

Simple Auctions with Simple Strategies

NIKHIL DEVANUR, Microsoft Research
JAMIE MORGENSTERN, Carnegie Mellon University
VASILIS SYRGKANIS, Microsoft Research
S. MATTHEW WEINBERG, Princeton University

We introduce single-bid auctions as a new format for combinatorial auctions. In single-bid auctions, each bidder submits a single real-valued bid for the right to buy items at a fixed price. Contrary to other simple auction formats, such as simultaneous or sequential single-item auctions, bidders can implement no-regret learning strategies for single-bid auctions in polynomial time. Price of anarchy bounds for correlated equilibria concepts in single-bid auctions therefore have more bite than their counterparts for auctions and equilibria for which learning is not known to be computationally tractable (or worse, known to be computationally intractable [Cai and Papadimitriou 2014; Dobzinski et al. 2015] this end, we show that for any subadditive valuations the social welfare at equilibrium is an $O(\log m)$ -approximation to the optimal social welfare, where m is the number of items. We also provide tighter approximation results for several subclasses. Our welfare guarantees hold for Nash equilibria and no-regret learning outcomes in both Bayesian and complete information settings via the smooth-mechanism framework. Of independent interest, our techniques show that in a combinatorial auction setting, efficiency guarantees of a mechanism via smoothness for a very restricted class of *cardinality* valuations extend, with a small degradation, to subadditive valuations, the largest complement-free class of valuations.

Categories and Subject Descriptors: F.2.0 [Analysis of Algorithms and Problem Complexity]: General
General Terms: Mechanism Design, Algorithms

Additional Key Words and Phrases: Mechanism design; auction design

1. INTRODUCTION

The design of combinatorial auctions has been a major topic of interest in economics as well as algorithmic game theory. The desiderata of efficiency (maximizing social welfare), incentive compatibility, computability (of both the mechanism outcome and equilibrium behavior for the participants), and simplicity are often not simultaneously achievable and hence different mechanisms offer different tradeoffs. The VCG mechanism achieves optimal efficiency and is incentive compatible but its outcome is, in general, not computable in polynomial time. The past decade has seen much success in the design of polynomial time computable truthful mechanisms that obtain a *fraction* of the optimal social welfare, and many obtain ratios that are quite good [Balcan et al. 2008; Dughmi et al. 2011; Dobzinski et al. 2006; Feige et al. 2011; Dobzinski et al. 2005; Lavi and Swamy 2005; Dobzinski 2007; Krysta and Vöcking 2012]. Unfortunately, most known mechanisms guaranteeing a $m^{o(1)}$ -approximation (where m

Morgenstern was supported in part by Simons Graduate Fellowship in Theoretical Computer Science, as well as NSF grants CCF-1101215 and CCF-1415460. Part of work done while an intern at Microsoft Research, Redmond. Syrgkanis supported in part by Simons Graduate Fellowship in Theoretical Computer Science. Part of work done while an intern at Microsoft Research, New England.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'15, June 15–19, 2015, Portland, OR, USA.

ACM 978-1-4503-3410-5/15/06 ...\$15.00.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

<http://dx.doi.org/10.1145/2764468.2764484>

is the number of items) are quite sophisticated, prohibiting them from broad practical use: mechanisms should be easy for designers to implement and transparent to its participants.

In practice, the desire for simplicity seems to trump the desire for truthfulness, and auctions such as *sequential item auctions*, where bidders submit sealed bids for each item one at a time, or *simultaneous item auctions*, where each bidder instead reports a separate sealed bid for each item all at once, are commonly used. A recent collection of exciting work has proven that several simple auctions allocate items approximately efficiently at various equilibrium concepts, further supporting their use [Bhawalkar and Roughgarden 2011; Leme et al. 2012; Feldman et al. 2013a; Syrgkanis and Tardos 2012]. However, a notable drawback of these auction formats is that *none* of these equilibrium concepts is known to be computationally tractable and some are even known to be computationally *intractable* [Cai and Papadimitriou 2014; Dobzinski et al. 2015]. In other words, while these auctions are indeed simple from a design perspective, they are still quite complex from a strategic one. This realization motivates the search for auctions with a *simple design* that are also *strategically simple* and have a *low price of anarchy*.

As a notion of being strategically simple, we propose the concept of *learnability*: suppose that the same auction is repeated many times, then a bidder should be able to *efficiently* run a *no-regret* algorithm on his strategy space. (If each bidder does so, then the joint empirical distribution of strategies converges to a correlated equilibrium.) This does not require the bidder to have any information about his opponents; this is in a sense a “zero information” setting, as opposed to the standard notions of complete and incomplete information settings. Both sequential and simultaneous auctions have strategy spaces of size \mathbb{R}^m which is exponential in the number of items. Therefore, the straightforward implementation of no-regret algorithms for these games would take exponential time and space, and it is not known whether computationally efficient no-regret algorithms exist for these auctions.¹

Our Results. In this paper, we introduce an extremely simple auction format that we call a *single-bid auction*. In a single-bid auction, bidders submit a single real number as their bid. They are then visited in decreasing order of bids, and each may pay their bid per item for any number of remaining items. Below is a formal description.

- (1) Initialize, for all $i \in [n]$, $S_i = \emptyset, P_i = 0$. The set of remaining items $I = [m]$.
- (2) Each bidder $i \in [n]$ submits a sealed bid b_i .
- (3) Sort bidders in decreasing order according to their bids. Break ties arbitrarily.
- (4) For $i = 1$ to n :
 - (5) Let i be the i^{th} highest bidder.
 - (6) Bidder i chooses any set $X_i \subseteq I$.
 - (7) Bidder i pays her bid for each item in X_i , i.e., $P_i = b_i |X_i|$.
 - (8) Update $I = I \setminus X_i$.
- (9) End For.

Importantly, the strategy space of single-bid auctions is simple enough so that one can efficiently deploy no-regret algorithms. The strategic choices in a single-bid auction consist of making a bid (Step 2) and selecting a set of remaining items to purchase (Step 6). The space of possible bids is just \mathbb{R} , but the space of possible sets to choose is still exponential in m . However, a bidder’s dominant strategy in Step 6 is extremely

¹Cai and Papadimitriou [2014] have some discussion of no-regret dynamics in simultaneous second price auctions that indicates that such algorithms may not exist.

simple: once the auction reaches this phase, bidder i faces an item pricing and may pay b_i for any item in I and has nothing to gain by selecting any set besides $X_i = \operatorname{argmax}_{X \subseteq I} \{v_i(X) - |X|b_i\}$.² In other words, when learning the effectiveness of different strategies, bidder i needs only to learn over different potential bids and not also over potential methods for choosing items to purchase.

The challenge, then, is to show that our auction achieves a good fraction of the optimal social welfare at outcomes of no-regret learning algorithms (or, equivalently, correlated equilibria). Such bounds are called bounds on the *price of anarchy* (PoA), which is the ratio of the optimal social welfare to the welfare at the worst possible equilibrium, for various equilibrium concepts. In a nutshell, we show that for subadditive (a.k.a. complement-free) valuations, the price of anarchy of single-bid auction w.r.t. correlated equilibria is at most $\frac{e}{e-1}H_m$ where m is the number of items and H_m is the m^{th} harmonic number. In comparison, for the same class of valuations, the best deterministic and randomized truthful auctions achieve approximation factors of $O(\sqrt{m})$ and $O(\log m \log \log m)$ respectively³, and simultaneous first price auctions have a price of anarchy of 2 w.r.t. Bayes Nash equilibria. On the other hand, *sequential* item auctions could have a price of anarchy of $\Omega(m)$ even for a much simpler class of valuations, namely the union of additive and unit-demand valuations, even for Nash equilibria in the complete information setting. This indicates that *identifying* the right auction is important in order to get a price of anarchy such as $O(\log m)$.

Our main technical results are stated below:

THEOREM 1.1 (INFORMAL). *There is a polynomial time no-regret learning algorithm for a bidder participating in single-bid auction.*

THEOREM 1.2. *The single-bid auction has a price of anarchy of at most $\frac{e}{e-1}H_m$ w.r.t. coarse correlated equilibria.*

We prove Theorem 1.2 by developing a reduction from proving price of anarchy bounds when bidders have subadditive valuations to proving PoA bounds when bidders have considerably simpler valuations that we call *constraint-homogeneous*. A bidder has constraint-homogeneous valuation if he has an interest set S and the same obtains value v per item in S and 0 per item not in S . This reduction itself may be of independent interest. The proof and formal statement of Theorem 1.2 can be found in Section 3.

We also provide stronger PoA bounds for restricted classes of valuations, such as unit-demand, concave-symmetric, and k -demand valuations in Appendix 4. In Section 3.1, we include a lower bound of $\Omega\left(\frac{\log m}{\log \log m}\right)$ on the possible PoA for single-bid auctions when we have additive bidders.

Finally, in Appendix A we provide PoA bounds for a sequential format of single-bid auctions which we call **draft auctions**. A draft auction proceeds in rounds: each bidder submits a bid in each round. The highest bidder in each round may pick any of the remaining items, and pays her bid for each item she picks. We show an $O(\log m)$ bound on the price of anarchy with subadditive bidders for draft auctions as well. This offers a significant advantage over *sequential* item auctions, for which the price

²For simplicity, we assume throughout that bidder i can find X_i in polynomial time. If not, then bidder i has no incentive not to select the best set of items to purchase that she can find computationally efficiently, and our approximation ratios degrade naturally based on how well the bidders can perform this optimization.

³If $\max_i \{v_i([m])\}$ is known, then the best randomized bound can be improved to $O(\log m)$ [Balcan et al. 2008].

of anarchy for even simple valuation classes such as the union of additive and unit-demand valuations is $\Omega(m)$ [Feldman et al. 2013b].

1.1. Related work

Truthful Auctions. The study of combinatorial auctions has long focused on the design of truthful auctions. Although the VCG mechanism is truthful and gives the socially optimal allocation, it is not computationally efficient. Within the AGT community, this computational barrier has spurred a lively line of research into designing truthful mechanisms that run in polynomial time and approximate the social welfare for various classes of valuations. The state-of-the-art for various instances are: an $O(\log m \log \log m)$ -approximation when bidders are subadditive [Dobzinski 2007], an $O(\log m)$ approximation when bidders are subadditive and $\max_i v_i([m])$ is known [Balcan et al. 2008], an $O(\log m)$ -approximation when bidders are fractionally subadditive [Krysta and Vöcking 2012], and an $\frac{e}{e-1}$ -approximation when bidders have coverage valuations [Dughmi et al. 2011]. These mechanisms (and others) are all quite impressive, but have some drawbacks preventing them from being used in practice, such as being algorithmically complex, having a high probability of achieving 0 welfare, or completely ignoring many participants. In the case of the random single price mechanism [Balcan et al. 2008], the auction requires knowledge (or approximate knowledge) of OPT, and can have probability at least $1 - O\left(\frac{1}{\log m}\right)$ of achieving 0 welfare.

Price of Anarchy. More recently, an alternate approach has been to analyze simple auctions that are commonly used in practice, by quantifying the inefficiency of equilibria via the price of anarchy [Christodoulou et al. 2008; Bhawalkar and Roughgarden 2011; Hassidim et al. 2011; Feldman et al. 2013a; Lucier and Borodin 2010; Leme and Tardos 2010; Lucier and Paes Leme 2011; Caragiannis et al. 2011; Syrgkanis and Tardos 2012; Leme et al. 2012; Feldman et al. 2013b]. The dominating theme here has been the emergence of a “smoothness” framework that captures many of the price of anarchy bounds, and allows these bounds to be extended to larger classes of equilibria: Roughgarden [2009] to outcomes of learning algorithms and Roughgarden [2012] and Syrgkanis [2012] to games of incomplete information. Syrgkanis and Tardos [2013] give a specialized smoothness framework for auctions with quasi-linear preferences, which we also use. The result most directly comparable to ours is that of Feldman et al. [2013a], which shows that simultaneous item auctions have a constant price of anarchy for subadditive bidders. The ratio is, of course, more desirable than ours. However it is unknown how to compute any of the equilibria for which their PoA guarantees hold (even approximately) in polynomial time. So, without further research, it is unclear whether one should expect bidders in simultaneous item auctions to play an (approximate) equilibrium. In contrast, bidders can reach equilibria of single-bid auctions in polynomial time via distributed no-regret learning, so it is quite reasonable to expect strategic play to approach equilibrium.

Equilibrium Computation. We conclude this section by briefly discussing positive and negative results related to equilibrium computation in simple auctions. Lehmann et al. [2001] showed how to efficiently compute a pure Nash equilibrium of simultaneous second-price auctions when bidders are submodular. Unfortunately, the equilibrium computed is quite unnatural: it selects a desired winner for each item and asks them to place a large bid on that item, and for all other bidders to bid 0. Even though their construction finds an equilibrium where the large bids are not “overbids,” it is still clear that this equilibrium is unnatural: it is carefully constructed by a centralized agent with a specific allocation in mind, and it asks bidders to play dominated strategies (why bid 0 if you have any positive value for adding an item?). To our knowledge,

there are no other positive results regarding equilibrium computation in simple auctions. On the negative side, [Cai and Papadimitriou \[2014\]](#) proved that it is PP-hard⁴ to find an exact Bayes-Nash equilibrium in simultaneous second-price auctions with sub-modular bidders, and that it is also NP-hard to find an ϵ -Bayes-Nash for some constant ϵ . They further extend their hardness to a notion of ϵ -Bayes-Coarse-Correlated equilibria, and show that this equilibrium is also NP-hard to find. [Dobzinski et al. \[2015\]](#) also show that computing pure Nash equilibria of simultaneous second-price item auctions requires exponential communication. This line of work suggests that simple auctions with strong PoA bounds which do not explicitly consider equilibrium computation may have less bite in a computationally-constrained world. Our work addresses this concern as bidders can run regret-minimization algorithms in polynomial time. A very recent paper of [Roughgarden \[2014\]](#) explores formal barriers to obtaining mechanisms with low price of anarchy. He shows, via a reduction to a communication complexity, that no mechanism requiring sub-doubly-exponential communication can have a price of anarchy better than 2 for subadditive bidders (achieved by simultaneous first price auctions [[Feldman et al. 2013a](#)]). It would be interesting to see if similar techniques can provide formal barriers to designing mechanisms “like” single-bid auctions with constant PoA.

1.2. Discussion and Future Work

Our paper contributes to the recent line of work addressing the design and analysis of simple combinatorial auctions with low price of anarchy. We propose coarse correlated equilibria of single-bid auctions as the first solution concept of a simple auction that both has a low price of anarchy and can be computed in polynomial time. Our paper also motivates several directions for future work. First, there are numerous questions related to the analysis of existing simple mechanisms, such as simultaneous item auctions. For example, *is there a poly-time no-regret algorithm for simultaneous item auctions?* It is conceivable that such algorithms exist despite the exponential strategy space and evidence found by [Cai and Papadimitriou \[2014\]](#). If not, *is there a different equilibrium concept for simultaneous item auctions that is well-motivated, poly-time computable, and has low price of anarchy?* Or more generally, *how should one expect bidders to behave in a simultaneous item auction? Can one bound the price of anarchy at this behavior?*

Also motivated is the design of new simple auctions with learnable equilibria and a constant price of anarchy. We note that doing so will require very different techniques than the present work in a formal sense. One significant generalization of single-bid auctions is the following: bidders first play an arbitrary game, where each bidder has $\text{poly}(m)$ possible strategies. Then, based on the strategies selected by each bidder, the bidders are visited sequentially (in an order determined by the strategies played) and each offered an item pricing over the remaining items (also determined by the strategies played). Single-bid auctions are an extremely special case of these mechanisms, where the game consists of each bidder simply making a bid, and the item pricing is uniform over the remaining items (and the order of bidders is determined by simply ranking their bids, and the uniform price offered is exactly their bid). Recent work of [Braverman et al. \[2015\]](#) shows that no auction of this general format can possibly have a price of anarchy $o(\log m / \log \log m)$ at any equilibrium concept. Therefore, modifications to single-bid auctions such as asking the bidders to report multiple bids, charging non-uniform prices, etc. cannot possibly improve the price of anarchy. We therefore hope that future research will be fruitful in designing novel simple mechanisms achieving low price of anarchy at learnable equilibrium concepts.

⁴PP is the class “BPP without the B,” and lies somewhere between the polynomial hierarchy and PSPACE.

2. PRELIMINARIES AND NOTATION

Learnability and correlated equilibria. We begin with a review of standard notions from the online learning literature. Suppose there are N actions and T rounds. An online algorithm A selects an action $a^t \in [N]$ (which is in general randomized and is drawn from a distribution, say x^t) in round t . An adversary selects a reward vector $r^t \in [0, h]^N$; r^t is chosen with the knowledge of x^t but not a^t . A receives reward $r_{a^t}^t$. In the *bandit* setting, this is all A learns, as opposed to the *experts* setting, where A learns the entire reward vector. We now define the *regret* of A .

DEFINITION 2.1. *We say that algorithm A achieves **regret** $R(T)$ with respect to an action sequence a'_1, \dots, a'_T if, for all reward vectors $r^1, \dots, r^T \in [0, h]^N$,*

$$\sum_{t=1}^T \mathbb{E}[r_{a'_t}^t - r_{a^t}^t] \leq R(T).$$

*If A achieves regret $R(T)$ with respect to all fixed action sequences ($a'_1 = a'_2 = \dots = a'_T$), we say that A achieves **external regret** of $R(T)$. If A achieves regret $R(T)$ with respect to all action sequences $f(a_1), f(a_2), \dots, f(a_T)$ for some $f : [N] \rightarrow [N]$, we say A achieves **swap regret** of $R(T)$.*

We say an algorithm is a **no-regret** algorithm if it achieves regret $R(T) = o(T)$. We say that a game is **learnable** if in the setting where the same game is repeated many times, each player has a polynomial time learning algorithm that achieves external/swap regret of $o(T^{1-\delta})$ ⁵ over the set of all his strategies.

The single-bid auction induces a multi-player simultaneous move game among all the bidders, where the strategy of bidder i is his bid b_i . A tuple of bids \mathbf{b} determines the outcome of the auction; player i 's utility is $u_i(\mathbf{b}) := v_i(S_i(\mathbf{b})) - P_i(\mathbf{b})$ where $S_i(\mathbf{b})$ is the set of items i wins and $P_i(\mathbf{b})$ is her total payment. Additionally, for any bid tuple \mathbf{b} , we denote with $p_j(\mathbf{b})$ the price that item j was sold at under \mathbf{b} . Finally, let $h = \max_i v_i([m])$ be the the maximum valuation any bidder has for the bundle of all goods (note that we do not require the designer to know h). Players may randomize their strategies, in which case the bids (and everything else that depends on the bids) are random variables.

The standard notion of equilibrium used in such games is that of Nash equilibrium, in which no player can unilaterally deviate from the equilibrium strategy and gain utility for himself. We consider the relaxed notion of *correlated equilibrium*: a central mediator suggests a strategy to each player, drawn jointly from some distribution. This distribution is a correlated equilibrium if each player, knowing the joint distribution and his suggestion but not other's suggestions, has no incentive to deviate.

DEFINITION 2.2. Correlated equilibrium *An α -correlated equilibrium is a joint distribution X over bid vectors \mathbf{b} such that, for each player i , following her suggestion b_i drawn from X is a best-response up to an additive error of α , in expectation over others' suggestions \mathbf{b}_{-i} , assuming everyone else plays according to their suggestion:*

$$\forall i, \forall b'_i, \mathbb{E}_{\mathbf{b} \sim X}[u_i(\mathbf{b}) \mid b_i] \geq \mathbb{E}_{\mathbf{b} \sim X}[u_i(b'_i, \mathbf{b}_{-i}) \mid b_i] - \alpha$$

*Note that the deviation b'_i is allowed to depend on the suggestion. In the event that b'_i is independent of b_i for all i , we call X an α -**coarse correlated equilibrium**.*

A correlated equilibrium is an equilibrium of the static game in the complete information setting. This means that, even if a player knows the types of all other players, and the joint distribution from which the suggestions are being drawn, he will not deviate

⁵We insist that $\delta = \Omega(1)$, so that convergence occurs in polynomially many rounds of running the no-regret algorithms.

from the suggested strategy. The following theorem relates an outcome in the repeated setting when each player employs a no-regret learning algorithm to a correlated equilibrium of the static game.

THEOREM 2.3 ([FOSTER AND VOHRA 1997], [HART AND MAS-COLELL 2000]).

Suppose that a game is repeated for T rounds and each player employs a no-regret learning algorithm with external regret (resp. swap regret) of at most R . Then the joint distribution over strategy tuples given by the empirical distribution of strategies played by the players in each of the T rounds is an R/T -coarse correlated equilibrium (resp. correlated equilibrium) of the static game.

Thus, for a learnable game, the empirical distribution over strategies when each player runs a no-regret algorithm converges to a correlated equilibrium.

No-regret learning algorithms. By Theorem 2.3, the rate of convergence to a correlated equilibrium will be governed by the regret achieved by each of the players' learning algorithms. In each round, each bidder can compute her payoff from the bid she chooses: it is her utility from that round. On the other hand, a bidder cannot know what items would be available for her if her bid caused her to be later in the ordering: she cannot compute her payoff for all bids. Therefore, we need algorithms which have low regret in the *bandit* setting, rather than in the experts setting. Previous work has given efficient algorithms with low external and swap regret in the bandit setting.

THEOREM 2.4 (AUER ET AL. [2003]; BLUM AND MANSOUR [2007]). *There exist efficient algorithms which achieve external regret (resp. swap regret) of at most $\sqrt{hNT \log N}$ (resp. $N\sqrt{hNT \log N}$) in a bandit setting.*

One option for each player is to employ the algorithms as given by the above lemmas over the $O(hm/\epsilon)$ experts in the discretized bid space $\{0, \frac{\epsilon}{m}, \frac{2\epsilon}{m}, \dots, \lfloor \frac{hm}{\epsilon} \rfloor \frac{\epsilon}{m}\}$.⁶ We state the convergence rate obtained from such a discretization of bids below.

COROLLARY 2.5. *If each player employs an algorithm with external regret (resp. swap regret) as given by Theorem 2.4 on the discretized bid space as mentioned above, then after $O\left(\frac{h^2 m}{\epsilon^3} \log \frac{hm}{\epsilon}\right)$ rounds (resp. $O\left(\frac{h^4 m^3}{\epsilon^5} \log \frac{hm}{\epsilon}\right)$ rounds), the players have reached an ϵ -approximate coarse correlated equilibrium (resp. correlated equilibrium).*

Notice that the above convergence rates are pseudopolynomial (they depend polynomially rather than polylogarithmically on h). Alternatively, each bidder could discretize their bid space as follows: $\{0, \frac{v_i([m])\epsilon}{nm}, \frac{2v_i([m])\epsilon}{nm}, \dots, \lfloor \frac{nm}{\epsilon} \rfloor \frac{v_i([m])\epsilon}{nm}\}$. This reduces the total number of bids in our discretization to $O\left(\frac{nm}{\epsilon}\right)$. Furthermore, each bid $b \in [0, v_i([m])]$ is within an additive $\frac{v_i([m])\epsilon}{nm}$ of some bid in the discretized bid-space, therefore this discretization allows us to approach a $\frac{h\epsilon}{n}$ -correlated equilibrium in polynomial time (as $h \geq v_i([m])$ for all i).

COROLLARY 2.6. *If each player employs an algorithm with external regret (resp. swap regret) as given by Theorem 2.4 on the discretized bid space as mentioned above, then after $O\left(\frac{n^3 m^2}{h\epsilon^3} \log \left(\frac{nm}{\epsilon}\right)\right)$ rounds (resp. $O\left(\frac{n^5 m^4}{h\epsilon^5} \log \left(\frac{nm}{\epsilon}\right)\right)$ rounds), the players have reached an $\frac{h\epsilon}{n}$ -approximate coarse correlated equilibrium (resp. correlated equilibrium) of the discretized bid space auction.*

⁶Note that bidder i never has incentive to report a bid larger than $v_i([m]) \leq h$, as such a bid would certainly lead to 0 utility.

The total error of the discretization and approximation to correlated equilibrium is additively $O\left(\frac{h\epsilon}{n}\right)$ per bidder. So, the difference in welfare between this approximate correlated equilibrium and an exact correlated equilibrium is at most $O(h\epsilon)$. Since $h \leq OPT$, this is at most $O(\epsilon OPT)$. Thus, any approximation guarantee we prove for exact correlated equilibria will extend to these learnable approximate equilibria, gaining at most an ϵ factor in the approximation guarantee.

Strategic Play and the Price of Anarchy. Strategic play in many auctions can lead to inefficient allocations of goods; furthermore, it is a priori quite difficult to predict what types of strategic play might arise. In recent years, focus has shifted towards the analysis of simple auctions via the *price of anarchy*: one proves claims of the form “as long as bidders use strategies that form a Nash/correlated/coarse correlated/Bayes-Nash equilibrium, the items are allocated approximately efficiently.” Formally, for a given valuation profile \mathbf{v} , let $SW(\text{OPT}(\mathbf{v}))$ be the optimal social welfare, which is the highest social welfare obtainable over all possible allocations of items to bidders. $SW(\text{OPT}(\mathbf{v})) := \max \left\{ \sum_{i \in [n]} v_i(S_i) : (S_i)_{i \in [n]} \text{ is a partition of } [m] \right\}$. Let T denote a particular set of equilibria, s an equilibrium in T and $SW(s)$ the social welfare at this equilibrium. Then the price of anarchy w.r.t equilibria in T is defined as

$$PoA(T) := \max_{s \in T} \frac{SW(\text{OPT}(\mathbf{v}))}{SW(s)}.$$

Smooth Mechanisms. Roughgarden [2009] introduced the notion of *smooth games*, which was later extended by Syrgkanis and Tardos [2013] to the notion of smooth mechanisms. The smooth mechanism framework provides a method by which to prove Price of Anarchy bounds that hold simultaneously for Nash and correlated equilibria in games of incomplete and complete information.

DEFINITION 2.7 ([SYRGKANIS AND TARDOS 2013]). *A mechanism is (λ, μ) -smooth for a class of valuations $\mathcal{V} = \times_i \mathcal{V}_i$ if for any valuation profile $\mathbf{v} \in \mathcal{V}$, there exists a mapping $b'_i : [0, h] \rightarrow \Delta([0, h])$ such that for all $\mathbf{b} \in [0, h]^n$:*

$$\sum_i \mathbb{E} [u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i)] \geq \lambda SW(\text{OPT}(\mathbf{v})) - \mu \sum_i P_i(\mathbf{b}) \quad (1)$$

THEOREM 2.8 ([SYRGKANIS AND TARDOS 2013]). *If a mechanism is (λ, μ) -smooth then the price of anarchy w.r.t. mixed Bayes-Nash equilibria of the incomplete information setting and correlated equilibria in the complete information setting is at most $\frac{\max\{1, \mu\}}{\lambda}$. Furthermore, if the mapping b'_i is independent of b_i , then this result holds for coarse correlated equilibria.*

3. PRICE OF ANARCHY BOUNDS FOR SUBADDITIVE VALUATIONS

To prove the upper bound on the price of anarchy of the single-bid auction for subadditive valuations, we will establish that the single-bid auction is a $\left(\frac{e-1}{e \cdot H_m}, 1\right)$ -smooth mechanism, where H_m is the m^{th} harmonic number. Our approach is the following: we first show that the mechanism is $\left(\frac{e-1}{e}, 1\right)$ -smooth for a very restricted class of valuations which we dub constraint-homogeneous valuations (CHV). Each CHV is additive, with value for each individual item either 0 or some value \hat{v} , common for all items. Then we show that smoothness of a mechanism for one class of valuations implies smoothness for a more general class, as long as the latter class can be *approximated* by the former within some factor (we use a non-standard notion of valuation approximation, which we precisely define in Lemma 3.3). Moreover, the smoothness property

degrades exactly by the factor of approximation. We conclude the proof by showing that subadditive valuations can be *approximated* by CHV within a factor of H_m .

DEFINITION 3.1 (CONSTRAINT-HOMOGENEOUS VALUATION). *A valuation on a set of items is constraint-homogeneous if it is defined via an interest set S and a per-unit value \hat{v} such that:*

$$\forall T \subseteq [m] : v(T) = \hat{v} \cdot |T \cap S| \quad (2)$$

LEMMA 3.2 (SMOOTHNESS FOR CONSTRAINT-HOMOGENEOUS). *The single-bid auction is a $(1 - \frac{1}{e}, 1)$ -smooth mechanism with constraint-homogeneous bidders.*

PROOF. Consider a constraint-homogeneous valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ and let S_i^* be the set of items allocated to player i in the welfare maximizing allocation for valuation profile \mathbf{v} . We will show that there exists a randomized deviation B'_i , which does not depend upon the behavior of other agents, such that for any bid profile \mathbf{b} :

$$\mathbb{E}[u_i(B'_i, \mathbf{b}_{-i})] \geq \left(1 - \frac{1}{e}\right) \hat{v} |S_i^*| - \sum_{j \in S_i^*} p_j(\mathbf{b}). \quad (3)$$

where $p_j(\mathbf{b})$ is the price at which item j is sold under bid profile \mathbf{b} , i.e. the bid of the player that acquires it under \mathbf{b} .

Suppose that player i deviates to some deterministic bid $t \in [0, \hat{v}]$. Then for any $j \in S_i^*$, if $t > p_j(\mathbf{b})$, it means that when player i gets to pick his set of items, item j is still available. Thus his utility from such a strategy is lower bounded by:

$$u_i(t, \mathbf{b}_{-i}) \geq \sum_{j \in S_i^*} (\hat{v} - t) \cdot \mathbf{1}\{t > p_j(\mathbf{b})\} \quad (4)$$

Thus if B'_i is distributed according to density function $f(t) = \frac{1}{\hat{v}-t}$ and support $[0, (1 - \frac{1}{e})\hat{v}]$ then:

$$\mathbb{E}[u_i(B'_i, \mathbf{b}_{-i})] \geq \sum_{j \in S_i^*} \int_{p_j(\mathbf{b})}^{(1-\frac{1}{e})\hat{v}} (\hat{v} - t) \cdot f(t) \cdot dt = \sum_{j \in S_i^*} \left(\left(1 - \frac{1}{e}\right) \hat{v} - p_j(\mathbf{b}) \right) \quad (5)$$

which is exactly the lower bound we wanted to show. Summing the latter lower bound for every player, we get the $(1 - \frac{1}{e}, 1)$ -smoothness property. \blacksquare

We will next show that smoothness for constraint-homogeneous valuations implies smoothness for a much larger class of valuations. We achieve this based on the following re-interpretation of the results in [Syrgkanis and Tardos 2013]⁷.

DEFINITION 3.3 (POINTWISE VALUATION APPROXIMATION). *A valuation class V is pointwise β -approximated by a valuation class V' , if for any valuation $v \in V$, and for any set $S \subseteq [m]$, there exists a valuation $v' \in V'$ such that: $\beta v'(S) \geq v(S)$ and for all $T \subseteq [m]$: $v(T) \geq v'(T)$.*

Note that, importantly, the valuation v' can depend on S . $\beta v'$ only needs to upper bound v at S , while v' needs to lower bound v everywhere else. This is much weaker than the related notion of approximation by a function class, where for every v we ask for a single v' such that v is sandwiched between $\beta v'$ and v' everywhere.

⁷[Hartline 2013] gives a special case of this re-interpretation for the mechanism defined by simultaneous single-item auctions, showing how smoothness for additive valuations implies smoothness for unit-demand (and XOS) valuations

LEMMA 3.4 (EXTENSION LEMMA). *If a mechanism for a combinatorial auction setting is (λ, μ) -smooth for the class of valuations V' and V is pointwise β -approximated by V' , then it is $(\frac{\lambda}{\beta}, \mu)$ -smooth for the class V .*

PROOF. Consider a valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ where each valuation v_i comes from valuation class V . For each player i let S_i^* be her optimal allocation under \mathbf{v} and let $\mathbf{v}^* = (v_1^*, \dots, v_n^*)$ be the valuation profile such that $v_i^* \in V'$ is the valuation that β -approximates v_i for set S_i^* : i.e. $\beta \cdot v_i^*(S_i^*) \geq v_i(S_i^*)$ and for all $T \subseteq [m]$: $v_i(T) \geq v_i^*(T)$. By the first property we get that $\beta \cdot SW(\text{OPT}(\mathbf{v}^*)) \geq SW(\text{OPT}(\mathbf{v}))$. By the second property we get that for all bid profiles \mathbf{b} : $u_i(\mathbf{b}; v_i) \geq u_i(\mathbf{b}; v_i^*)$. Let $b'_i : [0, h] \rightarrow \Delta([0, h])$ be the deviation mapping that is designated by the smoothness property of the mechanism under \mathbf{v}^* . Then for any bid profile \mathbf{b} :

$$\begin{aligned} \sum_i \mathbb{E} [u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i)] &\geq \sum_i \mathbb{E} [u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i^*)] \geq \lambda SW(\text{OPT}(\mathbf{v}^*)) - \mu \sum_i P_i(\mathbf{b}) \\ &\geq \frac{\lambda}{\beta} SW(\text{OPT}(\mathbf{v})) - \mu \sum_i P_i(\mathbf{b}) \end{aligned}$$

which implies the mechanism is smooth for the valuation class V . ■

To conclude the proof we show that sub-additive valuations can be H_m -approximated by constraint-homogeneous valuations.

LEMMA 3.5 (CONSTRAINT-HOMOGENEOUS H_m -APPROXIMATE SUBADDITIVE). *Subadditive valuations can be pointwise H_m -approximated by constraint-homogeneous valuations.*

PROOF. Consider a subadditive valuation v , some β , and some set of items $X \subseteq [m]$. Let h_S denote the constraint-homogeneous function $h_S(T) = \frac{v(X)}{|S|\beta} |T \cap S|$. It suffices to find S such that $\beta h_S(X) \geq v(X)$ and also $v(T) \geq h_S(T)$ for all T . We will either find such an S or find an upper bound on β .

Consider $S_1 = X$. Then, $\beta h_{S_1}(X) = \beta h_X(X) = v(X)$, so the first inequality holds. If $v(T) \geq h_{S_1}(T)$ holds for all T , then h_{S_1} pointwise β -approximates v at X . If not, there exists T_1 such that $v(T_1) < h_{S_1}(T_1)$. Then, since v is monotone, $v(T_1 \cap S_1) \leq v(T_1) < h_{S_1}(T_1) = h_{S_1}(T_1 \cap S_1)$.

Iteratively, consider $S_i = S_{i-1} \setminus T_{i-1}$. As above, $\beta h_{S_i}(X) = \frac{v(X)}{|S_i|\beta} |X \cap S_i| = v(X)$, so the first condition is satisfied by h_{S_i} for all i . If for some i , $v(T) \geq h_{S_i}(T)$ for all T , then h_{S_i} pointwise β -approximates v at X ; else, there exists some T_i such that $v(T_i) < h_{S_i}(T_i)$.

After $j \leq m$ iterations, we have either found an h_{S_j} which pointwise β -approximates v at X , or we have constructed a partition T_1, \dots, T_j of X such that for all i

$$v(T_i) < h_{S_i}(T_i) = \frac{v(X)}{\beta |S_i|} |S_i \cap T_i| \leq \frac{v(X)}{\beta |S_i|} |T_i| \quad (6)$$

Since v is subadditive: $v(X) \leq \sum_i v(T_i)$. Thus, combining this with Equation 6,

$$v(X) < \sum_i \frac{v(X)}{\beta |S_i|} |T_i| = \frac{v(X)}{\beta} \sum_i \frac{|T_i|}{|S_i|}.$$

Thus, $\beta < \sum_i \frac{|T_i|}{|S_i|}$. We need to upper-bound $\sum_i \frac{|T_i|}{|S_i|}$ to upper-bound β . Notice that

$$\frac{|T_i|}{|S_i|} = \sum_{t=0}^{|T_i|-1} \frac{1}{|S_i|} \leq \sum_{t=0}^{|T_i|-1} \frac{1}{|S_i|-t}$$

so we have as desired that $\beta < \sum_i \frac{|T_i|}{|S_i|} \leq \sum_i \sum_{t=0}^{|T_i|-1} \frac{1}{|S_i|-t} = \sum_{\ell=0}^{m-1} \frac{1}{|X|-\ell} = H_m$. ■

To draw connections to previous work, when the class V' is the set of general additive valuations, then whether V can be pointwise β -approximated by V' is equivalent to asking whether the class V is β -fractionally subadditive, i.e. whether there exist a set of additive valuations indexed by some index set \mathcal{L} such that for any S :

$$\max_{\ell \in \mathcal{L}} v^\ell(S) \leq v(S) \leq \beta \max_{\ell \in \mathcal{L}} v^\ell(S) \quad (7)$$

It is known that subadditive valuations are H_m -fractionally subadditive [Dobzinski 2007; Bhawalkar and Roughgarden 2011], or in other words, that subadditive valuations can be pointwise H_m -approximated by additive valuations. Hence, our Lemma 3.5 can be viewed as a strengthening of this result, stating that general additive valuations are not needed and only additive valuations with only one possible non-zero value for each individual item, suffices. This result can be of independent interest in algorithmic and mechanism design questions for sub-additive valuations.

Combining Lemma 3.5 with the smoothness of single-bid auctions for constraint-homogeneous valuations (Lemma 3.2) and the Extension Lemma (Lemma 3.4) we get that the single-bid auction is $\left(\frac{e-1}{e \cdot H_m}, 1\right)$ -smooth for subadditive valuations. Moreover, observing that in all our proofs the smoothness deviation was a fixed strategy and not a mapping, yields our main Theorem 1.2.

3.1. Almost Tight Lower Bound

This bound on the price of anarchy for single-bid auctions is nearly asymptotically tight, even when restricted to additive bidders.

THEOREM 3.6. *The price of anarchy of single-bid auctions at pure Nash equilibria is at least $\Omega\left(\frac{\log m}{\log \log m}\right)$, even when all bidders are additive.*

PROOF. For the sake of simplicity, assume that ties are broken lexicographically when determining bid order throughout this proof. Consider the following bidders, valuations, and items. Suppose there is a partition of the m items B_0, \dots, B_{k-1} . Let $|B_t| = k^t$. Let bidder 0 have valuation v_0 as follows. For each $j \in B_t$ and each t , $v_0(j) = k^{k-t}$; thus, $v_0(B_t) = k^k$ for each t and $v_0([m]) = k^{k+1}$.

Then, for each $j \in \{0, \dots, k-1\}$, let there be two bidders i_{j1}, i_{j2} with valuations $v_{i_{j1}}(i) = v_{i_{j2}}(i) = \frac{v_0(i)}{k}$ for all $i \in B_j$, and $v_{i_{j1}}(i) = v_{i_{j2}}(i) = 0$ for all $i \notin B_j$. Notice that if each of these “small” bidders bids $\frac{v_0(j)}{k}$, then they are both playing a deterministic best-response, irrespective of bidder 0’s bid.

Given that all of the “small” bidders are bidding $\frac{k^{k-t}}{k}$ for $t \in \{0, \dots, k-1\}$, bidder 0 will bid exactly one of these numbers in equilibrium. Suppose she bids $b_0 = \frac{k^{k-t^*}}{k}$. When she bids b_0 , all items in B_{t^*}, \dots, B_{k-1} will be available for her to purchase. Consider some item $j \in B_t$ for $t > t^* + 1$: it is clear that $v_0(j) = k^{k-t} < k^{k-t^*-1} = \frac{k^{k-t^*}}{k} = b_0$. Thus, bidder 0 will not choose to buy any item in $B_{t^*+1}, \dots, B_{k-1}$. Then, she will buy at most the sets B_{t^*-1}, B_{t^*} , obtaining value $v_0(B_{t^*-1} \cup B_{t^*}) = 2k^k$.

Suppose S_i is the set of items bidder i buys at this equilibrium. We just showed that $v_0(S_0) \leq 2k^k$. For all $i \neq 0$, $v_i(S_i) = \frac{v_0(S_i)}{k}$, so $\sum_i v_i(S_i) \leq 2k^k + (k-2)k^{k-1} \leq 3k^k$, while the optimal social welfare is k^{k+1} . Notice that $m = \sum_{t=0}^k k^t = \Theta(k^{k-1})$. Thus, the price of anarchy is at least $\Omega(k) = \Omega\left(\frac{\log m}{\log \log m}\right)$. ■

4. TIGHTER UPPER BOUNDS FOR SIMPLER VALUATIONS

In this section we show tighter price of anarchy bounds for two other important classes of valuations: unit-demand and symmetric valuations. A valuation is unit-demand if a player only wants one item and has no value for any extra item. Equivalently if it can be expressed as: $v_i(S) = \max_{j \in S} w_{ij}$ for some $w_{ij} \geq 0$. A valuation is symmetric, if it is a function of the number of items and not of the specific set, i.e. if all items are identical. We will consider the case of concave symmetric valuations, i.e., $v_i(S) = f_i(|S|)$ for some non-decreasing concave function $f_i: \mathbb{N} \rightarrow \mathbb{R}^+$.

We show that both unit-demand valuations and concave symmetric valuations can be pointwise 1-approximated by constraint-homogeneous valuations. As a corollary we get that the price of anarchy of single-bid auctions for this case is at most $\frac{e}{e-1}$.

THEOREM 4.1. *The class of concave symmetric valuations is pointwise 1-approximated by constraint-homogeneous valuations.*

PROOF. Consider a valuation profile v as described in the theorem (i.e. $v(S) = f(|S|)$). Consider a set $S \subseteq [m]$ and let v' be the constraint-homogeneous valuation with interest set S and per-unit valuation $\hat{v}' = \frac{f(|S|)}{|S|}$. By concavity of f and since $f(0) = 0$, we know that for any $y > x$, $\frac{f(y)}{y} \leq \frac{f(x)}{x}$. Thus we have for any $T \subseteq [m]$:

$$v'(T) = \hat{v}' \cdot |T \cap S| = \frac{f(|S|)}{|S|} \cdot |T \cap S| \leq f(|T \cap S|) \leq v(T)$$

Additionally, $v'(S) = f(|S|) = v(S)$. ■

LEMMA 4.2. *The class of unit-demand valuations is 1-approximated by constraint-homogeneous valuations.*

PROOF. For each set of items S , let $j(S) = \arg \max_{j \in S} w_{ij}$. Then consider the constraint-homogeneous valuation v' , with $\hat{v}' = w_{ij(S)}$ and interest set $S = \{j\}$. Then: $v'(T) = w_{ij(S)} \cdot \mathbf{1}\{j(S) \in T\} \leq \max_{j \in T} w_{ij}$ and $v'(S) = w_{ij(S)} = v(S)$. ■

LEMMA 4.3. *The class of k -demand valuations is H_k -approximated by constraint-homogeneous valuations.*

PROOF. Consider k -demand valuation v and interest set S . We will construct a constraint-homogeneous v' such that $v'(T) \leq v(T)$ for all T but $v'(S) \geq H_k v(S)$. Let $S' = \arg \max_{S'' \subseteq S, |S''|=k} v(S'')$. Then, $v(S') = v(S)$. Now, repeat the proof of Lemma 3.5, beginning with set S' instead of X . In the final line of the proof, $|T_i|$ and $|X|$ can be replaced with k , rather than m , implying H_k as an upper bound on β . ■

REFERENCES

- Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, January 2003.
- Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Item pricing for revenue maximization. In *Proceedings of the 9th ACM conference on Electronic Commerce*, pages 50–59. ACM, 2008.

- Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 700–709. SIAM, 2011.
- Avrim Blum and Yishay Mansour. From external to internal regret. *Journal of Machine Learning Research*, 8:1307–1324, December 2007.
- Mark Braverman, Jieming Mao, and S. Matthew Weinberg. Incentives in communication and applications to combinatorial auctions. *Manuscript*, 2015.
- Yang Cai and Christos H. Papadimitriou. Simultaneous bayesian auctions and computational complexity. In *ACM Conference on Economics and Computation*, pages 895–910, 2014.
- Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 81–90. ACM, 2011.
- George Christodoulou, Annamaria Kovacs, and Michael Schapira. Bayesian Combinatorial Auctions. In *ICALP '08 Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, 2008.
- Shahar Dobzinski. Two randomized mechanisms for combinatorial auctions. In *Proceedings of the 10th International Workshop on Approximation and the 11th International Workshop on Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 89–103, 2007.
- Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *Proceedings of the thirty-seventh Annual ACM Symposium on Theory of Computing*, pages 610–618. ACM, 2005.
- Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the thirty-eighth annual ACM Symposium on Theory of Computing*, pages 644–652. ACM, 2006.
- Shahar Dobzinski, Hu Fu, and Robert Kleinberg. On the complexity of computing an equilibrium in combinatorial auctions. *Proceedings of the twenty-fifth annual ACM-SIAM Symposium on Discrete Algorithms*, page To appear, 2015.
- Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. In *Proceedings of the forty-third annual ACM Symposium on Theory of Computing*, pages 149–158. ACM, 2011.
- Uriel Feige, Vahab S Mirrokni, and Jan Vondrak. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011.
- Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM Symposium on Theory of Computing*, pages 201–210. ACM, 2013a.
- Michal Feldman, Brendan Lucier, and Vasilis Syrgkanis. Limits of efficiency in sequential auctions. In *Proceedings of the 9th Workshop on Internet and Network Economics*, 2013b.
- Dean P Foster and Rakesh V Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1):40–55, 1997.
- Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.
- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- Jason Hartline. Lecture 16, lecture notes on price of anarchy, Northwestern University, 2013.
- Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 295–296. ACM, 2011.
- Piotr Krysta and Berthold Vöcking. Online mechanism design (randomized rounding on the fly). In *Automata, Languages, and Programming*, pages 636–647. Springer, 2012.
- R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 595–604, Oct 2005.
- Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 18–28. ACM, 2001.
- Renato Paes Leme and Eva Tardos. Pure and bayes-nash price of anarchy for generalized second price auction. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Sympo-*

- sium on, pages 735–744. IEEE, 2010.
- Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. Sequential auctions and externalities. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 869–886. SIAM, 2012.
- Brendan Lucier and Allan Borodin. Price of anarchy for greedy auctions. In *Proceedings of the twenty-first annual ACM-SIAM Symposium on Discrete Algorithms*, pages 537–553. Society for Industrial and Applied Mathematics, 2010.
- Brendan Lucier and Renato Paes Leme. Gsp auctions with correlated types. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 71–80. ACM, 2011.
- Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the forty-first annual ACM Symposium on Theory of Computing*, pages 513–522. ACM, 2009.
- Tim Roughgarden. The price of anarchy in games of incomplete information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 862–879. ACM, 2012.
- Tim Roughgarden. Barriers to near-optimal equilibria. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 71–80. IEEE, 2014.
- Vasilis Syrgkanis. Bayesian games and the smoothness framework. *CoRR*, abs/1203.5155, 2012.
- Vasilis Syrgkanis and Eva Tardos. Bayesian sequential auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 929–944. ACM, 2012.
- Vasilis Syrgkanis and Eva Tardos. Composable and efficient mechanisms. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 211–220. ACM, 2013.

A. DRAFT AUCTIONS

In this section, we formally define *draft* auctions, a sequential version of single-bid auctions, and prove draft auctions have similar smoothness guarantees as we proved for single-bid auctions. Draft auctions proceed in rounds: each round is a first-price auction in which each bidder submits a bid; the winner in each round chooses some subset of the remaining items, and pays her bid for *each* item:

- (1) Initialize, for all $i \in [n]$, $S_i = \emptyset, P_i = 0$. The set of remaining items $I = [m]$.
- (2) While $I \neq \emptyset$,
- (3) Each bidder $i \in [n]$ submits a sealed bid b_i and a set $X_i \subseteq I$.
- (4) Allocate X_{i^*} to one $i^* \in \arg \max_{i \in [n]} \{b_i\}$, i.e., $S_{i^*} = S_{i^*} \cup X_{i^*}$.
- (5) Bidder i^* pays her bid per item, i.e., $P_{i^*} = P_{i^*} + b_{i^*} |X_{i^*}|$.
- (6) The winner i^* , winning bid b_{i^*} and allocated bundle X_{i^*} is announced.
- (7) End While.

Suppose each bidder’s valuation $v_i \in V_i$ is drawn from a distribution: $v_i \sim D_i$. Bidder i knows v_i but only D_j (rather than v_j) for all $j \neq i$. Then, draft auctions form a sequential game of incomplete information (and, in the case that each D_i is a point mass, a sequential game of complete information). A *strategy* $s_i : \mathcal{V}_i \rightarrow \Delta(B_i)$ of bidder i is a function, from her valuation to a distribution over bid plans $b_i \in B_i$. Each bid plan b_i determines the bid b_{it} that a player makes at some round t and the set X_{it} of items he gets conditional on winning, based on the information h_{it} available to her up to that round. For any given valuation profile \mathbf{v} , a tuple of strategies $\mathbf{b} = \mathbf{s}(\mathbf{v}) = (s_i(v_i))_{i \in [n]}$ determines the outcome of the auction; let $u_i(\mathbf{b}; v_i)$ denote the utility, (or expected utility when b is a distribution over bid plans) obtained by bidder i as a function of the bid plans \mathbf{b} . Recall that for a deterministic profile the utility is $v_i(S_i(\mathbf{b})) - P_i(\mathbf{b})$ where $S_i(\mathbf{b})$ is the set of items i wins and $P_i(\mathbf{b})$ is her total payment. Additionally, for any bid plan \mathbf{b} , we denote with $p_j(\mathbf{b})$ the price that item j was sold at, under bid plan \mathbf{b} . Observe that a bid plan actually also contains information about *what might have happened*, i.e., they specify the result of possible deviations from the actual outcome, which becomes important in the definitions of equilibria. A Nash equilibrium in sequential games allows for *irrational threats*, where an equilibrium strategy of a bidder could be suboptimal beyond a certain round. A standard refinement of the Nash equilibrium for extensive form games is the *subgame perfect equilibrium*, that allows only for strategies that constitute an equilibrium of any subgame, conditional on any possible history of play (see [Fudenberg and Tirole 1991] for a formal definition and a more

comprehensive treatment). We discuss these different notions in the complete version of the paper, but mention that all of our results extend to complete-information correlated equilibria (a generalization of Nash equilibria).

A.1. Smoothness of Draft auctions

We now show that draft auctions are smooth mechanisms, which has the same implications on the price of anarchy as in Theorem 2.8. There are two main technical hurdles in extending the arguments of smoothness for single-bid auctions to draft auctions. Unlike single-bid auctions, draft auctions proceed in rounds. This means that strategies are functions that map history to bids in each round. Bidders' deviations need to aim for particular items *at their equilibrium prices*. So, a deviating bidder needs to behave as they do in equilibrium (to ensure she faces equilibrium prices) until the right moment, at which point they bid the "right bid", and procure the items they would get in the optimal allocation. The second difficulty is that, unlike in sequential item auctions, a player is not aware, without information about other bidders' strategies, at which step any item is going to be allocated, since this is endogenously chosen by one of his opponents. Thus, deviations of the form: "behave exactly as previously until the optimal item arrives and then deviate to acquire it", will not yield smoothness proofs in the case of draft auctions.⁸ Instead, our deviations for the unit-demand case have a player always attempt to get his optimal item, while it is still available, without changing the observed history when she loses. We show a deviation of the following form does just that: *At each time step, as long as your optimal item is still available, bid the maximum of your equilibrium bid and half your value for your optimal item. If you ever win, buy your optimal item.*

LEMMA A.1. *The draft auction for unit-demand bidders is $(\frac{1}{2}, 2)$ -smooth.*

Proof of Lemma A.1 : Consider a unit-demand valuation profile v (i.e. $v_i(S) = \max_{j \in S} v_{ij}$) and let j_i^* be the item assigned to player i in the optimal matching for valuation profile v . We will show that there exists a deviation mapping $b'_i : B_i \rightarrow B_i$ for each player i , such that for any bid profile b :

$$u_i(b'_i(b_i), \mathbf{b}_{-i}) \geq \frac{1}{2}v_{ij_i^*} - p_{j_i^*}(\mathbf{b}) - P_i(\mathbf{b}). \quad (8)$$

Consider the following b'_i : in every auction t , the player bids the maximum of her previous bid b_{it} (conditional on the history) and $\frac{v_{ij_i^*}}{2}$, until j_i^* gets sold. If she ever wins some auction, she picks j_i^* . Suppose that j_i^* was sold at some auction t^* under strategy profile \mathbf{b} . We consider the following two cases separately, which are exhaustive since i drops out after round t^* at most.

Case 1: i wins an auction $t' \leq t^$ in b'_i .* . If i wins with bid $b_{it'}$ then that must have been her payment under b_i as well, and $P_i(\mathbf{b}) = b_{it'}$. Otherwise it is $b_i^* = \frac{v_{ij_i^*}}{2}$. Therefore her utility is

$$u_i(b'_i, \mathbf{b}_{-i}) \geq v_{ij_i^*} - \max \left\{ \frac{v_{ij_i^*}}{2}, P_i(\mathbf{b}) \right\} \geq v_{ij_i^*} - \frac{v_{ij_i^*}}{2} - P_i(\mathbf{b}) \geq \frac{1}{2}v_{ij_i^*} - p_{j_i^*}(\mathbf{b}) - P_i(\mathbf{b}).$$

Case 2: i does not win any auction in b'_i . . In this case, it must be that $p_{j_i^*}(\mathbf{b}) \geq \frac{1}{2}v_{ij_i^*}$ since otherwise i would have won auction t^* . Her utility is zero in this case. Therefore (8) holds in this case as well.

Thus, the deviation b'_i always satisfies (8). The smoothness property follows by summing over all players and using that $\sum_i p_{j_i^*}(\mathbf{b}) = \sum_{j \in [m]} p_j(\mathbf{b}) = \sum_i P_i(\mathbf{b})$. ■

Thus, combining Lemma A.1 and Theorem 2.8, we have the following.

COROLLARY A.2. *The price of anarchy for draft auctions with unit-demand bidders is at most 4.*

⁸Even in the complete information setting, the time at which an item sells is defined by the strategies of other players: using this information to construct a deviation would not fit into the smoothness framework. In the case of mixed strategies, or incomplete information, the time an item sells is a random variable, so such a strategy is not even well-defined.

We now state that draft auctions are smooth for constraint-homogeneous valuations. This implies that the price of anarchy bounds stated for single-bid auctions hold for draft auctions as well. Just like in the case of single-bid auctions, the need to buy many items adds complexity to the proof of this corollary over the unit-demand setting.

LEMMA A.3. *Draft auctions are $(\frac{1}{4}, 2)$ -smooth for constraint-homogeneous bidders.*

Before proving Lemma A.3, we state a separate lemma, which shows there always exists a “good” deviation for constraint-homogeneous bidders in draft auctions.

LEMMA A.4 (CORE DEVIATION LEMMA FOR DRAFT AUCTIONS). *Suppose player i has a constraint-homogeneous valuation with interest set S and per-unit value \hat{v} . Then there is a deviation for draft auctions $b'_i : B_i \rightarrow B_i$ such that, for any strategy profile b :*

$$u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i) \geq \frac{1}{2} \frac{\hat{v} \cdot |S|}{2} - \sum_{j \in S} p_j - P_i(\mathbf{b}).$$

We now use Lemma A.4 to prove Lemma A.3.

Proof of Lemma A.3: Consider a constraint-homogeneous valuation profile v and a bid profile b . Let S_i^* be the units allocated to player i in the optimal allocation for profile v . Also let S_i be the interest set of each player and \hat{v}_i his per-unit value. Consider the alternative valuation profile where each player i has a constraint-homogeneous valuation v'_i with interest set $S'_i = S_i \cap S_i^*$ and per unit value $\hat{v}'_i = \hat{v}_i$.

Observe that for any $T \subseteq [m]$, $v_i(T) \geq v'_i(T)$ and $v_i(S_i^*) = v'_i(S_i^*)$. Thus, for any bid profile b : $u_i(b; v_i) \geq u_i(b; v'_i)$ and $SW(\text{OPT}(v')) \geq SW(\text{OPT}(v))$. Invoking Lemma A.4 on valuations v'_i , we get that there exists a deviation mapping $b'_i : B_i \rightarrow B_i$ for each player i such that for any strategy profile b :

$$\sum_i u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i) \geq \sum_i u_i(b'_i(b_i), \mathbf{b}_{-i}; v'_i) \geq \frac{1}{4} \text{OPT}(v') - 2 \sum_i P_i(\mathbf{b}) \geq \frac{1}{4} \text{OPT}(v) - 2 \sum_i P_i(\mathbf{b}),$$

where we have once again used the fact that $\sum_i p_{j_i^*}(\mathbf{b}) = \sum_{j \in [m]} p_j(\mathbf{b}) = \sum_i P_i(\mathbf{b})$. ■

Combining Lemma 3.5 with Lemma A.3 and Lemma 3.4, we get the following efficiency guarantee for draft auctions with subadditive valuations.

COROLLARY A.5. *The price of anarchy for draft auctions with subadditive bidders is at most $8H_m$.*

The Core Deviation for draft auctions is somewhat more complicated than the single-bid deviation, because it is multi-stage and needs to mimic a bidder’s equilibrium behavior. Just as in the case for single-bid auctions, the key deviation to prove smoothness for draft auctions is to bid the “right price”, half of her per-unit value, and then try to acquire the “right number” of those items, which is at least half the number of units in her optimal allocation. However, consider a round where her equilibrium bid is higher than the “right price”. If the bidder bids the right price, she may change the history for all the other players and sets the game down an off-equilibrium path, where the prices in the remaining off-equilibrium subgame are difficult to reason about. Thus, the deviations we consider have a player mimic her equilibrium play until she can acquire her optimal number of units at a good price.

To achieve this, the deviation bids the maximum of the original bid and the right price. If the original bid is higher, she follows the original strategy and picks the same set of items⁹. If the right price is higher, she then buys sufficient number of items to win the “right number” of units, and drops out of subsequent rounds.

DEFINITION A.6 (CORE DEVIATION FOR DRAFT AUCTIONS). *The core deviation for draft auctions b'_i for player i with a constraint-homogeneous valuation with interest set S and per-unit value \hat{v} is defined as follows. Let $b_i^* = \frac{\hat{v}}{2}$. In every auction t , she submits $b'_{it} = \max\{b_i^*, b_{it}\}$. If she wins with bid at most b_i^* , she buys $s^* - k_{i, < t}$ units of S and drops out. If she wins with a*

⁹If the deviation were for the bidder to buy all the right number of units when she won because of her equilibrium bid, she might pay too much for them.

bid of $b_{it} > b_i^*$, she buys what she did under b_i : k_{it} units together with any other items she was buying under strategy profile b_i at auction t . She continues to bid b'_{it} until she acquires s^* units or the number of units remaining are not sufficient for her to complete s^* units.

The crucial observation is this: as long as the player hasn't already acquired s^* units, she has not affected the game path created by strategy b_i in any way. From the perspective of the other bidders, she behaved exactly as under b_i , by winning at her price under b_i and getting the items she would have got under b_i . If she ever wins at a higher price, she acquires all the units needed to reach s^* units in that auction and then drops out. Thus the prices that she faces in all the auctions prior to having won s^* units are the same as the prices under strategy b_i .

The draft Core Deviation Lemma then follows from Lemmas A.7, A.8, and A.9.

LEMMA A.7. *If player i wins at least s^* units of S under b'_i , the Core Deviation for draft auctions, then*

$$u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i) \geq \frac{1}{2}s^*\hat{v} - P_i(\mathbf{b}).$$

PROOF. If player i wins at least s^* units of S under b'_i then the valuation for the items she wins is at least $s^*\hat{v}$. For the auctions in which she wins with a bid of b_{it} she pays at most $P_i(\mathbf{b})$ and for the (at most one) auction she wins with a bid of b_i^* she pays at most $s^*b_i^*$. So her total payment is at most $s^*b_i^* + P_i(\mathbf{b}) = s^*\frac{\hat{v}}{2} + P_i(\mathbf{b})$. ■

LEMMA A.8. *If player i wins fewer than s^* units of S under the Core Deviation for draft auctions b'_i then*

$$u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i) \geq \frac{1}{2}s^*\hat{v} - \sum_{j \in S} p_j - P_i(\mathbf{b}).$$

PROOF. Consider the auction under the original strategy profile \mathbf{b} . Let (by an abuse of notation) $p_1 \leq p_2 \leq \dots \leq p_{|S|}$ be the prices at which the items in S are sold under \mathbf{b} . This is not necessarily the order in which they are sold. We show in Lemma A.9 that, when bidder i wins fewer than s^* units under b'_i , it must be that $p_{s^*} \geq \frac{\hat{v}}{2}$. Thus,

$$\sum_{j \in S} p_j \geq \sum_{l=s^*}^{|S|} p_l \geq (|S| - s^* + 1)p_{s^*} \geq s^*p_{s^*} \geq \frac{\hat{v}}{2}s^*, \quad (9)$$

where we also used the simple observation that $s^* \leq \frac{|S|+1}{2}$.

The total payment of player i under b'_i in this case where she wins fewer than $|S|/2$ units of S is at most $P_i(\mathbf{b})$, therefore her utility is (trivially) at least $-P_i(\mathbf{b})$. The lemma now follows from adding the inequalities $u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i) \geq -P_i(\mathbf{b})$ and $0 \geq \frac{\hat{v}}{2}s^* - \sum_{j \in S} p_j$ (which holds by inequality (9)). ■

LEMMA A.9. *If player i wins fewer than s^* units of S under the Core Deviation b'_i then the s^* -th lowest price of the units in S under \mathbf{b} , is at least $\hat{v}/2$.*

PROOF. First, observe that if player i was obtaining at least s^* units under \mathbf{b} then she is definitely winning s^* units under b'_i , since she is always bidding at least as high. So, we can assume that under \mathbf{b} player i wins fewer than s^* units.

Recall that $p_1 \leq p_2 \leq \dots \leq p_{|S|}$ are the prices at which the units in S are sold under \mathbf{b} . Let P_t be the price of auction t (under \mathbf{b}). Let t^* be the first auction that was won at price $P_{t^*} \leq p_{s^*}$ under \mathbf{b} but not by bidder i . We know that such an auction must exist; under \mathbf{b} there are s^* units of S that are sold at a price at most p_{s^*} , and since player i wins fewer than s^* of them, some of them are not won by player i .

We now argue that player i is still bidding in auction t^* under b'_i . First of all, she has not won s^* units prior to t^* . The other condition needed for her to be active is that there are at least $s^* - k_{i, < t^*}$ units available for sale in that auction. This follows from the fact that for any auction $t < t^*$ for which $P_t \leq p_{s^*}$, we know that player i was winning under b_i . Thus every unit that was sold prior to t^* at a price of less than or equal to p_{s^*} was sold to player i . There are s^* units sold at a price $\leq p_{s^*}$ and the number of such units sold prior to t^* is at most the number of total

units won by bidder i prior to t^* . Thus the number of available units available at t^* is at least: $s^* - k_{i, < t^*}$.

Finally, we argue that $P_{t^*} \geq b_i^*$. Suppose for the sake of contradiction that $P_{t^*} < b_i^*$. Then player i wins auction t^* . Since she was not winning t^* under b_i , it must be that she is winning t^* with a bid of b_i^* . Thus in that auction she will buy every unit needed to reach s^* units. By the analysis in the previous paragraph, we know that there are still enough units available for sale to reach s^* . Thus in this case she will win s^* items, a contradiction with the main assumption of the Lemma. Therefore, $b_i^* \leq P_{t^*}$ and by definition, $P_{t^*} \leq p_{s^*}$ and $b_i^* = \frac{v}{2}$. ■

An easy corollary of the above core deviation lemma is that when all players have constraint-homogeneous valuations, the draft auction is a $(\frac{1}{4}, 2)$ -smooth mechanism, and thus has a price of anarchy of at most 8 for these valuations.