

Inefficiency of Standard Multi-unit Auctions

Bart de Keijzer¹, Evangelos Markakis^{3,*}, Guido Schäfer^{1,2},
and Orestis Telelis^{3,**}

¹ CWI Amsterdam, The Netherlands

² VU Amsterdam, The Netherlands

{b.de.keijzer,g.schaefer}@cwi.nl

³ Athens University of Economics and Business, Greece

{markakis,telelis}@gmail.com

Abstract. We study two standard multi-unit auction formats for allocating multiple units of a single good to multi-demand bidders. The first one is the Discriminatory Auction, which charges every winner his winning bids. The second is the Uniform Price Auction, which determines a uniform price to be paid per unit. Variants of both formats find applications ranging from the allocation of state bonds to investors, to online sales over the internet. For these formats, we consider two bidding interfaces: (i) standard bidding, which is most prevalent in the scientific literature, and (ii) uniform bidding, which is more popular in practice. In this work, we evaluate the economic inefficiency of both multi-unit auction formats for both bidding interfaces, by means of upper and lower bounds on the Price of Anarchy for pure Nash equilibria and mixed Bayes-Nash equilibria. Our developments improve significantly upon bounds that have been obtained recently for submodular valuation functions. Also, for the first time, we consider bidders with subadditive valuation functions under these auction formats. Our results signify near-efficiency of these auctions, which provides further justification for their use in practice.

1 Introduction

We study standard multi-unit auction formats for allocating multiple units of a single good to multi-demand bidders. Multi-unit auctions are one of the most widespread and popular tools for selling identical units of a good with a single auction process. In practice, they have been in use for a long time, one of their most prominent applications being the auctions offered by the U.S. and U.K. Treasuries for selling bonds to investors, see e.g., the U.S. treasury report [21]. In more recent years, they are also implemented by various online brokers [17].

* Supported by the project AGT of the research funding program THALES (co-financed by the EU and Greek national funds).

** Supported by project “DDCOD” (PE6-213). The project is implemented within the framework of the Action “Supporting Postdoctoral Researchers” of the Operational Program “Education and Lifelong Learning” (Action’s Beneficiary: General Secretariat for Research and Technology), and is co-financed by the European Social Fund (ESF) and the Greek State.

Table 1. Upper bounds on the (Bayesian) economic inefficiency of multi-unit auctions

Valuation Functions	Auction Format	
	(bidding: standard uniform)	
	<i>Discriminatory Auction</i>	<i>Uniform Price Auction</i>
<i>Submodular</i>	$e/(e - 1)$	3.1462
<i>Subadditive</i>	2 $2e/(e - 1)$	4 6.2924

In the literature, multi-unit auctions have been a subject of study ever since the seminal work of Vickrey [22] (although the need for such a market enabler was conceived even earlier, by Friedman, in [8]) and the success of these formats has led to a resurgence of interest in auction design.

There are three simple *Standard Multi-Unit Auction* formats that have prevailed and are being implemented; these are the *Discriminatory Auction*, the *Uniform Price Auction* and the *Vickrey Multi-Unit Auction*. All three formats share a common allocation rule and bidding interface and have seen extensive study in auction theory [12,15]. Each bidder under these formats is asked to issue a sequence of non-increasing marginal bids, one for each additional unit. For an auction of k units, the k highest marginal bids win, and each grants its issuing bidder a single unit. The formats differ in the way that payments are determined for the winning bidders. The Discriminatory Auction prescribes that each bidder pays the sum of his winning bids. The Uniform Price Auction charges the lowest winning or highest losing marginal bid per allocated unit. The Vickrey auction charges according to an instance of the Clarke payment rule (thus being a generalization of the well known single-item Second-Price auction).

Except for the Vickrey auction, which is truthful and efficient, the others suffer from a *demand reduction* effect [1], whereby bidders may have incentives to understate their value, so as to receive less units at a better price. This effect is amplified when bidders have non-submodular valuation functions, since the bidding interface forces them to encode their value within a submodular bid vector. Even worse, in many practical occasions bidders are asked for a *uniform bid* per unit together with an upper bound on the number of desired units. In such a setting, each bidder is required to “compress” his valuation function into a bid that scales linearly with the number of units. The mentioned allocation and pricing rules apply also in this *uniform bidding* setting, thus yielding different versions of Discriminatory and Uniform Price Auctions. Despite the volume of research from the economics community [1,16,6,3,4] and the widespread popularity of these auction formats, the first attempts of quantifying their economic efficiency are only very recent [14,19]. There has also been no study of these auction formats for non-submodular valuations, as noted by Milgrom [15].

Our Contributions. We study the inefficiency of the Discriminatory Auction and Uniform Price Auction under the standard and uniform bidding interfaces. Our main results are improved inefficiency bounds for bidders with submodular valuation functions and new bounds for bidders with subadditive valuation

functions.¹ The results are summarized in Table 1. Our bounds indicate that these auctions are nearly efficient, which, paired with their simplicity, provides further justification for their use in practice.

Our focus is on the inefficiency of Bayes-Nash equilibria; we refer the reader to the full version of the paper [11] for a discussion of pure Nash equilibria. For submodular valuation functions, we derive upper bounds of $\frac{e}{e-1}$ and $3.1462 < \frac{2e}{e-1}$ for the Discriminatory and the Uniform Price Auctions, respectively. These improve upon the previously best known bounds of $\frac{2e}{e-1}$ and $\frac{4e}{e-1}$ [19]. For the Uniform Price Auction, our bound is less than a factor 2 away from the known lower bound of $\frac{e}{e-1}$ [14]. We also prove lower bounds of $\frac{e}{e-1}$ and 2 for the Discriminatory Auction and Uniform Price Auction, with respect to the currently known proof techniques [19,7,5,2,10]. As a consequence, unless the upper bound of $\frac{e}{e-1}$ for the Discriminatory Auction is tight, its improvement requires the development of novel tools; the same holds for reducing the Uniform Price Auction upper bound below 2 (if $\frac{e}{e-1}$ from [14] is indeed worst-case). For subadditive valuations, we obtain bounds of $\frac{2e}{e-1}$ and $6.2924 < \frac{4e}{e-1}$ for Discriminatory and Uniform Price Auctions respectively, independent of the bidding interface. Further, for the standard bidding interface we are able to derive improved bounds of 2 and 4, respectively, by adapting a recent technique from [7]. We also give a lower bound of almost 2 for uniform pricing and subadditive valuations. In Section 4 we discuss further applications of our results in connection with the smoothness framework of [19]. In particular, some of our bounds carry over to simultaneous and sequential compositions of such auctions (Table 2).

Related Work. The multi-unit auction formats that we examine here present technical and conceptual resemblance to the *Simultaneous Auctions* format that has received significant attention recently [7,5,2,10,19]. However, upper bounds in this setting do not carry over to our format. Simultaneous auctions were first studied by Christodoulou, Kovacs and Schapira [5]. The authors proposed that each of a collection of distinct goods, with one unit available for each of them, is sold in a distinct *Second Price Auction*, simultaneously and independently of the other goods. Bidders in this setting may have combinatorial valuation functions over the subsets of goods, but they are forced to bid separately for each good. For bidders with fractionally subadditive valuation functions, they proved a tight upper bound of 2 on the mixed Bayesian Price of Anarchy of the Simultaneous Second Price Auction. Bhawalkar and Roughgarden [2] extended the study of inefficiency for subadditive bidders and showed an upper bound of $O(\log m)$ which was recently reduced to 4 by Feldman *et al.* [7]. For arbitrary valuation functions, Fu, Kleinberg and Lavi [9] proved an upper bound of 2 on the inefficiency of pure Nash equilibria, when they exist.

Hassidim *et al.* [10] studied *Simultaneous First Price Auctions*. They showed that pure Nash equilibria in this format are always efficient, when they exist. They proved constant upper bounds on the inefficiency of mixed Nash equilibria for (fractionally) subadditive valuation functions and $O(\log m)$ and $O(m)$ for

¹ To the best of our knowledge, for subadditive valuation functions our bounds provide the first quantification of the inefficiency of these auction formats.

the inefficiency of mixed Bayes-Nash equilibria for subadditive and arbitrary valuation functions. Syrgkanis showed in [20] that this format has Bayesian Price of Anarchy $\frac{e}{e-1}$ for fractionally subadditive valuation functions. Feldman *et al.* [7] proved an upper bound of 2 for subadditive ones.

Recently, Syrgkanis and Tardos [19] and Roughgarden [18] independently developed extensions of the *smoothness technique* for games of incomplete information. In [19], these ideas are further developed for analyzing the inefficiency of simultaneous and sequential *compositions* of simple auction mechanisms. They demonstrate applications of their techniques on welfare analysis of standard multi-unit auction formats *and* their compositions. For submodular valuation functions, they prove inefficiency upper bounds of $\frac{2e}{e-1}$ and $\frac{4e}{e-1}$ for the Discriminatory Auction and Uniform Price Auction, respectively. Here, we improve upon these results, also regarding simultaneous and sequential compositions.

2 Definitions and Preliminaries

We consider auctioning k units of a single good to a set $[n] = \{1, \dots, n\}$ of n bidders, indexed by $i = 1, \dots, n$. Every bidder $i \in [n]$ has a non-negative non-decreasing private valuation function $v_i : (\{0\} \cup [k]) \mapsto \mathbb{R}^+$ over quantities of units, where $v_i(0) = 0$. We denote by $\mathbf{v} = (v_1, \dots, v_n)$ the *valuation function profile* of bidders. We consider in particular (symmetric) *submodular* and *subadditive* functions:

- Definition 1.** A valuation function $f : (\{0\} \cup [k]) \mapsto \mathbb{R}^+$ is called:
- *submodular* iff for every $x < y$, $f(x) - f(x - 1) \geq f(y) - f(y - 1)$.
 - *subadditive* iff for every x, y , $f(x + y) \leq f(x) + f(y)$.

The class of submodular functions is strictly contained in the class of subadditive ones [13]. For any non-negative non-decreasing function $f : (\{0\} \cup [k]) \mapsto \mathbb{R}^+$ and any integers $x, y \in [k]$, $x < y$, the following are known to hold: If f is submodular, then $f(x)/x \geq f(y)/y$. If f is subadditive, then $f(x)/x \geq f(y)/(x + y)$.

Standard Multi-unit Auctions. The *standard format*, as described in auction theory [12,15], prescribes that each bidder $i \in [n]$ submits a vector of k non-negative non-increasing *marginal bids* $\mathbf{b}_i = (b_i(1), \dots, b_i(k))$ with $b_i(1) \geq \dots \geq b_i(k)$. We will often refer to these simply as *bids*. In the *uniform bidding format*, each bidder i submits only a single bid \bar{b}_i along with a quantity $q_i \leq k$, the interpretation being that i is willing to pay at most \bar{b}_i per unit for up to q_i units.

The allocation rule of standard multi-unit auctions grants the issuer of each of the k highest (marginal) bids a distinct unit per winning bid. The pricing rule differentiates the formats. Let $x_i(\mathbf{b})$ be the number of units won by bidder i under profile $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$. We study the following two pricing rules:

- (i) *Discriminatory Pricing.* Every bidder i pays for every unit a price equal to his corresponding winning bid, i.e., the utility of i is
- (ii) *Uniform Pricing.* Every bidder i pays for every unit a price equal to the *highest losing bid* $p(\mathbf{b})$, i.e., the utility of i is

$$u_i^{v_i}(\mathbf{b}) = v_i(x_i(\mathbf{b})) - \sum_{j=1}^{x_i(\mathbf{b})} b_i(j). \qquad u_i^{v_i}(\mathbf{b}) = v_i(x_i(\mathbf{b})) - x_i(\mathbf{b})p(\mathbf{b}).$$

For a bidding profile \mathbf{b} , the produced allocation $\mathbf{x}(\mathbf{b}) = (x_1(\mathbf{b}), x_2(\mathbf{b}), \dots, x_n(\mathbf{b}))$ has a *social welfare* equal to the bidders' total value: $SW(\mathbf{v}, \mathbf{b}) = \sum_{i=1}^n v_i(x_i(\mathbf{b}))$. The (pure) Price of Anarchy is the worst case ratio, over all pure Nash equilibrium profiles \mathbf{b} , of the optimal social welfare over $SW(\mathbf{v}, \mathbf{b})$.

Incomplete Information. Under the incomplete information model of Harsanyi, the valuation function \mathbf{v}_i of bidder i is drawn from a finite set V_i according to a discrete probability distribution $\pi_i : V_i \rightarrow [0, 1]$ (independently of the other bidders); we will write $\mathbf{v}_i \sim \pi_i$. The actual drawn valuation function of every bidder is *private*. A valuation profile $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathcal{V} = \times_{i \in [n]} V_i$ is drawn from a *publicly known distribution* $\pi : \mathcal{V} \rightarrow [0, 1]$, where π is the product distribution of π_1, \dots, π_n , i.e., $\pi(\mathbf{v}) \mapsto \prod_{i \in [n]} \pi_i(\mathbf{v}_i)$. Every bidder i knows his own valuation function but does not know the valuation function $\mathbf{v}_{i'}$ drawn by any other bidder $i' \neq i$. Bidder i may only use his knowledge of π to estimate \mathbf{v}_{-i} . Given the publicly known distribution π , the (possibly mixed) strategy of every bidder is a function of his own valuation \mathbf{v}_i , denoted by $B_i(\mathbf{v}_i)$. B_i maps a valuation function $\mathbf{v}_i \in V_i$ to a *distribution* $B_i(\mathbf{v}_i) = B_i^{v_i}$, over all possible bid vectors for i . In this case we will write $\mathbf{b}_i \sim B_i^{v_i}$, for any particular bid vector \mathbf{b}_i drawn from this distribution. We also use the notation $\mathbf{B}_{-i}^{\mathbf{v}_{-i}}$, to refer to the vector of randomized strategies of bidders other than i , under profile \mathbf{v}_{-i} .

A *Bayes-Nash equilibrium* (BNE) is a strategy profile $\mathbf{B} = (B_1, \dots, B_n)$ such that for every bidder i and for every valuation \mathbf{v}_i , $B_i(\mathbf{v}_i)$ maximizes the utility of i in expectation, over the distribution of the other bidders' valuations \mathbf{w}_{-i} given \mathbf{v}_i and over the distribution of i 's own and the other bidders' strategies, $\mathbf{B}(\mathbf{v}_i, \mathbf{w}_{-i})$, i.e., for every pure strategy \mathbf{c}_i of i :

$$\mathbb{E}_{\mathbf{w}_{-i}|\mathbf{v}_i, \mathbf{b} \sim \mathbf{B}(\mathbf{v}_i, \mathbf{w}_{-i})} [u_i^{v_i}(\mathbf{b})] \geq \mathbb{E}_{\mathbf{w}_{-i}|\mathbf{v}_i, \mathbf{b}_{-i} \sim \mathbf{B}^{\mathbf{w}_{-i}}} [u_i^{v_i}(\mathbf{c}_i, \mathbf{b}_{-i})]$$

where $\mathbb{E}_{\mathbf{v}}$ and $\mathbb{E}_{\mathbf{w}_{-i}|\mathbf{v}_i}$ denote expectation over the distributions π and $\pi(\cdot|\mathbf{v}_i)$.

Fix a valuation profile $\mathbf{v} \in \mathcal{V}$ and consider a (mixed) bidding configuration $\mathbf{B}^{\mathbf{v}}$ under \mathbf{v} . The Social Welfare $SW(\mathbf{v}, \mathbf{B}^{\mathbf{v}})$ under $\mathbf{B}^{\mathbf{v}}$ when the valuations are \mathbf{v} is defined as the expectation over the bidding profiles chosen by the bidders from their randomized strategies, i.e., $SW(\mathbf{v}, \mathbf{B}^{\mathbf{v}}) = \mathbb{E}_{\mathbf{b} \sim \mathbf{B}^{\mathbf{v}}} [\sum_i v_i(x_i(\mathbf{b}))]$. The *expected* Social Welfare in Bayes-Nash equilibrium $\mathbf{B}^{\mathbf{v}}$ is then $\mathbb{E}_{\mathbf{v} \sim \pi} [SW(\mathbf{v}, \mathbf{B}^{\mathbf{v}})]$. The socially optimum assignment under valuation profile $\mathbf{v} \in \mathcal{V}$ will be denoted by $\mathbf{x}^{\mathbf{v}}$. The *expected* optimum social welfare is then $\mathbb{E}_{\mathbf{v}} [SW(\mathbf{v}, \mathbf{x}^{\mathbf{v}})]$. Under these definitions, we will study the *Bayesian Price of Anarchy*, i.e., the worst case ratio $\mathbb{E}_{\mathbf{v}} [SW(\mathbf{v}, \mathbf{x}^{\mathbf{v}})] / \mathbb{E}_{\mathbf{v}} [SW(\mathbf{v}, \mathbf{B}^{\mathbf{v}})]$ over all possible product distributions π and Bayes-Nash equilibria \mathbf{B} for π .² Similarly to previous works, when analyzing the Uniform Price Auction we assume *no-overbidding*, i.e., each bidder never bids more than his value for every number of units; formally, for every $s \in [k]$, $\sum_{j \in [s]} b_i(j) \leq v_i(s)$. In our analysis, we will use $\beta_j(\mathbf{b})$ to refer to the j -th lowest winning bid under profile \mathbf{b} ; thus $\beta_1(\mathbf{b}) \leq \dots \leq \beta_k(\mathbf{b})$.

² As in previous works [5,7], we ensure existence of Bayes-Nash equilibria in our auction formats by assuming that bidders have bounded and finite strategy spaces, e.g., derived through discretization. Our bounds on the auctions' Bayesian inefficiency hold for sufficiently fine discretizations (see also Appendix D of [7]).

Due to space limitations, we omit several proofs from this extended abstract; all missing proofs can be found in the full version of the paper [11].

3 Bayes-Nash Inefficiency

Our main results concern the inefficiency of Bayes-Nash equilibria (we defer a discussion of pure Nash equilibria to the full version [11]). We derive bounds on the (mixed) Bayesian Price of Anarchy for the Discriminatory and the Uniform Price Auctions with submodular and subadditive valuation functions. For the latter class our bounds are the first results to appear in the literature of standard multi-unit auctions (see also the commentary in [15, Chapter 7]).

Theorem 1. *The Bayesian Price of Anarchy (under the standard or uniform bidding format) is at most*

- (i) $\frac{e}{e-1}$ and $\frac{2e}{e-1}$ for the Discriminatory Auction with submodular and subadditive valuation functions, respectively,
- (ii) $|W_{-1}(-1/e^2)| \approx 3.1462 < \frac{2e}{e-1}$ and $2|W_{-1}(-1/e^2)| \approx 6.2924 < \frac{4e}{e-1}$ for the Uniform Price Auction with submodular and subadditive valuation functions, respectively, W_{-1} being the lower branch of the Lambert W function.

This theorem improves on the currently best known upper bounds of $\frac{2e}{e-1}$ and $\frac{4e}{e-1}$ for the Discriminatory Auction and the Uniform Price Auction, respectively, with submodular valuation functions due to Syrgkanis and Tardos [19]. For the Uniform Price Auction, this further reduces the gap from the known lower bound of $\frac{e}{e-1}$ [14].

Syrgkanis and Tardos [19] obtained their bounds through an adaptation of the *smoothness framework* for games with incomplete information ([18,20]). The bounds of Theorem 1 and some additional results can also be obtained through this framework. We comment on this in more detail in Section 4.

For subadditive valuation functions and the standard bidding format, however, better bounds can be obtained by adapting a technique recently introduced by Feldman *et al.* [7], which does not fall within the smoothness framework. We were unable to derive these bounds via a smoothness argument and believe that this is due to the additional flexibility provided by this technique.

Theorem 2. *The Bayesian Price of Anarchy is at most 2 and 4 for the Discriminatory Auction and the Uniform Price Auction, respectively, with subadditive valuation functions under the standard bidding format.*

3.1 Proof Template for Bayesian Price of Anarchy

In order to present all our bounds from Theorem 1 and Theorem 2 in a self-contained and unified manner, we make use of a proof template, formalized in Theorem 3, below. Adaptations of it have been used in previous works [14,5,2].

Theorem 3. *Let V be a class of valuation functions. Suppose that for every valuation profile $\mathbf{v} \in V^n$, for every bidder $i \in [n]$, and for every distribution \mathcal{P}_{-i} over non-overbidding profiles \mathbf{b}_{-i} , there is a bidding profile \mathbf{b}'_i such that the following inequality holds for some $\lambda > 0$ and $\mu \geq 0$:*

$$\mathbb{E}_{\mathbf{b}_{-i} \sim \mathcal{P}_{-i}} \left[u_i^{\mathbf{v}_i}(\mathbf{b}'_i, \mathbf{b}_{-i}) \right] \geq \lambda \cdot v_i(x_i^{\mathbf{v}}) - \mu \cdot \mathbb{E}_{\mathbf{b}_{-i} \sim \mathcal{P}_{-i}} \left[\sum_{j=1}^{x_i^{\mathbf{v}}} \beta_j(\mathbf{b}_{-i}) \right]. \quad (1)$$

Then the Bayesian Price of Anarchy is at most

- (i) $\max\{1, \mu\}/\lambda$ for the Discriminatory Auction,
- (ii) $(\mu + 1)/\lambda$ for the Uniform Price Auction.

In this theorem we make no assumptions regarding the bidding interface; proving a bound for the uniform bidding interface only requires that we exhibit a uniform bidding strategy \mathbf{b}'_i for each bidder i and for any distribution \mathcal{P}_{-i} over uniform non-overbidding profiles \mathbf{b}_{-i} . In Section 3.3 we show that our bound of $\frac{e}{e-1}$ for the Discriminatory Auction is best possible, for the proof template of Theorem 3; this rules out achievement of better bounds via the techniques in [19,7].

3.2 Key Lemma and Proofs of Theorem 1 and Theorem 2

The following is our key lemma to prove Theorem 1. We point out that it applies to arbitrary valuation functions and to any multi-unit auction which is *discriminatory price dominated*, i.e., the total payment $P_i(\mathbf{b})$ of bidder i under profile \mathbf{b} satisfies $P_i(\mathbf{b}) \leq \sum_{j \in [x_i(\mathbf{b})]} b_i(j)$. Note that every multi-unit auction guaranteeing *individual rationality* must satisfy this condition.

Lemma 1 (Key Lemma). *Let \mathbf{v} be a valuation profile and suppose that the pricing rule is discriminatory price dominated. Define $\tau_i = \arg \min_{j \in [x_i^{\mathbf{v}}]} v_i(j)/j$ for every $i \in [n]$. Then for every bidder $i \in [n]$ and every bidding profile \mathbf{b}_{-i} there exists a randomized uniform bidding profile \mathbf{b}'_i such that for every $\alpha > 0$*

$$\mathbb{E}[u_i^{\mathbf{v}_i}(\mathbf{b}'_i, \mathbf{b}_{-i})] \geq \alpha \left(1 - \frac{1}{e^{1/\alpha}} \right) x_i^{\mathbf{v}} \frac{v_i(\tau_i)}{\tau_i} - \alpha \sum_{j=1}^{x_i^{\mathbf{v}}} \beta_j(\mathbf{b}_{-i}). \quad (2)$$

Proof. Define $B = (1 - e^{-1/\alpha})$ and let \mathbf{c}_i be the vector that is $v_i(\tau_i)/\tau_i$ on the first $x_i^{\mathbf{v}}$ entries, and is 0 everywhere else. Let t be a random variable drawn from $[0, B]$ with probability density function $f(t) = \alpha/(1 - t)$. Define the random deviation of bidder i as $\mathbf{b}'_i = t\mathbf{c}_i$. Note that \mathbf{b}'_i is always a uniform bid vector.

Let k^* be the number of items that bidder i would win in profile $(B\mathbf{c}_i, \mathbf{b}_{-i})$, i.e., the number of items won by i , when i would deviate to bid vector $B\mathbf{c}_i$. For $j = 0, \dots, k^*$, let γ_j refer to the infimum value in $[0, B]$ such that bidder i would win j items if he would deviate to bid vector $\gamma_j\mathbf{c}_i$. Note that this definition is equivalent to defining γ_j as the least value in $[0, B]$ that satisfies $\gamma_j v_i(\tau_i)/\tau_i = \beta_j(\mathbf{b}_{-i})$. For notational convenience, we define $\gamma_{k^*+1} = B$.

Let $x_i(\mathbf{b}'_i, \mathbf{b}_{-i})$ be the random variable that denotes the number of units allocated to bidder i under $(\mathbf{b}'_i, \mathbf{b}_{-i})$. It always holds that $x_i(\mathbf{b}'_i, \mathbf{b}_{-i}) \leq k^* \leq x_i^y$, because bidder i bids $b'_i(j) = 0$ for all $j = x_i^y + 1, \dots, k$. More precisely, we have $x_i(\mathbf{b}'_i, \mathbf{b}_{-i}) = j$ if $t \in (\gamma_j, \gamma_{j+1}]$ for $j = 0, \dots, k^*$. By assumption, the payment of bidder i under profile $(\mathbf{b}'_i, \mathbf{b}_{-i})$ is at most $tx_i(\mathbf{b}'_i, \mathbf{b}_{-i})v_i(\tau_i)/\tau_i$. Also note that, by definition of τ_i , it holds that $v_i(j) \geq jv_i(\tau_i)/\tau_i$ for $j \leq x_i^y$. Using these two facts, we can bound the expected utility of bidder i as follows:

$$\begin{aligned} \mathbb{E}[u_i^{v_i}(\mathbf{b}'_i, \mathbf{b}_{-i})] &\geq \sum_{j=1}^{k^*} \int_{\gamma_j}^{\gamma_{j+1}} \left(v_i(j) - tj \frac{v_i(\tau_i)}{\tau_i} \right) f(t) dt \\ &\geq \sum_{j=1}^{k^*} \int_{\gamma_j}^{\gamma_{j+1}} j \frac{v_i(\tau_i)}{\tau_i} (1-t) f(t) dt = \alpha \sum_{j=1}^{k^*} j \frac{v_i(\tau_i)}{\tau_i} \int_{\gamma_j}^{\gamma_{j+1}} 1 dt \\ &= \alpha \sum_{j=1}^{k^*} j \frac{v_i(\tau_i)}{\tau_i} (\gamma_{j+1} - \gamma_j) = \alpha B k^* \frac{v_i(\tau_i)}{\tau_i} - \alpha \sum_{j=1}^{k^*} \gamma_j \frac{v_i(\tau_i)}{\tau_i} \\ &= \alpha B k^* \frac{v_i(\tau_i)}{\tau_i} - \alpha \sum_{j=1}^{k^*} \beta_j(\mathbf{b}_{-i}) \geq \alpha B x_i^y \frac{v_i(\tau_i)}{\tau_i} - \alpha \sum_{j=1}^{x_i^y} \beta_j(\mathbf{b}_{-i}). \end{aligned}$$

The last inequality holds because $Bv_i(\tau_i)/\tau_i \leq \beta_j(\mathbf{b}_{-i})$, for $k^* + 1 \leq j \leq x_i^y$, by the definition of k^* . The above derivation implies (2). \square

The deviation \mathbf{b}'_i defined in Lemma 1 is a distribution on uniform bidding strategies. That is, the lemma applies to both the standard and the uniform bidding format. Observe also that \mathbf{b}'_i satisfies the no-overbidding assumption.

Proof (of Theorem 1). First consider the case of submodular valuation functions. In this case, $\tau_i = x_i^y$ for every $i \in [n]$, as explained in Section 2. Using our Key Lemma, we conclude that Theorem 3 holds for $(\lambda, \mu) = (\alpha(1 - e^{-1/\alpha}), \alpha)$. The stated bounds are obtained by choosing $\alpha = 1$ for the Discriminatory Auction and $\alpha = -1/(W_{-1}(-1/e^2) + 2) \approx 0.87$ for the Uniform Price Auction.

Next consider the case of subadditive valuation functions. The following lemma shows that subadditive valuation functions can be approximated by uniform ones, thereby losing at most a factor 2.

Lemma 2. *If v_i is subadditive, $\tau_i = \arg \min_{j \in [x_i^y]} \frac{v_i(j)}{j}$ yields $\frac{v_i(\tau_i)}{\tau_i} \geq \frac{1}{2} \frac{v_i(x_i^y)}{x_i^y}$.*

By combining Lemma 2 with our Key Lemma, it follows that Theorem 3 holds for $(\lambda, \mu) = (\frac{\alpha}{2}(1 - e^{-1/\alpha}), \alpha)$. The bounds stated in Theorem 1 are obtained by the same choices of α for the submodular valuation functions. \square

Next, consider subadditive valuations under the *standard* bidding format. We derive improved bounds of 2 and 4 for the Discriminatory and Uniform Price Auction, respectively. To this end, we adapt an approach recently developed by Feldman *et al.* [7] to establish an analog of our Key Lemma. The main idea is to construct the bid \mathbf{b}'_i by using the distribution \mathcal{P}_{-i} on the profiles \mathbf{b}_{-i} . Theorem 2 then follows from Theorem 3 in combination with Lemma 3 below.

Lemma 3. *Let V be the class of subadditive valuation functions. Then Theorem 3 holds true with $(\lambda, \mu) = (\frac{1}{2}, 1)$ for the Discriminatory and $(\lambda, \mu) = (\frac{1}{2}, 1)$ for the Uniform Price Auction (under the standard bidding format).*

3.3 Lower Bounds

A lower bound of approximately $\frac{e}{e-1}$ for Uniform Price Auctions with submodular bidders was given in [14]; our upper bound is less than a factor 2 away. For subadditive valuation functions, we prove a lower bound of almost 2:

Theorem 4. *The Price of Anarchy is at least $\frac{2k}{k+1}$ for the Uniform Price Auction with subadditive valuations (under the uniform bidding interface).*

No lower bound is known for the Discriminatory Auction, although *Demand Reduction* (which is responsible for welfare loss in this format) has been observed previously [12,1]. In light of this, we prove here an *impossibility result* showing that for the Discriminatory Auction no bound better than $\frac{e}{e-1}$ on the Price of Anarchy can be achieved via the proof template given in Theorem 3. Similarly, for the Uniform Price Auction we rule out that a bound better than 2 on the Price of Anarchy can be derived through this template.

Theorem 5. *There is a lower bound of $\frac{e}{e-1}$ and 2 on the Bayesian Price of Anarchy for the Discriminatory Auction and the Uniform Price Auction, respectively, with submodular valuation functions that can be derived through the proof template given in Theorem 3.*

Theorem 5 rules out the possibility of obtaining better bounds by means of the smoothness framework of [19], or by means of *any* approach aiming at identifying the \mathbf{b}'_i required by Theorem 3, including [7]. These are essentially the only known techniques for obtaining upper bounds on the Bayesian Price of Anarchy. Thus, any improvement on our upper bound for the Discriminatory Auction must use either specific properties of the (Bayes-Nash equilibrium) distribution \mathcal{D} , or a completely new approach altogether. The same holds for improvements of the upper bound for the Uniform Price Auction below 2 – and towards the only known lower bound of $\frac{e}{e-1}$ from [14] (should it be worst-case).

4 Smoothness and Its Implications

We elaborate on the connections of our results to the smoothness framework for auction mechanisms, which has very recently been developed by Syrgkanis and Tardos [19]. We first review the smoothness definitions introduced in [19] (adapted to our multi-unit auction setting). As introduced earlier, let $P_i(\mathbf{b})$ refer to the payment of bidder i under bidding profile \mathbf{b} .

Definition 2 ([19]). *A mechanism \mathcal{M} is (λ, μ) -smooth for $\lambda > 0$ and $\mu \geq 0$ if for any valuation profile \mathbf{v} and for any bidding profile \mathbf{b} there exists a randomized bidding profile $\mathbf{b}'_i = \mathbf{b}'_i(\mathbf{v}, \mathbf{b}_{-i})$ for each i such that*

$$\sum_{i \in [n]} \mathbb{E}[u_i^{\mathbf{v}_i}(\mathbf{b}'_i, \mathbf{b}_{-i})] \geq \lambda SW(\mathbf{v}, \mathbf{x}^{\mathbf{v}}) - \mu \sum_{i \in [n]} P_i(\mathbf{b}).$$

Table 2. Upper bounds on the Bayesian Price of Anarchy for compositions

Valuation Functions	Discriminatory Auction		Uniform Price Auction
	Simultaneous	Sequential	Simultaneous/Sequential
Submodular	$e/(e - 1)$	$2e/(e - 1)$	3.1462
Subadditive	$2e/(e - 1)$	$4e/(e - 1)$	6.2924

In [19] it is shown that if a mechanism is (λ, μ) -smooth, then several results follow automatically. One such result concerns upper bounds on the Price of Anarchy. Another result is that the smoothness property is retained under *simultaneous and sequential compositions*. In these compositions there are m mechanisms with separate allocation and payment rules. Every bidder specifies for each mechanism a bidding profile. In the simultaneous composition, these profiles are submitted simultaneously, while in the sequential composition, they are submitted sequentially. A bidder expresses his valuation for the m -tuples of outcomes of the mechanisms in a restricted way.³ We summarize the main composition results of Syrgkanis and Tardos [19] in the theorem below.

Theorem 6 (Theorems 4.2, 4.3, 5.1, and 5.2 in [19]).

- (i) If \mathcal{M} is (λ, μ) -smooth, then the correlated (or mixed Bayesian) Price of Anarchy of \mathcal{M} is at most $\max\{1, \mu\}/\lambda$.
- (ii) If \mathcal{M} is a simultaneous (respectively, sequential) composition of m (λ, μ) -smooth mechanisms, then \mathcal{M} is (λ, μ) -smooth (resp., $(\lambda, \mu + 1)$ -smooth).

By exploiting our Key Lemma, we can show that the Discriminatory Auction is smooth. Theorem 7 in combination with Theorem 6 leads to the composition results stated in Table 2 (these bounds are achieved for $\alpha = 1$).

Theorem 7. *The Discriminatory Auction is (λ, μ) -smooth (both in the standard and uniform bidding format) with*

- (i) $(\lambda, \mu) = (\alpha(1 - e^{-1/\alpha}), \alpha)$ for submodular valuation functions, and
- (ii) $(\lambda, \mu) = (\frac{\alpha}{2}(1 - e^{-1/\alpha}), \alpha)$ for subadditive valuation functions.

For auction mechanisms where one needs to impose a no-overbidding assumption, a different smoothness notion is introduced in [19]. Given a mechanism \mathcal{M} , define bidder i 's *willingness-to-pay* as the maximum payment he could ever pay conditional to being allocated x units, i.e., $B_i(\mathbf{b}_i, x) = \max_{\mathbf{b}_{-i}:x_i(\mathbf{b})=x} P_i(\mathbf{b})$.

Definition 3 ([19]). *A mechanism \mathcal{M} is weakly (λ, μ_1, μ_2) -smooth for $\lambda > 0$ and $\mu_1, \mu_2 \geq 0$ if for any valuation profile \mathbf{v} and for any bidding profile \mathbf{b} there exists a randomized bidding profile $\mathbf{b}'_i = \mathbf{b}'_i(\mathbf{v}, \mathbf{b}_i)$ for each bidder i such that*

$$\sum_{i \in [n]} \mathbb{E}[u_i^{\mathbf{y}^i}(\mathbf{b}'_i, \mathbf{b}_{-i})] \geq \lambda SW(\mathbf{v}, \mathbf{x}^{\mathbf{y}}) - \mu_1 \sum_{i \in [n]} P_i(\mathbf{b}) - \mu_2 \sum_{i \in [n]} B_i(\mathbf{b}_i, x_i(\mathbf{b})).$$

³ More precisely, in the simultaneous composition it is assumed that the valuation function of each bidder is *fractionally subadditive* across the m mechanisms (see [19] for formal definitions). In the sequential composition, the valuation function of each bidder is defined as the maximum of his valuations over these mechanisms.

Syrkkanis and Tardos [19] establish the following results.

Theorem 8 (Theorems 7.4, C.4 and C.5 in [19]).

- (i) If \mathcal{M} is (λ, μ_1, μ_2) -weakly smooth, then the correlated (or mixed Bayesian) Price of Anarchy of \mathcal{M} is at most $(\mu_2 + \max\{1, \mu_1\})/\lambda$.
- (ii) If \mathcal{M} is a simultaneous (resp., sequential) composition of m (λ, μ_1, μ_2) -weakly smooth mechanisms, then \mathcal{M} is (λ, μ_1, μ_2) -weakly smooth (resp., $(\lambda, \mu_1 + 1, \mu_2)$ -weakly smooth).

Using our Key Lemma, we can show that the Uniform Price Auction is weakly smooth. As a consequence, we obtain the composition results stated in Table 2 (these bounds are achieved for $\alpha = -1/(W_{-1}(-1/e^2) + 2) \approx 0.87$).

Theorem 9. *The Uniform Price Auction is weakly (λ, μ_1, μ_2) -smooth (both in the standard and uniform bidding format) with*

- (i) $(\lambda, \mu_1, \mu_2) = (\alpha(1 - e^{-1/\alpha}), 0, \alpha)$ for submodular valuation functions, and
- (ii) $(\lambda, \mu_1, \mu_2) = (\frac{\alpha}{2}(1 - e^{-1/\alpha}), 0, \alpha)$ for subadditive valuation functions.

Some additional results on mechanisms with *budgets* (see [19]) can be inferred from Theorems 7 and 9. We defer further details to the full version of the paper.

5 Conclusions

We derived inefficiency upper bounds in the incomplete information model for the widely popular Discriminatory and Uniform Price Auctions, when bidders have submodular or subadditive valuation functions. Notably, our bounds for subadditive valuation functions already improve upon the ones that were known for submodular bidders [14,19]. Moreover, for each of the two formats and valuation function classes we considered both the *standard* bidding interface [12,15] and a practically motivated *uniform* bidding interface.

To derive our results, we elaborated on several techniques from the recent literature on *Simultaneous Auctions* [19,7,5,2]. By the recent developments of [19], our bounds for submodular bidders yield improved inefficiency bounds for *simultaneous* and *sequential* compositions of the considered formats. In absence of an indicative lower bound in the incomplete information model, we showed that our upper bound of $\frac{e}{e-1}$ for the Discriminatory Auction with submodular valuation functions is best possible, w.r.t. the currently known proof techniques. Additionally, for the Uniform Price Auction (with submodular bidders), we showed that, proving an upper bound of less than 2, also requires novel techniques; this poses a particularly challenging problem, given the lower bound of $\frac{e}{e-1}$ from [14].

References

1. Ausubel, L., Cramton, P.: Demand Reduction and Inefficiency in Multi-Unit Auctions. Tech. rep., University of Maryland (2002)
2. Bhawalkar, K., Roughgarden, T.: Welfare Guarantees for Combinatorial Auctions with Item Bidding. In: Proc. of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 700–709 (2011)

3. Binmore, K., Swierzbinski, J.: Treasury auctions: Uniform or discriminatory? Review of Economic Design 5(4), 387–410 (2000)
4. Bresky, M.: Pure Equilibrium Strategies in Multi-unit Auctions with Private Value Bidders. Tech. Rep. 376, CERGE Economics Institute (2008)
5. Christodoulou, G., Kovács, A., Schapira, M.: Bayesian Combinatorial Auctions. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part I. LNCS, vol. 5125, pp. 820–832. Springer, Heidelberg (2008)
6. Engelbrecht-Wiggans, R., Kahn, C.M.: Multi-unit auctions with uniform prices. Economic Theory 12(2), 227–258 (1998)
7. Feldman, M., Fu, H., Gravin, N., Lucier, B.: Simultaneous Auctions are (almost) Efficient. In: Proc. of the 45th ACM Symposium on Theory of Computing (STOC), pp. 201–210 (2013)
8. Friedman, M.: A Program for Monetary Stability. Fordham University Press, New York (1960)
9. Fu, H., Kleinberg, R., Lavi, R.: Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In: Proc. of the 13th ACM Conference on Electronic Commerce, p. 586 (2012)
10. Hassidim, A., Kaplan, H., Mansour, Y., Nisan, N.: Non-price equilibria in markets of discrete goods. In: Proc. of the 12th ACM Conference on Electronic Commerce, pp. 295–296 (2011)
11. de Keijzer, B., Markakis, E., Schäfer, G., Telelis, O.: On the Inefficiency of Standard Multi-Unit Auctions. arXiv:1303.1646 [cs.GT] (2013)
12. Krishna, V.: Auction Theory. Academic Press (2002)
13. Lehmann, B., Lehmann, D.J., Nisan, N.: Combinatorial auctions with decreasing marginal utilities. Games and Economic Behavior 55(2), 270–296 (2006)
14. Markakis, E., Telelis, O.: Uniform Price Auctions: Equilibria and Efficiency. In: Serna, M. (ed.) SAGT 2012. LNCS, vol. 7615, pp. 227–238. Springer, Heidelberg (2012)
15. Milgrom, P.: Putting Auction Theory to Work. Cambridge University Press (2004)
16. Noussair, C.: Equilibria in a multi-object uniform price sealed bid auction with multi-unit demands. Economic Theory 5, 337–351 (1995)
17. Ockenfels, A., Reiley, D.H., Sadrieh, A.: Economics and Information Systems. In: Handbooks in Information Systems, ch. 12. Online Auctions, vol. 1, pp. 571–628. Elsevier (2006)
18. Roughgarden, T.: The price of anarchy in games of incomplete information. In: Proc. of the 13th ACM Conference on Electronic Commerce, pp. 862–879 (2012)
19. Syrgkanis, V., Tardos, E.: Composable and efficient mechanisms. In: Proc. of the 45th ACM Symposium on Theory of Computing (STOC), pp. 211–220 (2013)
20. Syrgkanis, V.: Bayesian games and the smoothness framework. arXiv:1203.5155 [cs.GT] (2012)
21. U.S. Department of Treasury: Uniform-Price Auctions: Update of the Treasury Experience (1998),
<http://www.treasury.gov/press-center/press-releases/Documents/upas.pdf>
22. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. Journal of Finance 16(1), 8–37 (1961)