

The Price of Anarchy in Auctions

Tim Roughgarden

*Computer Science Department, Stanford University,
Stanford, CA USA*

TIM@CS.STANFORD.EDU

Vasilis Syrgkanis

*Microsoft Research, 1 Memorial Drive,
Cambridge, MA USA*

VASY@MICROSOFT.COM

Éva Tardos

*Computer Science Department, Cornell University,
Ithaca, NY USA*

EVA@CS.CORNELL.EDU

Abstract

This survey outlines a general and modular theory for proving approximation guarantees for equilibria of auctions in complex settings. This theory complements traditional economic techniques, which generally focus on exact and optimal solutions and are accordingly limited to relatively stylized settings.

We highlight three user-friendly analytical tools: smoothness-type inequalities, which immediately yield approximation guarantees for many auction formats of interest in the special case of complete information and deterministic strategies; extension theorems, which extend such guarantees to randomized strategies, no-regret learning outcomes, and incomplete-information settings; and composition theorems, which extend such guarantees from simpler to more complex auctions. Combining these tools yields tight worst-case approximation guarantees for the equilibria of many widely-used auction formats.

1. Introduction

Many modern applications in computer science involve a number of self-interested participants, with objectives different from each other and from the application designer. *Auctions* are a canonical genre of such applications, ranging from the sale of antiques on eBay, to real-time and targeted Internet advertising, to the allocation of licenses for the wireless spectrum that constitutes much of the modern communication infrastructure.

Auctions have been studied by economists for over a half-century. But over the last several years, work in computer science has offered a fresh and relevant perspective. This new theory is the subject of this survey and it concerns *approximation guarantees* for the equilibrium performance of auctions, also known as “price of anarchy” bounds (Koutsoupias & Papadimitriou, 1999). There are a number of fundamental models that appear impossible to reason about without resorting to approximation.¹

For example, consider a seller with m different items for sale. There are n bidders, and the seller does not know what the bidders want. What should the seller do? One idea is to simply ask the bidders what they want, meaning ask them to bid on each of the possible

1. Computer science has also brought other important ideas to the table in auction design, including an emphasis on reasonable computational complexity and robustness to informational assumptions. These topics are outside the scope of this survey.

subsets of items that they might get. With a single item ($m = 1$), this is a practical idea, and is basically what happens in an eBay auction. In general, however, this idea requires soliciting 2^m bids from each bidder (one per subset of the items), which is a non-starter unless m is tiny.

So how are multiple items auctioned off in practice? One of the most common methods is to sell each item separately. That is, each bidder submits one bid for each item (m bids in all), and each item is awarded to the highest bidder on that item (e.g., for a price equal to the highest or second-highest bid). How good is this simple method — do bidders bid in a way that the items are allocated to those who want them the most?

For many decades, it has been known that simple auction formats do not generally result in the most efficient allocation of the items for sale, but conventional wisdom in economics states that the allocation should be “pretty good” provided bidders’ preferences over items are “sufficiently nice” (e.g. see Milgrom, 2004). Traditional economic tools appear inadequate for translating this empirical rule of thumb into a rigorous performance guarantee, for two reasons: (1) work in economics has focused on exact and optimal solutions, and for the most part has not considered approximation guarantees; (2) economic systems are traditionally studied by solving for equilibria and then analyzing them, while the equilibria of multi-item auctions are far too complex to characterize.

This survey outlines a fairly general and modular theory for proving rigorous performance guarantees for equilibria of auctions in complex settings. A representative consequence of the theory is: if multiple items are sold separately using first-price auctions and the willingness to pay of each bidder is a submodular function of the items that she receives (i.e., preferences are “sufficiently nice”), then every equilibrium of the auction achieves social welfare at least 63% of the maximum possible (i.e., is “pretty good”). Despite the complexity of the welfare-maximization problem and of the equilibria in such multi-item auctions, this guarantee follows from a user-friendly three-step recipe. Roughly, the first step is to consider only the very special case of a single item, with every bidder knowing the willingness to pay of every other bidder, and pure (i.e., deterministic) equilibria. The task in this special case is to use elementary arguments to translate the defining conditions of an equilibrium into a particular type of approximate welfare guarantee (a “smoothness-type inequality”). The second step is to apply an “extension theorem,” which extends the approximate welfare guarantee to single-item auctions in which there is uncertainty (i.e., randomness) both in what the bidders are willing to pay and in what bids they submit, or to outcomes reached by players using no-regret learning, which results in a form of correlation in this randomness. The third step is to apply a “composition theorem,” which extends the approximate welfare guarantee to equilibria of simultaneous auctions with any number of items.

1.1 Organization

Section 2 explores a simple but non-trivial example: the worst-case inefficiency of equilibria in first-price single-item auctions. In addition to introducing Bayes-Nash equilibrium analysis in a concrete and understandable setting, the efficiency analysis in this section already introduces the essence of many of the key ideas of the general framework. Section 3 considers a more complex example, the sale of multiple items via simultaneous single-item auctions.

After absorbing the analyses of these two examples, the general theory for approximate efficiency guarantees for smooth auctions via “extension theorems”, described in Section 4, follows naturally. Section 5 explains why guarantees for smooth auctions apply even when players fail to converge to an equilibrium, provided each achieves a “no-regret” property over repeated plays of the auction. Section 6 shows that the extension from single-item to simultaneous single-item auctions (Sections 2 and 3) is a general phenomenon, by proving a “composition theorem” for smooth auctions. Section 7 considers the limitations of simple auctions, and explains why they cannot enjoy good price-of-anarchy bounds with general bidder preferences. Section 8 offers pointers to the literature on related topics not covered in this survey, and Section 9 concludes with a number of open research directions.

2. A Simple Example: First-Price Single-Item Auctions

We begin our analysis of the equilibria of auctions with a simple but fundamental non-truthful auction, the first-price single-item auction. Consider a single item being auctioned off to one of n bidders (also called “players”). Each bidder i has some value v_i for winning the item — the maximum “willingness to pay” of the bidder — which is private information known only to her. If she wins the item and is asked to pay a price p_i , then her payoff is $v_i - p_i$. We refer to such payoffs by saying that players have quasi-linear preferences with respect to money.²

In a sealed-bid first-price auction, every player i simultaneously submits a bid b_i to the auctioneer. The player with the highest bid wins the item and pays her bid. Ties are broken arbitrarily.

2.1 First-Price Auctions and Bayes-Nash Equilibria

Bidding in a first-price auction is tricky.³ Certainly no player will bid her true valuation, as this would guarantee zero payoff. Instead, bidders will “shade” their bids, meaning that they will bid less than their values. By how much should a bidder shade her bid? The answer depends on the amount of competition she faces and on how other players behave. How can a bidder reason about what others will do when their valuations are unknown to her? The standard approach to modeling this issue is via (Bayesian) *games of incomplete information*. We assume that the valuation v_i of each player i is drawn independently from some distribution \mathcal{F}_i , and that these distributions are common knowledge to all the players. Intuitively, these distributions correspond to the common beliefs that players have about everyone’s valuations.

In this incomplete-information model, a *strategy* of a player is a function s_i that maps a value v_i in the support of \mathcal{F}_i to a bid $s_i(v_i)$. The semantics are: “when my valuation is v_i , I will bid $s_i(v_i)$.” The central equilibrium concept in Bayesian games is the *Bayes-Nash equilibrium*. By definition, a profile of strategies constitutes a Bayes-Nash equilibrium if for every player i and every valuation v_i that the player might have, the player chooses

2. Non-quasi-linear utility functions, such as those incorporating risk aversion or budgets, are also interesting. See also Sections 8 and 9.

3. By contrast, in a *second-price auction*, where the winning bidder pays the value of the second-highest bid, it is a weakly dominant strategy to bid one’s true value.

a bid $s_i(v_i)$ that maximizes her conditional expected utility. The expectation is over the valuations of other players, conditioned on bidder i 's valuation being v_i .

2.2 Symmetric Valuation Distributions

What do Bayes-Nash equilibria look like in a single-item first-price auction? To get a feel for this question, we begin with a simple example, with two bidders with valuations drawn independently and identically from the uniform distribution on $[0, 1]$.

EXAMPLE 2.1. (Two bidders with uniform $[0, 1]$ valuations) Let's "guess and check" a Bayes-Nash equilibrium for this example. First, since the setting is symmetric in the two bidders, it is natural to guess that Bayes-Nash equilibria are also symmetric, meaning that the two players use the same strategy $s(v)$. Let's also guess that the function $s(\cdot)$ is strictly increasing, continuous, and differentiable. Under these assumptions, the highest bidder is the bidder with the largest valuation. By symmetry, the probability that a bidder with valuation v wins is $\mathcal{F}(v) = v$. To check the Bayes-Nash equilibrium conditions, fix a player and condition on her valuation being v . We need to solve for the bid that maximizes the expected utility of the bidder. She could pretend to have value z and bid $s(z)$ for $z \in [0, 1]$. Her expected utility for such a bid is

$$g(z) = \underbrace{(v - s(z))}_{\text{utility of win}} \cdot \underbrace{\mathcal{F}(z)}_{\text{prob. of win}} = (v - s(z)) \cdot z.$$

To force the condition that the optimal bid of the form $s(z)$ is $s(v)$, as prescribed by the Bayes-Nash equilibrium conditions, we differentiate g with respect to z and set $s(\cdot)$ to force a zero derivative at v . This yields the condition

$$0 = g(z)'|_{z=v} = v - (s(z)z)'|_{z=v} = v - (v \cdot s(v))' \Leftrightarrow v \cdot s(v) = \frac{v^2}{2} + \text{constant}$$

on the function $s(\cdot)$. Setting $s(0) = 0$, we obtain the solution $s(v) = \frac{v}{2}$. This solution does indeed satisfy our initial assumptions of differentiability and monotonicity. It is also easy to check that it satisfies the Bayes-Nash equilibrium conditions for all bids, and not just for bids of the form $s(z)$ for some z .

This Bayes-Nash equilibrium turns out to be the unique equilibrium in this example. This equilibrium is fully efficient, in that the item is always allocated to the bidder with the higher valuation. ■

The argument in Example 2.1 generalizes to arbitrary settings in which players' valuations are drawn independently and identically from a distribution \mathcal{F} . Bayes-Nash equilibria also continue to be unique in this case (Chawla & Hartline, 2013). We refer to such settings as *symmetric first-price auctions*. For example, with n bidders with valuations drawn from the uniform distribution on $[0, 1]$, every player uses the strategy $s(v) = \frac{n-1}{n}v$ in the Bayes-Nash equilibrium. Thus, as competition increases, bidders shade their bids less at equilibrium.

Summarizing, there are two key take-aways about symmetric first-price auctions.

1. Bayes-Nash equilibria are relatively well understood.
2. Bayes-Nash equilibria are fully efficient, with the item always allocated to the bidder with the highest valuation.

2.3 Asymmetric Valuation Distributions

When there is information that distinguishes different bidders, for example the market shares of different companies, the symmetry assumption of Section 2.2 is no longer appropriate. Can we extend the results of that section to *asymmetric* first-price single-item auctions (Maskin & Riley, 2000), where bidders' valuations are drawn from different distributions?

This question has been extensively studied (see Krishna, 2002, Section 4.3). Because solving for a Bayes-Nash equilibrium in the asymmetric case is a daunting task and generally admits no closed-form solution, most papers in the area have considered only the case of two bidders and specific parametric distributions (Vickrey, 1961; Kaplan & Zamir, 2012). Even with two bidders with valuations drawn uniformly from $[0, 1]$ and $[0, 2]$, things get complicated.

EXAMPLE 2.2. (Two bidders with uniform $[0, 1]$ and uniform $[0, 2]$ distributions (Vickrey, 1961)) One can verify that the following bidding functions constitute an equilibrium in this example (see also Krishna, 2002):

$$s_1(v_1) = \frac{4}{3v_1} \left(1 - \sqrt{1 - \frac{3v_1^2}{4}} \right)$$

$$s_2(v_2) = \frac{4}{3v_2} \left(\sqrt{1 + \frac{3v_2^2}{4}} - 1 \right).$$

Both bidders bid in the range $[0, \frac{2}{3}]$, with the weaker bidder 1 bidding more aggressively than the stronger bidder 2 (i.e., $s_1(v) > s_2(v)$ for $v \in [0, 1]$). For intuition, recall from Example 2.1 that if the bidders had uniformly and identically distributed valuations, then at equilibrium both bid half their value. Recall also that the equilibrium bid of a player increases with the amount of competition she faces. In this example, from the first bidder's perspective, the second bidder represents stiffer competition than an identically distributed bidder, so she bids more aggressively than in the symmetric case. The opposite reasoning applies to the second bidder, who bids less aggressively than in the symmetric case. For this reason, the Bayes-Nash equilibrium is not fully efficient — there are valuation profiles in which the bidder with the lower valuation is the higher bidder and hence the winner. ■

Summarizing, even the simplest asymmetric first-price auctions are less well-behaved than symmetric ones, in two senses.

1. Solving for a Bayes-Nash equilibrium requires finding a solution to a system of partial differential equations, which in most cases has no closed-form solution.
2. Bayes-Nash equilibria are generally inefficient, in that the bidder with the highest valuation is not always the winner.

How inefficient can the Bayes-Nash equilibria of auctions be? The goal of this survey is to explain a number of general tools that are useful for answering this question, along with several representative applications. Specifically, for every auction format that we consider, we aim to show that the *price of anarchy* — the smallest ratio between the expected welfare of a Bayes-Nash equilibrium and the expected welfare of a welfare-maximizing allocation

— is at least some constant, independent of the parameters of the auctions (the number of bidders, the valuation distributions, etc.).

2.4 The Price of Anarchy of First-Price Auctions

The previous section demonstrated the futility of trying to characterize the Bayes-Nash equilibria of asymmetric first-price auctions in order to bound their inefficiency. Instead, our analysis will rely only on the fact that, in a Bayes-Nash equilibrium, every player is best responding to her opponents' strategies. We will then show that every strategy profile that satisfies this best response property is approximately efficient. We will prove that, in every (asymmetric) first-price auction, every Bayes-Nash equilibrium has expected welfare at least $\approx 63\%$ of the maximum possible. Thus non-trivial efficiency equilibrium guarantees do not require a detailed understanding of the structure of equilibria.

We use the following notation. For a bid profile $\mathbf{b} = (b_1, \dots, b_n)$, $x_i(\mathbf{b})$ denotes whether or not bidder i is the winner (1 or 0, respectively). We denote by $p(\mathbf{b}) = \max_{i \in \{1, \dots, n\}} b_i$ the selling price, which is the highest bid. We use $u_i(\mathbf{b}; v_i)$ to denote the utility of player i when her valuation is v_i and the bid profile is \mathbf{b} . Because we are analyzing the first-price auction, we can write

$$u_i(\mathbf{b}; v_i) = (v_i - b_i) \cdot x_i(\mathbf{b}). \quad (1)$$

Consider a strategy profile $\mathbf{s} = (s_1, \dots, s_n)$, where each strategy s_i is a function from the player's valuation to her bid. We use $\mathbf{s}(\mathbf{v})$ to denote the strategy vector resulting from the vector of valuations \mathbf{v} . For a vector \mathbf{w} , we use \mathbf{w}_{-i} to denote the vector \mathbf{w} with the i th component removed. For example, $\mathbf{s}_{-i}(\mathbf{v}_{-i})$ is the vector of bids of the players other than i when the valuations of these players are \mathbf{v}_{-i} . With this notation, a strategy profile $\mathbf{s} = (s_1, \dots, s_n)$ is a Bayes-Nash equilibrium if and only if

$$\mathbb{E}_{\mathbf{v}_{-i}} [u_i(\mathbf{s}(\mathbf{v}); v_i) \mid v_i] \geq \mathbb{E}_{\mathbf{v}_{-i}} [u_i(b'_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}); v_i) \mid v_i] \quad (2)$$

for every player i , every possible valuation v_i of the player, and every possible deviating bid b'_i . The expectations are over the valuations of the players other than i , according to the assumed prior distribution \mathcal{F} .

The *social welfare* of a bid profile \mathbf{b} when the valuation profile is $\mathbf{v} = (v_1, \dots, v_n)$ is

$$SW(\mathbf{b}; \mathbf{v}) = \sum_{i=1}^n v_i \cdot x_i(\mathbf{b}). \quad (3)$$

Note that the social welfare is the sum of the utilities of the players plus the revenue of the auctioneer. The maximum-possible social welfare in a single-item auction is

$$\text{OPT}(\mathbf{v}) = \sum_{i=1}^n v_i \cdot x_i^*(\mathbf{v}), \quad (4)$$

where $x_i^*(\mathbf{v})$ is the indicator variable for whether or not player i is the player with the highest valuation (with ties broken arbitrarily).

The *price of anarchy* of an auction, with valuation distribution \mathcal{F} , is the smallest value of the ratio

$$\frac{\mathbb{E}_{\mathbf{v}} [SW(\mathbf{s}(\mathbf{v}); \mathbf{v})]}{\mathbb{E}_{\mathbf{v}} [\text{OPT}(\mathbf{v})]},$$

ranging over all Bayes-Nash equilibria \mathbf{s} of the auction. Thus inefficiency is measured by the extent to which the price of anarchy is smaller than 1. The price of anarchy of an auction format is then the worst-case (i.e., smallest) price of anarchy of the auction in any setting, ranging over all choices n for the number of players and all valuation distributions \mathcal{F} .

Theorem 2.1 (Price of Anarchy of First-Price Single-Item Auctions). *The price of anarchy of the first-price single-item auction format is at least $1 - \frac{1}{e} \approx 0.63$.*

Proof. We prove a bound of $\frac{1}{2}$; the bound of $1 - \frac{1}{e}$ follows from an optimized version of the following argument (see Syrgkanis & Tardos, 2013).

Let \mathbf{s} be a Bayes-Nash equilibrium. By definition, every player i chooses a strategy that maximizes her expected utility, given her valuation v_i , the distribution on \mathbf{v}_{-i} , and the strategies $\mathbf{s}_{-i}(\mathbf{v}_{-i})$ used by the other players. In particular, if a bidder i deviates from the bid $s_i(v_i)$ that she uses in the equilibrium, to bidding half her value ($b_i^* = \frac{v_i}{2}$), then her expected utility can only go down.⁴ The choice of such hypothetical deviations b_i^* will be an important theme throughout this survey.

It is simple to bound from below the utility of a bidder that bids half her value ($b_i^* = \frac{v_i}{2}$). For every bid profile \mathbf{b} , we have

$$u_i(b_i^*, \mathbf{b}_{-i}; v_i) \geq \frac{1}{2}v_i - p(\mathbf{b}), \quad (5)$$

since the bidder either wins and obtains utility $v_i - b_i^* = \frac{1}{2}v_i \geq \frac{1}{2}v_i - p(\mathbf{b})$ or loses (in which case $\frac{1}{2}v_i < p(\mathbf{b})$) and obtains utility $0 \geq \frac{1}{2}v_i - p(\mathbf{b})$. Since the bid $b_i^* = \frac{v_i}{2}$ guarantees non-negative utility, we can also write

$$u_i(b_i^*, \mathbf{b}_{-i}; v_i) \geq \left(\frac{1}{2} \cdot v_i - p(\mathbf{b})\right) \cdot x_i^*(\mathbf{v}).$$

Summing this inequality over all bidders i , we obtain

$$\sum_{i=1}^n u_i(b_i^*, \mathbf{b}_{-i}; v_i) \geq \sum_{i=1}^n \left(\frac{1}{2} \cdot v_i - p(\mathbf{b})\right) \cdot x_i^*(\mathbf{v}) = \frac{1}{2}\text{OPT}(\mathbf{v}) - p(\mathbf{b}) \quad (6)$$

for every valuation profile \mathbf{v} and bid profile \mathbf{b} , where x_i^* is defined as in (4).

We now invoke the hypothesis that \mathbf{s} is a Bayes-Nash equilibrium: for every player i with valuation v_i ,

$$\mathbb{E}_{\mathbf{v}_{-i}} [u_i(\mathbf{s}(\mathbf{v}); v_i)] \geq \mathbb{E}_{\mathbf{v}_{-i}} [u_i(b_i^*, \mathbf{s}_{-i}(\mathbf{v}_{-i}); v_i)]. \quad (7)$$

Taking expectations over v_i , summing up the n inequalities of the form (7), and combining with inequality (6), we obtain

$$\sum_{i=1}^n \mathbb{E}_{\mathbf{v}} [u_i(\mathbf{s}(\mathbf{v}); v_i)] \geq \sum_{i=1}^n \mathbb{E}_{\mathbf{v}} [u_i(b_i^*, \mathbf{s}_{-i}(\mathbf{v}_{-i}); v_i)] \geq \mathbb{E}_{\mathbf{v}} \left[\frac{1}{2}\text{OPT}(\mathbf{v}) - p(\mathbf{s}(\mathbf{v})) \right].$$

4. To obtain the $1 - \frac{1}{e}$ bound, one needs to consider a randomized bid with support $[0, (1 - \frac{1}{e})v_i]$ in place of the deterministic bid $v_i/2$.

Last, observe that by the quasi-linear form of bidders' utilities, for every bid profile \mathbf{b} and valuation profile \mathbf{v} we have

$$\sum_{i=1}^n u_i(\mathbf{b}; v_i) = SW(\mathbf{b}; \mathbf{v}) - p(\mathbf{b}).$$

Combining the last inequality and equation yields

$$\mathbb{E}_{\mathbf{v}} [SW(\mathbf{s}(\mathbf{v}); \mathbf{v})] = \sum_{i=1}^n \mathbb{E}_{\mathbf{v}} [u_i(\mathbf{s}(\mathbf{v}); v_i)] + \mathbb{E}_{\mathbf{v}} [p(\mathbf{s}(\mathbf{v}))] \geq \frac{1}{2} \mathbb{E}_{\mathbf{v}} [\text{OPT}(\mathbf{v})],$$

which concludes the proof. \blacksquare

REMARK 2.1. It is natural to ask if the bound of $1 - 1/e$ in Theorem 2.1 is the best possible. All that is currently known is that no bound better than .87 is possible (Hartline, Hoy, & Taggart, 2014). Determining the precise worst-case price of anarchy of asymmetric first-price auctions is an interesting open question (Section 9). \blacksquare

Interestingly, the proof of Theorem 2.1 never used the assumption that bidders' valuations are independent. We therefore have an even stronger guarantee.

Theorem 2.2. *Every Bayes-Nash equilibrium of a first-price auction with correlated valuation distributions has expected social welfare at least $1 - 1/e$ times that of the expected optimal welfare.*

REMARK 2.2. For correlated valuation distributions, this bound of $1 - 1/e$ is tight (for an explicit example see Syrgkanis, 2014). \blacksquare

3. A More Complex Example: Simultaneous Single-Item First-Price Auctions

Before moving to a general auction setting, we explore a more complex example and introduce the idea of smooth auctions. In this section, we consider an auction for multiple heterogeneous items and quantify the equilibrium inefficiency of a simple decentralized auction format.

3.1 Multi-item Auctions

Consider a set of m items being auctioned off to a set of n bidders. Each bidder $i \in [n]$ has a value v_{ij} for each item $j \in [m]$ and only wants one item. If she happens to win a set of multiple items S , then her valuation for the set is the highest-valued item in the set, i.e., $v_i(S) = \max_{j \in S} v_{ij}$. Such bidders are often called *unit-demand* bidders. This setting has a long and distinguished history in the economics literature. It is a generalization of the *assignment model* analyzed by Shapley and Shubik (1971), where the notion of *core outcomes* was originally introduced, and it is also the setting considered by Demange, Gale, and Sotomayor (1986), who introduced the first ascending auctions for multi-item settings that implement welfare-optimal outcomes. From an algorithmic point of view, the welfare optimization problem in this setting is simply the maximum weighted bipartite matching

problem (Cook, Cunningham, Pulleyblank, & Schrijver, 1997), which also has a long and distinguished history in combinatorial optimization.

This section analyzes the efficiency of the simple auction format that sells each item j simultaneously and independently using a single-item first-price auction (Section 2). Each player i submits a bid b_{ij} for each item j . Each item j is awarded to the highest bidder for the item and each bidder is asked to pay her bid for each item that she won. The utility of a player is her value for the set of items she won minus her total payment. Thus, if we denote by $S_i(\mathbf{b})$ the set of items allocated to player i under a bid profile \mathbf{b} , then:

$$u_i(\mathbf{b}; v_i) = v_i(S_i(\mathbf{b})) - \sum_{j \in S_i(\mathbf{b})} b_{ij} = \max_{j \in S_i(\mathbf{b})} v_{ij} - \sum_{j \in S_i(\mathbf{b})} b_{ij} \quad (8)$$

The analysis of this auction game with unit-demand bidders dates back to the early work of Engelbrecht-Wiggans and Weber (1979) — when auction theory was still in its infant stages — who analyzed the very special case of n bidders and n items, with all players having a valuation of 1 for each item and constrained to bid on only one item. The optimal welfare is clearly n . Interestingly, they show that there exists a symmetric mixed Nash equilibrium whose expected welfare approaches $(1 - 1/e)n$ as $n \rightarrow \infty$. In other words, they showed that the price of anarchy of this auction game can be as bad as $1 - 1/e$. This result can be viewed as the first price-of-anarchy bound in the realm of simple auctions. Surprisingly, the techniques that we describe in this survey will show that the example proposed by Engelbrecht-Wiggans and Weber (1979) exhibits the worst-possible inefficiency, over all numbers of bidders and items, all choices of bidders' valuation distributions, and all Bayes-Nash equilibria.

Two decades later, Bikhchandani (1999) analyzed a generalization of the game where players can have arbitrarily complex valuations over the items and without the restriction of bidding on only one item. He focused on the special case of complete information, where all players' valuations are common knowledge, and of pure-strategy Nash equilibria, where each bidder deterministically chooses a single strategy. He showed that complete-information pure-strategy Nash equilibria are fully efficient. We are interested in the case of incomplete information and general Bayes-Nash equilibria; as we know from the previous section, these are not always fully efficient, even in the case of a single item.

The goal of this section is to prove an approximate efficiency guarantee for simultaneous first-price auctions with arbitrary independent distributions over unit-demand valuations.

Theorem 3.1. *Every Bayes-Nash equilibrium of the simultaneous first-price auction game with unit-demand bidders and independent valuations achieves expected social welfare at least $1 - 1/e$ times the expected optimal welfare.*

We start with the simple case where players' valuations are common knowledge. Unlike Bikhchandani (1999), we analyze the inefficiency of mixed-strategy Nash equilibria (where players can randomize). We will then see how our conclusions for mixed Nash equilibria of the complete information setting also extend to the incomplete information setting. In the remainder of the survey we will see how this type of extension is not specific to the game studied in this section, but applies more generally to many auction environments.

3.2 Warm-up: Complete Information

We begin with the case studied by Bikhchandani (1999), of pure-strategy complete information Nash equilibria. Fix a profile \mathbf{v} of unit-demand valuations. Fix a welfare-maximizing allocation, where without loss of generality each bidder receives at most one item, and let $j^*(i)$ denote the item awarded to bidder i in the allocation (if any).

At a pure Nash equilibrium, each player i submits a bid vector $b_i = (b_{ij})_{j \in [m]}$ such that, for all vectors b'_i :

$$u_i(\mathbf{b}; v_i) \geq u_i(b'_i, \mathbf{b}_{-i}; v_i). \quad (9)$$

Let $p_j(\mathbf{b})$ denote the price at which item j is sold. Since a player i does not gain from deviating to any other strategy, she does not gain by bidding infinitesimally above the current price on the item $j^*(i)$ and zero on all other items. Denote this deviating bid vector by b_i^* . By bidding b_i^* , player i definitely wins his optimal item and pays $p_{j^*(i)}(\mathbf{b})$, deriving utility

$$u_i(b_i^*, \mathbf{b}_{-i}; v_i) = v_{i j^*(i)} - p_{j^*(i)}(\mathbf{b})$$

Summing over the bidders i , we conclude that there exist deviations b_1^*, \dots, b_n^* such that

$$\sum_{i \in [n]} u_i(b_i^*, \mathbf{b}_{-i}; v_i) = \text{OPT}(\mathbf{v}) - \sum_{j \in [m]} p_j(\mathbf{b}). \quad (10)$$

Using the fact that the sum of the Nash equilibrium utilities is at least the sum of these deviating utilities (applying (9) with $b'_i = b_i^*$) and that the total utility equals the welfare minus the revenue, we can easily derive the following theorem.

Theorem 3.2 (Bikhchandani, 1999). *Every complete-information pure Nash equilibrium of the simultaneous first-price auction game with unit-demand players achieves the maximum-possible welfare.*

The above result may seem surprising — every equilibrium outcome corresponds to an optimal matching, the solution of a non-trivial combinatorial optimization problem. Thus decentralized optimization by competing bidders yields a globally optimal solution. But how robust is this result?

Unfortunately, the analysis in the proof of Theorem 3.2 breaks down when we try to apply it to a mixed Nash equilibrium, even in the case of complete information. Specifically, if the bid profile is random, then the price $p_j(\mathbf{b})$ of an item is a random variable. A player is not in a position to deviate to bidding $p_j(\mathbf{b})$ on an item, since she does not know the realization of $p_j(\mathbf{b})$ at the time of bidding. This issue cannot be mitigated by a different analysis, as there exist inefficient mixed Nash equilibria of the complete-information simultaneous first-price auction game.⁵

EXAMPLE 3.1. (Inefficiency of Mixed Nash Equilibria) Consider the case of two bidders and two items. Each bidder has a value of 1 for any of the items and bidders are restricted to bid on at most one item. It is relatively easy to show that the following is a symmetric mixed

5. As mentioned above, this was first observed by Engelbrecht-Wiggans and Weber (1979). The price of anarchy remains at most $1 - 1/e$ even if players can bid on multiple items, if players can have valuations slightly more general than unit-demand (Christodoulou, Kovács, Sgouritsa, & Tang, 2016b).

Nash equilibrium of the game: each bidder picks an item uniformly at random and then submits a bid x on that item drawn from a distribution with cumulative density function

$$F(x) = \frac{x}{1-x}$$

and support $[0, \frac{1}{2}]$. It is easy to check that the utility of a player from any bid between $[0, 1/2]$ on any item is equal to $1/2$ and is strictly lower for higher bids. The expected welfare of this equilibrium is equal to the expected number of items that are allocated to some bidder. Each item is allocated to some bidder with probability $3/4$. Thus the total expected welfare is equal to $3/2$, while the optimal welfare is equal to 2. If the example is extended to n bidders and n items, then the price of anarchy tends to $1 - 1/e$ as $n \rightarrow \infty$. ■

The primary obstacle to extending the full efficiency proof (Theorem 3.2) to mixed Nash equilibria is the dependence of the proposed deviating bid on the current bids of the other players. Any analysis that used deviations that depend on the realization of others' bids will only hold for the case of pure Nash equilibria.

One solution would be to prove an approximate efficiency result that makes use only of deviations that are independent of others' actions. But is this even possible?

The previous section provides an affirmative answer in the case of first-price single-item auctions. Specifically, in the proof of Theorem 2.1, the deviations used only require that a bidder bid half her value. This deviation did not depend on others' bids, and guaranteed a utility of at least half of the player's valuation minus the realization of the item's price, whatever the price may be (recall (5)).

We now extend this idea to multi-item auctions. Define bidder i 's deviation b_i^* as bidding half her value on the item $j^*(i)$ that she receives in some fixed optimal allocation (with at most one item per bidder), and 0 on all other items. Following the proof of Theorem 2.1, we have

$$u_i(b_i^*, \mathbf{b}_{-i}; v_i) \geq \frac{v_{ij^*(i)}}{2} - p_{j^*(i)}(\mathbf{b})$$

for every bid profile \mathbf{b} and hence

$$\sum_{i \in [n]} u_i(b_i^*, \mathbf{b}_{-i}; v_i) \geq \frac{1}{2} \text{OPT}(\mathbf{v}) - \sum_{j \in [m]} p_j(\mathbf{b}). \quad (11)$$

If \mathbf{b} is a pure Nash equilibrium, then the inequality (11) and the reasoning in the proof of Theorem 2.1 imply that the social welfare of \mathbf{b} is at least half of the maximum possible.

The key point is that, because the deviations b_i^* used in this derivation are independent of the bid profile \mathbf{b} , the approximation guarantee applies more generally to mixed Nash equilibria (and even beyond mixed Nash equilibria, see Section 5). To see this, note that each deviation b_i^* is now well defined even when others' bids are randomized. Taking the expectation of (11) over the mixed Nash equilibrium bid profile distribution (with the deviations b_i^* fixed) and using the fact that players are best-responding in expectation proves that the expected welfare of every mixed Nash equilibrium is at least $\frac{1}{2} \text{OPT}(\mathbf{v})$. The optimized deviations mentioned in Section 2 can be used to improve the bound from $\frac{1}{2}$ to $1 - 1/e$.

Theorem 3.3. *Every complete-information mixed Nash equilibrium of the simultaneous first-price auction game with unit-demand players achieves expected welfare at least $1 - 1/e$ times the maximum possible.*

3.3 Incomplete Information

The preceding section showed how to cope with randomness in bidders' strategies and prove efficiency guarantees for (complete-information) mixed Nash equilibria. We now consider the incomplete-information case, where the valuation of each bidder i is drawn independently from a distribution \mathcal{F}_i . In this setting, a bidder knows her own valuation and the \mathcal{F}_i 's, but not the realizations of others' valuations. Do our previous price-of-anarchy bounds continue to hold?

Coming up with well-defined deviations b_i^* is again the primary obstacle to extending our results. To understand the issue, recall how we defined b_i^* in the proof of Theorem 3.3, with bidder i bidding half her valuation on the item $j^*(i)$ (if any) that she receives in some fixed optimal allocation. The optimal allocation, and hence the identity of the item $j^*(i)$, depend on the full valuation profile \mathbf{v} . When valuations are commonly known, bidder i knows the identity of the item $j^*(i)$ and is in a position to execute the deviation b_i^* . If bidder i only knows a distribution over others' valuations and not the valuations themselves, then she only knows a distribution over the possible identities of $j^*(i)$. The previous deviation b_i^* is no longer well defined, derailing the argument.

This obstacle motivates defining each deviation b_i^* in a way that depends only on the player's own valuation. This seems rather restrictive! Interestingly, we show that there is, essentially, a black-box way to transform each price-independent deviation used in the proof of Theorem 3.3 in a way that it no longer depends on the valuations of other bidders, while at the same time implying the exact same efficiency guarantee! The key idea is to use the following (randomized) deviation: a bidder i samples valuations \mathbf{v}'_{-i} for the other bidders according to the (known) valuation distributions and uses \mathbf{v}'_{-i} as a surrogate for the true but unknown valuations \mathbf{v}_{-i} . That is, the bidder bids half her value on the item $j'(i)$ and 0 on the other items, where $j'(i)$ is the item i receives (if any) in an optimal allocation when bidder i has valuation v_i and the other bidders have valuations \mathbf{v}'_{-i} . Crucially, this randomized deviation depends only on bidder i 's valuation (and the distributions \mathcal{F}_{-i}), and not on anyone's bid nor on any other bidder's valuation. This idea originates by Christodoulou, Kovács, and Schapira (2016a) and its generality was made clear by Roughgarden (2012) and Syrgkanis (2012). It enables us to extend the argument in the proof of Theorem 3.3 to establish Theorem 3.1.

Proof of Theorem 3.1. We prove a price-of-anarchy bound of $\frac{1}{2}$. The improvement to $1 - 1/e$ follows similar lines as in the proof of Theorem 2.1.

Denote by $j^*(i, \mathbf{v})$ the item awarded to player i in the optimal allocation for the valuation profile \mathbf{v} . Define $b_i^*(\mathbf{v})$ as in the proof of Theorem 3.3, as the bid vector where player i bids half of her value on item $j^*(i, \mathbf{v})$ and zero on every other item. Inequality (11) from the proof of Theorem 3.3 implies that, for every valuation profile \mathbf{v} and every bid profile \mathbf{b} ,

$$\sum_{i \in [n]} u_i(b_i^*(\mathbf{v}), \mathbf{b}_{-i}; v_i) \geq \frac{1}{2} \text{OPT}(\mathbf{v}) - \mathcal{R}(\mathbf{b}), \quad (12)$$

where $\mathcal{R}(\mathbf{b}) = \sum_{j \in [m]} p_j(\mathbf{b})$ denotes the total revenue of the auction. Recall that $b_i^*(\mathbf{v})$ is not a valid deviation in the incomplete-information setting, as player i is not aware of the valuations \mathbf{v}_{-i} .

Consider a Bayes-Nash equilibrium profile of strategies where, conditional on her valuation being v_i , player i chooses a bid according to some distribution $D_i(v_i)$. For conciseness, denote by \mathcal{G}_i the distribution of player i 's bid at this equilibrium (drawing v_i from \mathcal{F}_i and then b_i from $D_i(v_i)$). Crucially, because players' valuations are independently distributed, the joint distribution of equilibrium bids is just the product distribution $\mathcal{G}_1 \times \dots \times \mathcal{G}_n$, which we denote by \mathcal{G} . In particular, the distribution \mathcal{G}_{-i} of the equilibrium bids of players other than i is unaffected by conditioning on player i 's valuation v_i . This would not be the case if players' valuations were correlated.

Now consider the following valid incomplete-information randomized deviation $b'_i \sim D'_i(v_i)$: player i first randomly samples a valuation profile $\mathbf{v}'_{-i} \sim \mathcal{F}_{-i}$ and then performs the deviation $b'_i(v_i, \mathbf{v}'_{-i})$. The Bayes-Nash equilibrium conditions imply that, for every player i and possible valuation v_i of the player,

$$\begin{aligned} \mathbb{E}_{b_i \sim D_i(v_i), \mathbf{b}_{-i} \sim \mathcal{G}_{-i}} [u_i(\mathbf{b}; v_i)] &\geq \mathbb{E}_{b'_i \sim D'_i(v_i), \mathbf{b}_{-i} \sim \mathcal{G}_{-i}} [u_i(b'_i, \mathbf{b}_{-i}; v_i)] \\ &= \mathbb{E}_{\mathbf{v}'_{-i} \sim \mathcal{F}_{-i}, \mathbf{b}_{-i} \sim \mathcal{G}_{-i}} [u_i(b'_i(v_i, \mathbf{v}'_{-i}), \mathbf{b}_{-i}; v_i)]. \end{aligned}$$

This inequality holds for every v_i , and hence also in expectation over v_i :

$$\begin{aligned} \mathbb{E}_{v_i \sim \mathcal{F}_i, b_i \sim D_i(v_i), \mathbf{b}_{-i} \sim \mathcal{G}_{-i}} [u_i(\mathbf{b}; v_i)] &\geq \mathbb{E}_{v_i \sim \mathcal{F}_i, \mathbf{v}'_{-i} \sim \mathcal{F}_{-i}, \mathbf{b}_{-i} \sim \mathcal{G}_{-i}} [u_i(b'_i(v_i, \mathbf{v}'_{-i}), \mathbf{b}_{-i}; v_i)] \\ &= \mathbb{E}_{\mathbf{v} \sim \mathcal{F}, \mathbf{b} \sim \mathcal{G}} [u_i(b'_i(\mathbf{v}), \mathbf{b}_{-i}; v_i)], \end{aligned} \quad (13)$$

where in the equation we have renamed \mathbf{v}'_{-i} as \mathbf{v}_{-i} and also used the fact that the utility of player i after deviating is independent of what she would bid at equilibrium. We emphasize that, on the right-hand side of (13), \mathbf{v} and \mathbf{b} are drawn independently from \mathcal{F} and \mathcal{G} , respectively.

On the other hand, taking the expectation of inequality (12) over $\mathbf{v} \sim \mathcal{F}$ and $\mathbf{b} \sim \mathcal{G}$ shows that the sum of the expected deviating utilities across players is

$$\sum_{i \in [n]} \mathbb{E}_{\mathbf{v} \sim \mathcal{F}, \mathbf{b} \sim \mathcal{G}} [u_i(b'_i(\mathbf{v}), \mathbf{b}_{-i}; v_i)] \geq \frac{1}{2} \mathbb{E}_{\mathbf{v} \in \mathcal{F}} [\text{OPT}(\mathbf{v})] - \mathbb{E}_{\mathbf{b} \sim \mathcal{G}} [\mathcal{R}(\mathbf{b})]. \quad (14)$$

Summing inequality (13) over the bidders i and combining it with inequality (14), we get

$$\mathbb{E}_{\mathbf{v} \sim \mathcal{F}, b_i \sim D_i(v_i) \forall i} \left[\sum_i u_i(\mathbf{b}; v_i) \right] \geq \frac{1}{2} \mathbb{E}_{\mathbf{v} \in \mathcal{F}} [\text{OPT}(\mathbf{v})] - \mathbb{E}_{\mathbf{b} \sim \mathcal{G}} [\mathcal{R}(\mathbf{b})].$$

Now adding $\mathbb{E}_{\mathbf{b} \sim \mathcal{G}} [\mathcal{R}(\mathbf{b})]$ to both sides shows that the expected welfare of the Bayes-Nash equilibrium is at least $\frac{1}{2}$ times the expected maximum welfare. ■

4. General Auctions and Smoothness

The examples in the previous sections portrayed how we can bypass the daunting task of characterizing the equilibria of a game of incomplete information and directly show that every equilibrium is approximately efficient. Our next goal is to develop a general framework for providing such efficiency guarantees, building on our previous arguments for single- and multi-item first-price auctions.

4.1 General Auction Mechanisms

We begin by formally defining a general mechanism design setting and a generic auction, and introducing some essential notation that is used throughout the survey. In a general mechanism design setting, the auctioneer solicits an action a_i from each player i from some action space \mathcal{A}_i (a bid in the case of the auctions considered in the previous two sections). Define $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$. Given the action profile $\mathbf{a} = (a_1, \dots, a_n) \in \mathcal{A}$, the auctioneer decides an outcome $o(\mathbf{a})$ among a set of feasible outcomes \mathcal{O} . Part of an outcome is also a payment $p_i(o)$ that the auctioneer receives from each player. We denote by $\mathcal{R}(o) = \sum_i p_i(o)$ the revenue of the auctioneer.

Each player derives some utility which is a function of the outcome and of a parameter v_i taking values in a parameter space \mathcal{V}_i , typically referred to as the *valuation* or *type* of the player. We denote by $u_i(o; v_i)$ the utility of a player with type $v_i \in \mathcal{V}_i$ in the outcome $o \in \mathcal{O}$. Write \mathcal{V} for the set $\mathcal{V}_1 \times \mathcal{V}_2 \times \dots \times \mathcal{V}_n$ of valuation profiles.

For a given auction, since the outcome is uniquely defined by the action profile, we overload notation and write $u_i(\mathbf{a}; v_i)$ for the utility of the player with type v_i for the outcome in the auction under an action profile \mathbf{a} . We also denote by $\mathcal{R}(\mathbf{a})$ the revenue of the auctioneer in action profile \mathbf{a} .

4.2 Additional Examples

Before formally defining equilibria and smoothness and proving price-of-anarchy bounds, we will give a few additional examples of auctions of interest that fit in the framework that was just defined. While the general definition allows the player utility $u_i(\mathbf{a}; v_i)$ to depend on the price paid in more complex ways, in all of our examples we use quasi-linear utility, where $u_i(\mathbf{a}; v_i) = v_i(\mathbf{a}) - p_i(\mathbf{a})$, for some value function $v_i(\mathbf{a})$ that depends only on the allocation of the auction, and $p_i(\mathbf{a})$ is the price paid by player i .

Our first additional example is again for the sale of a single item, but with the different payment rule in which *all* players pay their bid, and not only the winner. Such *all-pay auctions* arise naturally in contest settings where players have sunk costs that they have to pay irrespective of whether they win or not (e.g., where the work invested acts as the payment). There are many works in the economic literature (Amann & Leininger, 1996; Baye, Kovenock, & de Vries, 1996; Gneezy & Smorodinsky, 2006; Siegel, 2014) and a few recent ones in computer science (DiPalantino & Vojnovic, 2009; Chawla, Hartline, & Sivan, 2012) that analyze properties of all-pay auctions. As with the first-price auction, the asymmetric independent private values model has been under-studied because of the difficulty of characterizing its equilibria in a closed form.

EXAMPLE 4.1. (All-Pay Auction) Consider a setting where n players bid for a single item. Each player has a value v_i for the item, drawn from a commonly known distribution \mathcal{F}_i . Each player submits a bid b_i . The highest bidder wins the item and every player pays her bid (whether a winner or not). ■

A different example arises in the context of public goods, such as a bridge or a park. A group of people needs to decide whether to build such a public good, and how to share the cost. This scenario is especially tricky when the public good is non-excludable, meaning that everyone can use it, whether or not they contributed to its construction (as with most bridges and parks).

EXAMPLE 4.2. (Public Good) A group of n players bid on their share of a joint project that has a publicly known cost c . Each player has a value v_i for the completion of the project. Each player submits a bid b_i . If $\sum_i b_i \geq c$, the project is undertaken, and all players pay their bids, so bidder i 's utility for the outcome is $v_i - b_i$. If $\sum_i b_i < c$, then the project is abandoned and payments are not collected. ■

4.3 Incomplete Information and Equilibria

The type v_i of each player is private information and is drawn independently from a commonly known distribution \mathcal{F}_i . We denote by \mathcal{A}_i the possible actions of the player, and by $\Delta(\mathcal{A}_i)$ the probability distributions over these actions. The strategy of a player in this auction is a mapping $\sigma_i : \mathcal{V}_i \rightarrow \Delta(\mathcal{A}_i)$ from a type $v_i \in \mathcal{V}_i$ to a distribution over actions $\sigma_i(v_i) \in \Delta(\mathcal{A}_i)$.

Definition 4.1 (Bayes-Nash equilibrium (BNE)). *A strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is a Bayes-Nash equilibrium if: for each player i , for each type $v_i \in \mathcal{V}_i$, and for each action $a'_i \in \mathcal{A}_i$:*

$$\mathbb{E}_{\mathbf{v}_{-i} \sim \mathcal{F}_{-i}} [\mathbb{E}_{\mathbf{a} \sim \sigma(\mathbf{v})} [u_i(\mathbf{a}; v_i)]] \geq \mathbb{E}_{\mathbf{v}_{-i} \sim \mathcal{F}_{-i}} [\mathbb{E}_{\mathbf{a}_{-i} \sim \sigma_{-i}(\mathbf{v}_{-i})} [u_i(a'_i, \mathbf{a}_{-i}; v_i)]] .$$

Social welfare. We will be interested in analyzing the social welfare, defined as the utility of all participating parties (the bidders and the auctioneer):

$$SW(o; \mathbf{v}) = \sum_{i=1}^n u_i(o; v_i) + \mathcal{R}(o). \tag{15}$$

For players with quasi-linear utilities, as in (1) and (8), this definition coincides with that in (3). The definition in (15) makes sense with arbitrary player utility functions, not just quasi-linear utility functions. For a given valuation profile \mathbf{v} , we denote the optimal welfare of a feasible outcome by $\text{OPT}(\mathbf{v}) = \max_{o \in \mathcal{O}} SW(o; \mathbf{v})$.

We measure the inefficiency of Bayes-Nash equilibria of an auction with the *price of anarchy (PoA)*, defined as

$$\text{PoA} = \inf_{\mathcal{F}, \sigma \text{ is BNE}} \frac{\mathbb{E}_{\mathbf{v} \sim \mathcal{F}} [\mathbb{E}_{\mathbf{a} \sim \sigma(\mathbf{v})} [SW(\mathbf{a}; \mathbf{v})]]}{\mathbb{E}_{\mathbf{v} \sim \mathcal{F}} [\text{OPT}(\mathbf{v})]} .$$

The PoA is between 0 and 1, with numbers closer to 1 corresponding to more efficient equilibria.

4.4 Smooth Auctions

We begin our analysis of general auction settings by observing that all of the proofs of the efficiency guarantees in Sections 2 and 3 follow the exact same paradigm. The key step is to find an appropriate deviation (b_i^* , or more generally an action a_i^*) for each player, such that no matter what others' actions are, the utility achieved by the deviation can be bounded below by some fraction of the player's contribution to the optimal welfare, less some quantity that relates to the current revenue of the auction. (In fact, we only needed this in aggregate over all of the players.) This was crystalized by Equation (6) for the

case of Bayes-Nash equilibria of single-item first-price auctions, Equation (10) for complete-information pure Nash equilibria of simultaneous first-price auctions, and Equation (11) for mixed Nash equilibria of such auctions. All of these inequalities are variants of what we will call a *smoothness-type* inequality.

The only difference in the different versions of the argument is the informational requirements of the deviations a_1^*, \dots, a_n^* . For instance, in Equation (10) each b_i^* depends on the valuations and bids of all of the players; this restricts the analysis to complete-information pure Nash equilibria. In Equation (11), each b_i^* depends only on the valuations of players and not on their bids; this is sufficient to analyze mixed Nash equilibria and, after the application of a “black-box translation” to remove the dependence of a bidder’s deviation on others’ valuations, Bayes-Nash equilibria in the incomplete-information setting with independent valuation distributions. In Equation (6), each b_i^* depends on the player’s own valuation and on nothing else, and the consequent approximate efficiency bound applies even to Bayes-Nash equilibria of first-price single-item auctions with correlated bidder valuations.

Our primary goal in this section is to understand the efficiency of Bayes-Nash equilibria of auctions in the independent private values model. This motivates defining a *smooth* auction as one that satisfies a natural generalization of Equation (11), with the restriction that the deviations used do not depend on players’ actions (but can depend on all players’ valuations).

Definition 4.2 (Smooth auction). *For parameters $\lambda \geq 0$ and $\mu \geq 1$, an auction is (λ, μ) -smooth if for every valuation profile $\mathbf{v} \in \mathcal{V}$ there exist action distributions $D_1^*(\mathbf{v}), \dots, D_n^*(\mathbf{v})$ over $\mathcal{A}_1, \dots, \mathcal{A}_n$ such that, for every action profile \mathbf{a} ,*

$$\sum_i \mathbb{E}_{a_i^* \sim D_i^*(\mathbf{v})} [u_i(a_i^*, \mathbf{a}_{-i}; v_i)] \geq \lambda \text{OPT}(\mathbf{v}) - \mu \mathcal{R}(\mathbf{a}). \quad (16)$$

Definition 4.2 is due to Syrgkanis and Tardos (2013), inspired by previous definitions of “smooth games” in complete-information (Roughgarden, 2015) and incomplete-information (Roughgarden, 2012; Syrgkanis, 2012) settings. The realizable smoothness parameters of a fixed auction format generally depend on the class \mathcal{V} of permitted valuation profiles; this point is particularly important in Section 6.

On the surface, inequality (16) appears relevant only for the case of complete information (since \mathbf{v} is fixed in (16)) and pure Nash equilibria (since \mathbf{a} is similarly fixed). But the ideas we developed for simultaneous first-price auctions (Section 3) indicate how to extend efficiency guarantees when actions or valuations are randomized. For example, the proof of Theorem 3.3 generalizes to all smooth auctions without difficulty.

Theorem 4.3. *If an auction is (λ, μ) -smooth, then for every valuation profile $\mathbf{v} \in \mathcal{V}$, every complete-information mixed Nash equilibrium of the auction has expected welfare at least $\frac{\lambda}{\mu} \cdot \text{OPT}(\mathbf{v})$.*

Similarly, the proof of Theorem 3.1 can be generalized to all smooth auctions.

Theorem 4.4 (Extension to Incomplete Information). *If an auction is (λ, μ) -smooth, then for every profile $\mathcal{F}_1, \dots, \mathcal{F}_n$ of independent valuation distributions over $\mathcal{V}_1, \dots, \mathcal{V}_n$, every Bayes-Nash equilibrium of the auction has expected welfare at least $\frac{\lambda}{\mu} \cdot \mathbb{E}_{\mathbf{v} \sim \mathcal{F}}[\text{OPT}(\mathbf{v})]$.*

Results like Theorems 4.3 and 4.4 are sometimes called “extension theorems,” meaning that they extend an approximate efficiency guarantee from a restricted class of equilibria (like complete-information pure Nash equilibria) to a more general class (like Bayes-Nash equilibria). They free the analyst to focus on proving inequalities of the form (16), without concern for randomness in valuations or in actions. Proofs that establish inequalities of the form (16), for a suitable λ and μ , are sometimes called “smoothness proofs.”

4.5 Another Example Application of Smoothness: All-Pay Auctions

To portray the generality of the smoothness approach, we discuss how it applies to all-pay auctions (Example 4.1). The extension theorems are especially interesting in this case, as all-pay auctions do not have pure Nash equilibria. Despite this, our results allow the analyst to think only about a property of pure strategies, and then conclude an approximate efficiency guarantee for Bayes-Nash equilibria.

We show that the all-pay auction is $(1/2, 1)$ -smooth. Thus for every valuation profile \mathbf{v} , we require a deviation b_i^* for each player (possibly randomized) such that for every bid profile \mathbf{b} ,

$$\sum_{i \in [n]} \mathbb{E}_{b_i^*} [u_i(b_i^*, \mathbf{b}_{-i}; v_i)] \geq \frac{1}{2} \text{OPT}(\mathbf{v}) - \mathcal{R}(\mathbf{b}). \quad (17)$$

Without loss of generality, assume that player 1 has the highest valuation in the profile \mathbf{v} . Our deviating bids are: b_1^* is chosen uniformly at random from $[0, v_1]$, and $b_i^* = 0$ for all $i > 1$. Now fix a bid profile \mathbf{b} . Player 1 wins the item after deviating to b_1^* whenever this bid is above the highest bid in \mathbf{b}_{-1} , in which case she gets a value of v_1 . Her expected payment is exactly $v_1/2$ (recall it’s an all-pay auction). Thus her expected utility after deviating can be bounded below as follows:

$$\mathbb{E}_{b_1^*} [u_1(b_1^*, \mathbf{b}_{-1}; v_1)] = v_1 \cdot \Pr \left[b_1^* > \max_{j>1} b_j \right] - \frac{v_1}{2} \geq v_1 \cdot \frac{v_1 - \max_j b_j}{v_1} - \frac{v_1}{2} = \frac{v_1}{2} - \max_j b_j.$$

The utility of every other player from the deviation $b_i^* = 0$ is non-negative. Summing all these lower bounds and observing that $\text{OPT}(\mathbf{v}) = v_1$ and $\mathcal{R}(\mathbf{b}) \geq \max_j b_j$ verifies the inequality (16) with $\lambda = \frac{1}{2}$ and $\mu = 1$.

Combining this smoothness proof with Theorem 4.4, we conclude that every Bayes-Nash equilibrium of an asymmetric all-pay auction with players with independent types achieves expected welfare at least half of the expected optimal welfare.

4.6 Correlated Valuations

The guarantee in Theorem 4.6 is for independently drawn player valuations — what if valuations are correlated? Unfortunately, the guarantee no longer holds: there exist auctions that satisfy Definition 4.2 for constant λ and μ , but which have unbounded inefficiency with correlated valuations as the number of players and items in the market grows (Bhawalkar & Roughgarden, 2011; Feldman, Fu, Gravin, & Lucier, 2013).

On the positive side, Theorem 2.2 shows that the Bayes-Nash equilibria of single-item first-price auctions are approximately efficient even when players’ valuations are correlated. The crucial property that enables this result is that the single-item first-price auction satisfies a smoothness-type inequality where the deviating action depends only on the player’s

own valuation (and not on others' valuations). This type of smoothness property was defined by Lucier and Paes Leme (2011), who called it *semi-smoothness*. Since the only difference with Definition 4.2 is the independence of the deviating action from others' valuations, we refer to this property as *smoothness with private deviations*.

Definition 4.5 (Smooth Auction with Private Deviations). *For parameters $\lambda \geq 0$ and $\mu \geq 1$, an auction is (λ, μ) -smooth with private deviations if for every valuation profile $\mathbf{v} \in \mathcal{V}$ there exist action distributions $D_1^*(v_1), \dots, D_n^*(v_n)$ over $\mathcal{A}_1, \dots, \mathcal{A}_n$ such that, for every action profile \mathbf{a} ,*

$$\sum_i \mathbb{E}_{a_i^* \sim D_i^*(v_i)} [u_i(a_i^*, \mathbf{a}_{-i}; v_i)] \geq \lambda \text{OPT}(\mathbf{v}) - \mu \mathcal{R}(\mathbf{a}). \quad (18)$$

The requirement in Definition 4.5 is considerably stronger than that in Definition 4.2, but the reward is approximate efficiency guarantees with correlated valuations.

Theorem 4.6 (Extension to Correlated Valuations). *If an auction is (λ, μ) -smooth, then for every joint distribution \mathcal{F} over players' valuations, every Bayes-Nash equilibrium of the auction has expected welfare at least $\frac{\lambda}{\mu} \cdot \mathbb{E}_{\mathbf{v} \sim \mathcal{F}}[\text{OPT}(\mathbf{v})]$.*

Because of the stronger condition in Definition 4.5, the proof of Theorem 4.6 is simpler than that of Theorem 4.4 — one can just use each deviation $D_i^*(v_i)$ directly, and there is no need to randomly sample fictitious valuations for the other players (and hence no need for independent valuations). The proof follows the same lines as that of Theorem 2.1 (for details, see Lucier & Paes Leme, 2011; Roughgarden, 2012).

We note that the proof of smoothness in Section 4.5 for all-pay auctions does not use private deviations, hence does not extend automatically to correlated valuations. Determining the price of anarchy of single-item all-pay auctions with correlated values is an open problem.

5. No-Regret Learning

So far in this survey we have only analyzed auctions at equilibrium. But how do players arrive at this equilibrium? In many real-world applications of auction design, players do not participate in the auction only once and then vanish. Typically they participate in the auction repeatedly.

The Bayes-Nash equilibrium condition implicitly assumes that when the players arrive in a market, they have done their homework well: they have formed their beliefs about the competition and have computed a Bayes-Nash equilibrium of the market for these beliefs. On their arrival to the market, they simply invoke the equilibrium strategy that they have pre-computed for their realized value.

In many auction settings, the assumption that players are such diligent students is a strong one, especially in cases where entering players lack information about the environment, or when the problem of computing an equilibrium is computationally hard. The expectation of such diligent preparation is even more unreasonable in repeated auction environments where the stakes of each individual auction are small, with the aggregate payoff over time being of primary importance. A more reasonable assumption is that players experiment in the market and try to optimize their bid over time using their past experience

as a proxy for future rewards. For example, in Internet advertising auctions, there are a number of different (adaptive) bidding agents available. In general, the study of adaptive game playing is called *the theory of learning in games* (Fudenberg & Levine, 1998).

Can we bound the average efficiency of an auction when players use adaptive game-playing algorithms? Do the equilibrium efficiency guarantees that we provided thus far extend to such adaptive game playing?

One attractive model of adaptive game playing is *no-regret learning* (Freund & Schapire, 1999; Hart & Mas-Colell, 2000; Auer, Cesa-Bianchi, Freund, & Schapire, 1995), which dates back to the very early work of Hannan (1957). No-regret learning has a long and distinguished history even outside of game theory, when a single decision maker is facing a sequence of decisions among a fixed set of actions whose rewards at each time step are chosen by an adversary (for an extended survey see e.g. Cesa-Bianchi & Lugosi, 2006). It is easy to see the relevance of this model to a repeated game environment: instead of an adversary, the reward of each action at each iteration is affected by the actions of other players. When the other players are hard to predict, the player might as well treat them as adversaries.

A learning algorithm for a player satisfies the no-regret condition if, in the limit as the number of times the game is played goes to infinity, the average reward of the algorithm is at least as good as the average reward of the best fixed action in hindsight (assuming the sequence of actions of the other players remain unchanged). Many simple algorithms are known to achieve this property, including the *multiplicative weight updates* algorithm (Littlestone & Warmuth, 1994; Freund & Schapire, 1999) and the *regret matching* algorithm (Hart & Mas-Colell, 2000). See the work of Cesa-Bianchi and Lugosi (2006) for more general classes of no-regret algorithms.

This section addresses whether the price-of-anarchy guarantees of the previous sections extend to the average welfare of repeated auctions, assuming that each player uses a no-regret learning algorithm. We will argue that the price of anarchy bound of $\frac{\lambda}{\mu}$ for a (λ, μ) -smooth auction (Theorems 4.3–4.4) directly extends to such no-regret learning outcomes, thereby providing further robustness for the welfare properties of smooth auctions.

More formally, we consider an auction with n players that is repeated for T time steps. Each player i has some fixed valuation⁶ v_i and at each iteration t , she chooses to submit some action a_i^t which can depend on the history of play. After each iteration, each player observes the actions taken by the other players.⁷

If a player i uses a no-regret learning algorithm, then in hindsight her average regret for any alternative strategy a_i' (i.e. the difference between the average reward of always playing action a_i' and the average reward of the algorithm) goes to zero or becomes negative (as $T \rightarrow \infty$). When every player uses such an algorithm, the result is a *vanishing regret* sequence.

6. This valuation can be thought of as being drawn at the beginning of time from the distribution \mathcal{F}_i . Recent work of Hartline, Syrgkanis, and Tardos (2015) extends the results of this section to the case where player i 's valuation is drawn at each iteration from \mathcal{F}_i , rather than being fixed.

7. This assumption can be relaxed, and the theory also extends to a “bandit” model where each player only observes the utility of the action taken.

Definition 5.1 (Vanishing Regret). *A sequence of action profiles $\mathbf{a}^1, \mathbf{a}^2, \dots$, is a vanishing regret sequence if for every player i and action $a'_i \in \mathcal{A}_i$,*

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T (u_i(a'_i, \mathbf{a}_{-i}^t; v_i) - u_i(\mathbf{a}^t; v_i)) \leq 0. \quad (19)$$

We now argue that if an auction is (λ, μ) -smooth, then eventually the average welfare of every vanishing regret sequence is at least $\frac{\lambda}{\mu}$ times the optimal welfare. The proof is not hard and we sketch it here. For simplicity, suppose that each player after T time steps already has zero regret (or less) for each action. In particular, each player i has no regret with respect to the randomized action $a_i^* \sim D_i^*(\mathbf{v})$ prescribed by the smoothness property (Definition 4.2). This implies that

$$\frac{1}{T} \sum_{t=1}^T u_i(\mathbf{a}^t; v_i) \geq \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{a_i \sim D_i^*(v_i)} [u_i(a_i^*, \mathbf{a}_{-i}^t; v_i)].$$

Summing this inequality over all players and invoking the smoothness inequality (16) for each action profile \mathbf{a}^t , we can easily conclude that $\frac{1}{T} \sum_{t=1}^T SW(\mathbf{a}^t; \mathbf{v}) \geq \frac{\lambda}{\mu} \text{OPT}(\mathbf{v})$. The same reasoning straightforwardly yields the following result for vanishing regret sequences.

Theorem 5.2 (Extension to Vanishing Regret Sequences). *If an auction is (λ, μ) -smooth, then for every valuation profile \mathbf{v} , every vanishing regret sequence of the auction has expected welfare at least $\frac{\lambda}{\mu} \cdot \text{OPT}(\mathbf{v})$ as $T \rightarrow \infty$.*

6. Composability

“Most analyses of competitive bidding situations are based on the assumption that each auction can be treated in isolation. This assumption is sometimes unreasonable.” (Milgrom & Weber, 1982)

This section gives a general approach for analyzing the efficiency of multiple auctions that take place simultaneously, when players have valuations that are complex functions of the outcomes of the different auctions. Specifically, we will prove a “composition theorem” stating that, under a “complement-free” assumption on players’ utility functions, the simultaneous composition of smooth auctions is also smooth.

6.1 Simultaneous Composition

We consider a setting with n bidders and m auctions. Each auction j concerns its own set of items, its own feasible outcome space \mathcal{O}_j , for each player i its own action space \mathcal{A}_{ij} , and its own outcome function $o_j(\mathbf{a}_j)$. Similarly, we denote by $p_{ij}(o_j)$ the payment of player i in auction j in some outcome o_j and by $\mathcal{R}_j(o_j) = \sum_i p_{ij}(o_j)$ the revenue of auctioneer j . For example, if each auction is a first-price single-item auction, then each \mathcal{A}_{ij} is just the set of possible bids (i.e., \mathbb{R}^+), $o_j(\mathbf{a}_j)$ awards item j to the player who bid the highest for it, and $\mathcal{R}_j(o_j)$ equals the highest bid on j . The *simultaneous composition* of the m auctions is the auction in which each player i simultaneously picks an action a_{ij} for every auction j , resulting in the outcome $\mathbf{o} = (o_1(\mathbf{a}_1), \dots, o_m(\mathbf{a}_m))$.

Utilities and valuations. Importantly, a player’s utility is now a function $u_i(\mathbf{o}; v_i)$ of the outcomes $\mathbf{o} = (o_1, \dots, o_m)$ of all of the auctions, where $v_i \in \mathcal{V}_i$ is the player’s type. In the most commonly studied utility model, each player has a quasi-linear utility function and is indifferent about what happens to the other players. Mathematically, this translates to

$$u_i(\mathbf{o}; v_i) = v_i(S_i(\mathbf{o})) - \sum_j p_{ij}(o_j), \quad (20)$$

where $S_i(\mathbf{o})$ denotes the items awarded to player i in the outcome \mathbf{o} , and v_i is a *valuation function* that assigns a nonnegative value $v_i(S)$ to each subset S of items that player i might obtain. For example, if every auction is a single-item auction, then $S_i(\mathbf{o})$ corresponds to the auctions for which i is the winner.

No positive results are possible with arbitrary valuation functions (see Section 7), so in order to make progress we need to impose structure on players’ valuations. The simplest type of valuation function v_i is an *additive* function, where there are nonnegative “item valuations” v_{i1}, \dots, v_{im} such that, for every bundle S of items, the valuation $v_i(S)$ for the bundle is just the sum $\sum_{j \in S} v_{ij}$ of the corresponding item valuations. A player with an additive valuation can reason about each item separately. The *unit-demand* valuations introduced in Section 3 are another example. Here, for nonnegative item valuations v_{i1}, \dots, v_{im} , the corresponding unit-demand valuation is $v_i(S) = \max_{j \in S} v_{ij}$. Already in this case, the value of a player for some item j might depend on whether or not she wins a different item j' sold in a different auction.

Both additive and unit-demand valuations are special types of *monotone submodular* valuations, meaning that $v_i(S) \leq v_i(T)$ whenever $S \subseteq T$ and $v_i(T \cup \{j\}) - v_i(T) \leq v_i(S \cup \{j\}) - v_i(S)$ for every $S \subseteq T$ and item j also, referred to as having the decreasing marginal value property. Such functions model diminishing returns, and have been studied extensively in machine learning, optimization, and economics.

It is technically convenient to work with a still more general class of valuations, which we call *complement-free* valuations. Intuitively, the following definition prohibits complementarities between different items, where winning one item is valuable only if the bidder also wins some other item.⁸ Precisely, a valuation function v_i is *complement-free* if there exist additive valuations $\mathbf{w}_i^1, \dots, \mathbf{w}_i^r$ such that, for every subset S of items,

$$v_i(S) = \max_{\ell=1}^r \left\{ \sum_{j \in S} w_{ij}^\ell \right\}. \quad (21)$$

For example, if v_i is a unit-demand valuation with item values v_{i1}, \dots, v_{im} , then v_i is the maximum of m additive valuations $\mathbf{w}_i^1, \dots, \mathbf{w}_i^m$, where w_{ij}^ℓ equals v_{ij} if $\ell = j$ and equals 0 otherwise. More generally, every monotone submodular valuation is complement-free. The proof uses $m!$ additive valuations, one for each ordering π of the m items, and defines w_{ij}^π as $v_i(S_j^\pi \cup \{j\}) - v_i(S_j^\pi)$, where S_j^π denotes the set of items that precede j in π . The condition (21) holds: for every subset S of items, $\mathbf{w}_i^\pi(S) \leq v_i(S)$ for every π (by

8. In the literature, “complement-free” is often equated with subadditivity. What we are calling “complement-free” is often called “fractionally subadditive” or “XOS” (for “XOR-of-singletons”) in the literature. Valuations that are complement-free in the sense of this survey are always subadditive, but not conversely. See the work of Lehmann, Lehmann, and Nisan (2001) for further details.

submodularity), and there exists a π such that $\mathbf{w}_i^\pi(S) = v_i(S)$ (take a π where the items of S come first).

We can analogously define a utility function u_i over outcomes $\mathbf{o} \in \mathcal{O}$ as *additive* if it has the form

$$u_i(\mathbf{o}; v_i) = \sum_{j=1}^m u_{ij}(o_j; v_{ij})$$

for valuations of the form $v_i = (v_{i1}, \dots, v_{im})$, and *complement-free* if it is the maximum of additive utility functions:

$$u_i(\mathbf{o}; v_i) = \max_{\ell=1}^r \left\{ \sum_{j=1}^m u_{ij}(o_j; v_{ij}^\ell) \right\}. \quad (22)$$

For example, with single-item auctions and quasi-linear utilities (20), additive and complement-free valuations induce additive and complement-free utility functions, respectively (with $u_{ij}(o_j; v_{ij}^\ell)$ equal to $v_{ij}^\ell x_{ij}(o_j) - p_{ij}(o_j)$, where $x_{ij}(o_j)$ indicates whether or not i wins the j th item in o_j). The composition theorem below requires only that players' utilities are complement-free in the sense of (22), and does not assume quasi-linear utilities or single-item auctions.

6.2 The Composition Theorem

The main result of this section is the following composition theorem.

Theorem 6.1 (Composition Theorem). *If players have complement-free utility functions, then the simultaneous composition of (λ, μ) -smooth auctions is again a (λ, μ) -smooth auction.*

Proof. As a warm-up, suppose that every player i actually has an additive utility function, with valuation $v_i = (v_{i1}, \dots, v_{im})$. Intuitively, in this case the composition theorem follows just by adding up the smoothness inequalities for the constituent auctions.

Formally, write \mathbf{v}_j for the projection (v_{1j}, \dots, v_{nj}) of \mathbf{v} onto the valuation space for the j th auction. To save notation, we write $u_i(\mathbf{a}; v_i)$ to mean $u_i(o(\mathbf{a}); v_i)$, where $o(\mathbf{a})$ is the outcome resulting from the action profile \mathbf{a} . Since each auction j is (λ, μ) -smooth (Definition 4.2), there are action distributions $D_{1j}^*(\mathbf{v}_j), \dots, D_{nj}^*(\mathbf{v}_j)$ over $\mathcal{A}_{1j}, \dots, \mathcal{A}_{nj}$ such that, for every action profile $\mathbf{a}_j \in \mathcal{A}_{1j} \times \dots \times \mathcal{A}_{nj}$,

$$\sum_i \mathbb{E}_{\mathbf{a}_{ij}^* \sim D_{ij}^*(\mathbf{v}_j)} [u_{ij}(a_{ij}^*, \mathbf{a}_{-i,j}; v_{ij})] \geq \lambda \text{OPT}_j(\mathbf{v}_j) - \mu \mathcal{R}_j(\mathbf{a}_j), \quad (23)$$

where $\text{OPT}_j(\mathbf{v}_j)$ denotes the maximum social welfare (i.e., sum of players' utilities plus seller revenue) of any outcome of auction j with valuation profile \mathbf{v}_j . Now define $D_i^*(\mathbf{v}) = D_{i1}^*(\mathbf{v}_1) \times \dots \times D_{im}^*(\mathbf{v}_m)$. Using the additivity of players' utilities and summing up (23) over all j verifies the smoothness condition (16) for the composition of the auctions.

For the general case, fix a profile \mathbf{v} of valuations such that all players' have complement-free utility functions. The idea is to reduce to the additive case. To extract additive utility

functions from the given complement-free ones, let \mathbf{o}^* be a social welfare-maximizing outcome of the composition of auctions for valuations \mathbf{v} . For each player i , with valuation $v_i = (v_i^1, \dots, v_i^m)$ with each $u_i(\mathbf{o}; v_i^\ell)$ additive, define the *proxy valuation* v_i^* as a valuation v_i^ℓ that satisfies $u_i(\mathbf{o}^*; v_i^\ell) = u_i(\mathbf{o}^*; v_i)$. With the valuation profile \mathbf{v}^* , all players have additive utility functions. By the previous paragraph, we can define action distributions $D_1^*(\mathbf{v}^*), \dots, D_m^*(\mathbf{v}^*)$ such that, for every action profile $\mathbf{a} \in \mathcal{A}$,

$$\sum_i \mathbb{E}_{a_i^* \sim D_i^*(\mathbf{v}^*)} [u_i(a_i^*, \mathbf{a}_{-i}; v_i^*)] \geq \lambda \text{OPT}(\mathbf{v}^*) - \mu \mathcal{R}(\mathbf{a}). \quad (24)$$

The utilities on the left-hand side and social welfare on the right-hand side of (24) refer to the proxy valuations \mathbf{v}^* , not the true valuations \mathbf{v} . The complement-free condition (22) and choice of v_i^* imply that $u_i(\mathbf{o}^*; v_i) = u_i(\mathbf{o}^*; v_i^*)$ and $u_i(\mathbf{o}; v_i) \geq u_i(\mathbf{o}; v_i^*)$ for every player i and outcome $\mathbf{o} \in \mathcal{O}$. Hence, the left-hand side of (24) is at most $\sum_i \mathbb{E}_{a_i^* \sim D_i^*(\mathbf{v}^*)} [u_i(a_i^*, \mathbf{a}_{-i}; v_i)]$. Also, since \mathbf{o}^* is a feasible outcome, the optimal welfare with the proxy valuations \mathbf{v}^* is at least as large as it is with the true valuations \mathbf{v} :

$$\text{OPT}(\mathbf{v}^*) \geq \sum_i u_i(\mathbf{o}^*; v_i^*) + R(\mathbf{o}^*) = \sum_i u_i(\mathbf{o}^*; v_i) + R(\mathbf{o}^*) = \text{OPT}(\mathbf{v}),$$

and so the right-hand side of (24) is at least $\lambda \text{OPT}(\mathbf{v}) - \mu \mathcal{R}(\mathbf{a})$ (for every \mathbf{a}). Combining (24) with these two inequalities verifies the smoothness condition for the auction composition. ■

From simpler to more complex valuations. The heart of the proof of Theorem 6.1 extends a smoothness inequality from additive player utility functions to complement-free utility functions. The same argument shows more generally that whenever a mechanism is smooth for some class of utilities and valuations, it is equally smooth when the utility function of each player is a maximum over such utilities and valuations. Feige, Feldman, Immorlica, Izsak, Lucier, and Syrgkanis (2015) and Lucier and Syrgkanis (2015) give applications of this generalization.

6.3 A Promise Fulfilled

In the Introduction, we mentioned that the following result follows easily from the machinery developed in this survey.

Theorem 6.2. *Every Bayes-Nash equilibrium of the simultaneous first-price auction game with bidders with independent submodular valuations and quasi-linear utilities achieves expected social welfare at least $1 - 1/e$ times the expected optimal welfare.*

To review, Theorem 6.2 follows immediately from chaining together the following tools: the fact that the first-price single-item auction is $(1 - \frac{1}{e}, 1)$ -smooth (Theorem 2.1); the fact that submodular functions are complement-free and hence induce complement-free utilities for players with quasi-linear utility functions; the composition theorem, which implies that simultaneous first-price auctions are $(1 - \frac{1}{e}, 1)$ -smooth for such players (Theorem 6.1); and the extension theorem stating that every Bayes-Nash equilibrium of a (λ, μ) -smooth auction with independent player valuations has expected social welfare at least $\frac{\lambda}{\mu}$ times the maximum possible (Theorem 4.4).

Of course, this same machinery applies far more generally — whenever an auction can be expressed as the simultaneous composition of smooth auctions and players have complement-free utility functions over the outcomes of these auctions.

Relationship to Theorem 3.1. To understand how the proof of Theorem 6.1 generalizes that of Theorem 3.1, suppose that each of the m auctions is a first-price single-item auction and that each player has a unit-demand valuation. Recall that such a valuation $v_i = (v_{i1}, \dots, v_{im})$ can be written as the maximum of m additive valuations $\mathbf{w}_i^1, \dots, \mathbf{w}_i^m$, where w_{ij}^ℓ equals v_{ij} if $\ell = j$ and 0 otherwise. The proof of Theorem 6.1 defines proxy additive valuations \mathbf{v}^* by setting each v_i^* to $\mathbf{w}_i^{j^*(i)}$, where $j^*(i)$ is the item i receives in a welfare-maximizing allocation (or to the all-zero function, if $j^*(i)$ is not defined). The deviation $D_i^*(\mathbf{v}^*)$ of player i is then defined by deviating independently in each of single-item auctions j , as if its valuation for that item is $w_{ij}^{j^*(i)}$. The deviation used to prove smoothness of a first-price single-item auction is to bid half of one’s valuation (Theorem 2.1), so player i ’s deviation here is to bid $v_{ij^*(i)}/2$ on item $j^*(i)$ and 0 on every other item. These are exactly the deviations used in the proof of Theorem 3.1.

7. Impossibility Results

Can we improve over Theorem 6.2, either in the approximation guarantee or in the generality of players’ valuations? If not with simultaneous first-price auctions, then what about with some other “simple” auction format? This section outlines a general technique for ruling out good price-of-anarchy bounds for simple auctions.

Consider a setting with n bidders and m items. We now impose no restrictions on the valuation v_i of each bidder i , other than monotonicity (meaning $v_i(S) \leq v_i(T)$ whenever $S \subseteq T$) and normalization (meaning $v_i(\emptyset) = 0$). Such general valuations permit complementarities between items, where one item is valuable only in the presence of other items (e.g., the left and right shoes of a pair). An extreme example is a *single-minded* valuation v_i , which is defined by a subset S of items and nonnegative number w as

$$v_i(T) = \begin{cases} w & \text{if } T \supseteq S; \\ 0 & \text{otherwise.} \end{cases}$$

Can we extend the guarantee in Theorem 6.2 for simultaneous first-price auctions to general valuations, or at least to single-minded bidders? Hassidim, Kaplan, Mansour, and Nisan (2011) provided a negative answer, even for the complete-information case.

Theorem 7.1 (Hassidim et al., 2011). *With two bidders with general valuations, simultaneous first-price auctions can have mixed Nash equilibria with expected welfare arbitrarily smaller than the maximum possible (as $m \rightarrow \infty$).*

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. The examples in the proof of Theorem 7.1 make use of one single-minded bidder and one unit-demand bidder per item.

Can we overcome the negative result in Theorem 7.1 by designing a different auction? The rule of thumb in practice is that simple auctions can perform poorly with complementarities between items (see e.g. Milgrom, 2004). One reason for this is the “exposure

problem,” where a bidder with value only for a bundle of items risks acquiring only a strict subset of the bundle, and at a significant price. But how would we prove a rigorous version of this rule of thumb? For a fixed auction format, like simultaneous first-price auctions, the obvious way to rule out good price-of-anarchy bounds is via an explicit example (as in Theorem 7.1). How could we rule out good bounds for all simple auctions simultaneously?

The key idea is to proceed in two steps. The first step is to rule out good approximation algorithms for the computational problem of allocating the items to maximize the social welfare. The second step is to show that there is a simple auction with good equilibria only if there is a good approximation algorithm for the welfare-maximization problem.

The first step is implemented by Dobzinski, Nisan, and Schapira (2010). The formal statement is about the *communication complexity* of the welfare-maximization problem (Kushilevitz & Nisan, 1996; Roughgarden, 2016). Imagine a setup where each player i initially knows only her own valuation, and the players cooperate to (approximately) maximize the social welfare by exchanging information (in the form of bits) about their valuations. A *communication protocol* specifies how information is exchanged — who tells what to whom, as a function of the history-so-far. The *cost* of a protocol is the maximum number of bits exchanged over all possibilities for the players’ valuations. The communication complexity of a problem is the minimum cost of a communication protocol for it.

Theorem 7.2 (Dobzinski et al., 2010). *For every constant $\alpha > 0$, every communication protocol that achieves an α -approximation of the maximum social welfare for general player valuations has cost exponential in the number of items m .*

Theorem 7.2 is proved via a reduction from a canonical hard communication problem (a version of the “Disjointness” problem); further details are outside the scope of this survey.

The second step is to extend the lower bound in Theorem 7.2 from communication protocols to equilibria of “simple” auctions. For the following result, by “simple” we mean that the number of bids submitted by each player is subexponential in the number of items m . For example, simultaneous single-item auctions require only a linear number of bids per player. A “direct-revelation” mechanism, where each player submits her type, requires 2^m bids per player (one for each subset of items) and thus does not qualify as simple.

Theorem 7.3 (Roughgarden, 2014). *If every α -approximate communication protocol for the welfare-maximization problem has cost exponential in the number of items m , then no family of simple auctions guarantees equilibrium welfare at least α of the maximum possible.*

The proof of Theorem 7.3 effectively shows how to compute an equilibrium allocation of a simple auction without using too much communication. In the presence of a good price-of-anarchy bound, this constitutes a low-cost communication protocol with a good approximation guarantee. If the latter cannot exist, then neither can the former.

Technically, Theorem 7.3 (and Corollary 7.4 below) are only known to hold for ϵ -approximate equilibria — meaning that every player mixes only over strategies with expected utility within ϵ of a best response — where $\epsilon > 0$ can be taken arbitrarily small. Theorem 7.3 applies even to (approximate) mixed Nash equilibria in the complete-information setting.

Combining Theorems 7.2 and 7.3 makes precise the empirical rule of thumb that simple auctions cannot guarantee high welfare outcomes when there are complementarities between items.

Corollary 7.4. *With general bidder valuations, no family of simple mechanisms guarantees equilibrium welfare at least a constant fraction of the maximum possible (as $m \rightarrow \infty$).*

There are analogous results for restricted classes of valuations. For example, when players' valuations are subadditive (meaning $v_i(S \cup T) \leq v_i(S) + v_i(T)$ for all S, T), combining another result by Dobzinski et al. (2010) with Theorem 7.3 shows that no family of simple auctions guarantees equilibrium welfare more than 50% of the maximum possible. Since simultaneous first-price auctions do guarantee at least 50% of the maximum-possible welfare at equilibrium (Feldman et al., 2013), they are optimal simple auctions for subadditive player valuations in a precise sense. It is an open question to determine the best equilibrium welfare guarantee achievable by simple auctions for submodular player valuations (see Section 9).

8. Further Topics and a Guide to the Literature

This section surveys the many related topics not covered by this survey, with pointers for further reading.

8.1 Second-Price-Type Auctions

Throughout the survey we focused on first-price auctions, with a cameo from all-pay auctions. Some auctions used in practice have a different payment scheme. Examples include the generalized second price auction used in online ad auctions, and the uniform-price auction used in financial institutions (e.g., to sell treasury bonds). In these auctions, a player pays the minimum bid she would need to make to continue to win the same item(s).

The smoothness framework extends to such “second-price-type” auctions as well, with small technical modifications to the smoothness definition. The primary change needed in Definition 4.2 is that, in (16), the revenue of the bid profile \mathbf{b} is replaced by the sum of the winning bids in \mathbf{b} . To obtain an approximate welfare guarantee, one needs to assume that players do not bid above their valuations at equilibrium. This assumption implies that the sum of the winning bids is bounded above by the current welfare at equilibrium, and the resulting price-or-anarchy bound is $\lambda/(1 + \mu)$ (instead of λ/μ). Example applications include price-of-anarchy bounds for simultaneous second-price auctions (Christodoulou et al., 2016a; Bhawalkar & Roughgarden, 2011), the generalized second-price auction (Paes Leme & Tardos, 2010), and uniform-price auctions (Markakis & Telelis, 2012). Roughgarden (2012) and Syrgkanis and Tardos (2013) explain how to interpret the analyses in these papers as smoothness proofs.

One major technical issue in second-price-type auctions is the no-overbidding assumption. Even in a single-item second-price auction, where bidding truthfully is a weakly dominant strategy and leads to an efficient allocation, there are very inefficient equilibria if

people overbid.⁹ Thus the no-overbidding assumption is essential for equilibrium efficiency guarantees. In some cases, such as the generalized second-price auction (Paes Leme & Tardos, 2010) and the uniform-price auction (Markakis & Telelis, 2012), the no-overbidding assumption can be justified by proving that overbidding is a weakly dominated strategy. In more complex settings, such as simultaneous second-price auctions (Bhawalkar & Roughgarden, 2011), overbidding may be unnatural, but it is not always weakly dominated. For other auction formats, it is open if one can replace the no-overbidding assumption with the assumption that players do not play dominated strategies (for further discussion see Christodoulou et al., 2016a; Bhawalkar & Roughgarden, 2011; Roughgarden, 2012; Syrgkanis & Tardos, 2013; Feldman et al., 2013).

8.2 Further Applications of the Smoothness Framework

The smoothness framework was originally introduced in the context of general games (Roughgarden, 2015), and has been applied to many other games beyond auctions; for examples, see the recent book of Roughgarden (2016). There are also several other auction formats and mechanism design settings, beyond the ones discussed so far, where the smoothness framework has been employed to characterize the price of anarchy, or where in retrospect existing price-of-anarchy analyses can be cast as smoothness proofs. We have seen *first-price auctions* (Hassidim et al., 2011; Syrgkanis & Tardos, 2013) and *all-pay auctions* (Syrgkanis & Tardos, 2013; Christodoulou, Sgouritsa, & Tang, 2015b). Other auction formats include: variants of *uniform-price auctions* that are used widely in financial markets (Markakis & Telelis, 2012; de Keijzer, Markakis, Schfer, & Telelis, 2013), combinatorial auctions where the allocation of items is decided by a *greedy algorithm* (Lucier & Borodin, 2010), *position auctions* such as the generalized second-price auction and generalizations of it (Caragiannis, Kaklamani, Kanellopoulos, Kyropoulou, Lucier, Paes Leme, & Tardos, 2015; Paes Leme & Tardos, 2010; Syrgkanis & Tardos, 2013), the *proportional mechanism for bandwidth allocation* (Johari & Tsitsiklis, 2004; Syrgkanis & Tardos, 2013; Caragiannis & Voudouris, 2014; Christodoulou, Sgouritsa, & Tang, 2015a), the *Walrasian mechanism* for combinatorial markets, (Babaioff, Lucier, Nisan, & Paes Leme, 2014a), a particularly simple bidding system with a single bid for multiple items introduced by Devanur, Morgenstern, Syrgkanis, and Weinberg (2015), *relax-and-round mechanisms* (Dütting, Kesselheim, & Tardos, 2015), and auctions for *renewable energy markets* (Kesselheim, Kleinberg, & Tardos, 2015).

The price-of-anarchy bounds in all of these papers can be interpreted as showing that the auction is smooth (Definition 4.2) or, in some cases, even smooth with private deviations (Definition 4.5). As a rule of thumb, when the auction allows players to express their entire private valuation through the action space, then the auction is likely to satisfy smoothness with private deviations (e.g., the generalized second-price auction, relax-and-round mechanisms, and the Walrasian mechanism). If the action space is restricted and only allows the player to bid some proxy restricted valuation, then the auction tends to only satisfy the weaker smoothness condition that extends only to the independent private values setting

9. Think of two players, with player a having valuation ϵ and player b having valuation 1. Player a bidding 1 and the player b bidding 0 is an equilibrium of the auction with welfare only ϵ times the maximum possible.

(e.g., simultaneous single-item auctions). The key intuition is that in the latter auction formats, most price-of-anarchy proofs require some version of the argument in the proof of Theorem 6.1 that extends a smoothness equality from some restricted class of valuations to some more general class of valuations. However, this extension only establishes the weaker smoothness condition for general valuations, even if the auction is smooth with private deviations for the restricted class.

8.3 Price-of-Anarchy Bounds via Non-smooth Techniques

Not all of the known price-of-anarchy bounds for auctions are smoothness proofs. One nice example is given in the work of Feldman et al. (2013), who analyze simultaneous first- and second-price item auctions when players' valuations are subadditive (meaning $v_i(S \cup T) \leq v_i(S) + v_i(T)$ for all subsets S, T of items). Subadditive valuations are strictly more general than the complement-free valuations studied in Section 6. Feldman et al. (2013) prove that the price of anarchy is constant in this case, while the smoothness framework is only known to give a bound that degrades logarithmically with the number of items (Bhawalkar & Roughgarden, 2011; Hassidim et al., 2011; Syrgkanis & Tardos, 2013). Feldman et al. (2013) deviate from the smoothness framework by using a different set of deviating actions for each Bayes-Nash equilibrium. Dütting, Henzinger, and Starnberger (2013) propose a generalization of the smoothness framework that includes even the analysis of the work of Feldman et al. (2013) as a special case.

A new approach to bounding the price of anarchy was recently proposed by Kulkarni and Mirrokni (2015). For an arbitrary game, they formulate a convex program that has a strong connection with the welfare-maximization problem associated with the game. They prove that every vanishing regret sequence of action profiles (Definition 5.1) of the game can be associated with a solution to the Fenchel dual of this convex program, with the average welfare of the sequence being close to the value of the dual at this solution. Using duality, one can argue that the average welfare of every such sequence is close to the optimal welfare. The primary applications in the work of Kulkarni and Mirrokni (2015) concern routing, scheduling, and location games. Interesting open questions include whether or not this framework can prove better bounds than the smoothness framework for natural auction games, and whether or not it extends to incomplete-information settings.

Another line of results that fall outside of the smoothness framework are those that prove stronger bounds for restricted subclasses of equilibria. For example, Bikhchandani (1999) proves that pure Nash equilibria of simultaneous first-price auctions are fully efficient (Theorem 3.2). Since mixed-strategy Nash equilibria need not be fully efficient (Example 3.1), this result cannot be established by a smoothness proof. For another example, Bhawalkar and Roughgarden (2011) show that, assuming subadditive player valuations and no over-bidding, every pure Nash equilibrium of simultaneous second-price auctions has welfare at least 50% of the maximum possible. Again, this guarantee does not hold for mixed Nash equilibria and hence cannot be established via a smoothness proof. Further examples are provided by the works of Christodoulou et al. (2015a, 2015b), who prove better bounds for mixed Nash equilibria of complete-information first-price and all-pay auctions than are known for Bayes-Nash equilibria in the more general incomplete-information case.

8.4 Sequential Auctions

In all of the auctions considered in this survey, actions are chosen simultaneously by the players and the allocation and payments are decided in one shot. Many auctions in practice have a sequential component to them. The simplest-possible theoretical model is that of sequential item auctions, where items are sold one-by-one in some predefined order via single-item auctions. This model has a long history in economics, starting from the work of Milgrom and Weber (1982, 1999). Sequentiality leads to a host of new complications in price-of-anarchy analyses. The biggest problem is that when a player deviates at some stage of the game, the deviation can cause a ripple effect in subsequent stages, changing for instance the prices of subsequent items. To bypass this complication, Paes Leme, Syrgkanis, and Tardos (2012) and Syrgkanis and Tardos (2012) proposed a bluffing deviation, where a player pretends to play as in equilibrium of some random type, until the right moment arrives when she deviates to acquire some item. Such a deviation analysis, where the deviation involves simulating some current equilibrium behavior, extends to the incomplete-information setting in a black-box manner (Syrgkanis & Tardos, 2013). Technically, one can show that most of the smoothness theory extends even if the deviations a_i^* can depend on the current action a_i from which the player is deviating from. Such deviations unlock the ability to analyze sequential auctions from a price-of-anarchy point of view.

8.5 Budget Constraints

Most of price-of-anarchy analyses of auctions assume that players have quasi-linear preferences, meaning that a player's utility is her valuation for the items received less the payment made. This implicitly assumes that players are capable of paying an arbitrarily large amount. In many practical auction settings, players also have budget constraints. The simplest way to model a budget is to define the utility of a player as $-\infty$ if her payment exceeds her budget. Do the efficiency guarantees for smooth auctions extend to the setting where players' have finite budgets?

High prices can prevent a player with a small budget and a high valuation from obtaining the allocation she would get in a welfare-maximizing outcome. To address this issue, a sequence of papers have used an alternative benchmark to measure the efficiency of auctions in the presence of budgets. This benchmark is called optimal effective welfare (Syrgkanis & Tardos, 2013) or optimal liquid welfare (Dobzinski & Paes Leme, 2014), and it is defined as the maximum welfare achievable after capping each player's valuation for an allocation by her budget. Syrgkanis and Tardos (2013) show that, under minor additional assumptions, the equilibrium welfare of a (λ, μ) -smooth auction is at least $\frac{\lambda}{\mu}$ times the expected maximum effective welfare. This benchmark was further analyzed by Caragiannis and Voudouris (2014) for the proportional bandwidth allocation mechanism. One interesting open question is whether one can also lower bound the effective welfare at equilibrium, rather than just the actual welfare. Caragiannis and Voudouris (2014) provided such a result for the proportional bandwidth allocation mechanism.

8.6 The Price of Anarchy for Revenue

All of the price-of-anarchy guarantees covered in this survey concern the equilibrium welfare of auctions. Another quantity of primary importance is the revenue of the auctioneer. It is analytically intractable to explicitly characterize the expected revenue of most non-truthful auctions at equilibrium, even in simple settings such as asymmetric single-item first-price auctions (see e.g. Kirkegaard, 2014).

Remarkably, Hartline et al. (2014) showed that in “single-parameter” settings, where the private information of a player boils down to a single number, a variant of the smoothness framework can be used to prove lower bounds on the expected auction revenue at a Bayes-Nash equilibrium, relative to the maximum-possible expected revenue. They formulated a stronger version of smoothness, asserting a smoothness-type inequality on every player individually (rather than in aggregate). They showed that this stronger condition implies approximate revenue guarantees in several settings, for example when there is no bidder with a unique valuation distribution, and when the auction is augmented with the appropriate reserve price. This implies, for example, that the expected revenue at a Bayes-Nash equilibrium of an asymmetric first-price single-item auction with an appropriate reserve price is at least $\frac{1}{4}$ times the maximum expected revenue.

8.7 Algorithmic Characterizations of Smoothness

Smooth auctions have many advantages, including composition theorems, and extension theorems for incomplete-information and no-regret learning settings. Which allocation algorithms and payment rules yield smooth auctions?

Several papers provide sufficient (but not necessary) algorithmic conditions for smoothness. Lucier and Borodin (2010) showed that if the allocation is based on a large class of greedy c -approximation algorithms with a “loser-independence” property, then coupling it with a first- or second-price payment scheme leads to a smooth auction with price of anarchy $\Theta(c)$. Lucier and Syrgkanis (2015) gave a more general characterization along these lines: if an auction can be viewed as running a greedy algorithm in some abstract “element space,” subject to matroid or poly-matroid constraints, then coupling it with a first-price payment scheme leads to a smooth auction. Examples of auctions that fall into this characterization are simultaneous item auctions, uniform-price auctions and position auctions, even in the presence of externalities (Roughgarden & Tardos, 2012). Babaioff et al. (2014a) proved that in a combinatorial auction setting with gross substitute valuations, maximizing welfare with respect to the reported valuations and charging suitable payments yields a smooth auction. Dütting et al. (2015) showed that if an allocation algorithm is based on a technique in approximation algorithms known as “relax-and-round,” then coupling it with suitable payments yields a smooth auction.

Dütting and Kesselheim (2015) gave a necessary and sufficient condition for smoothness in single-parameter settings. This characterization enables impossibility results for smooth mechanisms that are independent of any computational concerns. Dütting and Kesselheim (2015) defined the *permeability* of an allocation algorithm and showed that, for every single-parameter setting, there exists a smooth auction with approximation guarantee c if and only if there exists an allocation algorithm for the setting with permeability $\Theta(c)$. Variants of permeability were used earlier as a sufficient condition for smoothness in the works of

Syrkkanis and Tardos (2013) and Hartline et al. (2014). Dütting et al. (2015) showed how their characterization yields impossibility results, for example for combinatorial auctions with single-minded bidders.

8.8 Valuations with Complementarities

Most of the results presented in this survey assume that the valuation function of each player exhibits no complementarities. The impossibility results in Section 7 show that some assumption of this type is necessary for good price-of-anarchy bounds. A natural goal is to prove welfare guarantees for auctions that degrade gracefully with the “degree of complementarity” of players’ valuations. Such an analysis was done for computationally efficient truthful mechanisms in the work of Abraham, Babaioff, Dughmi, and Roughgarden (2012), and for the equilibria of non-truthful auctions in the work of Feige et al. (2015). The results in the work of Feige et al. (2015) include, for example, an extension of the composition theorem in Section 6 (Theorem 6.1) beyond complement-free valuations. Feldman, Friedler, Morgenstern, and Reiner (2016) considered the simple, single bid auction framework of Devanur et al. (2015) in the context of player valuations with some degree of complementarity and showed that its inefficiency degrades smoothly with this degree.

8.9 The Price of Anarchy in Large Markets

A classic economic intuition is that as a market grows large, the effect of each player on the market and hence the opportunities for strategic behavior should diminish, leading to more efficient outcomes. Making this idea rigorous is non-trivial. Swinkels (2001) showed that in a uniform-price auction, where there is generally inefficiency in small markets, the equilibrium welfare converges to the optimal welfare as the number of players and units of the good go to infinity, under a noise assumption on the arrival of players or units. Feldman, Immorlica, Lucier, Roughgarden, and Syrgkanis (2016) extended these results to combinatorial markets with arbitrary bidder valuations across different goods and to the simultaneous uniform-price auction, offering an adaptation of the smoothness framework that can incorporate large market assumptions. Cole and Tao (2015) subsequently provided similar full-efficiency, large market results for the “Walrasian mechanism” with gross substitute valuations. One benefit of using a smoothness approach to prove such large market results is that it can prove full efficiency guarantees even when players do not behave straightforwardly in the limit (e.g., when players cannot fully express their valuations through their bids).

8.10 Complexity of Computing an Equilibrium

Another line of work addresses equilibrium computation in simple auctions. Cai and Papadimitriou (2014) analyzed the complexity of computing a Bayes-Nash equilibrium in simultaneous second-price auctions and showed that it is a computationally hard problem. Such computational intractability raises questions about the predictive power of the equilibrium concept.

When a Nash or Bayes-Nash equilibrium is hard to compute, no-regret learning often comes to the rescue. When the strategy space of each player is part of the input, the players can use no-regret algorithms (as discussed in Section 5) to approximately compute learning

outcomes, also called “coarse correlated equilibrium” (which is closely related to Definition 5.1) in polynomial time. The price-of-anarchy guarantees proved using smoothness arguments apply to such equilibria (Section 5).

With implicitly defined exponential-size strategy spaces, as in simultaneous item auctions with many items, it is an open question whether polynomial-time decentralized dynamics that lead to no-regret outcomes exist. On the negative side, Daskalakis and Syrgkanis (2016) showed that, under appropriate complexity assumptions, there are no polynomial-time algorithms that guarantee no regret in the worst case. Conceivably, one could dodge this hardness result by designing coordinated dynamics for all of the players (to avoid the hard instances of the learning problem). Daskalakis and Syrgkanis (2016) complemented their impossibility result by proposing a relaxed version of no-regret learning, called “no-envy learning.” They showed that there exist polynomial-time no-envy learning algorithms, and that most price-of-anarchy bounds proved via smoothness arguments extend to outcome sequences generated by no-envy learners.

Motivated by equilibrium tractability, Devanur et al. (2015) seek combinatorial auctions that simultaneously have good price-of-anarchy bounds and also a very small strategy space. They proposed an algorithm where the strategy space of each player is a single number, with equilibrium welfare at least an inverse logarithmic (in the number of items) fraction of the maximum possible. Braverman, Mao, and Weinberg (2016) showed that a large class of auctions cannot improve over this logarithmic bound.

Finally, Christodoulou et al. (2016a) and Dobzinski, Fu, and Kleinberg (2015) considered the complexity of computing a complete-information pure Nash equilibrium of simultaneous second-price auctions. In simultaneous first-price auctions, such equilibria correspond to Walrasian equilibria (Bikhchandani, 1999), and can therefore be computed using linear programming (whenever one exists).

8.11 No-Regret Learning in Dynamic Games

In discussing no-regret learning in games in Section 5 we analyzed learning outcomes in auctions, under the assumptions that all participants use no-regret learning, and that the underlying environment and player population are fixed. In typical online environments, however, bidders need to constantly adjust to changes in the environment, e.g., in the population of competitors. Lykouris, Syrgkanis, and Tardos (2016) studied learning outcomes in a repeated game with dynamically changing player population. They considered dynamic games where at each step a small fraction of the players is replaced by (arbitrary) new participants. They showed that learning players can achieve close to optimal social welfare in smooth mechanisms even in dynamically changing environments. To achieve this, the learning method used needs to guarantee low regret against a shifting comparator, meaning a sequence of bids that stays constant except for a small number of changes.

8.12 Bayesian No-Regret Learning and Bayes-Coarse Correlated Equilibria

Section 5 analyzed no-regret learning when players’ valuations remain fixed over time. Recent work of Hartline et al. (2015) showed that the efficiency guarantees of a smooth auction hold even when each player’s valuation is drawn anew at each iteration, independently from some distribution \mathcal{F}_i . The first step of the argument shows that if all players use no-regret

learning algorithms in such an environment, then the empirical distribution of joint play converges to an analog of the coarse correlated equilibrium for games of incomplete information (“strategic form Bayes-coarse correlated equilibrium”). The second step shows that the approximate efficiency guarantees of smooth auctions apply even to such equilibria. Caragiannis et al. (2015) had previously showed such a result only under the stronger condition of smoothness via private deviations (Definition 4.5).

8.13 Signaling and Bayes-Correlated Equilibria

This survey discussed only private value settings where each player knows her own valuation and nothing else. Caragiannis et al. (2015) show that for auctions that are smooth via private deviations, the approximate welfare guarantees hold even if players receive arbitrary signals about others’ valuations prior to bidding. The set of outcomes that can arise in such a setting is closely related to the notion of a Bayes-correlated equilibrium (Bergemann & Morris, 2011, 2013). Specifically, this result implies that in a (λ, μ) -smooth auction via private deviations, every Bayes-correlated equilibrium where players know at least their own valuation achieves expected welfare at least $\frac{\lambda}{\mu}$ times the expected optimal welfare.

9. Open Questions and Research Directions

We conclude with a dozen suggestions for future research.

9.1 Stronger Price-of-Anarchy Bounds for Common Auction Formats

This survey focused on auction formats for which the price of anarchy is relatively well understood, but some open questions remain even for these auctions.

1. The price of anarchy of first-price single-item auctions. What is the exact price of anarchy of asymmetric first-price single-item auctions with independent player valuations? The answer is at least $1 - 1/e$ (Theorem 2.1, due to Syrgkanis and Tardos (2013)) and at most .87 (Hartline et al., 2014).

2. Better simple auctions with submodular valuations. For subadditive player valuations, the communication complexity lower bounds in the work of Roughgarden (2014) (Section 7) imply that simultaneous first-price auctions have the best-possible price of anarchy of any simple auction. With the stronger assumption of submodular bidder valuations (Section 6), the price of anarchy of simultaneous first-price auctions is $1 - 1/e$ (Theorem 6.2), and this is tight in the worst case (Christodoulou et al., 2016b). The state-of-the-art in communication complexity only implies that there is no simple auction with price of anarchy better than $1 - 1/2e$ (Dobzinski & Vondrak, 2013). Intriguingly, there is a low-cost communication protocol that approximates the welfare to within a factor (slightly) larger than $1 - 1/e$ (Feige & Vondrak, 2010). Is there an analogously good simple auction? What is the best price of anarchy achievable by a simple auction with submodular player valuations?

3. Explicit impossibility results. In some non-auction settings, such as in routing and congestion games, there is a “generic” construction that always produces examples with the worst-possible price of anarchy (Roughgarden, 2015). Is there a similar generic construction for some family of auctions, for example simultaneous first-price auctions, that

generalizes the examples in the works of Christodoulou et al. (2015a) and Feldman et al. (2016)? Impossibility results derived from communication complexity (Section 7) do not seem to lead to such a construction.

9.2 Price-of-Anarchy Bounds for Other Auction Formats

There are a number of practically relevant auction formats that have been understudied from a price-of-anarchy perspective.

4. Procurement auctions. The price-of-anarchy literature has focused on forward auctions, where buyers with private valuations compete for resources in an auction. In many applications, such as the procurement of energy from power firms, the situation is the opposite: a set of sellers with private costs compete to provide some service to one or more buyers who are running an auction to obtain the service as cheaply as possible. Simple non-truthful procurement auctions are common in practice, and not much is known about their price of anarchy. See the work of Babaioff, Nisan, and Paes Leme (2014b) for some initial results in this direction.

5. Double auctions. Even less is known about simple double auctions, where an intermediary runs an auction that has both buyers and sellers with private valuations. The classic impossibility result of Myerson and Satterthwaite (1983) for worst-case instances suggests that assumptions are needed for positive results. A line of work in economics (Rustichini, Satterthwaite, & Williams, 1994; Satterthwaite & Williams, 2002; Cripps & Swinkels, 2006) showed convergence to full efficiency at equilibrium of simple auctions as the market grows large. Does the smoothness framework for large markets proposed in the work of Feldman et al. (2016) extend to large double auctions, ideally unifying all such previous results in a single analysis framework?

6. The combinatorial clock auction. Iterative auctions, which take place in rounds rather than in a single shot, pose a particular challenge for price-of-anarchy analyses. For example, the iterative combinatorial clock auction (Ausubel, Cramton, & Milgrom, 2006) is widely used in practice for selling wireless spectrum. This auction contains both sequential and simultaneous bidding elements, which complicates the analysis of its price of anarchy. Bousquet, Cai, Hunkenschrder, and Vetta (2016) recently provided a theoretical analysis of the efficiency of the auction when all players act truthfully. Extending this approximate welfare guarantee to equilibria appears highly non-trivial.

7. Mechanisms without money. In many of the “killer applications” of mechanism design, including kidney exchange and residency matching, monetary transfers are prohibited. These are examples of “mechanism design without money” (Schummer & Vohra, 2007; Procaccia & Tennenholtz, 2009). Can we quantify the equilibrium inefficiency of simple mechanisms in such settings, perhaps with a smoothness-type framework? In many of these applications, the first challenge is to identify a well-motivated objective function for which price-of-anarchy bounds might be possible.

9.3 Richer Utility Models

Most of the results in this survey concern bidders with independent private values and with quasi-linear utility functions. To what extent can these assumptions be relaxed?

8. Restricted correlation of private values. Equilibria of smooth auctions can have very low welfare when players' valuations are arbitrarily correlated (Bhawalkar & Roughgarden, 2011; Feldman et al., 2013). Are there natural forms of valuation correlation for which smooth auctions are still guaranteed to have near-optimal equilibria? For example, *affiliation* is a strong form of positive correlation that has unlocked several results in economics (Milgrom & Weber, 1982). Can we bound the inefficiency of simple multi-item auctions with some form of affiliated private valuations?

9. The price of anarchy with risk-averse players. Classical auction theory concerns risk-neutral bidders, while in practice many bidders are risk-averse (all else being equal, preferring low-variance outcomes to high-variance ones). Risk aversion poses a host of problems to the analysis of auctions, and only partial results are known in economics on understanding equilibria in this case (see Maskin & Riley, 1984, and several follow-up works). Can we at least bound the equilibrium efficiency of a first-price single-item auction with risk-averse bidders? See the work of Fu, Hartline, and Hoy (2013) for some results on revenue of such auctions and the work of Lianas, Nikolova, and Stier-Moses (2015) for some results along these lines in non-auction domains, such as selfish routing networks.

9.4 Generalizing the Smoothness Framework

The smoothness framework presented in this survey is already rather general, but good researchers are greedy and always want more.

10. Beyond no-regret dynamics. The smoothness approach directly extends to adaptive game-playing when the vanishing regret condition holds (Section 5). There are a number of interesting adaptive game-playing algorithms that are not guaranteed to achieve this condition. One famous example is *fictitious play* (Brown, 1951), which is only known to converge to an equilibrium (and therefore satisfy the vanishing regret condition) in some special cases, such as zero-sum games (Robinson, 1951), potential games (Monderer & Shapley, 1996), and two-player games where one player has only two strategies (Berger, 2005). Does fictitious play achieve the vanishing regret property in smooth auctions? Is the average welfare of fictitious play over time guaranteed to be close to the maximum welfare in smooth auctions? These questions are also relevant for many other forms of adaptive game-playing. Recently, Dütting and Kesselheim (2017) provided results for best-response dynamics in the context of simultaneous item auctions. See also the work of Fudenberg and Levine (1998) for a starting point for other dynamics.

11. Characterizations of smoothness in multi-dimensional domains. Dütting and Kesselheim (2015) provided the first algorithmic characterization of smoothness in single-dimensional mechanism design domains. An obvious open question is whether or not there is an analogous characterization for multi-dimensional domains like multi-item auctions. A related direction is to understand the extent to which the approximation achievable by

low-cost communication protocols (Section 7) characterizes the best-possible guarantee of a simple smooth mechanism.

12. The price of anarchy for revenue in multi-dimensional settings. Hartline et al. (2014) adapted the smoothness framework to prove bounds on the revenue of simple auctions in single-dimensional environments. Can we obtain interesting guarantees for the revenue of simple mechanisms for multi-dimensional environments? Key to the results of Hartline et al. (2014) is the equivalence of expected revenue and expected “virtual welfare” in single-parameter settings (Myerson, 1981). While this equivalence does not hold in multi-parameter environments, there are cases where a version of virtual welfare well-approximates the revenue of an auction (Chawla, Hartline, Malec, & Sivan, 2010). Similar connections between revenue maximization and virtual welfare maximization were also given in the works of Haghpanah and Hartline (2014) and Cai, Devanur, and Weinberg (2016) for more general multidimensional settings. Can we use these connections and an adaptation of smoothness to bound the revenue of other simple auctions?

References

- Abraham, I., Babaioff, M., Dughmi, S., & Roughgarden, T. (2012). Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC '12*, pp. 3–16, New York, NY, USA. ACM.
- Amann, E., & Leininger, W. (1996). Asymmetric all-pay auctions with incomplete information: The two-player case. *Games and Economic Behavior*, 14(1), 1 – 18.
- Auer, P., Cesa-Bianchi, N., Freund, Y., & Schapire, R. E. (1995). Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science, FOCS '95*, pp. 322–, Washington, DC, USA. IEEE Computer Society.
- Ausubel, L., Cramton, P., & Milgrom, P. (2006). The clock-proxy auction: A practical combinatorial auction design. In Cramton, P., Shoham, Y., & Steinberg, R. (Eds.), *Combinatorial Auctions*, chap. 5, pp. 115–138. MIT Press.
- Babaioff, M., Lucier, B., Nisan, N., & Paes Leme, R. (2014a). On the efficiency of the Walrasian mechanism. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation, EC '14*, pp. 783–800, New York, NY, USA. ACM.
- Babaioff, M., Nisan, N., & Paes Leme, R. (2014b). Price competition in online combinatorial markets. In *Proceedings of the 23rd International Conference on World Wide Web, WWW '14*, pp. 711–722, New York, NY, USA. ACM.
- Baye, M., Kovenock, D., & de Vries, C. (1996). The all-pay auction with complete information. *Economic Theory*, 8(2), 291–305.
- Bergemann, D., & Morris, S. (2011). Correlated Equilibrium in Games with Incomplete Information. Cowles foundation discussion papers 1822, Cowles Foundation for Research in Economics, Yale University.
- Bergemann, D., & Morris, S. (2013). Robust predictions in games with incomplete information. *Econometrica*, 81(4), 1251–1308.

- Berger, U. (2005). Fictitious play in $2 \times n$ games. *Journal of Economic Theory*, 120(2), 139–154.
- Bhawalkar, K., & Roughgarden, T. (2011). Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pp. 700–709. SIAM.
- Bikhchandani, S. (1999). Auctions of heterogeneous objects. *Games and Economic Behavior*, 26(2), 193 – 220.
- Bousquet, N., Cai, Y., Hunkenschrder, C., & Vetta, A. (2016). On the economic efficiency of the combinatorial clock auction. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'16, pp. 1407–1423.
- Braverman, M., Mao, J., & Weinberg, S. M. (2016). Interpolating between truthful and non-truthful mechanisms for combinatorial auctions. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pp. 1444–1457.
- Brown, G. W. (1951). Iterative solution of games by fictitious play. *Activity analysis of production and allocation*, 13(1), 374–376.
- Cai, Y., Devanur, N. R., & Weinberg, S. M. (2016). A duality based unified approach to bayesian mechanism design. In *Symposium of the Theort of Computing (STOC 2016)*, pp. 926–939.
- Cai, Y., & Papadimitriou, C. (2014). Simultaneous Bayesian auctions and computational complexity. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pp. 895–910, New York, NY, USA. ACM.
- Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M., Lucier, B., Paes Leme, R., & Tardos, É. (2015). On the efficiency of equilibria in generalized second price auctions. *Journal of Economic Theory*, 156, 343–388.
- Caragiannis, I., & Voudouris, A. (2014). Welfare guarantees for proportional allocations. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory*, SAGT '14, pp. 206–217.
- Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA.
- Chawla, S., & Hartline, J. D. (2013). Auctions with unique equilibria. In *ACM Conference on Electronic Commerce*, pp. 181–196.
- Chawla, S., Hartline, J. D., Malec, D. L., & Sivan, B. (2010). Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, STOC '10, pp. 311–320, New York, NY, USA. ACM.
- Chawla, S., Hartline, J. D., & Sivan, B. (2012). Optimal crowdsourcing contests. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 856–868. SIAM.
- Christodoulou, G., Kovács, A., & Schapira, M. (2016a). Bayesian combinatorial auctions. *Journal of the ACM*, 63(2), 11.

- Christodoulou, G., Kovács, A., Sgouritsa, A., & Tang, B. (2016b). Tight bounds for the price of anarchy of simultaneous first price auctions. *ACM Transactions on Economics and Computation*, 4(2), 9.
- Christodoulou, G., Sgouritsa, A., & Tang, B. (2015a). On the efficiency of all-pay mechanisms. In *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pp. 349–360.
- Christodoulou, G., Sgouritsa, A., & Tang, B. (2015b). On the efficiency of the proportional allocation mechanism for divisible resources. In *Algorithmic Game Theory - 8th International Symposium, SAGT 2015, Saarbrücken, Germany, September 28-30, 2015, Proceedings*, pp. 165–177.
- Cole, R., & Tao, Y. (2015). When does the price of anarchy tend to 1 in large Walrasian auctions and Fisher markets?. *CoRR*, abs/1508.07370.
- Cook, W. J., Cunningham, W. H., Pulleyblank, W. R., & Schrijver, A. (1997). *Combinatorial Optimization*. John Wiley & Sons, Inc.
- Cripps, M. W., & Swinkels, J. M. (2006). Efficiency of large double auctions. *Econometrica*, 74(1), 47–92.
- Daskalakis, C., & Syrgkanis, V. (2016). Regret is hard, envy is easy. In *57th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*.
- de Keijzer, B., Markakis, E., Schfer, G., & Telelis, O. (2013). Inefficiency of standard multi-unit auctions. In Bodlaender, H., & Italiano, G. (Eds.), *Algorithms ESA 2013*, Vol. 8125 of *Lecture Notes in Computer Science*, pp. 385–396. Springer Berlin Heidelberg.
- Demange, G., Gale, D., & Sotomayor, M. (1986). Multi-item auctions. *Journal of Political Economy*, 94(4), pp. 863–872.
- Devanur, N., Morgenstern, J., Syrgkanis, V., & Weinberg, S. M. (2015). Simple auctions with simple strategies. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pp. 305–322, New York, NY, USA. ACM.
- DiPalantino, D., & Vojnovic, M. (2009). Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pp. 119–128. ACM.
- Dobzinski, S., Nisan, N., & Schapira, M. (2010). Approximation algorithms for combinatorial auctions with complement-free bidders. *Mathematics of Operations Research*, 35(1), 1–13.
- Dobzinski, S., & Vondrak, J. (2013). Communication complexity of combinatorial auctions with submodular valuations. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1205–1215.
- Dobzinski, S., Fu, H., & Kleinberg, R. D. (2015). On the complexity of computing an equilibrium in combinatorial auctions. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pp. 110–122.
- Dobzinski, S., & Paes Leme, R. (2014). Efficiency guarantees in auctions with budgets. In Esparza, J., Fraigniaud, P., Husfeldt, T., & Koutsoupias, E. (Eds.), *Automata*,

- Languages, and Programming*, Vol. 8572 of *Lecture Notes in Computer Science*, pp. 392–404. Springer Berlin Heidelberg.
- Dütting, P., Henzinger, M., & Starnberger, M. (2013). Valuation compressions in VCG-based combinatorial auctions. In *Web and Internet Economics - 9th International Conference, WINE 2013, Cambridge, MA, USA, December 11-14, 2013, Proceedings*, pp. 146–159.
- Dütting, P., & Kesselheim, T. (2015). Algorithms against anarchy: Understanding non-truthful mechanisms. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pp. 239–255.
- Dütting, P., & Kesselheim, T. (2017). Best-response dynamics in combinatorial auctions with item bidding. In *Proceeding of 28th ACM-SIAM Symposium on Discrete Algorithms (SODA 2017)*. SIAM.
- Dütting, P., Kesselheim, T., & Tardos, É. (2015). Algorithms as mechanisms: The price of anarchy of relax-and-round. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pp. 187–201.
- Engelbrecht-Wiggans, R., & Weber, R. J. (1979). An example of a multi-object auction game. *Management Science*, 25(12), pp. 1272–1277.
- Feige, U., Feldman, M., Immorlica, N., Izsak, R., Lucier, B., & Syrgkanis, V. (2015). A unifying hierarchy of valuations with complements and substitutes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 872–878.
- Feige, U., & Vondrak, J. (2010). The submodular welfare problem with demand queries. *Theory of Computing*, 6(1), 247–290.
- Feldman, M., Friedler, O., Morgenstern, J., & Reiner, G. (2016). Simple mechanisms for agents with complements. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC'16)*, pp. 251–267.
- Feldman, M., Fu, H., Gravin, N., & Lucier, B. (2013). Simultaneous auctions are (almost) efficient. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC '13*, pp. 201–210, New York, NY, USA. ACM.
- Feldman, M., Immorlica, N., Lucier, B., Roughgarden, T., & Syrgkanis, V. (2016). The price of anarchy in large games. In *Proceedings of the 48th symposium on Theory of Computing, STOC '16*, pp. 963–976.
- Freund, Y., & Schapire, R. E. (1999). Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1), 79–103.
- Fu, H., Hartline, J. D., & Hoy, D. (2013). Prior-independent auctions for risk-averse agents. In *ACM Conference on Electronic Commerce, EC '13, Philadelphia, PA, USA, June 16-20, 2013*, pp. 471–488.
- Fudenberg, D., & Levine, D. K. (1998). *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- Gneezy, U., & Smorodinsky, R. (2006). All-pay auctions—an experimental study. *Journal of Economic Behavior & Organization*, 61(2), 255–275.

- Haghpanah, N., & Hartline, J. D. (2014). Multi-dimensional virtual values and second-degree price discrimination. *CoRR*, *abs/1404.1341*.
- Hannan, J. (1957). Approximation to Bayes risk in repeated play. In Dresher, M., Tucker, A. W., & Wolfe, P. (Eds.), *Contributions to the Theory of Games*, Vol. 3, pp. 97–139. Princeton University Press.
- Hart, S., & Mas-Colell, A. (2000). A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, *68*(5), 1127–1150.
- Hartline, J., Hoy, D., & Taggart, S. (2014). Price of anarchy for auction revenue. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pp. 693–710, New York, NY, USA. ACM.
- Hartline, J., Syrgkanis, V., & Tardos, E. (2015). No-regret learning in repeated bayesian games. In *Proceedings of the 29th Annual Symposium on Neural Information Processing Systems*, NIPS '15, pp. 3061–3069.
- Hassidim, A., Kaplan, H., Mansour, Y., & Nisan, N. (2011). Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, pp. 295–296, New York, NY, USA. ACM.
- Johari, R., & Tsitsiklis, J. N. (2004). Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, *29*(3), pp. 407–435.
- Kaplan, T., & Zamir, S. (2012). Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory*, *50*(2), 269–302.
- Kesselheim, T., Kleinberg, R., & Tardos, E. (2015). Smooth online mechanisms: A game-theoretic problem in renewable energy markets. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pp. 203–220, New York, NY, USA. ACM.
- Kirkegaard, R. (2014). Ranking asymmetric auctions: Filling the gap between a distributional shift and stretch. *Games and Economic Behavior*, *85*, 60–69.
- Koutsoupias, E., & Papadimitriou, C. H. (1999). Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, Vol. 1563 of *Lecture Notes in Computer Science*, pp. 404–413.
- Krishna, V. (2002). *Auction Theory*. Academic Press.
- Kulkarni, J., & Mirrokni, V. (2015). Robust price of anarchy bounds via LP and Fenchel duality. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pp. 1030–1049. SIAM.
- Kushilevitz, E., & Nisan, N. (1996). *Communication Complexity*. Cambridge.
- Lehmann, B., Lehmann, D., & Nisan, N. (2001). Combinatorial auctions with decreasing marginal utilities. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, EC '01, pp. 18–28, New York, NY, USA. ACM.
- Lianas, T., Nikolova, E., & Stier-Moses, N. E. (2015). Asymptotically tight bounds for inefficiency in risk-averse selfish routing. arXiv:1510.02067.

- Littlestone, N., & Warmuth, M. K. (1994). The weighted majority algorithm. *Information and computation*, 108(2), 212–261.
- Lucier, B., & Borodin, A. (2010). Price of anarchy for greedy auctions. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, pp. 537–553, Philadelphia, PA, USA. Society for Industrial and Applied Mathematics.
- Lucier, B., & Paes Leme, R. (2011). GSP auctions with correlated types. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, pp. 71–80, New York, NY, USA. ACM.
- Lucier, B., & Syrgkanis, V. (2015). Greedy algorithms make efficient mechanisms. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pp. 221–238, New York, NY, USA. ACM.
- Lykouris, T., Syrgkanis, V., & Tardos, É. (2016). Learning and efficiency in games with dynamic population. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 120–129. SIAM.
- Markakis, E., & Telelis, O. (2012). Uniform price auctions: Equilibria and efficiency. In Serna, M. (Ed.), *Algorithmic Game Theory*, Lecture Notes in Computer Science, pp. 227–238. Springer Berlin Heidelberg.
- Maskin, E., & Riley, J. (1984). Optimal auctions with risk averse buyers. *Econometrica*, 52(6), pp. 1473–1518.
- Maskin, E., & Riley, J. (2000). Asymmetric auctions. *The Review of Economic Studies*, 67(3), pp. 413–438.
- Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press.
- Milgrom, P. R., & Weber, R. J. (1982). A Theory of Auctions and Competitive Bidding. *Econometrica*, 50(5), 1089–1122.
- Milgrom, P. R., & Weber, R. J. (1999). A Theory of Auctions and Competitive Bidding, II. In Klemperer, P. (Ed.), *The Economic Theory of Auctions*. Edward Elgar Publishing.
- Monderer, D., & Shapley, L. S. (1996). Fictitious play property for games with identical interests. *Journal of Economic Theory*, 68(1), 258–265.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1), 58–73.
- Myerson, R. B., & Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2), 265 – 281.
- Paes Leme, R., Syrgkanis, V., & Tardos, E. (2012). Sequential auctions and externalities. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pp. 869–886. SIAM.
- Paes Leme, R., & Tardos, E. (2010). Pure and Bayes-Nash price of anarchy for generalized second price auction. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, FOCS '10, pp. 735–744, Washington, DC, USA. IEEE Computer Society.

- Procaccia, A. D., & Tennenholtz, M. (2009). Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pp. 177–186. ACM.
- Robinson, J. (1951). An iterative method of solving a game. *Annals of Mathematics*, 54(2), 296–301.
- Roughgarden, T. (2015). Intrinsic robustness of the price of anarchy. *Journal of the ACM*, 62(5), 32.
- Roughgarden, T. (2016). Communication complexity (for algorithm designers). *Foundations and Trends in Theoretical Computer Science*, 11(3-4), 217–404.
- Roughgarden, T. (2012). The price of anarchy in games of incomplete information. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC '12*, pp. 862–879.
- Roughgarden, T. (2014). Barriers to near-optimal equilibria. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pp. 71–80.
- Roughgarden, T. (2016). *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press.
- Roughgarden, T., & Tardos, E. (2012). Do externalities degrade GSP’s efficiency?. In *Eighth Ad Auctions Workshop*.
- Rustichini, A., Satterthwaite, M. A., & Williams, S. R. (1994). Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 62(5), pp. 1041–1063.
- Satterthwaite, M. A., & Williams, S. R. (2002). The optimality of a simple market mechanism. *Econometrica*, 70(5), pp. 1841–1863.
- Schummer, J., & Vohra, R. V. (2007). Mechanism design without money. In Nisan, N., Roughgarden, T., Tardos, É., & Vazirani, V. (Eds.), *Algorithmic Game Theory*, chap. 10, pp. 243–299. Cambridge University Press.
- Shapley, L. S., & Shubik, M. (1971). The assignment game I: The core. *International Journal of Game Theory*, 1(1), 111–130.
- Siegel, R. (2014). Asymmetric all-pay auctions with interdependent valuations. *Journal of Economic Theory*, 153, 684 – 702.
- Swinkels, J. M. (2001). Efficiency of large private value auctions. *Econometrica*, 69(1), pp. 37–68.
- Syrngkanis, V. (2012). Bayesian games and the smoothness framework. *CoRR*, abs/1203.5155.
- Syrngkanis, V. (2014). *Efficiency of Mechanisms in Complex Markets*. Ph.D. thesis, Cornell University.
- Syrngkanis, V., & Tardos, E. (2012). Bayesian sequential auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC '12*, pp. 929–944, New York, NY, USA. ACM.

- Syrnganis, V., & Tardos, E. (2013). Composable and efficient mechanisms. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC '13*, pp. 211–220, New York, NY, USA. ACM.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1), 8–37.