Towards a Characterization of Worst Case Equilibria in the Discriminatory Price Auction

Evangelos Markakis[1], Alkmini Sgouritsa[2], and Artem Tsikiridis[1]

[1] Athens University of Economics and Business, Dept. of Informatics, Athens, Greece
{markakis, artem}@aueb.gr
[2] University of Liverpool, Dept. of Computer Science, Liverpool, UK
asgouritsa2@liverpool.ac.uk

Abstract. We study the performance of the discriminatory price auction under the uniform bidding interface, which is one of the popular formats for running multi-unit auctions in practice. We undertake an equilibrium analysis with the goal of characterizing the inefficient mixed equilibria that may arise in such auctions. We consider bidders with capped-additive valuations, which is in line with the bidding format, and we first establish a series of properties that help us understand the sources of inefficiency. Moving on, we then use these results to derive new lower and upper bounds on the Price of Anarchy of mixed equilibria. For the case of two bidders, we arrive at a complete characterization of inefficient equilibria and show an upper bound of 1.1095, which is also tight. For multiple bidders, we show that the Price of Anarchy is strictly worse, improving the best known lower bound for submodular valuations. We further present an improved upper bound of 4/3 for the special case where there exists a "high" demand bidder. Finally, we also study Bayes-Nash equilibria, and exhibit a separation result that had been elusive so far. Namely, already with two bidders, the Price of Anarchy for Bayes-Nash equilibria is strictly worse than that for mixed equilibria. Such separation results are not always true (e.g., the opposite is known for simultaneous second price auctions) and reveal that the Bayesian model here introduces further inefficiency.

1 Introduction

Multi-unit auctions form a popular transaction means for selling multiple units of a single good. They have been in use for a long time, and there are by now several practical implementations across many countries. Some of the most prominent applications involve government sales of treasury securities to investors [6], as well as electricity auctions (for distributing electrical energy) [18]. Apart from governmental use, they are also run in other financial markets, and they are being deployed by various online brokers [16]. In the economics literature, multi-unit auctions have been a subject of study ever since the seminal work of Vickrey [23], and some formats were conceived even earlier, by Friedman [10].

The focus of our work is on the welfare performance of the *discriminatory price auction*, which is also referred to as pay-your-bid auction. In particular, we study the *uniform bidding interface*, which is the format most often employed in practice. Under this format, each bidder submits two parameters, a monetary per-unit bid, along with an upper bound on the number of units desired. Hence, each bidder is essentially asked to declare a *capped-additive* curve (a special case of submodular functions). The auctioneer then allocates the units by satisfying first the demand of the bidder with the highest monetary bid, then moving to the second highest bid, and so on, until there are no units left. As a price, each winning bidder pays his bid multiplied by the number of units received.

It is easy to see that the discriminatory price auction is not a truthful mechanism, and the same holds for other formats used in practice. Consequently, in the more recent years, a series of works have studied the social welfare guarantees that can be obtained at equilibrium. The outcome of these works is quite encouraging for the discriminatory price auction. Namely, pure Nash equilibria are always efficient, whereas for mixed and Bayes-Nash equilibria, the Price of Anarchy is bounded by 1.58 [13] for submodular valuations. These results suggest that simple auction formats can attain desirable guarantees and provide theoretical grounds for the overall success in practice.

Despite these positive findings, there has been no progress on further improving the current Price of Anarchy bounds. The known lower bound of 1.109 by [8] is quite far from the upper bounds derived by the commonly used smoothness-based approaches, [13, 22], which however do not seem applicable for producing further improvements. We believe the main difficulty in getting tighter results is that one needs to delve more deeply into the properties of Nash equilibria. But obtaining any form of characterization results for mixed or Bayesian equilibria is a notoriously hard problem. Even with two bidders it is often difficult to describe how the set of equilibria looks like. This is precisely the focus of our work, where we manage to either partially or fully characterize equilibrium profiles towards obtaining improved Price of Anarchy bounds, as we outline below.

1.1 Contribution

Motivated by the previous discussion, in Section 3 we initiate an equilibrium analysis for mixed equilibria. We consider bidders with capped-additive valuations, which is a subclass of submodular valuations, and consistent with the bidding format. Our results can be seen as a partial characterization of inefficient mixed equilibria, and our major highlights include both structural properties on the demand profile (see Theorem 3), as well as properties on the distributions of the mixed strategies (see Corollary 2, Theorem 4 and Lemma 7).

In Section 4, we use these results to derive new lower and upper bounds on the Price of Anarchy for mixed equilibria. For two bidders, we arrive at a complete characterization of inefficient equilibria and show an upper bound of 1.1095, which is tight.[1] For multiple bidders, we show that the Price of Anarchy is

[1] In [8] there is a lower bound of 1.109 that applies to our setting with two bidders and three units. The lower bound we provide here is just slightly better, but most importantly, it is tight and can be seen as a generalization of the instance in [8] to many units.

strictly worse, which also improves the best known lower bound for submodular valuations [8]. We further present an improved upper bound of $4/3$ for the special case where there exists a "high" demand bidder. We believe these latter instances are representative of the worst-case inefficiency that may arise, and refer to the relevant discussion in Section 4.2. To summarize, our results show that in several cases, the Price of Anarchy is even lower than the previous bound of [13] and strengthen the perception that such auctions can work well in practice.

Finally, in Section 5, we also study Bayes-Nash equilibria, and we exhibit a separation result that had been elusive so far: already with two bidders, the Price of Anarchy for Bayes-Nash equilibria is strictly worse than for mixed equilibria. Such separation results, though intuitive, do not hold for all auction formats. For example, in simultaneous second price auctions with submodular valuations [7], the known tight bounds for mixed equilibria extend to the Bayesian model via smoothness arguments [19]. This reveals that the Bayesian model in our setting introduces a further source of inefficiency. Note that to obtain this result, we transform the underlying optimization of social welfare at equilibrium to a well-posed variational calculus problem. This technique may be of independent interest and have other applications in mechanism design.

1.2 Related Work

The work of [1] was among the first ones that studied the sources of inefficiency in multi-unit auctions. For the discriminatory price auction, the Price of Anarchy was later studied in [22], and for bidders with submodular valuations, the currently best upper bound of $e/(e-1) \approx 1.58$ has been obtained by [13] (both for mixed and for Bayes-Nash equilibria). These results exploit the smoothness-based techniques, developed by [19,22]. One can also obtain slightly worse upper bounds for subadditive valuations, by using a different methodology, based on [9]. As for lower bounds, the only construction known for submodular valuations is by [8], yielding a bound of at least 1.109. In parallel to these results, there has been a series of works on the inefficiency of many other auction formats, ranging from multi-unit to combinatorial auctions, see among others, [4,5,7,9].

Apart from social welfare guarantees, several other aspects or properties of equilibrium behavior have been studied. Recently in [17], a characterization of equilibria is given for a model where the supply of units can be drawn from a distribution. In the past, several works have focused on revenue equivalence results between the discriminatory price and the uniform price auction, see e.g. [2,20]. On a different direction, comparisons from the perspective of the bidders are carried out in [3].

For a more detailed exposition on multi-unit auctions and their earlier applications, we refer the reader to the books [14] and [15]. For more recent applications, we refer to [6,11,18], for treasury bonds, carbon licence auctions, and electricity auctions, respectively.

2 Notation and Definitions

We consider a discriminatory price multi-unit auction, involving the allocation of k identical units of a single item, to a set $\mathcal{N} = \{1, \ldots, n\}$ of bidders. Each bidder $i \in \mathcal{N}$ has a private value $v_i > 0$, which reflects her value per unit and a private demand $d_i \in \mathbb{Z}_+$ which reflects the maximum number of units bidder i requires. Therefore, if the auction allocates $x_i \leq k$ units to bidder i, her total value will be $\min\{x_i, d_i\} \cdot v_i$. We note that this class of valuations is a subclass of submodular valuations, and includes all additive vectors (when $d_i = k$). We will refer to them as *capped-additive* valuations.

We focus on the following simple format for the discriminatory price auction, which is known as the *uniform bidding* interface. The auctioneer asks each bidder $i \in \mathcal{N}$ to submit a tuple (b_i, q_i), where $b_i \geq 0$, is her monetary bid per unit (not necessarily equal to v_i), and q_i is her demand bid (not necessarily equal to d_i). We denote by $\mathbf{b} = (b_1, \ldots, b_n)$ the monetary bidding vector, and similarly $\mathbf{q}$ will be the declared demand vector. For a bidding profile $(\mathbf{b}, \mathbf{q})$, the auctioneer allocates the units by satisfying first the demand of the bidder with the highest monetary bid, then moving to the second highest bid, and so on, until there are no units left. Hence, all the winners have their reported demand satisfied, except possibly for the one selected last, who may be partially satisfied. Moreover, we assume that in case of ties, a deterministic tie-breaking rule is used, which does not depend on the input bids submitted by the players to the auctioneer (e.g., a fixed ordering of the players suffices).

For every bidding profile $(\mathbf{b}, \mathbf{q})$, we let $x_i(\mathbf{b}, \mathbf{q})$ be the number of units allocated to bidder i, where obviously $x_i(\mathbf{b}, \mathbf{q}) \leq q_i$. In the discriminatory auction, the auctioneer requires each bidder i to pay b_i per allocated unit, hence a total payment of $b_i \cdot x_i(\mathbf{b}, \mathbf{q})$. The utility function of bidder $i \in \mathcal{N}$, given a bidding profile $(\mathbf{b}, \mathbf{q})$, is: $u_i(\mathbf{b}, \mathbf{q}) = \min\{x_i(\mathbf{b}, \mathbf{q}), d_i\}v_i - x_i(\mathbf{b}, \mathbf{q})b_i$.

Viewed as games, these auctions have an infinite pure strategy space, and we also allow bidders to play mixed strategies, which are probability distributions over their set of pure strategies. When each bidder $i \in \mathcal{N}$ uses a mixed strategy G_i, she *independently* draws a bid (b_i, q_i) from G_i. We refer to $\mathbf{G} = \times_{i=1}^{n} G_i$ as the product distribution of bids. Under mixed strategies, the expected utility of a bidder i is $\mathbb{E}_{(\mathbf{b},\mathbf{q})\sim\mathbf{G}}[u_i(\mathbf{b}, \mathbf{q})]$.

Definition 1. *We say that* $\mathbf{G}$ *is a mixed Nash equilibrium when for all* $i \in \mathcal{N}$*, all* $b'_i \geq 0$ *and all* $q'_i \in \mathbb{Z}_+$

$$\mathop{\mathbb{E}}_{(\mathbf{b},\mathbf{q})\sim\mathbf{G}}[u_i(\mathbf{b}, \mathbf{q})] \geq \mathop{\mathbb{E}}_{(\mathbf{b}_{-i},\mathbf{q}_{-i})\sim\mathbf{G}_{-i}}[u_i((b'_i, \mathbf{b}_{-i}), (q'_i, \mathbf{q}_{-i}))] \,.$$

We note that in any equilibrium, if a bidder i declares with positive probability a bid that exceeds v_i, she should not be allocated any unit, since such strategies are strictly dominated by bidding the actual value v_i.

Fact 1 *Let* $\mathbf{G}$ *be a mixed Nash equilibrium. The probability that a bidder* i *is allocated some units, conditioned that she bids higher than* v_i*, is 0.*

In the sequel, we focus on equilibria, where the monetary bids never exceed the value per unit.

Given a valuation profile $(\mathbf{v}, \mathbf{d})$, we denote by $OPT(\mathbf{v}, \mathbf{d})$ the optimal social welfare (which can be computed very easily by running the allocation algorithm of the auction with the true value and demand vector). We also denote by $SW(\mathbf{G})$ the expected social welfare of a mixed Nash equilibrium $\mathbf{G}$, i.e., equal to $\mathbb{E}_{(\mathbf{b},\mathbf{q})\sim\mathbf{G}}[\sum_i \min\{x_i(\mathbf{b}, \mathbf{q}), d_i\}v_i]$. The Price of Anarchy is the worst-case ratio $\frac{OPT(\mathbf{v},\mathbf{d})}{SW(\mathbf{G})}$, over all valuation profiles $(\mathbf{v}, \mathbf{d})$, and all equilibria $\mathbf{G}$.

We refer to an equilibrium as *inefficient* when its social welfare is strictly less than the optimal.

3 Towards a Characterization of Inefficient Mixed Equilibria

In this section, we derive a series of important properties, that help us understand better how can inefficient equilibria arise. These properties will help us analyze the Price of Anarchy in Section 4.

3.1 Mixed Nash Equilibria with Demand Revelation

Our first result is that it suffices to focus on equilibria where bidders truthfully reveal their demand, resulting therefore in a single-parameter strategy space for the bidders (Theorem 1). We further argue that the inefficiency in equilibria appears only when the total demand exceeds k (Lemma 1) and therefore this is what we assume for the rest of the paper.

Theorem 1. *Let $(\mathbf{v}, \mathbf{d})$ be a valuation profile, and $\mathbf{G}$ be a mixed Nash equilibrium. Then, for every $i \in \mathcal{N}$, and in every pure strategy profile $(b_i, q_i) \sim G_i$, we can replace q_i by d_i so that the resulting distribution remains a mixed Nash equilibrium with the same social welfare.*

Lemma 1. *If $\sum_i d_i \leq k$ then the social welfare of any mixed Nash equilibrium is optimal.*

3.2 Existence of Non-empty-handed Bidders

For the rest of the paper we consider only strategy profiles where the bidders' demand bid matches their true demand. The main goal of this subsection is to derive Theorem 3, where we show that in any inefficient mixed equilibrium, there always exists a bidder such that the total demand of the other winners is strictly less than k, meaning that at least one item is allocated to him for sure (with probability one). This is a crucial property for understanding the formation of inefficient mixed equilibria. To proceed, we give first some further notation to be used in this and the following sections.

Further notation. Given Theorem 1, instead of using distributions on tuples (b_i, q_i), we suppose that each bidder $i \in \mathcal{N}$ independently draws only a monetary bid b_i from a distribution B_i and we refer to $\mathbf{B} = \times_{i=1}^{n} B_i$ as the product distribution of monetary bids or just bids from now on. For a bidding profile $\mathbf{b}$, the utility of a bidder i will simply be denoted as $u_i(\mathbf{b})$, instead of $u_i(\mathbf{b}, \mathbf{d})$. Definition 1 is also simplified, and we say that $\mathbf{B}$ is an equilibrium if $\mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})] \geq \mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i((b_i', \mathbf{b}_{-i}))]$, for any i and any $b_i' \geq 0$. Similarly, the social welfare of a mixed Nash equilibrium $\mathbf{B}$ is given by just $SW(\mathbf{B})$.

For a mixed strategy bidding profile $\mathbf{B}$, we denote by $W(\mathbf{B})$ the set of bidders with positive expected utility, i.e., $W(\mathbf{B}) = \{j : \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_j(\mathbf{b})] > 0\}$, and let $\mathbf{B}_W = \times_{i \in W(\mathbf{B})} B_i$. Moreover, the support of a bidder i in $\mathbf{B}$ is the domain of the distribution B_i, that i plays under $\mathbf{B}$, denoted by $Supp(B_i)$. We denote by $\ell(B_i), h(B_i)$ the leftmost and rightmost points, respectively, in the support of bidder i. In particular, if the rightmost part of the domain of B_i is a mass point b or an interval in the form $[a, b]$, then $h(B_i) = b$, and similarly for $\ell(B_i)$. In cases of distributions over intervals, we can safely assume that the domain contains only *closed* intervals, because the endpoints are chosen with zero probability. We further denote by $\ell(\mathbf{B}_W), h(\mathbf{B}_W)$ the leftmost and rightmost points, respectively, of the union of the supports of $W(\mathbf{B})$.

For $i = 1, \ldots, n$ we denote by F_i the CDF of B_i and by f_i their PDF. Moreover, given a profile $\mathbf{b}$, it is often useful in the analysis to consider the vector of bids (thresholds) that a bidder i competes against, denoted by $\beta(\mathbf{b})_{-i} = (\beta_1(\mathbf{b}_{-i}), \ldots, \beta_k(\mathbf{b}_{-i}))$. Here, $\beta_j(\mathbf{b}_{-i})$ is the j-th lowest winning bid of the profile $\mathbf{b}_{-i}$, for $j = 1, \ldots, k$, so that $\beta(b)_{-i}$ describes the winning bids if i didn't participate. This implies that, under profile $\mathbf{b}$, bidder i is allocated $j = 1, \ldots, k-1$ units capped by d_i, when $\beta_j(\mathbf{b}_{-i}) < b_i < \beta_{j+1}(\mathbf{b}_{-i})$ and d_i units, when $\beta_k(\mathbf{b}_{-i}) < b_i$. We note that because we focus on the uniform bidding interface, some consecutive β_j values may coincide and be equal to the bid of the same bidder. When $\mathbf{b}_{-i} \sim \mathbf{B}_{-i}$, for $i = 1, \ldots, n$, we denote the CDF of the random variable $\beta_j(\mathbf{b}_{-i})$ as $\hat{F}_{ij}$, for $j = 1, \ldots, k$. In the next fact, we express the expected allocation of any bidder i for bidding some $\alpha > 0$, in terms of the values $\hat{F}_{ij}(\alpha)$.

Fact 2 *Let $\mathbf{B}_{-i}$ be a product distribution of bids. Then for all $\alpha \geq 0$, where no bidder other than (possibly) i has a mass point,* $\mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[x_i(\alpha, \mathbf{b}_{-i})] = \sum_{j=1}^{d_i} \hat{F}_{ij}(\alpha)$.

Given a bidding profile $\mathbf{B}$, for any bidder i we define $\hat{F}_i^{avg}(x) = \frac{\sum_{j=1}^{d_i} \hat{F}_{ij}(x)}{d_i}$, to be the average CDF of the winning bids that bidder i competes against. Note that $\hat{F}_i^{avg}$ is a CDF since it is the average of a number of CDFs.

Remark 1. The $\hat{F}_{ij}$ functions are right continuous, as they are CDFs, and moreover, if the F_i functions have no mass point, the same holds for the $\hat{F}_{ij}$ functions. Additionally, if for any j, the $\hat{F}_{ij}$ functions are continuous, so is $\hat{F}_i^{avg}$, as the average of continuous functions.

We start by ruling out certain scenarios that cannot occur at inefficient equilibria. First, we can safely ignore bidders with zero expected utility, since in any *inefficient* mixed Nash equilibrium they do not receive any units.

Lemma 2. *Any mixed Nash equilibrium* $\mathbf{B}$ *with at least one bidder with zero expected utility, but positive expected number of allocated units, is efficient.*

Next, we show that to have inefficiency at an equilibrium, there must exist at least two bidders with positive expected utility.

Lemma 3. *Let* $(\mathbf{v}, \mathbf{d})$ *be a valuation profile and* $\mathbf{B}$ *be an inefficient mixed Nash equilibrium. Then,* $|W(\mathbf{B})| \geq 2$.

The next warm-up properties involve the expected utility of a bidder under an equilibrium $\mathbf{B}$, conditioned that she bids within a certain interval or at a single point. We start with Fact 3, which is a straightforward implication of the equilibrium definition, and proceed by arguing that no two bidders may bid on the same point with positive probability. Theorem 2 concludes by stating the main property regarding the utility of bidders when bidding in their support.

Fact 3 *Let* $\mathbf{B}$ *be an equilibrium. For a bidder* i*, consider a partition of* $Supp(B_i)$ *(or of a subset of it) into smaller disjoint sub-intervals, say* $I_1, \ldots, I_\ell$*, such that* B_i *has a positive probability on each sub-interval (mass points may also be considered as sub-intervals). Then, it should hold that* $\mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b}) \mid b_i \in I_r] = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})]$*, for every* $r = 1, \ldots, \ell$.

Based on Fact 3, we can obtain the following point-wise version. Variations of the version below have also appeared in related works, see e.g., [8].

Theorem 2. *Given a mixed Nash equilibrium* $\mathbf{B}$*, bidder* i *and* $z \in Supp(B_i)$*, where no other bidder has a mass point on* z*,* $\mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(z, \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})]$.

We further give the following observation regarding the existence of mass points on $\ell(\mathbf{B}_W)$.

Observation 1 *In any inefficient mixed Nash equilibrium* $\mathbf{B}$*, there can be no bidders* $i, j \in W(\mathbf{B})$ *such that both* $Pr[b_i = \ell(\mathbf{B}_W)] > 0$ *and* $Pr[b_j = \ell(\mathbf{B}_W)] > 0$.

The main theorem of this section follows, stating the existence of a special bidder, who always receives at least one unit, and is referred to as *non-empty-handed.*

Theorem 3. *Let* $(\mathbf{v}, \mathbf{d})$ *be a valuation profile, and let* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium. Then, there exists a bidder* $i \in W(\mathbf{B})$*, such that*

$$\sum_{j \in W(\mathbf{B}) \setminus \{i\}} d_j \leq k - 1\,.$$

Proof. On the contrary, suppose that for every $i \in W(\mathbf{B})$, $\sum_{j \in W(\mathbf{B}) \setminus \{i\}} d_j \geq k$. Let i be some bidder with $\ell = \ell(\mathbf{B}_W) \in Supp(B_i)$. We distinguish two cases.
Case 1: There exists an interval in the form $[\ell, \ell+\epsilon]$, on which B_i has a positive probability mass and on which the bidders of $W(\mathbf{B}) \setminus \{i\}$ have a zero mass. We note that we also allow $\epsilon = 0$, i.e., that i has a mass point on ℓ and the other bidders do not. This means that when bidder i bids within $[\ell, \ell+\epsilon]$, all the other bidders from $W(\mathbf{B})$ are above him. Since we assumed that the total demand of $W(\mathbf{B}) \setminus \{i\}$ is at least k, bidder i does not win any units in this case. Since i bids with positive probability in $[\ell, \ell+\epsilon]$, by Fact 3, we have $\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] = 0$, which contradicts the fact that $i \in W(\mathbf{b})$.
Case 2: Note that by Observation 1, it cannot happen that both bidder i and at least one bidder $j \in W \setminus \{i\}$, have a mass point on ℓ. Hence, the only remaining case to consider is that any mass point that may exist by the bidders is at some $x > \ell$, and there is also no interval starting from ℓ that is used only by bidder i. Thus, there exists an interval I in the form $I = [\ell, \ell+\epsilon]$ for some small enough $\epsilon > 0$, and a bidder $j \in W(\mathbf{B}) \setminus \{i\}$, such that both B_i and B_j contain I in their support, and have positive probability mass on I without mass points.

By Theorem 2, we obtain that $\mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(\ell, \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] > 0$. This is a contradiction, because by bidding ℓ, bidder i ranks lower than all other bidders of $W(\mathbf{B})$ with probability one. By our assumption that $\sum_{j \in W(\mathbf{B}) \setminus \{i\}} d_j \geq k$, there are no units left for i when she ranks last among $W(\mathbf{B})$, and therefore, $\mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(\ell, \mathbf{b}_{-i})] = 0$. □

The property above already implies the following interesting corollary, that if all bidders have unit demand, any mixed Nash equilibrium is efficient.

Corollary 1. *Let* $(\mathbf{v}, \mathbf{d})$ *be a valuation profile with only unit-demand bidders, i.e.,* $d_i = 1$ *for all* i. *Then any mixed Nash equilibrium* $\mathbf{B}$ *is efficient.*

3.3 The support and the CDFs of Mixed Nash Equilibria

The existence of a non-empty-handed bidder (Theorem 3) helps us to establish further properties that characterize the structure of inefficient mixed Nash equilibria. These properties (and especially Theorem 4) will be important to establish the inefficiency results that follow. We start with an observation regarding the highest bid of any bidder $i \in W(\mathbf{B})$, which should be strictly less than v_i.

Observation 2 *For any bidder* $i \in W(\mathbf{B})$, $h(B_i) < v_i$.

The next lemma shows that at any equilibrium $\mathbf{B}$, bidders who are not non-empty-handed cannot have higher bids in their support than the support of the non-empty-handed bidders. Moreover, any bidder who is non-empty-handed does not have a reason to use bids that are higher than the maximum bid of all other winning bidders. The reason is that if such differences existed, then there would be incentives to win the same number of units by lowering one's bid. Then, Lemma 5 shows that no bidder will bid alone at any point or interval, and Lemma 6 specifies that no mass points may exist apart from one case.

Lemma 4. *Let* $(\mathbf{v}, \mathbf{d})$ *be a valuation profile and* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium. Then, for any non-empty-handed bidder* i*, it holds that* $h(B_i) = h(\mathbf{B}_{W\setminus\{i\}}) = h(\mathbf{B}_W)$*.*

Lemma 5. *Let* $(\mathbf{v}, \mathbf{d})$ *be any valuation profile and* $\mathbf{B}$ *be any mixed Nash equilibrium. For all* $i \in W(\mathbf{B})$*, it holds that* $Supp(B_i) \subseteq \bigcup_{j\in W(\mathbf{B})\setminus\{i\}} Supp(B_j)$*.*

Lemma 6. *Let* $(\mathbf{v}, \mathbf{d})$ *be a valuation profile and* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium.*
1) There exists no bidder $i \in W(\mathbf{B})$ *and no point* $z \in Supp(B_i) \setminus \{\ell(\mathbf{B}_W)\}$*, with* $F_i(z) > \lim_{z\to z^-} F_i(z)$*, i.e., there are no mass points among the bidders of* $W(\mathbf{B})$*, except possibly the leftmost endpoint of all bidders' distributions.*
2) At most one bidder $i \in W(\mathbf{B})$ *may have a mass point on* $\ell(\mathbf{B}_W)$*, in which case,* i *is a non-empty-handed bidder.*

By combining Theorem 2 and Lemma 6 we get the following Corollary.

Corollary 2. *For any inefficient mixed Nash equilibrium* $\mathbf{B}$*, the following hold:*
1) For any bidder i *and* $z \in Supp(B_i) \setminus \{\ell(\mathbf{B}_W)\}$*,* $\mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(z, \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})]$*.*
2) If there exists a bidder i *with* $Pr[b_i = \ell(\mathbf{B}_W)] > 0$*, then* i *is a non-empty-handed bidder and* $\mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(\ell(\mathbf{B}_W), \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})]$*.*
3) If no non-empty-handed bidder exists with mass point on $\ell(\mathbf{B}_W)$*, for any bidder* i *with* $\ell(\mathbf{B}_W) \in Supp(B_i)$*,* $\mathbb{E}_{\mathbf{b}_{-i}\sim\mathbf{B}_{-i}}[u_i(\ell(\mathbf{B}_W), \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})]$*.*

Observation 3 *For any inefficient mixed Nash equilibrium* $\mathbf{B}$*, either there exists a non-empty-handed bidder* $i \in W(\mathbf{B})$ *with a mass point on* $\ell(\mathbf{B}_W)$*, or there are at least two non-empty-handed bidders with* $\ell(\mathbf{B}_W)$ *in their support.*

Given any (inefficient) equilibrium, the next theorem specifies the average CDF of the winning bids that bidder i competes against, i.e., $\hat{F}_i^{avg}$, in i's support.

Theorem 4. *Let* $(\mathbf{v}, \mathbf{d})$ *be any valuation profile and* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium. Then, for* $i \in W(\mathbf{B})$*, the CDF* $\hat{F}_i^{avg}$ *satisfies*

$$\hat{F}_i^{avg}(z) = \frac{u_i}{d_i(v_i - z)}\,, \quad \forall z \in Supp(B_i)\,,$$

where $u_i = \mathbb{E}_{\mathbf{b}\sim\mathbf{B}}[u_i(\mathbf{b})] > 0$*.*

A corollary of Theorem 4 is that the union of the support of the winners is an interval.

Corollary 3. *Let* $(\mathbf{v}, \mathbf{d})$ *be any valuation profile and* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium. Then, for every bidder* $i \in W(\mathbf{B})$*,* $\bigcup_{j\in W(\mathbf{B})\setminus\{i\}} Supp(B_j) = [\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$*.*

The final lemma of this section shows that the rightmost point in the support of $\mathbf{B}$ is a function of the parameters of certain non-empty-handed bidders.

Lemma 7. *Let* $(\mathbf{v}, \mathbf{d})$ *be any valuation profile and* $\mathbf{B}$ *be any inefficient mixed Nash equilibrium. Let* $i \in W(\mathbf{B})$ *be the non-empty-handed bidder such that* $Pr[b_i = \ell(\mathbf{B}_W)] > 0$*, or if no such bidder exists, then let* i *be any non-empty-handed bidder with* $\ell(\mathbf{B}_W)$ *in his support. We have*

$$h(\mathbf{B}_W) = h(B_i) = v_i - (k - \sum_{j \in W(\mathbf{B}) \setminus \{i\}} d_j) \frac{v_i - \ell(\mathbf{B}_W)}{d_i} .$$

4 Price of Anarchy for mixed equilibria

We can now exploit the properties derived so far for mixed equilibria, in order to analyze the inefficiency of the discriminatory price auction. Since we focus on inefficient equilibria, we assume that in any valuation profile considered in this section, there are at least two bidders with a different value per unit.

4.1 The case of two bidders

We pay particular attention to the case of $n = 2$. This is a setting where we can fully characterize in closed form the distributions of the inefficient mixed Nash equilibria, and derive valuable intuitions for the worst-case instances with respect to the Price of Anarchy, that are helpful also for auctions with multiple bidders. The main result of this subsection is the following theorem, showing that the inefficiency is quite limited.

Theorem 5. *For* $k \geq 2$*,* $n = 2$ *and capped additive valuation profiles, the Price of Anarchy of mixed equilibria is at most* 1.1095*, and this is tight as* k *goes to infinity.*

We postpone the proof of Theorem 5, as we first need to establish some properties regarding the form of inefficient mixed Nash equilibria with two bidders. For $n = 2$, a capped-additive valuation profile can be described as $(\mathbf{v}, \mathbf{d}) = ((v_1, d_1), (v_2, d_2))$. Recall also that it is sufficient to focus our attention only on profiles where $d_1 + d_2 > k$, since otherwise, by Lemma 1 any mixed equilibrium is efficient. We start our analysis by characterizing the support of inefficient mixed Nash equilibria.

Lemma 8. *Let* $(\mathbf{v}, \mathbf{d}) = ((v_1, d_1), (v_2, d_2))$ *be any capped-additive valuation profile of two bidders, and* $\mathbf{B} = (B_1, B_2)$ *be any inefficient mixed Nash equilibrium. Then:*

1. *Supp*$(B_1) =$ *Supp*$(B_2) = [\ell(B_1), h(B_1)]$*, and* $\ell(B_1) = 0$.
2. $h(B_1)$ *takes one of the following values*

$$h(B_1) = v_1 \frac{d_1 + d_2 - k}{d_1} \quad or \quad h(B_1) = v_2 \frac{d_1 + d_2 - k}{d_2} .$$

The following theorem specifies the cumulative distribution functions that comprise any inefficient mixed Nash equilibrium, along with a necessary condition for the existence of such equilibria. For a bidder i below, we use the notation v_{-i} and d_{-i} to denote the value and demand of the other bidder.

Theorem 6. *Let* $(\mathbf{v}, \mathbf{d}) = ((v_1, d_1), (v_2, d_2))$ *be a capped-additive valuation profile of two bidders, and* $\mathbf{B} = (B_1, B_2)$ *be any inefficient mixed Nash equilibrium.*

1. *The cumulative distribution function of bidder i, for $i = 1, 2$, is*

$$F_i(z) = \frac{1}{d_1 + d_2 - k} \left(\frac{d_{-i}(v_{-i} - h(B_i))}{v_{-i} - z} - (k - d_i) \right). \tag{1}$$

2. *Furthermore, for i being the non-empty-handed bidder with a mass point at 0, or if no such bidder exists, being any non-empty-handed bidder, it holds that* $\frac{v_{-i}}{v_i} \geq \frac{d_{-i}}{d_i}$,

Remark 2. By Lemma 8 and Theorem 6, we can see that there can be at most two inefficient equilibria, depending on how the interval of the support was determined.

We are now ready to prove Theorem 5.

Proof sketch of Theorem 5. The properties established so far imply a full characterization of instances that have inefficient equilibria. To establish Theorem 5, we will group instances into three appropriate classes and we will solve an appropriately defined optimization problem that approximates the Price of Anarchy for each subclass to arbitrary precision.

Suppose without loss of generality that we are given a value profile $(\mathbf{v}, \mathbf{d}) = ((v_1, d_1), (v_2, d_2))$ of k units, such that $d_1 \geq d_2 > 0$. Let $\bar{d}_1 := \frac{d_1}{k}$ and $\bar{d}_2 = \frac{d_2}{k}$, be the *normalized* demands of the bidders. Essentially, we intend to use $v_1, v_2, \bar{d}_1$ and $\bar{d}_2$ as the variables of the optimization problem mentioned before.

Let $\mathbf{B}$ be any inefficient mixed Nash equilibrium. With a slight abuse of notation we view the term $h(B_i)$ as a function of the valuation profile parameters, as established by Lemma 8, and define the functions $h_i(\mathbf{v}, \bar{\mathbf{d}}) = v_i \frac{\bar{d}_1 + \bar{d}_2 - 1}{\bar{d}_i}$ for $i = 1, 2$. Our goal now is to express the social welfare of $\mathbf{B}$, solely in terms of the value profile parameters, $(\mathbf{v}, \mathbf{d})$ and k, and without dependencies on the underlying equilibrium distributions. To proceed, we define first two auxiliary functions; namely, for $i = 1, 2$, we let $S_i(\mathbf{v}, \bar{\mathbf{d}})$ be equal to:

$$\bar{d}_{-i}(v_{-i} - v_i) \left(1 - \int_0^{h_i(\mathbf{v}, \bar{\mathbf{d}})} \frac{1}{\bar{d}_1 + \bar{d}_2 - 1} \left(\frac{\bar{d}_i(v_i - h_i(\mathbf{v}, \bar{\mathbf{d}}))}{v_i - z} - (1 - \bar{d}_{-i}) \right) \frac{v_{-i} - h_i(\mathbf{v}, \bar{\mathbf{d}})}{(v_{-i} - z)^2} dz \right) + v_i .$$

With these expressions in mind, the following lemma allows us to obtain the social welfare in a form that we can later exploit for producing our upper bound. The lemma follows by Theorem 6, which tells us what the equilibrium CDFs are, in terms of the valuation profile.

Lemma 9. *Let i be a non-empty handed bidder with a mass point at 0. Then,* $SW(\mathbf{B}) = kS_i(\mathbf{v}, \bar{\mathbf{d}})$*. If no such bidder exists, then either* $SW(\mathbf{B}) = kS_1(\mathbf{v}, \bar{\mathbf{d}})$ *or* $SW(\mathbf{B}) = kS_2(\mathbf{v}, \bar{\mathbf{d}})$.

To conclude the proof of the upper bound, we solve a sequence of optimization problems as determined by the cases arising in the statement of Lemma 9, and by the ordering of the values v_1, v_2. By solving these problems numerically, we found out that in the worst case instance $v_1 = 1, v_2 \approx 0.526, \bar{d}_1 = 1, \bar{d}_2 \approx 0.357$. It is not hard to convert the variables to the underlying worst case instance, which we present in the next paragraph.

Tight Example. Consider an instance of the discriminatory auction for $k \geq 4$ units and $n = 2$ bidders. Bidder 1 has value $v_1 = 1$ and $d_1 = k$, whereas bidder 2 has a value $v_2 = 0.526$ and $d_2 = \lceil 0.357k \rceil$ units. Let B_1, B_2 be two distributions supported in $[0, \frac{d_2}{k}]$. Note that $v_2 > \frac{d_2}{k}$. In accordance to Equation (1), the cumulative distribution functions of B_1 and B_2 are

$$F_1(z) = \frac{v_2 - \frac{d_2}{k}}{v_2 - z}, \qquad F_2(z) = \frac{k - d_2}{d_2} \frac{z}{1 - z}.$$

It is easy to verify that $\mathbf{B} = (B_1, B_2)$ is indeed a mixed equilibrium. The optimal allocation is for bidder 1 to obtain all k units and the expected social welfare of $\mathbf{B}$, by Lemma 9, is $SW(\mathbf{B}) = kS_1(\mathbf{v}, \bar{\mathbf{d}})$, since $F_1(0) > 0$. The worst case inefficiency ratio occurs as k grows and is approximately 1.1095. □

4.2 Multiple Bidders

Inspired by the construction in the previous section, we move to instances with more than two bidders and provide first a lower bound on the Price of Anarchy. This bound shows a separation between $n = 2$ and $n > 2$, in the sense that equilibria can be more inefficient with a higher number of bidders. It also improves the best known lower bound of the discriminatory price auction for the class of submodular valuations, which was 1.109, by [8]. The improvement however is rather small.

Theorem 7. *For $n > 2$, and for the class of mixed strategy Nash equilibria, the Price of Anarchy is at least* 1.1204.

The above bound is the best lower bound we have been able to establish, even after some extensive experimentation (driven by the results in the remainder of this section). It is natural to wonder if there is a matching upper bound, which would establish that the Price of Anarchy remains very small even for a large number of bidders. Recall that from [13], we know already a bound of $e/(e-1) \approx 1.58$. Although we have not managed to settle this question, we will provide an improved upper bound for a special case, for which there is evidence that it captures worst-case scenarios of inefficiency. At the same time, we will be able to characterize the format of such worst case equilibria.

To obtain some intuition, it is instructive to look at the proofs of our two lower bounds, in Theorem 5 and in Theorem 7. One can notice that the main source of inefficiency is the fact that the auctioneer accepts multi-unit demand declarations. When this does not occur, we have already shown in Corollary 1 that mixed Nash equilibria attain optimal welfare. When multi-demand bidders are present, Theorem 5 shows that in the case of two bidders, the most inefficient mixed Nash equilibrium occurs when a participating bidder declares a demand for all the units, whereas the opponent requires a much smaller fraction of the supply. In the proof of Theorem 7 above, we have extended this paradigm for multiple bidders with an arbitrary demand structure, but under the assumption that one of the bidders requires all the units (the additive bidder). Such a setting, of one large-demand bidder facing competition by multiple small-demand bidders has also been discussed in [3]. Furthermore, there exist other auction formats that also needed such a demand profile at their worst case instances, see e.g., [5] for the uniform price auction. To summarize, it seems unlikely that the worst instances involve only bidders with low demand or small variation on their demands.

Given the above, we will analyze the family of instances where there exists an additive bidder (with demand equal to k), and where she also has the highest value per unit. In fact, the latter assumption is needed only for the Price of Anarchy analysis but not for the characterization of the worst-case demand profile and the equilibrium strategies. We strongly believe that this class is representative of the most inefficient mixed Nash equilibria (which is true already for the case of two bidders).

The main result of this section is the following.

Theorem 8. *Consider the class of valuation profiles, where there exists an additive bidder α with the highest value, and an equilibrium $\mathbf{B}$, such that $\alpha \in W(\mathbf{B})$. Then, the Price of Anarchy is at most $4/3$.*

The proof of the theorem is by following a series of steps. The existence of the additive bidder helps in the analysis, because a direct corollary of Theorem 3 is that the additive bidder is the sole non-empty-handed bidder (everyone else faces competition for all the units).

Corollary 4 (by Theorem 3). *Consider a valuation profile $(\mathbf{v}, \mathbf{d})$ with an additive bidder α, that admits an equilibrium $\mathbf{B}$, such that $\alpha \in W(\mathbf{B})$. Then, bidder α is the unique non-empty-handed bidder under $\mathbf{B}$, thus, $\sum_{i \in W(\mathbf{B}) \setminus \{\alpha\}} d_i \leq k-1$.*

To proceed, we ensure that for the instances described by Theorem 8, it suffices to analyze the equilibria where bidder α belongs to $W(\mathbf{B})$, i.e., there cannot exist a more inefficient equilibrium $\mathbf{B}'$ of these instances with $\alpha \notin W(\mathbf{B}')$. This is addressed by the following lemma.

Lemma 10. *Consider a valuation profile, and suppose that it admits two distinct inefficient equilibria, $\mathbf{B}$ and $\mathbf{B}'$. If $i \in W(\mathbf{B})$ is a non-empty-handed bidder in $\mathbf{B}$, then $i \in W(\mathbf{B}')$.*

Using Lemma 10 and Corollary 4, from now on, we fix a bidder α and an inefficient equilibrium $\mathbf{B}$, so that α is additive and $\alpha \in W(\mathbf{B})$.

Corollary 4 already gives us an insight about the competition in such an equilibrium $\mathbf{B}$. While bidder α will have to compete against the other bidders of $W(\mathbf{B})$ to win extra units, in addition to those that she is guaranteed to obtain, each bidder in $W(\mathbf{B}) \setminus \{\alpha\}$ only competes against α. Each of them is not guaranteed any units, unless she outbids α (bidder α is the only cause of externality for bidders in $W(\mathbf{B}) \setminus \{\alpha\}$, and anyone bidding lower than α cannot get any units). If bidder α did not exist, the other winners could be automatically granted the demand they are requesting since, in total, it is smaller than k and hence, there is no competition among them.

Observation 4 *$\hat{F}_i^{avg}(z) = F_\alpha(z)$, for every $i \in W(\mathbf{B}) \setminus \{\alpha\}$, where F_α is the CDF of bidder α.*

We continue with further properties on the support of the mixed strategies.

Lemma 11. *For the equilibrium $\mathbf{B}$ under consideration, it is true that:*

1. *$Supp(B_\alpha) = [\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$.*
2. *For any two bidders $i, j \in W(\mathbf{B}) \setminus \{\alpha\}$ such that $v_i \neq v_j$, the set $Supp(B_i) \cap Supp(B_j)$ is of measure 0 (intersection points can occur only at endpoints of intervals).*

Lemma 11 suggests that we can group the bidders according to their values (since only bidders with the same value can overlap in their support). Let $r \leq |W(\mathbf{B}) \setminus \{\alpha\}|$ represent the number of distinct values $v_1, \ldots, v_r$, that bidders in $W(\mathbf{B}) \setminus \{\alpha\}$ have. We can partition the bidders of $W(\mathbf{B}) \setminus \{\alpha\}$ into r groups $W_1(\mathbf{B}), \ldots, W_r(\mathbf{B})$, such that, for $j = 1, \ldots, r$, the bidders in group $W_j(\mathbf{B})$ have value v_j. Similarly, we split the support of the winning bidders $[\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$ into r intervals, i.e., $[\ell(\mathbf{B}_W), h(\mathbf{B}_W)] = \bigcup_{j=1}^{r} I_j(\mathbf{B})$, where each interval $j \in \{1, \ldots, r\}$ is formed as $I_j(\mathbf{B}) = \bigcup_{i \in W_j(\mathbf{B})} Supp(B_i)$. The following is a direct corollary of Lemma 11.

Corollary 5. *For every $s, t \in \{1, \ldots, r\}$ with $s \neq t$, the set $I_s(\mathbf{B}) \cap I_t(\mathbf{B})$ is of measure 0.*

When all bidders in $W(\mathbf{B}) \setminus \{\alpha\}$ have distinct values there are precisely $|W(\mathbf{B}) \setminus \{\alpha\}|$ intervals, whereas when they all have a common value, they must be bidding on the entire interval $[\ell(W(\mathbf{B})), h(W(\mathbf{B}))]$ (the equilibrium in the 2-bidder case when $d_1 = k$, in Section 4.1, is one such example). We sometimes denote as $I_0(\mathbf{B})$ the interval of losing bidders $[0, \ell(\mathbf{B}_W)]$, i.e., for the bidders in $\mathcal{N} \setminus W(\mathbf{B})$. Note that given $\mathbf{B}$, the only criterion for the membership of the support of a bidder i in an interval $I_s(\mathbf{B})$ is their value.

The next step is quite crucial in simplifying the extraction of our upper bound. We show that the worst case demand structure for the bidders in $W(\mathbf{B}) \setminus \{\alpha\}$ is when they all have unit demand.

Theorem 9. *For the value profile $(\mathbf{v}, \mathbf{d})$ and the equilibrium $\mathbf{B}$ under consideration, there exists another value profile $(\mathbf{v}', \mathbf{d}')$ and a product distribution $\mathbf{B}'$ such that*

1. *$\alpha \in W(\mathbf{B}')$ is an additive bidder and for every bidder $i \in W(\mathbf{B}') \setminus \{\alpha\}$, it holds that $d'_i = 1$.*
2. *$\mathbf{B}'$ is a mixed Nash equilibrium for $(\mathbf{v}', \mathbf{d}')$.*
3. *$\frac{OPT(\mathbf{v},\mathbf{d})}{SW(\mathbf{B})} = \frac{OPT(\mathbf{v}',\mathbf{d}')}{SW(\mathbf{B}')}$.*

For the remainder of the section, it suffices to analyze valuation profiles, that possess equilibria where the members of $W(\mathbf{B})$ are either additive or unit-demand. Recall, that due to Corollary 4, there must be a unique additive bidder. Hence, we fix an instance given by a valuation profile $(\mathbf{v}, \mathbf{d})$, so that at the equilibrium $\mathbf{B}$, the set $W(\mathbf{B})$ consists of n unit-demand bidders plus the additive bidder α, i.e., $n = |W(\mathbf{B}) \setminus \{\alpha\}|$. Moreover, due to the following observation we may assume, without loss of generality, that the support of each unit-demand bidder has no overlapping intervals with other bidders from $W(\mathbf{B}) \setminus \{\alpha\}$.

Lemma 12. *Let $(\mathbf{v}, \mathbf{d})$ be a value profile, and let $\mathbf{B}$ be any mixed Nash equilibrium, such that the members of $W(\mathbf{B})$ are all unit-demand bidders aside from one additive bidder. Then, there exists a mixed Nash equilibrium $\mathbf{B}'$ with disjoint support intervals such that $SW(\mathbf{B}) = SW(\mathbf{B}')$.*

Therefore, by Corollary 5 and the discussion preceding it, the support of each bidder $i = 1, \ldots, n$ is $[\ell(B_i), h(B_i)]$. Note that due to Lemma 11, the unit-demand bidders must cover the entire interval $[\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$. Hence, for a unit-demand bidder $i = 1, \ldots, n$, it must be that $\ell(B_i) = h(B_{i-1})$, assuming for convenience that $h(B_0) = \ell(\mathbf{B}_W)$.

The next theorem provides a more complete understanding of the support intervals and the distributions of the equilibrium $\mathbf{B}$.

Theorem 10. *For the value profile $(\mathbf{v}, \mathbf{d})$ under consideration, the following properties hold:*

1. *For bidder α, we have $h(B_\alpha) = h(B_n) = h(\mathbf{B}_W) = v_\alpha - (k-n)\frac{v_\alpha - \ell(B_\alpha)}{k}$. Moreover, for every unit-demand bidder $i = 1, \ldots, n-1$ it holds that*

$$\ell(B_{i+1}) = h(B_i) = v_\alpha - \frac{(k-n)(v_\alpha - \ell(B_\alpha))}{k-n+i}.$$

2. *The CDF F_α of bidder α, is a branch function, so that for $i = 1, \ldots, n$, $F_\alpha(z) = F_\alpha^i(z)$ for every $z \in [h(B_{i-1}), h(B_i)]$ with*

$$F_\alpha^i(z) = \prod_{j=i+1}^{n} \left(\frac{v_j - h(B_j)}{v_j - h(B_{j-1})} \right) \frac{v_i - h(B_i)}{v_i - z}.$$

Before proving our upper bound, we present two additional lemmas. The first is a straightforward inequality, that is a direct consequence of the definition of a mixed equilibrium, and the second is an expression for the social welfare. Both of these are useful for obtaining our final Price of Anarchy upper bound.

Lemma 13. *Consider a value profile* $(\mathbf{v}, \mathbf{d})$*, and any inefficient mixed Nash equilibrium* $\mathbf{B}$*, with* $W(\mathbf{B})$ *consisting only of additive or unit-demand bidders. Then, for* $i = 2, \ldots, n$*,* $m = 1, \ldots, i-1$*, and every* $z \in [h(B_{m-1}), h(B_m)]$*,*

$$\prod_{j=m+1}^{i-1} \frac{v_j - h(B_j)}{v_i - h(B_{j-1})} \leq \frac{v_m - z}{v_m - h(B_m)} \frac{v_i - h(B_{i-1})}{v_i - z} \,. \tag{2}$$

Lemma 14. *Consider a value profile* $(\mathbf{v}, \mathbf{d})$*, and any inefficient mixed Nash equilibrium* $\mathbf{B}$*, with* $W(\mathbf{B})$ *consisting only of additive or unit-demand bidders. The expected social welfare is*

$$kv_\alpha - (k-n)(v_\alpha - \ell(B_\alpha)) \sum_{i=1}^{n} \prod_{j=i+1}^{n} \left(\frac{v_j - h(B_j)}{v_j - h(B_{j-1})} \right) \int_{h(B_{i-1})}^{h(B_i)} \frac{v_i - h(B_i)}{v_i - z} \frac{v_\alpha - v_i}{(v_a - z)^2} dz \,.$$

Proof of Theorem 8. For brevity, we denote $\ell(B_a)$ as ℓ and for $j = 1, \ldots, n$, we denote $h(B_j)$ as h_j. Moreover, by assumption $v_a \geq v_n$. To simplify the calculations, we assume that $v_a = 1$ by rescaling all values in the instance.

Given a mixed Nash equilibrium $\mathbf{B}$, we lower bound the expected social welfare $SW(\mathbf{B})$ described in the equation of Lemma 14 as

$$\begin{aligned}
SW(\mathbf{B}) &= k - (k-n)(1-\ell) \sum_{i=1}^{n} \prod_{j=i+1}^{n} \left(\frac{v_j - h_j}{v_j - h_{j-1}} \right) \int_{h_{i-1}}^{h_i} \frac{v_i - h_i}{v_i - z} \frac{1 - v_i}{(1-z)^2} dz \\
&= k - (k-n)(1-\ell) \sum_{i=1}^{n} \prod_{j=i+1}^{n} \left(\frac{v_j - h_j}{v_j - h_{j-1}} \right) \\
&\qquad \left(\int_{h_{i-1}}^{h_i} \frac{v_i - h_i}{v_i - z} \frac{1}{(1-z)} dz - \int_{h_{i-1}}^{h_i} \frac{v_i - h_i}{(1-z)^2} dz \right) \\
&> k - (k-n)(1-\ell) \sum_{i=1}^{n} \prod_{j=i+1}^{n} \left(\frac{v_j - h_j}{v_j - h_{j-1}} \right) \int_{h_{i-1}}^{h_i} \frac{v_i - h_i}{v_i - z} \frac{1}{(1-z)} dz \\
&\geq k - (k-n)(1-\ell) \int_{\ell}^{h_n} \frac{v_n - h_n}{(v_n - z)(1-z)} dz \\
&\geq k - (k-n)(1-\ell) \int_{\ell}^{h_n} \frac{1 - h_n}{(1-z)^2} dz \geq k - (k-n)(1-\ell) \\
&= k - (k-n)(h_n - \ell) = k - (k-n) \left(\frac{n}{k}(1-\ell) \right) \geq k - \frac{(k-n)n}{k} \geq \frac{3}{4} k \,.
\end{aligned}$$

The first inequality is true since for all bidders $i = 1, \ldots, n$, it holds that $v_i > h_i$ by Observation 2. The second one is an application of the mixed Nash equilibrium property encoded by Equation (2) of Lemma 13. The next two inequalities occur by observing that the respective functions are increasing in terms of v_n (which, by assumption, we upper bound with $v_n \leq 1$) and ℓ (which we lower bound with $\ell \geq 0$). The last inequality follows by setting $x = \frac{n}{k}$ and minimizing the function $s(x) = 1 - x + x^2$ for $x \in (0, 1)$. The theorem follows by observing that the optimal welfare is k, since the additive bidder has the highest value. $\square$

5 A Separation between Mixed and Bayesian Cases

In this section we explore the more general solution concept of Bayes Nash equilibrium. We consider the following incomplete information setting. Let (v_i, d_i) be the type of bidder $i \in \mathcal{N}$. We suppose that the private value v_i of a bidder i is drawn independently from a distribution V_i. The second part of bidder i's type is his demand d_i; for the purposes of this section (we only construct a lower bound instance), we assume d_i to be deterministic private information.

Each bidder i is aware of her own value per unit v_i and the product distribution formed by the V_j's, and decides a strategy $(b_i, q_i) \sim G_i(v_i)$ for each value $v_i \sim V_i$. The bidding strategy is in general a mixed strategy. In the special case that bidder i chooses a single bid $(b_i(v_i), q_i)$ for each drawn value v_i, he submits a pure strategy, where q_i is not necessarily d_i.

Definition 2. *Given $\mathbf{V} = \times_{i=1}^{n} V_i$ and $\mathbf{d}$, a profile $\mathbf{G}(\mathbf{v})$ is a Bayes Nash equilibrium if for all $i \in \mathcal{N}$, v_i in V_i's domain, $b_i' \geq 0$ and $q_i' \in \mathbb{Z}_+$ it holds that*

$$\mathop{\mathbb{E}}_{\mathbf{v}_{-i} \sim \mathbf{V}_{-i}} \left[\mathop{\mathbb{E}}_{(\mathbf{b},\mathbf{q}) \sim \mathbf{G}(\mathbf{v})} [u_i^{v_i}(\mathbf{b}, \mathbf{q})] \right] \geq \mathop{\mathbb{E}}_{\mathbf{v}_{-i} \sim \mathbf{V}_{-i}} \left[\mathop{\mathbb{E}}_{(\mathbf{b}_{-i},\mathbf{q}_{-i}) \sim \mathbf{G}_{-i}(\mathbf{v}_{-i})} [u_i^{v_i}((b_i', q_i'), (\mathbf{b}_{-i}, \mathbf{q}_{-i}))] \right],$$

where $u_i^{v_i}(\cdot)$ stands for bidder i's utility when his value is v_i.

We can define the Bayesian Price of Anarchy in the same way as before, by comparing against the expected optimal welfare, over the value distributions.

Although in a few other auction formats, the inefficiency does not get worse when one moves to incomplete information games, we exhibit that this is not the case here. We present a lower bound on the Bayesian Price of Anarchy of 1.1204, with two bidders. For mixed equilibria and two bidders, Theorem 5 showed that the Price of Anarchy is at most 1.1095. Although this difference is small, it shows that the Bayesian model is more expressive and can thus create more inefficiency. In particular, we stress that the bound obtained here for two bidders is inspired by the same bound of 1.1204 for mixed equilibria in Theorem 7, where we had to use a large number of bidders.

Theorem 11. *For $n = 2$, $k \geq 2$, and capped additive valuation profiles, the Price of Anarchy of Bayes Nash equilibria is at least* 1.1204.

Remark 3. When $k = 1$, there is a lower bound of 1.15 in [12] for the first price auction. However this requires a very large number of bidders. There is a simpler construction with two bidders in [21] but it only yields a lower bound of 1.06.

Acknowledgements. This research was supported by the Hellenic Foundation for Research and Innovation (H.F.R.I.). The first author was supported under the "1st Call for H.F.R.I. Research Projects to support faculty members and researchers and the procurement of high-cost research equipment" grant (Project Number: HFRI-FM17-3512), and the third author was supported under the H.F.R.I. PhD Fellowship grant (Fellowship Number: 289).

References

1. Ausubel, L., Cramton, P.: Demand Reduction and Inefficiency in Multi-Unit Auctions. Tech. rep., University of Maryland (2002)
2. Ausubel, L., Cramton, P., Pycia, M., Rostek, M., Weretka, M.: Demand Reduction and Inefficiency in Multi-Unit Auctions. The Review of Economic Studies **81**, 1366–1400 (2014)
3. Baisa, B., Burkett, J.: Large Multi-Unit Auctions With a Large Bidder. J. Econ. Theory **174**, 1–15 (2018)
4. Bhawalkar, K., Roughgarden, T.: Welfare Guarantees for Combinatorial Auctions with Item Bidding. In: ACM-SIAM Symposium on Discrete Algorithms, SODA 2011. pp. 700–709 (2011)
5. Birmpas, G., Markakis, E., Telelis, O., Tsikiridis, A.: Tight Welfare Guarantees for Pure Nash Equilibria of the Uniform Price Auction. Theory Comput. Syst. **63**(7), 1451–1469 (2019)
6. Brenner, M., Galai, D., Sade, O.: Sovereign Debt Auctions: Uniform or Discriminatory? Journal of Monetary Economics **56**(2), 267–274 (2009)
7. Christodoulou, G., Kovács, A., Schapira, M.: Bayesian combinatorial auctions. J. ACM **63**(2), 11:1–11:19 (2016)
8. Christodoulou, G., Kovács, A., Sgouritsa, A., Tang, B.: Tight Bounds for the Price of Anarchy of Simultaneous First-Price Auctions. ACM TEAC **4**(2), 9:1–9:33 (2016)
9. Feldman, M., Fu, H., Gravin, N., Lucier, B.: Simultaneous Auctions Without Complements are (almost) Efficient. Games and Econ. Behavior **123**, 327–341 (2020)
10. Friedman, M.: A Program for Monetary Stability. Fordham University Press, New York, NY (1960)
11. Goldner, K., Immorlica, N., Lucier, B.: Reducing Inefficiency in Carbon Auctions with Imperfect Competition. In: Innovations in Theoretical Computer Science, ITCS 2020. pp. 15:1–15:21 (2020)
12. Hartline, J., Hoy, D., Taggart, S.: Price of Anarchy for Auction Revenue. In: ACM Conference on Economics and Computation, EC 2014. pp. 693–710 (2014)
13. de Keijzer, B., Markakis, E., Schäfer, G., Telelis, O.: Inefficiency of Standard Multi-unit Auctions. In: European Symposium on Algorithms, ESA 2013. pp. 385–396
14. Krishna, V.: Auction Theory. Academic Press (2002)
15. Milgrom, P.: Putting Auction Theory to Work. Cambridge University Press (2004)
16. Ockenfels, A., Reiley, D.H., Sadrieh, A.: Economics and Information Systems, chap. 12. Online Auctions, pp. 571–628 (2006)
17. Pycia, M., Woodward, K.: Auctions of Homogeneous Goods: A Case for Pay-as-Bid. In: ACM Conference on Economics and Computation, EC 2021 (2021)
18. Rio, P.D.: Designing Auctions for Renewable Electricity Support. Best Practices From Around the World. Energy for Sustainable Development **41**, 1–13 (2017)
19. Roughgarden, T.: The Price of Anarchy in Games of Incomplete Information. In: ACM Conference on Economics and Computation, EC 2012. pp. 862–879 (2012)
20. Swinkels, J.: Efficiency of Large Private Value Auctions. J. Econ. Theory **69**(1), 37–68 (2001)
21. Syrgkanis, V.: Efficiency of Mechanisms in Complex Markets. Ph.D. thesis, Cornell University (2014)
22. Syrgkanis, V., Tardos, E.: Composable and Efficient Mechanisms. In: ACM Symposium on Theory of Computing, STOC 2013. pp. 211–220 (2013)
23. Vickrey, W.: Counterspeculation, Auctions, and Competitive Sealed Tenders. Journal of Finance **16**(1), 8–37 (1961)