data (the samples) and the bidders (a fresh sample from the same distributions). See also Figure 2.

3.6 Positive Results

The hope is that our positive results for the single-bidder problem (Section 3.4) carry over to single-item auctions with multiple bidders. First, provided F_1, \dots, F_n are regular distributions, it is still possible to get a coarse but non-trivial approx-

bidders with valuations drawn i.i.d. from a regular distribution F is at least that of an optimal auction (i.e., the monopoly price) for a single bidder with valuation drawn from F .

imation (namely, 25%) with a single sample — this follows from a generalization of the Bulow-Klemperer theorem given by Hartline and Roughgarden [2009]. But what about very close approximations, like 90% or 99%?

In the special case where bidders are homogeneous — meaning have identically distributed valuations — the positive results for a single bidder continue to hold. Intuitively, the reason is that the form of the optimal auction is independent of the number of bidders — it is simply a second-price auction with a reserve set to the monopoly price for the distribution F . Since a single statistic about the distribution F determines the optimal auction for an arbitrary number of homogeneous bidders, it makes sense that the sample complexity of approximating this optimal auction is independent of n .

Thus, in these cases, the amount of data — the granularity of knowledge about the valuation distributions — necessary to achieve near-optimal revenue is relatively modest, and does not depend on the number of bidders.

3.7 Negative Results

The approximately optimal mechanism design paradigm identifies a qualitative difference between the cases of homogeneous and heterogeneous bidders. When bidders are heterogeneous and we seek a close approximation of the optimal revenue, the sample complexity depends fundamentally on the number of bidders.

THEOREM 3.1 [COLE AND ROUGHGARDEN 2014]. *There is a constant $c > 0$ such that, for every sufficiently small $\epsilon > 0$ and every $n \geq 2$, there is no auction that depends on at most $cn/\sqrt{\epsilon}$ samples and has expected revenue at least $1 - \epsilon$ times the Myerson benchmark for every profile F_1, \dots, F_n of regular distributions.*

The valuation distributions used in the proof of Theorem 3.1 are not pathological — exponential distributions, truncated at different maximum values, already yield the lower bound.

The proof of Theorem 3.1 shows more generally that every auction that fails to implicitly learn all bidders' virtual valuation functions (recall (1)) up to small error is doomed to having expected revenue less than $1 - \epsilon$ times the Myerson benchmark in some cases. In this sense, detailed knowledge of the valuation distributions is an unavoidable feature of every near-optimal single-item auction with heterogeneous bidders.⁷

4. CASE STUDY: DO GOOD COMBINATORIAL AUCTIONS REQUIRE COMPLEX BID SPACES?

In this section we switch gears and study the problem of allocating multiple items to bidders with private valuations to maximize the social welfare. We instantiate the approximately optimal mechanism design paradigm to identify conditions on bidders' valuations that are necessary and sufficient for the existence of simple combinatorial auctions. The take-away from this section is that rich bidding

⁷There is also a converse to Theorem 3.1: for every $\epsilon > 0$ and $n \geq 1$, and for an arbitrary number of bidders with n distinct valuation distributions, a polynomial number (in n and ϵ^{-1}) of samples is sufficient to achieve a $(1 - \epsilon)$ -approximation of the Myerson benchmark [Cole and Roughgarden 2014].

spaces are an essential feature of every good combinatorial auction when items are complements, while simple auctions can perform well when bidders' valuations are complement-free.

4.1 The VCG Mechanism

We adopt the standard setup for allocating multiple items via a combinatorial auction. There are n bidders and m non-identical items. Each bidder has, in principle, a different private valuation $v_i(S)$ for each bundle S of items it might receive. Thus, each bidder has 2^m private parameters. In this section, we assume that the objective is to determine an allocation S_1, \dots, S_n that maximizes the social welfare $\sum_{i=1}^n v_i(S_i)$.

The Vickrey auction can be extended to the case of multiple items; this extension is the Vickrey-Clarke-Groves (VCG) mechanism [Vickrey 1961; Clarke 1971; Groves 1973]. The VCG mechanism is a direct-revelation mechanism, so each bidder i reports a valuation $b_i(S)$ for each bundle of items S . The mechanism then computes an allocation that maximizes welfare with respect to the reported valuations. As in the Vickrey auction, suitable payments make truthful revelation a dominant strategy for every bidder.

Even with a small number of items, the VCG mechanism is a non-starter in practice, for a number of reasons [Ausubel and Milgrom 2006]. We focus here on the first step. Every direct revelation mechanism, including the VCG mechanism, solicits 2^m numbers from each bidder. This is an exorbitant number: roughly a thousand parameters when $m = 10$, roughly a million when $m = 20$.

4.2 Motivating Question

In this case study, we apply the approximately optimal mechanism design paradigm to study the following question.

Does a near-optimal combinatorial auction require rich bidding spaces?

Thus, as in the previous case study, we seek conditions under which “simple auctions” can “perform well.” This time, our design space of “simple auctions” consists of mechanism formats in which the dimension of every player’s bid space is growing polynomially with the number m of items (say m or m^2), rather than exponentially with m as in the VCG mechanism.

“Performing well” means, as usual, achieving objective function value (here, social welfare) close to that of a benchmark. We use the *VCG benchmark*, meaning the welfare obtained by the best arbitrarily complex mechanism (the VCG mechanism), which is simply the maximum-possible social welfare.

This case study contributes to the debate about whether or not package bidding is an important feature of combinatorial auctions, a topic over which much blood and ink has been spilled over the past twenty years. We can identify auctions with no or limited package bidding with low-dimensional mechanisms, and those that support rich package bidding with high-dimensional mechanisms. With this interpretation, our results make precise the intuition that flexible package bidding is crucial when items are complements, but not otherwise.

4.3 A Simple Auction: Selling Items Separately

Our goal is to understand the power and limitations of the entire design space of low-dimensional mechanisms. To make this goal more concrete, we begin by examining a specific simple auction format.

The simplest way of selling multiple items is by selling each separately. Several specific auction formats implement this general idea. We analyze one such format, simultaneous first-price auctions [Bikhchandani 1999]. In this auction, each bidder submits simultaneously one bid per item — only m bidding parameters, compared with its 2^m private parameters — and each item is sold in parallel using a first-price auction.

When do we expect simultaneous first-price auctions to have reasonable welfare at equilibrium? Not always. With general bidder valuations, and in particular when items are complements, we might expect severe inefficiency due to the “exposure problem” (e.g., [Milgrom 2004]). For example, consider a bidder in an auction for wireless spectrum licenses that has large value for full coverage of California but no value for partial coverage. When items are sold separately, such a bidder has no vocabulary to articulate its preferences, and runs the risk of obtaining a subset of items for which it has no value, at a significant price.

Even when there are no complementarities amongst the items, we expect inefficiency when items are sold separately (e.g., [Krishna 2010]). The first reason is “demand reduction,” where a bidder pursues fewer items than it truly wants, in order to obtain them at a cheaper price. Second, if bidders’ valuations are drawn independently from different valuation distributions, then even with a single item, Bayes-Nash equilibria are not always fully efficient.

4.4 Valuation Classes

Our discussion so far suggests that simultaneous first-price auctions are unlikely to work well with general valuations, and suffer from some degree of inefficiency even with simple bidder valuations. To parameterize the performance of this auction format, we introduce a hierarchy of bidder valuations (Figure 3); the literature also considers more fine-grained hierarchies [Feldman et al. 2014; Lehmann et al. 2006].

The biggest set corresponds to general valuations, which can encode complementarities among items. The other three sets denote different notions of “complement-free” valuations. In this survey, we focus on the most permissive of these, *subadditive valuations*. Such a valuation v_i is monotone ($v_i(T) \subseteq v_i(S)$ whenever $T \subseteq S$) and satisfies $v_i(S \cup T) \leq v_i(S) + v_i(T)$ for every pair S, T of bundles. This class is significantly larger than the well-studied classes of gross substitutes and submodular valuations.⁸ In particular, subadditive valuations can have “hidden complements” — meaning two items become complementary given that a third item has already been acquired — while submodular valuations cannot [Lehmann et al. 2006].

⁸Submodularity is the set-theoretic analog of “diminishing returns:” $v_i(S \cup \{j\}) - v_i(S) \leq v_i(T \cup \{j\}) - v_i(T)$ whenever $T \subseteq S$ and $j \notin S$. The gross substitutes condition — which states that a bidder’s demand for an item only increases as the prices of other items rise — is strictly stronger and guarantees the existence of Walrasian equilibria [Kelso, Jr. and Crawford 1982; Gul and Stacchetti 1999].

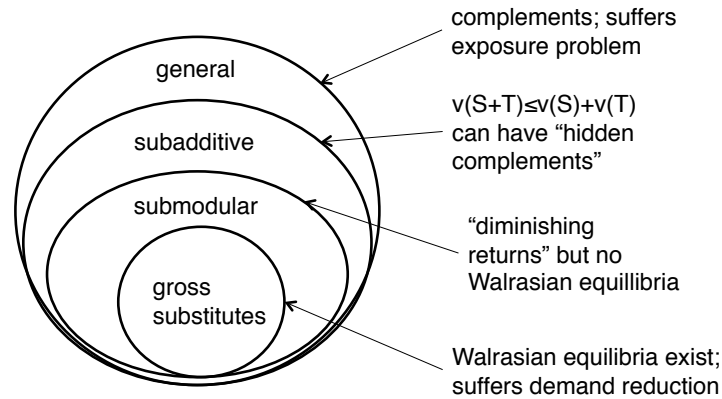


Fig. 3. A hierarchy of valuation classes.

4.5 When Do Simultaneous First-Price Auctions Work Well?

Our intuition about the performance of simultaneous first-price auctions translates nicely into rigorous statements. First, for general valuations, selling items separately can be a disaster.

THEOREM 4.1 [HASSIDIM ET AL. 2011]. *With general bidder valuations, simultaneous first-price auctions can have mixed-strategy Nash equilibria with expected welfare arbitrarily smaller than the VCG benchmark.*

For example, equilibria of simultaneous first-price auctions need not obtain even 1% of the maximum-possible welfare when there are complementarities between many items.

On the positive side, even for the most permissive notion of complement-free valuations — subadditive valuations — simultaneous first-price auctions suffer only bounded welfare loss.

THEOREM 4.2 [FELDMAN ET AL. 2013]. *If every bidder's valuation is drawn independently from a distribution over subadditive valuations, then the expected welfare obtained at every Bayes-Nash equilibrium of simultaneous first-price auctions is at least 50% of the expected VCG benchmark value.*

In Theorem 4.2, the valuation distributions of different bidders do not have to be identical, just independent. The guarantee improves to roughly 63% for the special case of submodular bidder valuations [Syrkkanis and Tardos 2013].

Taken together, Theorems 4.1 and 4.2 suggest that simultaneous first-price auctions should work reasonably well if and only if there are no complementarities among items.

4.6 Digression on Approximation Ratios

Before proceeding to our final set of technical results, we pause to emphasize how worst-case approximation results like Theorems 4.1 and 4.2 should be interpreted.

Many researchers have a tendency to fixate unduly on and take too literally such approximation guarantees.

Both of the primary motivations for applying the approximately optimal mechanism design paradigm strive for qualitative insights, not fine-grained performance predictions (recall Section 2.4). The first goal is to identify mechanisms or mechanism features that are potentially useful in practice. The auction formats implicitly recommended by our case studies, such as selling items separately with first-price auctions provided bidders' valuations are sufficiently simple, corroborate well with folklore beliefs. The second goal of the approximation paradigm is to quantify the cost of a side constraint like “simplicity” on the mechanism design space. In our case studies, we are coarsely classifying such constraints as “tolerable” or “intolerable” according to whether or not imposing the constraint reduces the achievable performance by a modest constant factor. This viewpoint leads to interesting and sensible conclusions in both of our case studies: complexity is unavoidable in near-optimal revenue-maximizing single-item auctions if and only if bidders are heterogeneous, and complexity is unavoidable in near-optimal welfare-maximizing auctions for selling multiple items if and only if there are complementarities among the items.

To the reader who insists on interpreting approximation guarantees literally, against our advice, we offer a few observations. First, in most applications of the approximately optimal mechanism design framework, the benchmark is constructed so that there is no mechanism in the design space that always achieves 100% of the benchmark. When 100% is unachievable, the best-possible approximation is going to be some number bounded below 100% — it cannot be arbitrarily close to 100% when nothing is tending to infinity.⁹ Examples that demonstrate mechanism sub-optimality are often “small” in some sense, which translates to impossibility results for worst-case approximation guarantees better than relatively modest fractions like 50% or, if you're lucky, 75%. Finally, remember that the benchmark being approximated — for example, the performance of a mechanism so complex as to be unrealizable — is generally not an option on the table. The benchmark represents a utopia that exists only in the analyst's mind — like your favorite baseball team winning 162 games, or receiving referee reports on your journal submission in less than six months.

Of course, like any general analysis framework, the approximation paradigm can be abused and should be applied with good taste. In settings where the approximately optimal mechanism design paradigm does not give meaningful results, the approach should be modified — by defining a different benchmark, changing the notion of benchmark approximation, or using a completely different analysis framework.

⁹In some cases it makes sense to speak of the asymptotic optimality of a mechanism, such as the sample complexity results of Section 3 and in large markets (e.g., [Segal 2003; Swinkels 2001]). Asymptotic results are clearly interesting, but are applicable to only a fraction of the problems that we want to reason about.

4.7 Negative Results

We now return to the question of when simple mechanisms, meaning mechanisms with low-dimensional bid spaces, can achieve non-trivial welfare guarantees. Section 4.5 considered the special case of simultaneous first-price auctions; here we consider the full design space.

First, the poor performance of simultaneous first-price auctions with general bidder valuations is not an artifact of the specific format: *every* simple mechanism is vulnerable to arbitrarily large welfare loss when there are complementarities among items. This impossibility result argues forcefully for a rich bidding language, such as flexible package bidding, in such environments.

THEOREM 4.3 [ROUGHGARDEN 2014]. *With general bidder valuations, no family of simple mechanisms guarantees equilibrium welfare at least a constant fraction of the VCG benchmark.*

In Theorem 4.3, the mechanism family is parameterized by the number of items m ; “simple” means that the number of dimensions in each bidder’s bid space is bounded above by some polynomial function of m . The theorem states that for every such family and constant $c > 0$, for all sufficiently large m , there is a valuation profile and a full-information mixed Nash equilibrium of the mechanism with expected welfare less than c times the maximum possible.¹⁰

We already know from Theorem 4.2 that, in contrast, simple auctions can have non-trivial welfare guarantees with complement-free bidder valuations. Our final result states that no simple mechanism outperforms simultaneous first-price auctions with these bidder valuations.

THEOREM 4.4 [ROUGHGARDEN 2014]. *With subadditive bidder valuations, no family of simple mechanisms guarantees equilibrium welfare more than 50% of the VCG benchmark.*

5. CONCLUSIONS

5.1 Motivating Questions Revisited

To close the circle, we return to the motivating questions of our case studies and review the answers provided by the approximately optimal mechanism design paradigm. The first question was:

Does a near-optimal single-item auction require detailed distributional knowledge?

To answer this question, we took the design space to be auctions with limited knowledge of the valuation distributions — in the form of s i.i.d. samples — and studied the number of samples necessary and sufficient to achieve a $(1 - \epsilon)$ -approximation of the Myerson benchmark. We discovered that the amount of knowledge (i.e., samples) required scales linearly with the number of distinct valuation distributions represented in the bidder population. Thus, detailed distributional knowledge is

¹⁰Technically, Theorem 4.3 proves this statement for an ϵ -approximate Nash equilibrium — meaning every player mixes only over strategies with expected utility within ϵ of a best response — where $\epsilon > 0$ can be made arbitrarily small. The same comment applies to Theorem 4.4.

required for near-optimal revenue maximization if and only if the bidders are heterogeneous.

The second motivating question was:

Does a near-optimal combinatorial auction require rich bidding spaces?

Here, we defined the design space to be families of mechanisms for which the number of parameters in a bidder’s bid space grows polynomially with the number m of items. We adopted the VCG benchmark, which equals the maximum-possible social welfare. We discovered that high-dimensional bid spaces are fundamental to non-trivial welfare guarantees when there are complementarities among items, but not otherwise. We also learned that, in some cases, selling items separately with first-price auctions achieves the best-possible worst-case approximation guarantee of any family of simple mechanisms.

5.2 Further Discussion

We showed how the approximately optimal mechanism design paradigm yields basic insights about two fundamental problems. Moreover, it is not clear how to glean these insights without resorting to an analysis framework that incorporates approximation. Our first case study fundamentally involved suboptimality — the less knowledge the seller has, the less revenue it can obtain. Similarly, inefficiency was an unavoidable aspect of our second case study, since simple mechanisms are suboptimal even in very simple settings (e.g., due to demand reduction or bidder asymmetry). An approximation framework is the obvious way to reason about and compare different degrees of suboptimality.

An alternative idea, given a design space and an objective function, is to simply identify the mechanism in the design space with the “best” objective function value. The fundamental issue here is how to meaningfully compare two different mechanisms, which will generally have incomparable performance. For example, for two different single-item auctions that depend on the valuation distributions F_1, \dots, F_n only through s samples (Section 3), typically either one can have higher expected revenue than the other, depending on the choice of F_1, \dots, F_n . Similarly, for two different combinatorial auctions with low-dimensional bid spaces, one generally has higher welfare for some valuation profiles, and the other for other valuation profiles. The traditional approach in mechanism design to resolving such trade-offs is to impose a prior on the unknown information and maximize expected performance with respect to the prior. But this approach would return us to the very bind we intended to escape, of uninformatively complex optimal mechanisms that require detailed distributional knowledge.

Are our insights surprising? The presented results both confirm some existing intuitions — which we view as important sanity checks for the theory — and go beyond them. For example, in single-item auctions, the result that modest data is sufficient for near-optimal revenue-maximization with homogeneous bidders is natural given that the optimal auction depends only on the valuation distribution’s monopoly price. While revenue-maximization with heterogeneous bidders can only be a more complex problem, it is not clear a priori how such complexity scales with bidder heterogeneity, or even how “complexity” should be defined. The fact that the sample complexity scales linearly with the number of distinct valuation distri-

butions is a satisfying and non-obvious formalization of the idea that “heterogeneity matters.”

For the case study of selling multiple items, the high-level take-aways of our analysis are in line with prevailing intuition — simple auctions enjoy reasonable performance when there are no complementarities among items, but not otherwise. One pleasant surprise of the analysis, which deserves further investigation, is that the positive results for simple auctions hold even for the most general notion of “complement-free valuations,” well beyond the more well-studied special cases of gross substitutes and submodular valuations.

5.3 Open Questions

This survey presented two recent applications of the approximately optimal mechanism design paradigm. There have been dozens of other applications over the past fifteen years (Section 2.6), and there is still much to do.

For example, the sample-complexity formalism of Section 3 shows promise of deepening our understanding of Bayesian-optimal mechanism design. Proving that modest distributional knowledge suffices for near-optimal mechanism performance is an important step in arguing the practical relevance of a theoretically optimal design. Upper bounds on the number of samples needed (Section 3.4) generally suggest interesting methods of incorporating data, such as past bidding data, into designs. Thus far, only the simple settings of single-item auctions (Section 3) and single-bidder multi-item mechanisms [Dughmi et al. 2014] have been studied from this perspective.

For welfare-maximization with multiple items, results like those in Section 4 give preliminary insights into which auction designs might work well, as a function of bidders’ preferences. An important direction for future work is to draw sharper distinctions between different plausibly useful formats. For example, there is ample empirical evidence that ascending auctions for multiple items perform better than their sealed-bid counterparts. Can this observation be made formal using the approximately optimal mechanism design paradigm?

REFERENCES

- AUSUBEL, L. M. AND MILGROM, P. 2006. The lovely but lonely Vickrey auction. In *Combinatorial Auctions*, P. Cramton, Y. Shoham, and R. Steinberg, Eds. MIT Press, Boston, MA, USA, Chapter 1, 57–95.
- BALIGA, S. AND VOHRA, R. 2003. Market research and market design. *Advances in Theoretical Economics* 3.
- BIKHCHANDANI, S. 1999. Auctions of heterogeneous objects. *Games and Economic Behavior* 26, 193–220.
- BULOW, J. AND KLEMPERER, P. 1996. Auctions versus negotiations. *American Economic Review* 86, 1, 180–194.
- CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., KYROPOULOU, M., LUCIER, B., LEME, R. P., AND TARDOS, É. 2012. On the efficiency of equilibria in generalized second price auctions. arXiv:1201.6429.
- CHAWLA, S., HARTLINE, J. D., AND KLEINBERG, R. D. 2007. Algorithmic pricing via virtual valuations. In *EC*. 243–251.
- CLARKE, E. H. 1971. Multipart pricing of public goods. *Public Choice* 11, 1, 17–33.
- COLE, R. AND ROUGHGARDEN, T. 2014. The sample complexity of revenue maximization. In *STOC (to appear)*. ACM.

- DHANGWATNOTAI, P., ROUGHGARDEN, T., AND YAN, Q. 2010. Revenue maximization with a single sample. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*. 129–138.
- DUGHMI, S., HAN, L., AND NISAN, N. 2014. Sampling and representation complexity of revenue maximization. Working paper.
- FELDMAN, M., FEIGE, U., IMMORLICA, N., IZSAK, R., LUCIER, B., AND SYRGGANIS, V. 2014. A unified approach to valuations with restricted complements. Working paper.
- FELDMAN, M., FU, H., GRAVIN, N., AND LUCIER, B. 2013. Simultaneous auctions are (almost) efficient. In *45th ACM Symposium on Theory of Computing (STOC)*. 201–210.
- GOLDBERG, A. V., HARTLINE, J. D., KARLIN, A., SAKS, M., AND WRIGHT, A. 2006. Competitive auctions. *Games and Economic Behavior* 55, 2, 242–269.
- GROVES, T. 1973. Incentives in teams. *Econometrica* 41, 4, 617–631.
- GUL, F. AND STACCHETTI, E. 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* 87, 1, 95–124.
- HARTLINE, J. D. 2013. Mechanism design and approximation. Book draft.
- HARTLINE, J. D. AND ROUGHGARDEN, T. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*. 225–234.
- HASSIDIM, A., KAPLAN, H., MANSOUR, M., AND NISAN, N. 2011. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*. 295–296.
- HUANG, Z., MANSOUR, Y., AND ROUGHGARDEN, T. 2014. Making the most of few samples. Working paper.
- KELSO, JR., A. S. AND CRAWFORD, V. P. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica* 50, 6, 1483–1504.
- KRISHNA, V. 2010. *Auction Theory*, Second ed. Academic Press.
- LEHMANN, B., LEHMANN, D., AND NISAN, N. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55, 2, 270–296.
- LEHMANN, D., O'CALLAGHAN, L. I., AND SHOHAM, Y. 2002. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM* 49, 5, 577–602.
- LUCIER, B. AND BORODIN, A. 2010. Price of anarchy for greedy auctions. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 537–553.
- MILGROM, P. 2004. *Putting Auction Theory to Work*. Cambridge.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1, 58–73.
- NEEMAN, Z. 2003. The effectiveness of English auctions. *Games and Economic Behavior* 43, 2, 214–238.
- NISAN, N. 2014. Algorithmic mechanism design: Through the lens of multi-unit auctions. In *Handbook of Game Theory*. North-Holland. Forthcoming.
- NISAN, N. AND RONEN, A. 2001. Algorithmic mechanism design. *Games and Economic Behavior* 35, 1/2, 166–196.
- NISAN, N. AND SEGAL, I. 2006. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory* 129, 1, 192–224.
- OSTROVSKY, M. AND SCHWARZ, M. 2009. Reserve prices in internet advertising auctions: A field experiment. Working paper.
- ROUGHGARDEN, T. 2014. Barriers to near-optimal equilibria. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. To appear.
- SEGAL, I. 2003. Optimal pricing mechanisms with unknown demand. *American Economic Review* 93, 3, 509–529.
- SEGAL, I. 2006. The communication requirements of combinatorial allocation problems. In *Combinatorial Auctions*, P. Cramton, Y. Shoham, and R. Steinberg, Eds. MIT Press, Chapter 11.
- SWINKELS, J. M. 2001. Efficiency of large private value auctions. *Econometrica* 69, 1, 37–68.
- SYRGGANIS, V. AND TARDOS, É. 2013. Composable and efficient mechanisms. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*. 211–220.
- VALIANT, L. G. 1984. A theory of the learnable. *Communications of the ACM* 27, 11, 1134–1142.

VICKREY, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16, 1, 8-37.