



**PROJECT: COMRADE – Combinatorial Reverse Auctions:
Design of Efficient Mechanisms for Crowdsourcing Markets**

Funded under the 1st Call for HFRI Research Projects to Support Faculty Members and Researchers and Procure High-Value Research Equipment

WP5: Design and analysis of mechanisms for combinatorial reverse auctions

***Deliverable D.5.2:** Technical report with our research results and conclusions on the design of mechanisms for combinatorial reverse auctions*

Prepared by the team members:

Michail Fasoulakis, Evangelos Markakis, Thanassis Papaioannou, Georgios Papasotiropoulos, Panagiotis Tsamopoulos



1 Introduction

This deliverable concerns WP5, and in particular, it is the last deliverable of WP5. It was completed by Month 34, in accordance to the amended timetable of COMRADE. The goal of this report is to provide a final overview on our research concerning the design of mechanisms for reverse auctions.

In WP4, we examined a class of fundamental algorithmic problems, concerning task assignment in procurement, which can be modeled as combinatorial optimization problems. In WP5, the goal has been to enhance these algorithms by taking into account the strategic behavior of the agents. I.e., agents may have incentives to misreport their actual costs so as to be eligible to receive higher payments. This is an extremely important concern in designing auction mechanisms, and we are therefore interested in algorithms that apart from determining the allocation of tasks to agents, also discourage agents from reporting false preferences and false costs. The remainder of this deliverable is structured as follows. In Section 2, we describe the two basic class of mathematical models that capture the problems of interest. In Sections 3 and 4, we then exhibit our research results on the design of truthful mechanisms for these models. In Section 5, we also report complementary research results on non-truthful mechanisms. Even though such mechanisms may provide incentives in worst case for strategic behavior, they have other advantages that make them popular in real applications. Finally in the Appendix, we provide our relevant publications on this topic.

2 Basic Mathematical Models for Auctions

For the sake of completeness, we recall the general setup and the algorithmic questions we have been interested in.

We consider a procurement auction setting, and let $A = [n] = \{1, 2, \dots, n\}$ denote a set of n agents who are willing to participate in the reverse auction. The auctioneer is interested in having some predefined set of tasks completed, possibly subject to certain constraints, and is therefore aiming at hiring some of the participating agents. There are various possibilities for modeling the capabilities of each agent, the requirements of the auctioneer, and the bidding interface so that the agents can make offers on the tasks they can complete. E.g., the auctioneer may simply want to assign each task to a single agent, or assign tasks multiple times (for quality assurance). Each agent may be capable of performing a single task out of the tasks posted by the auctioneer, or she may be able to bid for performing a subset of the tasks. The interface may also allow agents to make multiple bids on different combinations of tasks, or to make bids only for single tasks. The choices made for these aspects can lead to a variety of different models for task assignment.

Based on the current state of the art, and on our Deliverable D.2.1 from WP2, which performed a literature survey, we decided to work on three types of optimization problems that capture different aspects of task assignment in reverse auctions. We describe them separately below.

2.1 Utility maximization under budget constraints

We consider a generalization of a class of models that were originally introduced by Singer [8]. In these problems, the auctioneer asks each agent i to submit one parameter, denoted by c_i to the mechanism, corresponding to the cost for executing the tasks she is capable

of. This assumption can capture several different scenarios. E.g., for an agent i , her cost c_i could simply be the cost incurred to her for executing a single task, among the tasks advertised by the auctioneer. In other cases, it can denote the cost for executing a collection of tasks that she is capable of doing (and in which case the agent should declare along with c_i the list of these tasks). Regardless of the interpretation of the cost c_i , all these settings can be treated in a unified manner.

Furthermore, to evaluate the selected outcome, the auctioneer is also associated with a valuation function $v : P(A) \rightarrow \mathbb{R}$, where $P(A)$ is the powerset of the set of agents, A . Hence, $v(S)$ denotes the value (which can be thought of as the level of profit or happiness) derived by the auctioneer if she selects the set S of agents as the auction winners. It is also assumed that when a set S of agents is selected as the winners, these agents are committed to executing the tasks corresponding to their offer (this could be enforced by some penalty mechanism and it is beyond the scope of our work). Finally, the auctioneer also has a hard budget constraint, in the sense that the payments made by the auctioneer to the agents should not exceed a fixed parameter $B > 0$ that specifies the budget.

To summarize, the auctioneer collects the offers of the agents, and then needs to select a set $S \subseteq A$, as the winning agents, along with a payment to each winner. In WP4 we are interested only in the non-strategic version of these problems, where the agents are honest about their cost c_i , and therefore, we only care about the question of determining the allocation of the tasks, i.e., the set of winning agents. In doing so, a natural goal for the auctioneer under this model, is to design an algorithm for maximizing her attained value under the given constraints. In other words we are interested in finding a set of winning agents S , so as to maximize the value $v(S)$, subject to the budget constraints.

In our study, given the progress that has been made on the problem since the work of [8], we want to focus on two extensions of the original auction setting, where it is possible to hire an agent partially or to hire an agent for a particular quality of service.

Agents with Multiple Levels of Service: In this setting, each agent offers a service that consists of multiple levels. We can think of the levels as corresponding to different qualities of service. Hence, the auctioneer can choose not to hire an agent, or hire the first x number of levels of an agent, for some integer x , or hire the agent entirely, i.e., for all the levels that she is offering. Furthermore, the valuation function associated with each agent is concave, meaning that the marginal value of each level of service is non-increasing. This setting was first introduced by Chan and Chen [3] in the context of each agent offering multiple copies of the same good and each additional copy having a smaller marginal value. In their work it is assumed that the cost of a single level is arbitrary, meaning that it is plausible that the auctioneer can only afford to hire a single level of service of a single agent. The crucial difference with our setting is that we assume that the auctioneer's budget is big enough to afford any single individual agent entirely, which is in line with the indivisible setting in which the auctioneer can afford to hire any individual agent.

Divisible Agents: Another relevant setting is the setting in which agents are offering a divisible service, e.g., offering their time. In this case, it is reasonable to assume that the auctioneer can hire each agent for any fraction of the service that they are offering. Again, the valuation function associated with each agent is assumed to be concave, meaning that the marginal gain is non-increasing in the fraction of the acquired service. Note that this problem is the fractional relaxation of the problem introduced by Singer [8], when it is assumed that the auctioneer can afford to hire any individual agent entirely. Anari et al. [1] were the first to study the divisible setting.

2.2 Cost minimization in spatial crowdsourcing models

The last problem is motivated by spatial models in reverse auctions. There are three main differences compared to the previous two subsections. The first one is that we assume that the tasks lie in some geometric space, i.e., they may correspond to transportation tasks within a Euclidean space. The simplest example is to suppose that the tasks are ordered on a line. The second difference is that we now want to have all the tasks completed. This is unlike Section 2.1, where not all tasks are necessarily assigned, either due to the budget restrictions or due to the maximization of social welfare. The third difference is that each task has a demand, on the quality of its completion, and this implies that it may have to be assigned to multiple agents, to achieve its demand constraint. In applications where it is important to complete all the tasks listed by the auctioneer, it is more natural to focus more on producing a feasible solution of minimum cost. This gives rise to a cost minimization problem that we describe below.

Cost Minimization Interval Cover (CMIC): Consider a set of tasks, say $\{1, \dots, m\}$, that are ordered on a line, and a set of available bidders $A = \{1, \dots, n\}$. An instance of CMIC is determined by a tuple $(\mathbf{b}, \mathbf{q}, \mathbf{d})$, where:

- The vector $\mathbf{b} = (b_1, \dots, b_n)$ is the bidding profile. For each bidder $j \in A$, $b_j = (c_j, [s_j, f_j])$, where $[s_j, f_j] = I_j$ is the interval of contiguous tasks $\{s_j, s_j+1, \dots, f_j\} \subseteq [m]$, that j is able to contribute to, and $c_j \in \mathbb{R}_{\geq 0}$ is the cost incurred to her, if she is selected to contribute. We often denote $\mathbf{b}$ as $(\mathbf{c}, \mathbf{I})$ where $\mathbf{c}$ is the cost vector and $\mathbf{I}$ is the vector of the intervals.
- The vector $\mathbf{q} = (q_1, \dots, q_n)$ is the contribution vector, so that $q_j \in \mathbb{R}_{>0}$ denotes the contribution that bidder j can make to the tasks that belong to $[s_j, f_j]$.
- The vector $\mathbf{d} = (d_1, \dots, d_m)$ specifies the demand $d_j \in \mathbb{R}_{>0}$ of a task j .

The goal is to select a set of bidders $S \subseteq A$ of minimum cost, that satisfies

$$\sum_{i \in \mathcal{N}_j(\mathbf{I}) \cap S} q_i \geq d_j, \quad \forall j \in [m], \quad (1)$$

where for $j \in [m]$, $\mathcal{N}_j(\mathbf{I}) := \{i \in A \mid j \in [s_i, f_i]\}$, i.e., it is the set of bidders who can contribute to task j .

3 Mechanisms for utility maximization under budget constraints

Following the discussion in Section 2.1, we will present truthful mechanisms for the two settings that are defined there. We start with the scenario of agents with multiple levels of service.

3.1 Agents with multiple levels of service

We consider the following multiple service level model as a natural extension of the standard model introduced above (see also [3]). We refer to this model as the *k-level model* for short: suppose each agent $i \in N$ offers $k \geq 1$ levels of service and has an associated valuation function $v_i : \{0, \dots, k\} \rightarrow \mathbb{R}_{\geq 0}$ which is public information. Here, $v_i(j)$ denotes

the value that the j th level of service of agent i has to the auctioneer. Observe that in this setting, the total cost of agent i is linear, i.e., the cost of using $x_i = j$ service levels of agent i is $j \cdot c_i$.

Valuation Functions: Without loss of generality, we assume that each v_i is normalized such that $v_i(0) = 0$. We study the general class of *concave* valuation functions, i.e., for each agent i , $v_i(j) - v_i(j-1) \geq v_i(j+1) - v_i(j)$ for all $j = 1, \dots, k-1$. We also define the j -th marginal valuation of agent i as $m_i(j) := v_i(j) - v_i(j-1)$ for $j \in \{1, \dots, k\}$. Given a profile $\mathbf{c}$, the total value that the auctioneer derives from an allocation $\mathbf{x}$ is defined by the separable concave function $V(\mathbf{x}(\mathbf{c})) = \sum_{i \in N} v_i(x_i(\mathbf{c}))$.

Benchmark: The performance of a mechanism is measured by comparing $V(\mathbf{x}(\mathbf{c}))$ with the underlying (non-strategic) combinatorial optimization problem, which is commonly referred to as the *k-Bounded Knapsack Problem*¹ (see, e.g., [7] for a classification of knapsack problems):

$$\text{OPT}_I^k(\mathbf{c}, B) := \max \sum_{i=1}^n v_i(x_i), \quad \text{s.t.} \quad \sum_{i=1}^n c_i x_i \leq B, \quad x_i \in \{0, \dots, k\}, \forall i \in N. \quad (2)$$

The k -Bounded Knapsack Problem is NP-hard in general. We say that an algorithm is α -approximate with $\alpha \geq 1$ if $V(\mathbf{x}(\mathbf{c})) \geq \frac{1}{\alpha} \text{OPT}_I^k(\mathbf{c}, B)$. We also consider the *Fractional k-Bounded Knapsack Problem* that follows from the k -Bounded Knapsack Problem in (2) by relaxing the integrality constraint:

$$\begin{aligned} \text{OPT}_F^k(\mathbf{c}, B) &:= \max \sum_{i=1}^n v_i(x_i) + m_i(x_i)(x_i - x_i), \\ &\text{such that} \quad \sum_{i=1}^n c_i x_i \leq B, \quad x_i \in [0, k] \forall i \in N. \end{aligned}$$

Naturally, it also holds that $\text{OPT}_F^k(\mathbf{c}, B) \geq \text{OPT}_I^k(\mathbf{c}, B)$. Note that $\text{OPT}_F^1(\mathbf{c}, B)$ is the fractional relaxation of the well-known *Knapsack Problem*. It is not hard to see that $\text{OPT}_F^k(\mathbf{c}, B)$ inherits the well-known properties of its one-dimensional analogue, including the fact that an optimal solution can be computed by a simple greedy algorithm in polynomial time.

To summarize, an instance I of the k -level model will be denoted by a tuple $I = (N, \mathbf{c}, B, k, (v_i)_{i \in N})$. Whenever part of the input is clear from the context, we omit its explicit reference for conciseness (e.g., often we refer to instance simply by its corresponding cost vector $\mathbf{c}$).

3.1.1 A greedy mechanism

We propose a greedy mechanism for the k -level model, where the main underlying idea is as follows: If there is an agent i^* whose maximum valuation $v_{i^*}(k)$ is valuable enough (in a certain sense), then we simply pick all levels of service for this agent. Otherwise, we compute an allocation using the following greedy procedure: We first compute an optimal allocation $\mathbf{x}^*(\mathbf{c})$ to the corresponding Fractional k -Bounded Knapsack Problem (which can be done in polynomial time) and use the integral part of this solution as an initial allocation. Note that this allocation is close to the optimal fractional solution because $\mathbf{x}^*(\mathbf{c})$ has at most one fractional component. We then repeatedly discard the worst service level (in terms of marginal value-per-cost) of an agent from this allocation,

¹Note that for $k = 1$, the problem reduces to the well-known *0-1 Knapsack Problem*.

until the total value of our allocation would drop below an α -fraction of the optimal solution.

We need some more notation for the formal description of our mechanism: Given an allocation $\mathbf{x}$, we denote by $\ell(\mathbf{x})$ the agent whose $x_{\ell(\mathbf{x})}$ -th level of service is the least valuable in $\mathbf{x}$, in terms of their marginal value-per-cost ratio. Notice that due to the fact that the valuation functions are concave, the worst case marginal value-per-cost ratio indeed corresponds to the $x_{\ell(\mathbf{x})}$ -th ratio of agent $\ell(\mathbf{x})$. When $\mathbf{x}$ is clear from the context, we refer to this agent simply as ℓ . A detailed description of our greedy mechanism is given in Mechanism 1.

Our main result for this algorithm is the following theorem:

Theorem 3.1. *Mechanism SORT-&-REJECT(k) with $\alpha = \frac{1}{2+\sqrt{3}}$ is truthful, budget-feasible and $(2+\sqrt{3})$ -approximate for instances of the k -level model, and runs in time polynomial in n and k .*

Proof. The polynomial running time for the allocation is straightforward. So it suffices to establish the claimed approximation guarantee. For a cost profile $\mathbf{c}$ we will prove the claimed guarantee against the optimal solution to the fractional relaxation of the bounded knapsack instance, i.e., we will show that $V(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}_F^k(\mathbf{c})$. As mentioned earlier, this establishes our guarantee since $\text{OPT}_F^k(\mathbf{c}) \geq \text{OPT}_I^k(\mathbf{c})$. Again we distinguish the following two cases:

Case 1: $v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \cdot \text{OPT}_F^k(\mathbf{c}_{-i^*})$. We directly have

$$v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \cdot \text{OPT}_F^k(\mathbf{c}_{-i^*}) \geq \frac{\alpha}{1-\alpha} \left(\text{OPT}_F^k(\mathbf{c}) - v_{i^*}(k) \right),$$

where the second inequality follows from the fact that $\text{OPT}_F^k(\mathbf{c}_{-i}) + v_i(k) \geq \text{OPT}_F^k(\mathbf{c})$ for all $i \in N$. Rearranging terms leads to $v_{i^*}(k) \geq \alpha \text{OPT}_F^k(\mathbf{c})$ which concludes the proof since $v_{i^*}(k) = V(\mathbf{x}(\mathbf{c}))$.

Case 2: $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \cdot \text{OPT}_F^k(\mathbf{c}_{-i^*})$. In this case, it follows from the **while** condition of SORT-&-REJECT(k) that $V(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}_F^k(\mathbf{c})$. Note that for the final solution $\mathbf{x}(\mathbf{c})$, we have that $\mathbf{x}(\mathbf{c}) \neq \mathbf{0}$, since the **while** condition evaluates to *False* when a single agent is hired for a single level of service. $\square$

3.2 The model of divisible agents

We now show how to exploit the previous result on the k -level model to obtain an algorithm for the case where it is possible to hire an agent partially up to any percentage. Recall that in the divisible agent model, we have $\mathbf{x}(\mathbf{c}) \in [0, 1]^n$ and concave non-decreasing valuation functions $\bar{v}_i : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ with $\bar{v}_i(0) = 0$ for all $i \in N$. Throughout this section, we assume that all valuation functions are L -regular for some $L \geq 1$.

There is a natural correspondence between the setting with $k \geq 1$ levels of service and the setting with divisible agents: If we subdivide the $[0, 1]$ interval into k chunks of length $\frac{1}{k}$ and evaluate the $\bar{v}_i(\cdot)$'s at $\frac{1}{k}, \frac{2}{k}, \dots, \frac{k}{k}$, then this can be interpreted as the value of hiring 1, 2, $\dots$, k levels of service, respectively. We can then obtain results for the setting with divisible agents by applying this discretization, using SORT-&-REJECT(k) from Section 3.1 and letting k grow. Our CHUNK-&-SOLVE mechanism basically exploits this idea. A detailed description is given in Mechanism 2.

The following is the main result of this section:

Algorithm 1 SORT-&-REJECT(k)

▷ **Input:** A profile $\mathbf{c}$ and a parameter $\alpha \in (0, 1)$ Set $i^* = \arg \max_{i \in N} v_i(k) / \text{OPT}_F^k(\mathbf{c}_{-i})$

if $v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \cdot \text{OPT}_F^k(\mathbf{c}_{-i^*})$ **then**

 | set $x_{i^*} = k$ and $x_i = 0$ for all $i \in N \setminus \{i^*\}$

else

 Compute an optimal allocation $\mathbf{x}^*(\mathbf{c})$ of $\text{OPT}_F^k(\mathbf{c})$.

 Initialize $\mathbf{x} = (x_1^*(\mathbf{c}), \dots, x_n^*(\mathbf{c}))$.

for $i \in N$ and $j = 1, \dots, x_i$ **do**

 | add the marginal value-per-cost ratio $m_i(j)/c_i$ to a list $\mathcal{L}$.

 Sort $\mathcal{L}$ in non-increasing order and let ℓ be the index of the last agent of $W(\mathbf{x})$ in $\mathcal{L}$.

while $(\mathbf{x}) - m_\ell(x_\ell) \geq \alpha \text{OPT}_F^k(\mathbf{c})$ **do**

 | Set $x_\ell = x_\ell - 1$.

 | Remove the last element from $\mathcal{L}$ and update ℓ .

Allocate $\mathbf{x}$ and determine the payments $\mathbf{p}$ according to Myerson's lemma (described in Deliverable D.5.1).

Algorithm 2 CHUNK-&-SOLVE

▷ **Input:** A profile $\mathbf{c}$ and a positive integer k

 Initialize for each $i \in N$, $v_i : \{0, \dots, k\} \rightarrow \mathbb{R}_+$ with $v_i(0) = 0$.

for $i \in N$ and $j \in \{1, \dots, k\}$ **do** set $v_i(j) := \bar{v}_i(j/k)$;

 Set $\alpha(L, k) = (3k + L - \sqrt{5k^2 + 6kL + L^2})/2k$.

 Let $\tilde{I} = (N, \mathbf{c}, kB, k, (v_i)_{i \in N})$ denote the resulting discretized instance of the k -level model.

 Compute $\mathbf{x}$ by running SORT-&-REJECT(k)($\mathbf{c}, \alpha(L, k)$) on $\tilde{I}$. **for** $i \in N$ **do** set $\bar{x}_i = x_i/k$;

 Allocate $\bar{\mathbf{x}}$ and determine the payments $\bar{\mathbf{p}}$ according to Myerson's lemma.

Theorem 3.2. *Mechanism CHUNK-&-SOLVE is truthful, budget-feasible and $(L(1 + \frac{1}{k})/\alpha(L, k))$ -approximate for L -regular instances of the divisible agent model.*

It is a matter of using simple calculus to show that $\lim_{k \rightarrow \infty} \alpha(L, k) = \frac{1}{\phi+1}$, and thus the approximation ratio of Theorem 3.2 goes to $(\phi + 1)L$. Given that the running time of SORT-&-REJECT(k) is polynomial in k , a reasonable question is whether we can have a good approximation guarantee for 'small' k when L is $O(1)$. Again, it is a matter of calculations to show that using $k = O(L)$ suffices. For instance, taking $k = 11L$ implies an approximation ratio of $3L$ for any $L \geq 1$. Qualitatively, this means that, for $L \in O(1)$, CHUNK-&-SOLVE achieves a constant approximation ratio in polynomial time.

4 Mechanisms for spatial crowdsourcing

In the literature, the currently best algorithm for the problem CMIC that we are interested in, has an approximation ratio of 4, following from [2]. However, this algorithm is not convenient for designing a truthful mechanism. Instead, what we are after is algorithms that satisfy the monotonicity property that is identified in Deliverable D.5.1. So far, the only algorithm that has been identified to have such properties for avoiding strategic manipulation is described in [5], and achieves an approximation ratio of $2 \cdot \max_{i \in N} \{q_i\}$. Note that this approximation ratio is dependent on the contribution parameters of the

workers, which can become arbitrarily large, and not bounded by any function of n and m .

Our main result in this section is a truthful mechanism, established via a greedy algorithm that employs the local-ratio technique. In particular, we get an improvement, dependent on the maximum number of workers that contribute to any given task, which for an instance $((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$ of CMIC, is $\Delta(\mathbf{I}) := \max_{j=1, \dots, m} |N_j(\mathbf{I})|$. We denote it simply by Δ when $\mathbf{I}$ is clear from context. Obviously, Δ is always upper bounded by the number of workers, n . Our result can be summarized as follows:

Theorem 4.1. *There exists a polynomial-time, truthful mechanism, that achieves a Δ -approximation for the CMIC problem.*

Our research on this topic has already been published in a conference, and our research article is included at the Appendix of this deliverable. In the remainder of this section, we only provide a sketch of how the proof of Theorem 4.1 is derived. The main component of the proof is an approximation-preserving reduction to a particular job scheduling problem for a single machine [2], defined below. This also reveals an interesting connection between our problem and a family of job scheduling problems.

Loss Minimization Interval Scheduling (LMIS): We are given a limited resource whose amount may vary over a time period, which WLOG, is defined by the integral time instants $\{1, \dots, m\}$. We are also given a set of activities $J = \{1, \dots, n\}$, each of which requires the utilization of the resource, for an interval of time instants. An instance of LMIS is determined by a tuple $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$, where:

- The vector $\mathbf{p} = (p_1, \dots, p_n)$ specifies a penalty $p_j \in \mathbb{R}_{>0}$, for each activity $j \in J$, reflecting the cost that is incurred by not scheduling the activity.
- For each activity $j \in J$, we are given an interval $T_j = [s_j, f_j]$, such that $s_j, f_j \in \{1, \dots, m\}$ are the start and finish times of j respectively. Let $\mathbf{T} = (T_1, \dots, T_n)$, be the vector of all activity intervals.
- The vector $\mathbf{r}$ contains, for each activity $j \in J$, the width $r_j \in \mathbb{R}_{>0}$, reflecting how much resource the activity requires, i.e., this means that activity j requires r_j units of resource at every integral time instant of its interval T_j .
- The vector $\mathbf{D} = (D_1, \dots, D_m)$ specifies the amount of available resource $D_i \in \mathbb{R}_{>0}$, at each integral time instant $i \in [m]$.

Let $J_i(\mathbf{T}) := \{j \in J \mid i \in T_j\}$. The goal in LMIS is to select a set of activities $S \subseteq J$ to schedule, that meet the resource constraint

$$\sum_{j \in S \cap J_i(\mathbf{T})} r_j \leq D_i, \quad i = 1, \dots, m, \quad (3)$$

and such that $\sum_{j \in J \setminus S} p_j$ is minimized, i.e., we want to minimize the sum of the penalties for the non-scheduled activities.

Our work highlights an interesting connection between LMIS and CMIC. This can be seen via the reduction provided by algorithm $\hat{A}$ below, where for an instance $((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$, we let $\mathbf{Q}(\mathbf{I}) = (Q_1(\mathbf{I}), \dots, Q_m(\mathbf{I}))$ and $Q_j(\mathbf{I}) := \sum_{i \in N_j(\mathbf{I})} q_i, \forall j \in [m]$.

Theorem 4.2. *Algorithm $\hat{A}$ converts any α -approximation algorithm for LMIS to an α -approximation algorithm for CMIC.*

Algorithm 3 $\hat{A}(\mathbf{b})$

▷ **Input:** A bidding profile $\mathbf{b} = (\mathbf{c}, \mathbf{I})$ of a CMIC instance $((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$

Construct the LMIS instance $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D}) = (\mathbf{c}, \mathbf{I}, \mathbf{q}, \mathbf{Q}(\mathbf{I}) - \mathbf{d})$, with $J = N$.

Run an approximation algorithm for the LMIS instance, and let S be the set of scheduled activities.

return $N \setminus S$

Approximation Guarantee and Monotonicity of $\hat{A}$

We will consider an algorithm for LMIS, which is a tailor made modification of the algorithm by [2] for our purposes. This is presented as Algorithm 4 below. We can then plug in this algorithm as the component of $\hat{A}$ for solving LMIS.

We introduce first a notion that will be useful both for the statement of Algorithm 4 and for our analysis. Consider an instance $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$ of LMIS. Given a set of jobs $S \subseteq J$, and a time instant $i = 1, \dots, m$, we define

$$R_i(S, \mathbf{T}, \mathbf{D}) := \sum_{\ell \in S \cap J_i(\mathbf{T})} r_\ell - D_i.$$

The quantity $R_i(S, \mathbf{T}, \mathbf{D})$ measures how much (if at all), the resource constraint of Equation (3), for the i -th time instant is violated when scheduling all the activities in S . Accordingly, define $R^*(S, \mathbf{T}, \mathbf{D}) := \max_{i=1, \dots, m} R_i(S, \mathbf{T}, \mathbf{D})$. Note that a schedule S is feasible if and only if $R^*(S, \mathbf{T}, \mathbf{D}) \leq 0$.

Algorithm 4 constructs a feasible schedule S as follows: Initially, it checks if the entire set of activities $S = J$ constitutes a feasible schedule. If not, the algorithm iteratively removes one activity per iteration from S , in a greedy fashion, until S becomes feasible. The algorithm determines the time instant t^* with the most violated feasibility constraint, by computing $R^*(S, \mathbf{T}, \mathbf{D})$, and considers all activities from S whose interval contains t^* . Then, it removes from S one of these activities that minimizes a certain ratio, dependent on the current penalties and resource requirements, while it simultaneously decreases the penalty of all other activities that contain t^* .

Algorithm 4 LMIS-LR($\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D}$)

▷ **Input:** An instance $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$ of LMIS

Initialize $S = J, k = 0$, and $\mathbf{p}_k = (p_{i,k})_{i \in [n]} = \mathbf{p}$.

while $R^*(S, \mathbf{T}, \mathbf{D}) > 0$ **do**

 Let $t^* \in [m]$ be a maximizer of $R^*(S, \mathbf{T}, \mathbf{D})$.

$S_k = S \cap_{t^*}(\mathbf{T})$

$\varepsilon_k = \min_{i \in S_k} \frac{p_{i,k}}{\min\{R^*(S, \mathbf{T}, \mathbf{D}), r_i\}}$

 For $i = 1, \dots, m$ let

$$p_{i,k+1} = \begin{cases} p_{i,k} - \varepsilon_k \min\{R^*(S, \mathbf{T}, \mathbf{D}), r_i\}, & \text{if } i \in S_k, \\ p_{i,k}, & \text{o/w.} \end{cases}$$

 Let $j^* \in S_k$ be a minimizer of ε_k .

 Set $S = S \setminus \{j^*\}$, and $k = k + 1$.

return S

Formally, the analysis of this algorithm yields the following result, which concludes

also the approximation analysis for CMIC and the proof of Theorem 4.1.

Theorem 4.3. *Algorithm 4 achieves a Δ -approximation for the LMIS problem, where $\Delta = \max_{j=1,\dots,m} |N_j(\mathbf{I})|$, and the analysis of its approximation is tight.*

Finally, one also needs to analyze the monotonicity of the whole algorithm for CMIC. This is established in the following theorem, and we defer the full analysis in the Appendix.

Theorem 4.4. *Algorithm $\hat{A}$ is monotone, when using Algorithm 4 as a black box for solving LMIS.*

5 Non-truthful mechanisms for restricted auction models

In this last technical section of the report, we comment on the possibility of employing non-truthful mechanisms. The notion of truthfulness is undoubtedly a highly desirable concept in mechanism design. Nevertheless, truthful mechanisms are often hard to obtain or they may be more complex to run, and harder to understand for the auction participants. As a consequence, it occurs quite often in practice to also utilize non-truthful mechanisms. Although such mechanisms may provide incentives for strategic behavior, in some cases one can quantify the extent to which this may happen, and produce provable guarantees on the quality of the auction outcomes.

A quite popular example that falls in the above category is the discriminatory price multi-unit auction, involving the allocation of k identical units to a set $\mathcal{N} = \{1, \dots, n\}$ of bidders. This is a special case of auctions where all goods are the same. In the context of reverse auctions, this corresponds to the special case where we have k identical tasks to assign, and all bidders make offers on how many of these tasks they are willing to carry out.

In more detail, each bidder $i \in \mathcal{N}$ has a private value $v_i > 0$, which reflects her cost per task and a private demand $d_i \in \mathbb{Z}_+$ which reflects the maximum number of tasks bidder i requires (or is willing to undertake). 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 studied 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 $\mathbf{b} = (b_1, \dots, b_n)$ 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 cheapest cost per task, then moving to the second cheapest 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). Finally, as a price, each winning bidder pays his bid multiplied by the number of tasks received.

Multi-unit auctions form a popular transaction means for selling multiple units of a single good or for assigning multiple identical tasks/services to interested participants. It is easy to see however that the discriminatory price auction is not a truthful mechanism. Consequently, in the more recent years, a series of works have studied the social welfare

guarantees that can be obtained at equilibrium. This corresponds to a Price of Anarchy analysis, as demonstrated in the works of [6] [4] 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 analysis on the evaluation of Nash equilibria and the state of the art Price of Anarchy bounds. The known lower bounds are quite far from the upper bounds derived by the commonly used smoothness-based approaches, [6], 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 was 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.

5.1 Our Contribution

Motivated by the previous discussion, we study the welfare performance of the discriminatory price auction and 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, as well as properties on the distributions of the mixed strategies. We then 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. For multiple bidders, we show that the Price of Anarchy is strictly worse, which also improves the best known lower bound for submodular valuations [4]. 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. To summarize, our results show that in several cases, the Price of Anarchy is even lower than the previous bound of [6] and strengthen the perception that such auctions can work well in practice.

Our results from this work can be found in our research article , which already led to a conference publication and is included in the Appendix.

6 Conclusions and resulting publications

We have presented an overview of mechanisms for various auction settings that are relevant for procurement (reverse) auctions. We plan to submit an extended version of the results from Section 3, within the coming months, to an international conference. An expanded version of the results of Section 4 has already been presented as a conference publication at the 15th International Symposium on Algorithmic Game Theory (SAGT 2022). Moreover, the results of Section 5 have already been presented at the 16th International Conference on Web and Internet Economics (WINE 2021).

References

- [1] Nima Anari, Gagan Goel, and Afshin Nikzad. Budget feasible procurement auctions. *Operations Research*, 66(3):637–652, 2018.
- [2] Amotz Bar-Noy, Reuven Bar-Yehuda, Ari Freund, Joseph Naor, and Baruch Schieber. A unified approach to approximating resource allocation and scheduling. *Journal of the ACM*, 48(5):1069–1090, 2001.
- [3] Hau Chan and Jing Chen. Truthful multi-unit procurements with budgets. In *Web and Internet Economics: 10th International Conference, WINE 2014, Proceedings*, pages 89–105, 2014.
- [4] G. Christodoulou, A. Kovács, A. Sgouritsa, and B. Tang. Tight Bounds for the Price of Anarchy of Simultaneous First-Price Auctions. *ACM TEAC*, 4(2):9:1–9:33, 2016.
- [5] Pankaj Dayama, Balakrishnan Narayanaswamy, Dinesh Garg, and Y. Narahari. Truthful interval cover mechanisms for crowdsourcing applications. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, page 1091–1099, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.
- [6] B. de Keijzer, E. Markakis, G. Schäfer, and O. Telelis. Inefficiency of Standard Multi-unit Auctions. In *European Symposium on Algorithms, ESA 2013*, pages 385–396.
- [7] Silvano Martello and Paolo Toth. *Knapsack problems: algorithms and computer implementations*. John Wiley & Sons, Inc., 1990.
- [8] Yaron Singer. Budget feasible mechanisms. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*,, pages 765–774, 2010.

Appendix

A Resulting publications

In the next pages, we include 2 conference publications from our research team, which are already mentioned in the previous sections of this report. The first one concerns the results of Section 4, which have been presented as a conference publication at the 15th International Symposium on Algorithmic Game Theory (SAGT 2022). The second one concerns the results of Section 5, which have already been presented at the 16th International Conference on Web and Internet Economics (WINE 2021).

On Improved Interval Cover Mechanisms for Crowdsourcing Markets

Evangelos Markakis, Georgios Papasotiropoulos, Artem Tsikiris

Athens University of Economics and Business

Abstract. We study a covering problem motivated by spatial models in crowdsourcing markets, where tasks are ordered according to some geographic or temporal criterion. Assuming that each participating bidder can provide a certain level of contribution for a subset of consecutive tasks, and that each task has a demand requirement, the goal is to find a set of bidders of minimum cost, who can meet all the demand constraints. Our focus is on truthful mechanisms with approximation guarantees against the optimal cost. To this end, we obtain two main results. The first one, is a truthful mechanism that achieves a bounded approximation guarantee. This mechanism improves the state of the art, which is a mechanism with an arbitrarily large factor in worst case. The second one, concerns a class of instances that generalizes the minimum knapsack problem. Namely, we consider inputs with a constant number of tasks, for which we provide a truthful FPTAS. Finally, we also highlight connections of our problem with other well-studied optimization problems, out of which, we infer further conclusions on its (in)approximability.

1 Introduction

We consider a mechanism design problem, that emerges under certain spatial models for crowdsourcing and labour matching markets. It is instructive to start with an example, so as to introduce the main aspects of the model. Imagine a set of cities, located geographically in a consecutive order, and consider a company that has opened a store in each of these cities. In Figure 1, we see an instance with 5 cities, named A to E. The company needs to meet a demand constraint, i.e., a lower bound on the volume of goods that need to be transported to each store, based on consumption and future planning. To achieve this goal, it chooses to hire other firms, or single individuals, that can make deliveries, via a reverse (procurement) auction (for an exposition on auctions for transportation routes, see [15]). Suppose that every participating entity, referred to as a bidder or a worker, can only cover a certain interval of contiguous cities, at a cost that she specifies, and furthermore, she can only accommodate a certain volume of goods, assumed to be the same for each city in her declared interval (among others, dependent on the transportation means that she owns). The problem boils down to selecting a set of winning bidders who can jointly cover all the demand constraints at minimum cost, and in a way that prevents the workers from misreporting their true preferences.




Fig. 1: An example with 5 tasks (A to E) and 5 workers.

In the crowdsourcing jargon, we can view the store in each city as a task with a demand requirement. For instance, in labeling/classification tasks, the demand could correspond to the number of people who should execute the task in order to acquire a higher confidence on the outcome. In other cases, it could be interpreted as a (not necessarily integral) volume coverage requirement.

vehicle	contribution	cost
truck	8	10
van	6	5
pickup	5	4
motorbike	1	2
car	3	4

Table 1: Contribution and costs of the workers of Figure 1.

Coming back to the example of Figure 1, say that the demand of cities A to E are given by the vector $(6, 2, 9, 1, 3)$. The contribution per task, as well as the cost of the workers are given by Table 1. The interval covered by each worker is visible in Figure 1. Obviously, hiring all the workers is a feasible solution. Selecting the workers with the van, the pickup and the car also forms a feasible solution since the demands of all cities are indeed satisfied. Notice that the optimal solution is to hire the workers with the truck and the motorbike at a total cost of 12.

From an algorithmic viewpoint, there has been already significant progress on the problem. As it can be easily seen to be NP-hard, the main results on this front include constant factor approximation algorithms, with further improvements for special cases. However, in the context of mechanism design, one needs to consider strategic aspects as well. Bidders may attempt to report higher costs in order to achieve higher payments, or they can lie about the subset of tasks they can actually fulfil. It would be therefore desirable to have mechanisms that, on the one hand, achieve competitive approximation guarantees, and on the other, deter bidders from misreporting. To our knowledge, the currently best result on this direction is by [18], where however the approximation ratio of their truthful mechanism is unbounded in terms of the number of tasks and workers (and depends on numerical parameters of each instance). It has remained open since then, if truthfulness can be compatible with bounded approximation guarantees.

Contribution

We focus on truthful mechanisms and their approximation guarantees against the optimal cost. To this end, we provide two main results. The first one, in Section 3, is a truthful Δ -approximation mechanism, where Δ is the maximum number of workers that are willing to work on the same task (the maximum

being taken over all tasks). This mechanism, improves significantly the state of the art, coming to a guarantee that is polynomially bounded in terms of the input size. Apart from the improvement, we note also that our result is based on the local-ratio technique from approximation algorithms [4], a technique that has not been used very often for building truthful mechanisms (for an exception, see [20]). Moving on, our second result, in Section 4, concerns the class of instances with a constant number of tasks, which generalizes MIN-KNAPSACK, an NP-hard variant of KNAPSACK that corresponds to the single-task case of our problem. For this class we provide a truthful FPTAS, by mainly exploiting and adapting the framework of [8]. In doing so, we also identify a flaw in a previous attempt for designing a truthful FPTAS for MIN-KNAPSACK. Finally, in Section 5, we discuss some further implications and extensions. Namely, we bring to light some interesting connections with well-studied optimization problems, related to unsplittable flow problems and caching. By exploiting known results for these problems, we rule out a FPTAS for the general case, and also identify a special case where a PTAS is likely to exist.

Related Work

As already discussed, the work most related to ours is [18], which introduced the model in the context of crowdsourcing. They provide a truthful optimal mechanism when the workers have an identical contribution, whereas for the general case, they present a truthful approximate mechanism, the ratio of which is dependent on the contribution parameters of the workers. Furthermore, [41] studies the case of unit-demand tasks and unit-contribution workers and identifies a truthful optimal mechanism. For the same setting, they also propose a mechanism when workers can submit multiple bids, for which they attain a logarithmic factor. Additionally, in [42] the authors studied the prominent special case of a single task (MIN-KNAPSACK) and provided a randomized, truthful-in-expectation mechanism that achieves an approximation factor of 2.

Regarding the purely algorithmic problem, without the constraint of truthfulness, it has appeared under the name of (0-1)RESOURCE ALLOCATION, and a 4-approximation was presented in [10]. The currently best known algorithm achieves a factor of 3, in [33]. For the MIN-KNAPSACK problem, a PTAS is implied by [17] and a FPTAS is given in [29].

There are quite a few problems that can be viewed as generalizations of what we study here, such as *general scheduling problem* [2], *multidimensional min-knapsack*, *column restricted covering integer programs* [11]. Moreover, several problems in discrete optimization can be seen as related variants, but are neither extensions nor special cases of ours. Indicatively: *bandwidth allocation* [12], *multiset multicover* [7,38], *geometric knapsack* [23], *capacitated network design* [9], *admission control* [36].

Finally, for general spatio-temporal models appearing in the crowdsourcing literature, we refer to two recent surveys [25,39], which cover to a big extent the relevant results.

2 Preliminaries

In this section, we first define formally the problem that we study, together with some additional necessary notation. In the sequel, we discuss the relevant definitions for the design of truthful mechanisms. We note that all the missing proofs from the following sections can be found in the full version of our work.

Problem Statement

We are interested in the optimization problem defined below. For motivating applications, we refer the reader to Section 1.2 in [18].

Cost Minimization Interval Cover (CMIC): Consider a set of tasks, say $\{1, \dots, m\}$, that are ordered in a line, and a set of available bidders $\mathcal{N} = \{1, \dots, n\}$. We will interchangeably use the term bidder or worker in the sequel. An instance of CMIC is determined by a tuple $(\mathbf{b}, \mathbf{q}, \mathbf{d})$, where:

- The vector $\mathbf{b} = (b_1, \dots, b_n)$ is the bidding profile. For each bidder $j \in \mathcal{N}$, $b_j = (c_j, [s_j, f_j])$, where $[s_j, f_j] = I_j$ is the interval of contiguous tasks $\{s_j, s_j + 1, \dots, f_j\} \subseteq [m]$, that j is able to contribute to, and $c_j \in \mathbb{R}_{\geq 0}$ is the cost incurred to her, if she is selected to contribute. We can assume positive costs since workers of zero cost are trivially included in the solution. We often denote $\mathbf{b}$ as $(\mathbf{c}, \mathbf{I})$ where $\mathbf{c}$ is the cost vector and $\mathbf{I}$ is the vector of the intervals.
- The vector $\mathbf{q} = (q_1, \dots, q_n)$ is the contribution vector, so that $q_j \in \mathbb{R}_{> 0}$ denotes the contribution that bidder j can make to the tasks that belong to $[s_j, f_j]$.
- The vector $\mathbf{d} = (d_1, \dots, d_m)$ specifies the demand $d_j \in \mathbb{R}_{> 0}$ of a task j .

The goal is to select a set of bidders $S \subseteq \mathcal{N}$ of minimum cost, that satisfies

$$\sum_{i \in \mathcal{N}_j(\mathbf{I}) \cap S} q_i \geq d_j, \quad \forall j \in [m], \quad (1)$$

where for $j \in [m]$, $\mathcal{N}_j(\mathbf{I}) := \{i \in \mathcal{N} \mid j \in [s_i, f_i]\}$, i.e. is the set of bidders who can contribute to task j .

The problem belongs to the broad family of problems described by covering integer programs. It may also seem reminiscent of multicover variants of the SET COVER problem on an interval universe. We note however that in CMIC, each worker is allowed to be picked at most once, and moreover, the coverage requirement is not necessarily integral, which is a substantial difference.

Given a feasible solution $S \subseteq \mathcal{N}$, we refer to its total cost, $\sum_{i \in S} c_i$, as the derived *social cost*. Throughout all our work, we focus on deterministic allocation algorithms that use a consistent deterministic tie-breaking rule. Given an instance $P = (\mathbf{b}, \mathbf{q}, \mathbf{d})$, and an allocation algorithm A , we denote by $\mathcal{W}_A(\mathbf{b})$ the set of bidders selected by A , when $\mathbf{q}, \mathbf{d}$ are clear from the context¹. In the same spirit, we let also $C(A, \mathbf{b}) := \sum_{i \in \mathcal{W}_A(\mathbf{b})} c_i$ be the *social cost* derived by algorithm A on input P . Finally, we use $OPT(\mathbf{b})$ to denote the cost of an optimal solution.

¹ It is convenient to highlight the dependence on $\mathbf{b}$, especially when arguing about truthful mechanisms in the remaining sections.

Truthful Mechanisms

We move now to the strategic scenario, where bids can be private information. A mechanism for a reverse auction, like the ones we study here, is a tuple $M = (A, \pi)$, consisting of an *allocation algorithm* A and a *payment rule* π . Initially, each bidder $j \in \mathcal{N}$ is asked to submit a bid $b_j := (c_j, [s_j, f_j])$, which may differ from her actual cost and interval. Then, given a bidding profile $\mathbf{b} = (\mathbf{c}, \mathbf{I})$, the allocation algorithm $A(\mathbf{b})$ selects the set of winning bidders, i.e., this is a binary setting where the allocation decision for each bidder is whether she is included in the solution or not. Finally, the mechanism computes a vector of payments $\pi(\mathbf{b})$, so that $\pi_i(\mathbf{b})$ is the amount to be paid to bidder i by the auctioneer. Naturally, we will consider that non-winning bidders do not receive any payment.

Note that we consider the contribution q_j , for $j \in \mathcal{N}$, to be public information, available to the auctioneer. The reason is that such a parameter could be estimated by past statistical information on the performance or capacity of the worker. As an example, we refer to [18], where for labeling tasks, it is explained how q_j can be computed as a function of a worker's quality (i.e., the probability for a worker to label correctly), that can be available in a crowdsourcing platform with rating scores or reviews for the workers.

Our setting corresponds to what is usually referred to as a pseudo-2-parameter environment (or almost-single-parameter [5]), since each bidder has two types of private information, the monetary cost and the interval. And in particular, our model can be seen as a special case of *single-minded* bidders, who are interested in a single subset each, but for reverse auctions. As in [18], when the true type of a worker i is $(c_i, [s_i, t_i])$, and she declares her true interval or any non-empty subset of it, then she enjoys a utility of $\pi_i - c_i$, if she is selected as a winner by the mechanism (with π_i being her payment). If she declares any other interval that contains any task $j \notin [s_i, t_i]$, and she is selected as a winner, then the bidder has a utility of $-\infty$, or equivalently has an infinite cost. This simply models the fact that the worker may not be capable of executing or does not desire to be assigned to any task outside her true interval (and therefore would have no incentive for such deviations).

The previous discussion allows us to exploit the sufficient conditions proposed by [31] (for forward auctions), to obtain truthful mechanisms, as an extension of the seminal result by [34]. For reverse auctions, the same framework is also applicable, implying that as long as an allocation algorithm is exact (each selected bidder is assigned her declared interval), the crucial property we need to enforce is *monotonicity*. Monotonicity of an algorithm means that if a winning bidder declares a more competitive bid, she should still remain a winner. To be more precise, we define first the following partial order on possible bids.

Definition 1. Let $b_i = (c_i, [s_i, f_i])$ and $b'_i = (c'_i, [s'_i, f'_i])$ be two bids of bidder $i \in \mathcal{N}$. We say that $b'_i \succeq b_i$, if $c_i \geq c'_i$ and $[s_i, f_i] \subseteq [s'_i, f'_i]$.

Definition 2. An allocation algorithm A is monotone if for every bidding profile $\mathbf{b}$, for any bidder $i \in \mathcal{W}_A(\mathbf{b})$ and any bid $b'_i \succeq b_i$, it holds that $i \in \mathcal{W}_A(b'_i, \mathbf{b}_{-i})$.

Theorem 1 (cf. [31], Theorem 9.6). *In settings with single-minded bidders, given a monotone and exact algorithm A , there exists an efficiently computable payment rule π , such that $\mathcal{M} = (A, \pi)$ is a truthful mechanism.*

Finally, we stress that all our proposed algorithms are exact by construction, which means that bidders will only be assigned to the set of tasks which they asked for, as in [8], and hence we only need to care about monotonicity.

3 An Improved Truthful Approximation Mechanism

As already stated in Section 1, the currently best algorithm for CMIC has an approximation ratio of 3 [33], based on refining the 4-approximation algorithm by [10]. However, as the next proposition shows, these algorithms are not monotone². More generally, it has been noted in [18] that primal-dual algorithms with a “delete phase” at the end, are typically non-monotone, without, however, providing an example. For the sake of completeness, in the full version of our work, we provide a concrete example that proves the following statement:

Proposition 1. *The current state of the art constant factor approximation algorithms for CMIC by [10, 33] are not monotone.*

So far, the only truthful mechanism that has been identified [18], achieves an approximation ratio of $2 \cdot \max_{i \in \mathcal{N}} \{q_i\}$. Note that this approximation ratio is dependent on the contribution parameters of the workers, which can become arbitrarily large, and not bounded by any function of n and m . The main result of this section is the following theorem, established via a greedy, local-ratio algorithm, which reduces the gap between truthful and non-truthful mechanisms. In particular, this gap is related to the maximum number of workers that contribute to any given task, which for an instance $((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$ of CMIC, is $\Delta(\mathbf{I}) := \max_{j=1, \dots, m} |\mathcal{N}_j(\mathbf{I})|$. We denote it simply by Δ when $\mathbf{I}$ is clear from context. Obviously, Δ is always upper bounded by the number of workers, n .

Theorem 2. *There exists a truthful, polynomial-time mechanism, that achieves a Δ -approximation for the CMIC problem.*

The rest of the section is devoted to the proof of Theorem 2. The main component of the proof is an approximation-preserving reduction to a particular job scheduling problem for a single machine [3], defined as follows:

Loss Minimization Interval Scheduling (LMIS): We are given a limited resource whose amount may vary over a time period, which WLOG, is defined by the integral time instants $\{1, \dots, m\}$. We are also given a set of activities $\mathcal{J} = \{1, \dots, n\}$, each of which requires the utilization of the resource, for an interval of time instants. An instance of LMIS is determined by a tuple $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$, where:

² For a similar reason, the 40-approximation for CMIC by [11], which uses as a subroutine a primal-dual algorithm involving a “delete phase”, is non-monotone as well.

- The vector $\mathbf{p} = (p_1, \dots, p_n)$ specifies a penalty $p_j \in \mathbb{R}_{>0}$, for each activity $j \in \mathcal{J}$, reflecting the cost that is incurred by not scheduling the activity.
- For each activity $j \in \mathcal{J}$, we are given an interval³ $T_j = [s_j, f_j]$, such that $s_j, f_j \in \{1, \dots, m\}$ are the start and finish times of j respectively. Let $\mathbf{T} = (T_1, \dots, T_n)$, be the vector of all activity intervals.
- The vector $\mathbf{r}$ contains, for each activity $j \in \mathcal{J}$, the width $r_j \in \mathbb{R}_{>0}$, reflecting how much resource the activity requires, i.e., this means that activity j requires r_j units of resource at every integral time instant of its interval T_j .
- The vector $\mathbf{D} = (D_1, \dots, D_m)$ specifies the amount of available resource $D_i \in \mathbb{R}_{>0}$, at each integral time instant $i \in [m]$.

Let $\mathcal{J}_i(\mathbf{T}) := \{j \in \mathcal{J} \mid i \in T_j\}$. The goal in LMIS is to select a set of activities $S \subseteq \mathcal{J}$ to schedule, that meet the resource constraint

$$\sum_{j \in S \cap \mathcal{J}_i(\mathbf{T})} r_j \leq D_i, \quad i = 1, \dots, m, \quad (2)$$

and such that $\sum_{j \in \mathcal{J} \setminus S} p_j$ is minimized, i.e., we want to minimize the sum of the penalties for the non-scheduled activities.

Our work highlights an interesting connection between LMIS and CMIC. This can be seen via the reduction provided by algorithm $\hat{A}$ below, where for an instance $((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$, we let $\mathbf{Q}(\mathbf{I}) = (Q_1(\mathbf{I}), \dots, Q_m(\mathbf{I}))$ and $Q_j(\mathbf{I}) := \sum_{i \in \mathcal{N}_j(\mathbf{I})} q_i, \forall j \in [m]$.

Algorithm 1: $\hat{A}(\mathbf{b})$

▷ **Input:** A bidding profile $\mathbf{b} = (\mathbf{c}, \mathbf{I})$ of a CMIC instance $((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$

- 1 Construct the LMIS instance $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D}) = (\mathbf{c}, \mathbf{I}, \mathbf{q}, \mathbf{Q}(\mathbf{I}) - \mathbf{d})$, with $\mathcal{J} = \mathcal{N}$.
 - 2 Run an approximation algorithm for the LMIS instance, and let S be the set of scheduled activities.
 - 3 **return** $\mathcal{N} \setminus S$
-

Theorem 3. *Algorithm $\hat{A}$ converts any α -approximation algorithm for LMIS to an α -approximation algorithm for CMIC.*

Theorem 3 is based on the lemma below, which shows the connection between the feasible solutions of the two problems.

Lemma 1. *Consider a CMIC instance $P = ((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$. Let also P' be the LMIS instance defined by $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D}) = (\mathbf{c}, \mathbf{I}, \mathbf{q}, \mathbf{Q}(\mathbf{I}) - \mathbf{d})$, with $\mathcal{J} = \mathcal{N}$. Then, for every feasible solution S of P , it holds that $\mathcal{J} \setminus S$ is a feasible solution for P' with the same cost, and vice versa.*

³ Originally, the problem was defined using a semi-closed interval for each activity, but it is easy to see that defining it using a closed one instead, is equivalent and more convenient for our purposes.

Approximation Guarantee and Monotonicity of $\hat{A}$

If we only cared about the approximation ratio of $\hat{A}$, it would suffice to use as a black box any algorithm for LMIS. And in fact, the best known algorithm for LMIS achieves a 4-approximation, and was obtained in [3], using the local-ratio framework. Plugging in this algorithm however, does not ensure that we will end up with a truthful mechanism for CMIC. Instead, we will consider an appropriate modification of the algorithm by [3], which enforces monotonicity of $\hat{A}$, but at the price of a higher approximation ratio. This is presented as Algorithm 2 below.

We introduce first a notion that will be useful both for the statement of Algorithm 2 and for our analysis. Consider an instance $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$ of LMIS. Given a set of jobs $S \subseteq \mathcal{J}$, and a time instant $i = 1, \dots, m$, we define

$$R_i(S, \mathbf{T}, \mathbf{D}) := \sum_{\ell \in S \cap \mathcal{J}_i(\mathbf{T})} r_\ell - D_i.$$

The quantity $R_i(S, \mathbf{T}, \mathbf{D})$ measures how much (if at all), the resource constraint of Equation (2), for the i -th time instant is violated when scheduling all the activities in S . Accordingly, define $R^*(S, \mathbf{T}, \mathbf{D}) := \max_{i=1, \dots, m} R_i(S, \mathbf{T}, \mathbf{D})$. Note that a schedule S is feasible if and only if $R^*(S, \mathbf{T}, \mathbf{D}) \leq 0$.

Algorithm 2 constructs a feasible schedule S as follows: Initially, it checks if the entire set of activities $S = \mathcal{J}$ constitutes a feasible schedule. If not, the algorithm iteratively removes one activity per iteration from S , in a greedy fashion, until S becomes feasible. The algorithm determines the time instant t^* with the most violated feasibility constraint, by computing $R^*(S, \mathbf{T}, \mathbf{D})$, and considers all activities from S whose interval contains t^* . Then, it removes from S one of these activities that minimizes a certain ratio, dependent on the current penalties and resource requirements, while it simultaneously decreases the penalty of all other activities that contain t^* .

Algorithm 2 LMIS-LR($\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D}$)

▷ **Input:** An instance $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$ of LMIS

- 1 Initialize $S = \mathcal{J}$, $k = 0$, and $\mathbf{p}_k = (p_{i,k})_{i \in [n]} = \mathbf{p}$.
 - 2 **while** $R^*(S, \mathbf{T}, \mathbf{D}) > 0$ **do**
 - 3 Let $t^* \in [m]$ be a maximizer of $R^*(S, \mathbf{T}, \mathbf{D})$.
 - 4 $S_k = S \cap \mathcal{J}_{t^*}(\mathbf{T})$
 - 5 $\varepsilon_k = \min_{i \in S_k} \frac{p_{i,k}}{\min\{R^*(S, \mathbf{T}, \mathbf{D}), r_i\}}$
 - 6 For $i = 1, \dots, m$ let

$$p_{i,k+1} = \begin{cases} p_{i,k} - \varepsilon_k \min\{R^*(S, \mathbf{T}, \mathbf{D}), r_i\}, & \text{if } i \in S_k, \\ p_{i,k}, & \text{o/w.} \end{cases}$$
 - 7 Let $j^* \in S_k$ be a minimizer of ε_k .
 - 8 Set $S = S \setminus \{j^*\}$, and $k = k + 1$.
 - 9 **return** S
-

Remark 1. A variation of Algorithm 2 for LMIS is stated in [3]. The main difference is that the algorithm of [3] has an additional step to ensure that the solution returned is maximal. The extra step helps in improving the approximation ratio, but it destroys the hope for monotonicity of $\hat{A}$. This can be demonstrated using the same example that we used for the primal-dual algorithms of Proposition 1. Furthermore, we note that the algorithm of [3] is presented using the local-ratio jargon. We have chosen to present Algorithm 2 in a self-contained way for ease of exposition but for its analysis, we do make use of the local-ratio framework.

Theorem 4. *Algorithm 2 achieves a Δ -approximation for the LMIS problem, where $\Delta = \max_{j=1,\dots,m} |\mathcal{N}_j(\mathbf{I})|$, and the analysis of its approximation is tight.*

It remains to be shown that $\hat{A}$, with Algorithm 2 as a subroutine, becomes a monotone allocation algorithm for CMIC. For establishing the monotonicity, according to Definition 2, we have to examine the ways in which a winning worker i can deviate from the truth with a bid $b'_i \succeq b_i$, where b_i is her initial bid. By Definition 1, this means that under b'_i , a lower or equal cost and a larger or the same interval are declared, compared to b_i . But to argue about $\hat{A}$, we first have to understand how such deviations from the truth affect the outcome of Algorithm 2, when it is called by $\hat{A}$. Lowering the cost at a CMIC instance corresponds to lowering the penalty of an activity at the LMIS instance that $\hat{A}$ constructs. The lemma below examines precisely what happens when we lower the penalty of a non-scheduled activity in a LMIS instance.

Lemma 2. *For an instance $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$ of LMIS, let S be the schedule returned by Algorithm 2, and $j \in \mathcal{J} \setminus S$. Then, for any $p'_j \leq p_j$, it holds that $j \in \mathcal{J} \setminus S'$, where S' is the schedule returned by Algorithm 2 for $((p'_j, \mathbf{p}_{-j}), \mathbf{T}, \mathbf{r}, \mathbf{D})$.*

The next lemma examines enlarging the interval of an activity. This needs a different argument from the previous lemma because the deviation causes an activity to participate in more time instants.

Lemma 3. *For an instance $(\mathbf{p}, \mathbf{T}, \mathbf{r}, \mathbf{D})$ of LMIS, let S be a schedule returned by Algorithm 2 and $j \in \mathcal{J} \setminus S$. For any interval $T'_j \supseteq T_j$, consider the instance $P' = (\mathbf{p}, (T'_j, \mathbf{T}_{-j}), \mathbf{r}, \mathbf{D}')$, where $\mathbf{D}' = (D'_1, \dots, D'_m)$ such that:*

$$D'_\ell = \begin{cases} D_\ell + r_j, & \text{if } \ell \in T'_j \setminus T_j, \\ D_\ell, & \text{o/w.} \end{cases}$$

Then, $j \in \mathcal{J} \setminus S'$, where S' is the schedule returned by Algorithm 2 for P' .

Combining Lemma 2 and Lemma 3 we get the following:

Theorem 5. *Algorithm $\hat{A}$ is monotone, when using Algorithm 2 as a black box for solving LMIS.*

Proof. Fix a bidding profile $\mathbf{b}$ and a winning worker $i \in \mathcal{W}_{\hat{A}}(\mathbf{b})$. Let $b_i = (c_i, [s_i, f_i])$, and consider an arbitrary deviation of i , say $b''_i = (c'_i, [s'_i, f'_i])$, such

that $b_i'' \succeq b_i$. We need to show that $i \in \mathcal{W}_{\hat{A}}(b_i'', \mathbf{b}_{-i})$ and we will do this in two steps. First, we consider the deviation $b_i' = (c_i', [s_i, f_i])$. Since b_i' differs from b_i only with respect to the declared cost and the bids of the remaining workers remain the same, we can directly use Lemma 2 to conclude that $i \in \mathcal{W}_{\hat{A}}(b_i', \mathbf{b}_{-i})$. Having established that i is still a winner under $(b_i', \mathbf{b}_{-i})$, consider now the deviation from b_i' to b_i'' . Note that $(b_i'', \mathbf{b}_{-i})$ differs from $(b_i', \mathbf{b}_{-i})$ only with respect to the declared interval of bidder i . Recall also that Algorithm $\hat{A}$ calls Algorithm 2 with input the tuple $(\mathbf{c}, \mathbf{I}, \mathbf{q}, \mathbf{Q}(\mathbf{I}) - \mathbf{d})$. This means that under b_i'' , the vector $\mathbf{Q}(\mathbf{I}) - \mathbf{d}$ in the constructed LMIS instance changes only for time instants that belong to $[s_i', f_i'] \setminus [s_i, f_i]$ (where we simply add q_i). But then, this can be handled by Lemma 3, and obtain that $i \in \mathcal{W}_{\hat{A}}(b_i'', \mathbf{b}_{-i})$.

To conclude, it is trivial that Algorithm $\hat{A}$ runs in polynomial time, when using Algorithm 2 for solving LMIS, and hence, the monotonicity of $\hat{A}$, together with Theorems 3 and 4 complete the proof of Theorem 2.

4 A Truthful FPTAS for a Small Number of Tasks

At what follows, we investigate whether we can have truthful mechanisms with a better approximation ratio for special cases of restricted problem size. An instance of CMIC with a constant number of workers can be optimally solved in polynomial time by a brute force algorithm, which, together with the VCG payment scheme, results in a truthful mechanism. On the other hand, the story is different when we have a small number of tasks, since CMIC is NP-hard even for one task [18]. Building upon this negative result, we provide a truthful mechanism that achieves the best possible approximation factor, for the case of a constant number of tasks, and our main result of this section is the following:

Theorem 6. *There exists a truthful FPTAS for CMIC, when the number of tasks is constant.*

We would like first to pay attention to the special case of a single task, which corresponds to the MIN-KNAPSACK problem, the minimization version of KNAPSACK, where items have costs (instead of values) and there is a covering requirement (instead of a capacity constraint) for the selected items. The work of [8], which proposes a truthful FPTAS for the classic maximization version of KNAPSACK, claims (without providing a proof) that the analogous result holds for MIN-KNAPSACK too. To our knowledge, the only published work that explicitly attempts to extend [8] and describe a truthful FPTAS for MIN-KNAPSACK is [13]. However, we found that the analysis of truthfulness there is flawed and we refer to our full version for a counterexample, establishing the following claim.

Proposition 2. *The FPTAS for MIN-KNAPSACK, proposed by [13], is not monotone.*

Therefore, our Theorem 6 helps to resolve any potential ambiguities for MIN-KNAPSACK. Finally, for two or more tasks, we are not aware of any truthful mechanism attaining any bound better than the one provided by Theorem 2.

Remark 2. As in other approaches for constructing a truthful FPTAS, e.g. [8], we also make the assumption from here onwards, that $c_i \geq 1$ for every worker i , i.e., the workers will not be allowed to declare a cost lower than 1. In the full version of our work, we prove that we can adjust the assumption to $c_i \geq \delta$ for any arbitrarily small δ , but for convenience, here we stick to $\delta = 1$.

Furthermore, we bring in some additional notation that we use in this section. Given a bidding profile $\mathbf{b}$, let $c_{sum}(\mathbf{b}) := \sum_{i \in \mathcal{N}} c_i$. Similarly let $c_{min}(\mathbf{b})$ (resp. $c_{max}(\mathbf{b})$), be the minimum (resp. maximum) cost by the bidders.

4.1 A Pseudopolynomial Dynamic Programming Algorithm

The first step towards designing the FPTAS, is a pseudopolynomial dynamic programming algorithm that returns the optimal solution for the case of a constant number of tasks. For simplicity, we focus on describing the algorithm for the case of two tasks. The generalization is rather obvious (and discussed briefly at the end of this subsection).

Given an instance with two tasks, let d_1, d_2 be the demand requirements of the tasks. Note that $\mathcal{N}$ can be partitioned into three sets, W_0, W_1, W_2 , since we can have at most three types of workers: W_0 is the set of workers who can contribute to both tasks, and for $\ell \in \{1, 2\}$, W_ℓ is the set of workers who are capable of contributing only to task ℓ .

We define a 3-dimensional matrix $Q[\ell, i, c]$, where for $\ell = 0, 1, 2$, for $i = 0, 1, \dots, n$, and for $c = 0, 1, \dots, c_{sum}(\mathbf{b})$, $Q[\ell, i, c]$ denotes the maximum possible contribution that can be jointly achieved by any set of workers in $W_\ell \cap \{1, \dots, i\}$ with a total cost of exactly c . For our purposes, we assume⁴ that for $i \in [n]$, each c_i is an integer, so that c also takes only integral values. Our algorithm is based on computing the values of the cells of Q and we claim that this can be done by exploiting the following recursive relation:

$$Q[\ell, i, c] = \begin{cases} 0, & \text{if } i = 0 \\ Q[\ell, i-1, c], & \text{if } i > 0 \text{ and either } i \notin W_\ell \text{ or } c_i > c \\ \max\{Q[\ell, i-1, c], Q[\ell, i-1, c-c_i] + q_i\}, & \text{o/w} \end{cases} \quad (3)$$

Observe that for a feasible solution S , the workers who contribute to the demand of task 1 (resp. 2) are those from $S \cap W_0$ and $S \cap W_1$ (resp. $S \cap W_0$ and $S \cap W_2$). Hence, for $\ell \in \{1, 2\}$, it should hold that $\sum_{j \in S \cap W_0} q_j + \sum_{j \in S \cap W_\ell} q_j \geq d_\ell$. Our algorithm then can work as follows: After computing the values of Q , according to Equation (3), return the set of workers that minimize $c^{(0)} + c^{(1)} + c^{(2)}$, subject to $Q[0, n, c^{(0)}] + Q[\ell, n, c^{(\ell)}] \geq d_\ell$, for $\ell \in \{1, 2\}$. This can be done by enumerating all possible options, for breaking down the final cost as a sum of 3 values, $c^{(0)}$, $c^{(1)}$ and $c^{(2)}$. The formal statement can be found below.

⁴ It becomes clear in the next subsection, that the dynamic programming procedure is only needed for integral cost values.

Algorithm 3: DP(**b**) (presented for two tasks)

▷ **Input:** A bidding profile $\mathbf{b} = (\mathbf{c}, \mathbf{I})$ of a CMIC instance $(\mathbf{b}, \mathbf{q}, \mathbf{d})$ with $m = 2$

```

1 for  $\ell \in \{0, 1, 2\}$  do
2   for  $i \in \{0, 1, \dots, n\}$  do
3     for  $c \in \{0, 1, \dots, c_{sum}(\mathbf{b})\}$  do
4        $\lfloor$  Compute  $Q[\ell, i, c]$  using Equation (3)
5 return the set of workers that minimize  $c^{(0)} + c^{(1)} + c^{(2)}$ , s.t.
    $Q[0, n, c^{(0)}] + Q[\ell, n, c^{(\ell)}] \geq d_\ell, \forall \ell \in \{1, 2\}$  (or  $+\infty$ , if  $(\mathbf{b}, \mathbf{q}, \mathbf{d})$  has no solution)

```

The optimality of the DP algorithm is straightforward from the preceding discussion. Furthermore, its running time is pseudopolynomial, since the size of the table Q is $3 \cdot (|\mathcal{N}| + 1) \cdot (c_{sum}(\mathbf{b}) + 1)$ and to find the optimal solution we need to check at most $\binom{c_{sum}(\mathbf{b})}{3}$ different combinations for the decomposition of the total cost in three terms, as described earlier. It is easy to extend these ideas, for more tasks given the interval structure of the problem, i.e., the first dimension of Q will have a range of $O(m^2)$ and the enumeration part of the algorithm will require an order of $\binom{c_{sum}}{m^2}$ steps. Finally, we note that since this is an optimal mechanism and we use a deterministic, consistent tie-breaking rule, it will trivially be monotone.

Henceforth, we will be referring to this pseudopolynomial dynamic programming algorithm for any constant number of tasks, as the DP algorithm.

Theorem 7. *Given an instance of CMIC on a profile $\mathbf{b}$, with a constant number of tasks and integer costs, Algorithm DP($\mathbf{b}$) is optimal, monotone, and runs in pseudopolynomial time, i.e. polynomial in the input size and in c_{sum} .*

4.2 The FPTAS

In order to convert the DP algorithm to a truthful FPTAS, we adapt the framework of [8]. To that end, we define, for every integer k , an algorithm $A_k(\mathbf{b}, \epsilon)$, that uses the DP algorithm as a subroutine, on a subset of the initial set of bidders, with rounded costs, as follows:

Algorithm 4: $A_k(\mathbf{b}, \epsilon)$

▷ **Input:** A bidding profile $\mathbf{b} = (\mathbf{c}, \mathbf{I})$ of a CMIC instance $(\mathbf{b}, \mathbf{q}, \mathbf{d})$, $\epsilon \in (0, 1)$

Let $\mathcal{L}_k(\mathbf{c}) = \{i \in \mathcal{N} : c_i \leq 2^{k+1}\}$

```

1  $a_k = \frac{n}{\epsilon 2^k}$ 
2 for  $i \in \mathcal{L}_k(\mathbf{c})$  do
3    $\bar{c}_i = \lceil a_k \cdot c_i \rceil$ 
4  $\bar{\mathbf{b}} = (\bar{c}_i, [s_i, f_i])_{i \in \mathcal{L}_k(\mathbf{c})}$ 
5 return DP( $\bar{\mathbf{b}}$ )

```

Lemma 4. *Let $0 < \epsilon < 1$. For a bidding profile $\mathbf{b}$ and $k \geq 0$, the algorithm $A_k(\mathbf{b}, \epsilon)$ runs in time polynomial in the input size and in $\frac{1}{\epsilon}$, and if $2^k \leq OPT(\mathbf{b}) < 2^{k+1}$, it computes a solution of cost at most $(1 + \epsilon)OPT(\mathbf{b})$.*

Using Lemma 4, we can achieve the desired approximation by checking all possible values for k . However, to provide a polynomial-time mechanism, we can only test polynomially many such values, and hope that we compute the same outcome as if we were able to test all such algorithms. We will show that Algorithm 5, that tests all values up to a certain threshold, is what we need.

Algorithm 5: $A_{\text{FPTAS}}(\mathbf{b}, \epsilon)$

▷ **Input:** A bidding profile $\mathbf{b} = (\mathbf{c}, \mathbf{I})$ of a CMIC instance $((\mathbf{c}, \mathbf{I}), \mathbf{q}, \mathbf{d})$, $\epsilon \in (0, 1)$

- 1 **for** $k = 0, \dots, \lceil \log \left(\frac{nc_{\max}(\mathbf{b})}{\epsilon} \right) \rceil$ **do**
- 2 Run $A_k(\mathbf{b}, \epsilon)$ and store the winning set and its cost.
- 3 **return** the set of workers that achieve the minimum cost among the above, breaking ties in favor of the algorithm with the lowest index

Let $k^*(\mathbf{b}) := \lceil \log \left(\frac{nc_{\max}(\mathbf{b})}{\epsilon} \right) \rceil$ (or simply k^* when the bidding profile is clear from the context). To see that the algorithm is well-defined, recall that $c_{\max}(\mathbf{b}) \geq 1$ and since also $\frac{n}{\epsilon} \geq 1$, we have that $k^* \geq 0$. To establish that this is indeed a FPTAS, we prove in the following lemma that one cannot find a better solution by running an algorithm A_k for a value of k higher than k^* . In combination with Lemma 4, this directly establishes that A_{FPTAS} is a FPTAS for CMIC.

Lemma 5. *Given an instance of CMIC and an $\epsilon \in (0, 1)$, it holds that $\mathcal{W}_{A_k(\mathbf{b})} = \mathcal{W}_{A_{k^*}(\mathbf{b})}$, for every $k > k^*(\mathbf{b})$.*

Monotonicity of A_{FPTAS}

To establish monotonicity, we will make use of the following operator:

Definition 3. *Let $\mathcal{A} = \{A_0, A_1, \dots\}$ be the set of all allocation algorithms A_k . For a profile $\mathbf{b}$ and a finite collection of algorithms $S \subseteq \mathcal{A}$, let $\text{MIN}(S, \mathbf{b}) := \arg \min_{A \in S} C(A, \mathbf{b})$, with ties broken in favor of the lowest index.*

Given a bidding profile $\mathbf{b}$, the algorithm A_{FPTAS} can be expressed as $\text{MIN}\{A_0, \dots, A_{k^*(\mathbf{b})}\}$. Hence, the next step is to determine when is the MIN operator monotone. The framework of [8] defines a set of sufficient conditions, for maximization objectives. We adapt these properties below, and we note that they are sufficient conditions for minimization problems as well.

Definition 4. *A monotone allocation algorithm A is bitonic w.r.t. the social cost function C if for any bidding profile $\mathbf{b}$ and any worker i , the following hold:*

1. $i \in \mathcal{W}_A(\mathbf{b}) \Rightarrow C(A, \mathbf{b}) \geq C(A, (b'_i, \mathbf{b}_{-i})) \quad \forall b'_i \succeq b_i$
2. $i \notin \mathcal{W}_A(\mathbf{b}) \Rightarrow C(A, \mathbf{b}) \geq C(A, (b'_i, \mathbf{b}_{-i})) \quad \forall b'_i \preceq b_i$

Lemma 6. *For $k \geq 0$, Algorithm 4 is monotone and bitonic w.r.t. the social cost function C .*

Lemma 6 will be used in conjunction with the following Lemma.

Lemma 7. *(implied by [8]) For a bidding profile $\mathbf{b}$, $\text{MIN}(\{A_0, \dots, A_\ell\}, \mathbf{b})$ is monotone if $A_0, \dots, A_\ell$ are monotone algorithms that are additionally bitonic w.r.t. the social cost function C .*

Before we proceed, we would like to comment on a subtle point regarding the MIN operator. We stress that the set of algorithms A_k that are called by A_{FPTAS} , depends on the input profile $\mathbf{b}$. There is no a priori fixed set of algorithms that are run in every profile, but instead, this is determined by the quantity $k^*(\mathbf{b})$. As a result, Lemma 7 does not suffice on its own. For the monotonicity of A_{FPTAS} , we also need to consider the case that a winning worker declares a lower cost that changes $c_{\max}(\mathbf{b})$, and decreases the number of algorithms that A_{FPTAS} runs. This is where Lemma 5 comes to rescue, as explained in the proof of Theorem 8, and this is where the flaw in [13] is located (i.e., the algorithm of [13] does not perform enough iterations).

Given the above discussion, by performing a suitable case analysis and by using Lemmas 6 and 7, we can prove the following theorem, which also completes the proof of Theorem 6 and concludes this section.

Theorem 8. *For the domain of bidding profiles $\mathbf{b}$, such that $c_{\min}(\mathbf{b}) \geq 1$, the algorithm A_{FPTAS} is monotone.*

5 Further Implications and Extensions

Apart from truthful mechanisms, it would be interesting to also consider CMIC from the purely algorithmic angle. In particular, the existence of a FPTAS (or other types of approximation schemes) could be possible beyond the case studied in Section 4. Furthermore, it is natural to examine whether the mechanisms presented in the previous sections could be used in more general crowdsourcing scenarios. The following subsections shed more light towards these two questions.

5.1 Relation to Unsplittable Flow Problems and (In)approximability

At what follows, we first establish a connection between CMIC and other relevant algorithmic problems, namely UFP, weighted UFP-COVER, and FAULT-CACHING. These problems concern (a) the selection of subpaths of a path-graph so as to satisfy a demand (resp. capacity) constraint on every edge, at a minimum (resp. maximum) total cost and (b) the minimization of the total cache misses in a fault model with a cache of fixed size and requests for non-uniform size pages. We refer to the full version of the work for the definitions of the problems. Such connections are interesting in their own right, and are summarized in the following theorem.

Theorem 9. *CMIC is equivalent to weighted UFP-COVER. Furthermore UFP and FAULT-CACHING are special cases of CMIC.*

We can now exploit some known results about these problems and obtain analogous implications for CMIC. The following table lists the most important results that concern the polynomial solvability and the approximability of (weighted) UFP-COVER, UFP and FAULT-CACHING that also apply for CMIC, due to Theorem 9. As an example, we can now rule out a FPTAS for the general case of CMIC (and in fact for its simpler cases, as stated in the first two lines of Table 2), unless $P=NP$. We note that before our work, there were no results about the hardness of CMIC, other than the obvious NP-hardness result in [18]. We comment further on Table 2 in Section 6.

Restriction	Result for CMIC	Reduction
$d_i = 1 \quad \forall i \in [m]$ $q_i \in \{1, 2, 3\} \quad \forall i \in \mathcal{N}$ $c_i \in \mathbb{Z}$	strongly NP-hard $\xrightarrow{[40]} \nexists$ FPTAS *	UFP [6]
$d_i = 1 \quad \forall i \in [m]$ $c_i = 1 \quad \forall i \in \mathcal{N}$	strongly NP-hard $\xrightarrow{[40]} \nexists$ FPTAS *	FAULT-CACHING [14]
$c_i = 1 \quad \forall i \in \mathcal{N}$	W[1]-hard $\xrightarrow{[21]} \nexists$ EPTAS $^\odot$	UFP-COVER [16]
$q_i : \text{quasi-poly} \quad \forall i \in \mathcal{N}$	QPTAS $\xrightarrow{[1]} \nexists$ APX-hardness $^\diamond$	weighted UFP-COVER [26]

Table 2: Results concerning the (in)approximability of CMIC implied by Theorem 9 together with existing works on UFP-COVER, UFP and FAULT-CACHING. We have used the symbols $\star, \odot, \diamond$ to denote the complexity assumptions $P \neq NP$, $W[1] \neq FPT$ and $QP \neq NP$ respectively. For the W[1]-hardness result, it is assumed that UFP-COVER is parameterized by the cardinality of the optimal solution.

5.2 Generalization to Non-Interval Structures

The focus of our work has been largely on highlighting the difference on the approximation ratio between truthful and non-truthful algorithms. We conclude our work by showing that one can get tighter results, when moving to more general scenarios. A direct generalization of CMIC is to drop the linear arrangement of the tasks, and allow each worker $j \in \mathcal{N}$ to declare an arbitrary subset of tasks $I_j \subseteq [m]$. The rest of the input remains the same (costs, contributions and demands), and we refer to this problem as Cost Minimization Demand Cover (CMDC). Notice that CMDC with $q_i = 1$ for every worker i , and $d_j = 1$ for every task j , is nothing but the famous SET COVER problem.

Further extensions of CMDC have been studied under various names in a series of works regarding covering IPs (e.g., [28,38]) and approximation algorithms that match the factor of Algorithm 1 exist (e.g. [9,22,30,37]). However, the focus of these works was not about monotonicity and it is unclear if any of these algorithms are monotone (in fact some of them are certainly non-monotone).

Towards obtaining a monotone algorithm, a careful inspection of the proofs of Section 3, suffices to deduce that Algorithm 1 can be used for this more general setting as well (after defining first the appropriate generalization of LMIS) and it continues to yield the Δ factor for CMDC. Furthermore, under this setting, Algorithm 1 yields essentially a tight result, according to the following Proposition. The proof of Proposition 3 is straightforward due to the hardness results

for k -UNIFORM HYPERGRAPH VERTEX COVER [27,19], which is a special case of SET COVER, and therefore a special case of CMDC.

Proposition 3. *For the class of instances where Δ is constant, CMDC is $\Delta-1-\epsilon$ inapproximable, unless $P = NP$, and $\Delta-\epsilon$ inapproximable assuming the Unique Games Conjecture.*

Finally, we note that the algorithms of Section 4 can in principle also be applied for CMDC with no loss in the approximation factor, but at the expense of a much higher running time (doubly exponential in m , which still remains polynomial, as long as the number of tasks is a constant).

6 Discussion and Open Problems

From a mechanism design viewpoint, the most important question for future research is to design truthful mechanisms with better approximation guarantees for CMIC, as there is still a large gap between the non-truthful 3-approximation and our truthful Δ -approximation of Section 3. Moreover, exploring the approximability of well-motivated special cases of CMIC, other than the restriction on a constant number of tasks that we examined in Section 4, is also an intriguing topic (Table 2 in Section 5.1 suggests such possible restrictions). In the context of crowdsourcing, some of these special cases are also meaningful to study in a 2-dimensional model, where workers can cover circular areas of a given radius, or other geometric shapes. Apart from positive results, we believe it is very interesting to investigate the existence of lower bounds on the worst case performance of polynomial time truthful mechanisms. After all, it is conceivable that there may be a strict separation on the approximability by monotone and non-monotone algorithms, as, e.g., for combinatorial public project problems [35].

From a purely algorithmic viewpoint, we would like to point out some interesting open-problems for CMIC, that emerge from the results of Table 2. In particular, it has been known by [18], that we have polynomial solvability when all the workers have the same contribution parameter. On the contrary, by the first row of Table 2, we have a hardness result when the workers can be partitioned in three distinct groups by their contribution. It still remains open to determine what happens when we have two distinct groups for the contribution of the workers. Interestingly the last row of Table 2 implies that the existence of a PTAS is likely (at least for the case of quasi-polynomially bounded contributions) and, yet, such a PTAS still remains to be found. Finally, even though the problem is $W[1]$ -hard (parameterized by the number of workers in the optimal solution) for unit costs, it belongs to FPT for unit costs and integral numerical values [32,24], and hence an EPTAS may well exist for this case.

Acknowledgements. This research was supported by the Hellenic Foundation for Research and Innovation. The first two authors were supported by the “1st Call for HFRI Research Projects to support faculty members and researchers and the procurement of high-cost research equipment” (Project Num. HFRI-FM17-3512) and the third author by the HFRI PhD Fellowship grant (Fellowship Num. 289).

References

1. Bansal, N., Chakrabarti, A., Epstein, A., Schieber, B.: A quasi-PTAS for unsplittable flow on line graphs. In: *Proceedings of the 38th ACM Symposium on Theory of Computing*. pp. 721–729 (2006)
2. Bansal, N., Pruhs, K.: The geometry of scheduling. *SIAM Journal on Computing* **43**(5), 1684–1698 (2014)
3. Bar-Noy, A., Bar-Yehuda, R., Freund, A., Naor, J., Schieber, B.: A unified approach to approximating resource allocation and scheduling. *Journal of the ACM* **48**(5), 1069–1090 (2001)
4. Bar-Yehuda, R., Bendel, K., Freund, A., Rawitz, D.: The local ratio technique and its application to scheduling and resource allocation problems. In: *Graph Theory, Combinatorics and Algorithms*. Springer (2005)
5. Blumrosen, L., Nisan, N.: *Algorithmic Game Theory. Introduction to mechanism design*, Cambridge University Press (2007)
6. Bonsma, P., Schulz, J., Wiese, A.: A constant-factor approximation algorithm for unsplittable flow on paths. *SIAM Journal on Computing* **43**(2), 767–799 (2014)
7. Brederbeck, R., Faliszewski, P., Niedermeier, R., Skowron, P., Talmon, N.: Mixed integer programming with convex/concave constraints: Fixed-parameter tractability and applications to multicovering and voting. *Theoretical Computer Science* (2020)
8. Briest, P., Krysta, P., Vöcking, B.: Approximation techniques for utilitarian mechanism design. *SIAM Journal on Computing* **40**(6), 1587–1622 (2011)
9. Carr, R.D., Fleischer, L.K., Leung, V.J., Phillips, C.A.: Strengthening integrality gaps for capacitated network design and covering problems. In: *Proceedings of the 11th ACM-SIAM Symposium on Discrete Algorithms*. p. 106–115 (2000)
10. Chakaravarthy, V.T., Kumar, A., Roy, S., Sabharwal, Y.: Resource allocation for covering time varying demands. In: *European Symposium on Algorithms* (2011)
11. Chakrabarty, D., Grant, E., Könemann, J.: On column-restricted and priority covering integer programs. In: *International Conference on Integer Programming and Combinatorial Optimization*. pp. 355–368. Springer (2010)
12. Chen, B., Hassin, R., Tzur, M.: Allocation of bandwidth and storage. *IIE Transactions* **34**(5), 501–507 (2002)
13. Chen, J., Ye, D., Ji, S., He, Q., Xiang, Y., Liu, Z.: A truthful FPTAS mechanism for emergency demand response in colocation data centers. In: *Proceedings of the IEEE International Conference on Computer Communications-INFOCOM* (2019)
14. Chrobak, M., Woeginger, G.J., Makino, K., Xu, H.: Caching is hard—even in the fault model. *Algorithmica* **63**(4), 781–794 (2012)
15. Cramton, P., Shoham, Y., Steinberg, R., et al.: *Combinatorial auctions*. Tech. rep., University of Maryland (2006)
16. Cristi, A., Mari, M., Wiese, A.: Fixed-parameter algorithms for unsplittable flow cover. *Theory of Computing Systems* (2021)
17. Csirik, J.: Heuristics for the 0-1 min-knapsack problem. *Acta Cybernetica* (1991)
18. Dayama, P., Narayanaswamy, B., Garg, D., Narahari, Y.: Truthful interval cover mechanisms for crowdsourcing applications. In: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems* (2015)
19. Dinur, I., Guruswami, V., Khot, S., Regev, O.: A new multilayered PCP and the hardness of hypergraph vertex cover. *SIAM Journal on Computing* **34**(5), 1129–1146 (2005)

20. Elkind, E., Goldberg, L.A., Goldberg, P.W.: Frugality ratios and improved truthful mechanisms for vertex cover. In: Proceedings of the 8th ACM Conference on Electronic Commerce (2007)
21. Feldmann, A., Karthik, C., Lee, E., Manurangsi, P.: A survey on approximation in parameterized complexity: Hardness and algorithms. *Algorithms* **13**(6) (2020)
22. Fujito, T., Yabuta, T.: Submodular integer cover and its application to production planning. In: International Workshop on Approximation and Online Algorithms. pp. 154–166. Springer (2004)
23. Gálvez, W., Grandoni, F., Ingala, S., Heydrich, S., Khan, A., Wiese, A.: Approximating geometric knapsack via L-packings. *ACM Transactions on Algorithms* **17**(4), 1–67 (2021)
24. Gavenciak, T., Knop, D., Koutecký, M.: Integer programming in parameterized complexity: three miniatures. In: 13th International Symposium on Parameterized and Exact Computation. pp. 21:1–21:16 (2019)
25. Gummidi, S.R.B., Xie, X., Pedersen, T.B.: A survey of spatial crowdsourcing. *ACM Transactions on Database Systems* **44**(2), 1–46 (2019)
26. Höhn, W., Mestre, J., Wiese, A.: How unsplittable-flow-covering helps scheduling with job-dependent cost functions. *Algorithmica* (2018)
27. Khot, S., Regev, O.: Vertex cover might be hard to approximate to within $2 - \epsilon$. *Journal of Computer and System Sciences* **74**(3), 335–349 (2008)
28. Kolliopoulos, S., Young, N.: Approximation algorithms for covering/packing integer programs. *Journal of Computer and System Sciences* **71**(4), 495–505 (2005)
29. Kothari, A., Parkes, D., Suri, S.: Approximately-strategyproof and tractable multi-unit auctions. *Decision Support Systems* **39**(1), 105–121 (2005)
30. Koufogiannakis, C., Young, N.: Greedy δ -approximation algorithm for covering with arbitrary constraints and submodular cost. *Algorithmica* (2013)
31. Lehmann, D., O’Callaghan, L.I., Shoham, Y.: Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM* **49**(5), 577–602 (2002)
32. Lenstra, H.: Integer programming with a fixed number of variables. *Mathematics of Operations Research* (1983)
33. Mondal, S.: Improved algorithm for resource allocation problems. *Asia-Pacific Journal of Operational Research* **35**(01), 1–23 (2018)
34. Myerson, R.: Optimal auction design. *Mathematics of Operations Research* **6**(1), 58–73 (1981)
35. Papadimitriou, C., Schapira, M., Singer, Y.: On the hardness of being truthful. In: Proceedings of the 49th Foundations of Computer Science. pp. 250–259 (2008)
36. Phillips, C.A., Uma, R., Wein, J.: Off-line admission control for general scheduling problems. *Journal of Scheduling* **3**(6), 365–381 (2000)
37. Pritchard, D., Chakrabarty, D.: Approximability of sparse integer programs. *Algorithmica* **61**(1), 75–93 (2011)
38. Rajagopalan, S., Vazirani, V.: Primal-dual RNC approximation algorithms for (multi)-set (multi)-cover and covering integer programs. In: Proceedings of the 34th Foundations of Computer Science (1993)
39. Tong, Y., Zhou, Z., Zeng, Y., Chen, L., Shahabi, C.: Spatial crowdsourcing: a survey. *The VLDB Journal* **29**(1), 217–250 (2020)
40. Vazirani, V.: Approximation algorithms. Springer (2001)
41. Xu, J., Xiang, J., Yang, D.: Incentive mechanisms for time window dependent tasks in mobile crowdsensing. *IEEE Transactions on Wireless Communications* (2015)
42. Zhang, L., Ren, S., Wu, C., Li, Z.: A truthful incentive mechanism for emergency demand response in colocation data centers. In: Proceedings of the IEEE International Conference on Computer Communications-INFOCOM (2015)

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

Evangelos Markakis¹, Alkmini Sgouritsa², and Artem Tsikiris¹

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 $\mathcal{N} = \{1, \dots, n\}$ 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 $\mathbf{b} = (b_1, \dots, b_n)$ 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 $q'_i \in \mathbb{Z}_+$

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

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 $i = 1, \dots, n$ 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 $\beta(\mathbf{b})_{-i} = (\beta_1(\mathbf{b}_{-i}), \dots, \beta_k(\mathbf{b}_{-i}))$. Here, $\beta_j(\mathbf{b}_{-i})$ is the j -th lowest winning bid of the profile $\mathbf{b}_{-i}$, for $j = 1, \dots, k$, 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 $j = 1, \dots, k-1$ 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 $i = 1, \dots, n$, we denote the CDF of the random variable $\beta_j(\mathbf{b}_{-i})$ as $\hat{F}_{ij}$, for $j = 1, \dots, k$. 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 $\text{Supp}(B_i)$ (or of a subset of it) into smaller disjoint sub-intervals, say $I_1, \dots, I_\ell$, 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 $r = 1, \dots, \ell$.*

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 $z \in \text{Supp}(B_i)$, 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 $\ell = \ell(\mathbf{B}_i) \in \text{Supp}(B_i)$. 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, $\mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(\ell, \mathbf{b}_{-i})] = 0$. $\square$

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 $\text{Supp}(B_i) \subseteq \bigcup_{j \in W(\mathbf{B}) \setminus \{i\}} \text{Supp}(B_j)$.

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 $z \in \text{Supp}(B_i) \setminus \{\ell(\mathbf{B}_W)\}$, with $F_i(z) > \lim_{z \rightarrow z^-} F_i(z)$, 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 $z \in \text{Supp}(B_i) \setminus \{\ell(\mathbf{B}_W)\}$, $\mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(z, \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})]$.
- 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 $\ell(\mathbf{B}_W) \in \text{Supp}(B_i)$, $\mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(\ell(\mathbf{B}_W), \mathbf{b}_{-i})] = \mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})]$.

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

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

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 $i \in W(\mathbf{B})$, $\bigcup_{j \in W(\mathbf{B}) \setminus \{i\}} \text{Supp}(B_j) = [\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$.

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

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

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*

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

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, 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

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

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. $\square$

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. $\text{Supp}(B_\alpha) = [\ell(\mathbf{B}_W), h(\mathbf{B}_W)]$.
2. For any two bidders $i, j \in W(\mathbf{B}) \setminus \{\alpha\}$ such that $v_i \neq v_j$, the set $\text{Supp}(B_i) \cap \text{Supp}(B_j)$ 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 $v_1, \dots, v_r$, that bidders in $W(\mathbf{B}) \setminus \{\alpha\}$ have. We can partition the bidders of $W(\mathbf{B}) \setminus \{\alpha\}$ into r groups $W_1(\mathbf{B}), \dots, W_r(\mathbf{B})$, such that, for $j = 1, \dots, r$, 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 $j \in \{1, \dots, r\}$ is formed as $I_j(\mathbf{B}) = \bigcup_{i \in W_j(\mathbf{B})} \text{Supp}(B_i)$. The following is a direct corollary of Lemma 11.

Corollary 5. *For every $s, t \in \{1, \dots, r\}$ 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 $i = 1, \dots, n$ 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 $i = 1, \dots, n$, 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 $i = 1, \dots, n - 1$ 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 $i = 1, \dots, n$, $F_\alpha(z) = F_\alpha^i(z)$ 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 $i = 2, \dots, n$, $m = 1, \dots, i-1$, and every $z \in [h(B_{m-1}), h(B_m)]$,

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

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

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

Proof of Theorem 8. For brevity, we denote $\ell(B_a)$ as ℓ and for $j = 1, \dots, n$, 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

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

The first inequality is true since for all bidders $i = 1, \dots, n$, 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

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

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., Tsikiris, 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 (2013)
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)