Budget Feasible Mechanisms for Procurement Auctions with Divisible Agents

Sophie Klumper[1,2] and Guido Schäfer[1,3]

[1] Networks and Optimization, Centrum Wiskunde & Informatica (CWI)
[2] Department of Mathematics, Vrije Universiteit Amsterdam
[3] Institute for Logic Language and Computation, University of Amsterdam
The Netherlands
{sjk,g.schaefer}@cwi.nl

Abstract. We consider budget feasible mechanisms for procurement auctions with additive valuation functions. For the divisible case, where agents can be allocated fractionally, there exists an optimal mechanism with approximation guarantee $e/(e-1)$ under the small bidder assumption. We study the divisible case without the small bidder assumption, but assume that the true costs of the agents are bounded by the budget. This setting lends itself to modeling economic situations in which the goods represent time and the agents' true costs are not necessarily small compared to the budget. Non-trivially, we give a mechanism with an approximation guarantee of 2.62, improving the optimal result of 3 for the indivisible case. Additionally, we give a lower bound on the approximation guarantee of 1.25. We then study the problem in more competitive markets and assume that the agents' value over cost efficiencies are bounded by some θ. For $\theta \leq 2$, we give a mechanism with an approximation guarantee of 2 and a lower bound of 1.25. Finally, we extend our results to different agent types with concave additive valuation functions.

Keywords: Mechanism design · Procurement auction · Budget feasible mechanism · Divisible agents · Knapsack auction · Additive valuations

1 Introduction

We consider procurement auctions in which the auctioneer has a budget that limits the total payments that can be paid to the agents. In this setting, we are given a set of agents $A = [n]$ offering some service (or good), where each agent $i \in A$ has a privately known cost c_i and a publicly known valuation v_i. The auctioneer wants to do business with these agents and needs to decide with which agents $S \subseteq A$ to do so. In order for each agent $i \in S$ to comply, the auctioneer will have to make a payment p_i to this agent i. The auctioneer has a total budget B available for these payments. The goal of the auctioneer is to find a subset of agents $S \subseteq A$ that (approximately) maximizes the total value, while the total payment is at most B; such payments are said to be *budget feasible*. As the costs of the agents are assumed to be private, agents may misreport their actual costs. The goal is to design a mechanism that computes this subset and budget feasible payments, such that the agents have no incentive to misreport their actual costs (i.e., the agents are truthful).

Basically, there are two standard approaches in the literature to study this problem: (i) in the *Bayesian setting* it is assumed that the distributions of the agents' true costs are known, and (ii) in the *prior-free setting* it is assumed that nothing is known about these distributions. In this paper, we focus on the prior-free setting as it may not be possible to extract representative distributions of the agents' true costs.

Notably, most previous studies focus on the *indivisible case*, where the services (or goods) offered by the agents need to be allocated integrally. In contrast, the *divisible case*, where the agents can also be allocated fractionally, received much less attention and has been studied only under the so-called *small bidder assumption*, i.e., when the agents' costs are much smaller than the available budget. While this assumption is justified in certain settings (e.g., for large markets), it is less appropriate in settings where the agents' costs may differ vastly or be close to the budget. This motivates the main question that we address in this paper: Can we derive budget feasible truthful mechanisms with attractive approximation guarantees if the agents are divisible?

Singer [1] initiated the study of budget feasible mechanisms in the prior-free setting with indivisible agents and designed a deterministic mechanisms with an approximation guarantee of 5. Later, Chen et al. [2] improved this to $2+\sqrt{2}$. Almost a decade later, Gravin et al. [3] improved the approximation guarantee to 3 and showed that no mechanism can do better than 3 (when compared to the fractional optimal solution). All three papers also provide some additional results such as randomized algorithms with improved approximation guarantees and mechanisms for the more general setting of submodular functions.

To the best of our knowledge, the divisible case of knapsack procurement auctions has not been studied so far. When an agent is offering a service, which can be interpreted as offering time, it is suitable to model this as a divisible good. As an example, consider the situation where an auctioneer has a budget available and wants to organize a local comedy show. Agents are offering to perform and for each agent the auctioneer has a valuation, which reflects the amusement when this agent performs. The agents have costs related to their performance which could consist of the invested time in their performance, attributes needed, travel costs, etc. In this example it makes sense for the auctioneer to have the option to select agents fractionally. After selecting one agent, the budget left might only be enough to let some agent perform half of what they are offering. Or the auctioneer might want to have at least three performances and must select agents fractionally to achieve this due to the budget constraint.

As mentioned, the prior-free setting with divisible goods has been studied by Anari et al. [4] under the small bidder assumption. More formally, if $c_{\max} = \max_{i\in A}\{c_i\}$ and $r = c_{\max}/B$, then the results are analysed for $r \to 0$. Anari et al. [4] give optimal mechanisms for both the divisible (deterministic) and indivisible (randomized) case, both with an approximation guarantee of $e/(e-1)$. They also mention that in the divisible case, no truthful mechanism with a finite approximation guarantee exists without the small bidder assumption. This can already be shown by an instance with one agent with value v and true cost $c > B$ and budget B. If the agent is not strategic, an optimal fractional value of $v\frac{B}{c}$ can be achieved. A mechanism must select the agent fractionally, say to the extent of x, in order to achieve some positive value. Additionally, it must be that $c \cdot x \leq p \leq B$, with p the payment of the agent, for the mechanism to be individually rational and budget feasible. As we do not know c or any bound on c, there is no positive value of x for which we can guarantee this, and therefore no finite approximation guarantee exists.

However, if we assume that the true costs of the agents are bounded by the budget, there is still a lot we can do without the small bidder assumption and there are settings in which it makes sense to have this assumption. If we revisit our example of an auctioneer wanting to organize a local comedy show, an internationally famous comedian with a true cost exceeding the budget will most likely not be one of the agents offering to perform. On the other hand, there might be a national well-known comedian offering to perform. This comedian might have a true cost that is smaller than the budget, but not much smaller than the available budget, which is what the small bidder assumption requires.

We note that we can derive some mechanisms with finite approximation guarantees if the true costs of the agents are bounded by γB for some constant $\gamma > 1$. However, we omit the details of these results from this extended abstract and defer further discussion to the full version of the paper.

Our Contributions. For the knapsack procurement auction with divisible goods and true costs bounded by the budget, we give a mechanism with an approximation guarantee of 2.63 in Section 3. Additionally, we proof that no truthful mechanism can achieve an approximation guarantee better than 1.25. Although the divisible case gives more freedom in designing an allocation rule, improving the approximation guarantee compared to the indivisible case is non-trivial. In particular, it remains difficult to bound the threshold payments and the complexity of the payment rule increases compared to the indivisible case. Additionally, if an agent is selected fractionally one needs to determine this fraction exactly while remaining budget feasible. A natural method is to determine this fraction based on the agents' declared costs, but then one must limit the influence that this method gives to the agents.

Proving the above lower bound of 1.25 on the approximation guarantee requires that the true costs of the agents differ significantly ($\epsilon \ll B$), a property that is used more often when proving lower bounds. Therefore, in Section 4, we introduce a setting in which the agents' efficiencies (i.e., value over cost ratios) are bounded by some $\theta \geq 1$. It is reasonable to assume that the efficiencies of the agents is somewhat bounded, as agents will cease to exist if they cannot compete with the other agents in terms of efficiency. Another interpretation of this setting is some middle ground between the prior-free and the Bayesian setting. It might be impossible to find representative distributions of the agents' true

Assumptions		**Origin**	**Approximation Guarantee**	
			Upper bound	Lower bound
Indivisible		[1]	5	2
		[2]	≈ 3.41	≈ 2.41
		[3]	3	3^*
Divisible ($c_i \leq B\ \forall i$)	none	Section 3	≈ 2.62	1.25
	$\theta \leq 2$	Section 4	2	≈ 1.18
	$\big((t_i),(l_j)\big)$	Section 5	≈ 2.62	1.25
	$\theta \leq 2$ and $\big((t_i),(l_j)\big)$	Section 5	2	≈ 1.18

* When compared to the fractional knapsack problem.

Table 1. Overview of the results obtained in this paper.

costs and address the problem in the Bayesian setting, but the auctioneer could have information about the minimum and maximum efficiencies in the market by previous experiences or market research. We give a truthful mechanism with an approximation guarantee of 2 when the efficiency of the agents is bounded by a factor $\theta \leq 2$. For this case we also proof a lower bound of 1.18 on the approximation guarantee and generalize this for different values of θ.

In Section 5, we extend our results to a more general model in which agents may have different types. Our new model allows us to capture more complex settings. If we revisit the example of hosting a comedy show, this translates to each agent having a certain type of comedy that they preform. In order to set up a nice diverse program, the auctioneer wishes the jokes of a certain type to be limited. The auctioneer knows that, if at some point in time too many jokes of the same type are preformed, the additional value added to the show decreases. This is modeled by a concave non-decreasing valuation function for each type. We prove that for this setting the mechanisms of Section 3 and Section 4 still hold with the same approximation guarantee when slightly altered.

Related Work. As mentioned earlier, settings with different valuation functions have been studied for the indivisible case. In the case of submodular valuation functions, Singer [1] gave a randomized mechanism with an approximation guarantee of 112. Again, Chen et al. [2] improved this to 7.91 and gave a deterministic exponential time mechanism with an approximation guarantee of 8.34. Later Jalaly and Tardos [5] improved this to 5 and 4.56 respectively. For subadditive valuation functions, Dobzinski et al. [10] gave a randomized and deterministic mechanism with an approximation guarantee of $O(\log^2 n)$ and $O(\log^3 n)$ respectively. Bei et al. [11] improved this to $O(\log n/(\log\log n))$ with a polynomial time randomized mechanism and gave a randomized exponential time mechanism with an approximation guarantee of 768 for XOS valuation functions. Leonardi et al. [7] improved the latter to 436 by tuning the parameters of the mechanism.

Related settings have also been studied for the indivisible case. Leonardi et al. [7] consider the problem with an underlying matroid structure, where each element corresponds to an agent and the auctioneer can only select an independent set. Chan and Chen [8] studied the setting in which agents offer multiple units of their good. They regard concave additive and subadditive valuation functions. The setting in which the auctioneer wants to get a set of heterogeneous tasks done and where each task requires the performing agent to have a certain skill has been studied by Goel et al. [6]. They give a randomized mechanism with an approximation guarantee of 2.58, which is truthful under the small bidder assumption. Jalaly and Tardos [5] match this result with a deterministic mechanism. The results of Goel et al. [6] can be extended to settings in which tasks can be done multiple times and agents can perform multiple tasks. Related to this line of work is also the strategic version of matching and coverage, in which edges and subsets represent strategic agents, that Singer [1] also studied. Chen et al. [2] also studied the knapsack problem with heterogeneous items, were items are divided in groups and at most one item from each group can be selected. Amanatidis et al. [12] give randomized and deterministic mechanisms for a subclass of XOS problems. For the mechanism design version of the budgeted max weighted matching problem, they give a randomized (deterministic) mechanism with an approximation guarantee of 3 (4) and they generalize their results to problems with a similar combinatorial structure.

2 Preliminaries

More formally the problem is given declared costs $\mathbf{b} = (b_i)_{i\in A}$, non-negative valuations $\mathbf{v} = (v_i)_{i\in A}$ and non-negative budget B, to design a mechanism that computes an allocation vector $\mathbf{x} = (x_i)_{i\in A}$ and a payment vector $\mathbf{p} = (p_i)_{i\in A}$. The allocation vector should satisfy $x_i \in [0,1]$ for all $i \in A$ with x_i the fraction of agent i selected. An element p_i of the payment vector corresponds to the payment of agent i. We want to (approximately) maximize the value of the allocation vector and the total of the payments to be within the budget. We assume that the true costs $\mathbf{c} = (c_i)_{i\in A}$ are positive and bounded by the budget, i.e. $0 < c_i \leq B$ for all $i \in A$. We also assume that the cost incurred by the agents are linear, i.e. given allocation vector $\mathbf{x}$ the cost incurred by agent i is equal to $c_i x_i$. Therefore the utility u_i of agent $i \in A$ is equal to $u_i = p_i - c_i x_i$. The auctioneer has an additive valuation function, i.e. given allocation vector $\mathbf{x}$ the value derived by the auctioneer is $v(\mathbf{x}) = \sum_{i\in A} v_i x_i$. The goal of each player $i \in A$ is to maximize their utility u_i, and as the costs of the agents are assumed to be private, they may misreport their actual costs to achieve this. If the agents are not strategic, i.e. the costs are publicly known, the above setting naturally corresponds to the fractional knapsack problem. As common in mechanism design, we seek mechanisms that are

1. *Truthful:* for every agent reporting their true cost is a dominant strategy;
2. *Individually rational:* every agent has a non-negative utility;
3. *Budget Feasible:* the sum of all the payments is smaller than the budget;
4. *Approximation Guarantee:* the mechanism has an approximation guarantee of $\alpha \geq 1$ if for each instance of the problem it holds that $\alpha \cdot v(\mathbf{x}) \geq \texttt{opt}$, with $\texttt{opt}$ the value of the fractional knapsack problem;

5. *Computationally Efficient:* the allocation and payment vector can be computed in polynomial time.

In order to design truthful mechanisms we will use Theorem 1 (stated below) and come up with mechanisms that have a *monotone* allocation rule. Given declared costs $\mathbf{b} = (b_i)_{i\in A}$, let x_i be the fraction of agent $i \in A$ selected. Then the allocation rule is monotone if agent i is selected to at least the extent of x_i if they would have declared a lower cost and everything else remains unchanged.

Theorem 1. *(Archer and Tardos [9]) A monotone allocation rule $x(\mathbf{b})$ admits a truthful payment rule that is individually rational if and only if for all $i, \mathbf{b}_{-i}$: $\int_0^\infty x_i(u, \mathbf{b}_{-i})du < \infty$. In this case we can take the payment rule $p(\mathbf{b})$ to be*

$$p_i(b_i, \mathbf{b}_{-i}) = b_i \cdot x_i(b_i, \mathbf{b}_{-i}) + \int_{b_i}^{\infty} x_i(u, \mathbf{b}_{-i})du, \quad \forall i. \tag{1}$$

Next we introduce some notation that is used throughout the paper. We say agent i is a *winning* agent if $x_i > 0$. We define the cost of an allocation vector $\mathbf{x}$ as $c(\mathbf{x}) = \sum_{i\in A} c_i x_i$. We define the *efficiency* of an agent i as the ratio of value over (declared) cost ($\frac{v_i}{c_i}$). Agents with a high efficiency are preferable as, compared to agents with a lower efficiency, they relatively contribute a higher value per unit (declared) cost. Whenever we order the agents according to decreasing efficiencies, we assume that ties are broken arbitrarily but consistently. We define $\texttt{opt}(A, \mathbf{c})$ as the value of the optimal solution regarding the set of agents A, costs $\mathbf{c} = (c_i)_{i\in A}$ and values $\mathbf{v} = (v_i)_{i\in A}$. $\texttt{opt}_{-i}(A, \mathbf{c})$ is defined similarly, only regarding the set of agents in $A \setminus \{i\}$. For the sake of readability, we omit the valuations $\mathbf{v}$ as an argument as these are publicly known. Also, we omit the remaining arguments if they are clear from the context and simply use $\texttt{opt}$.

3 Divisible Agents

The Mechanism. It is well-known that designing a truthful and individually rational mechanism comes down to designing a monotone allocation rule and implementing it together with the payment rule of Theorem 1. When selecting an agent i in this case, the first term of (1) is equal to cost incurred by agent i and the second term can be interpreted as the amount agent i is overpaid. It is beneficial for this second term to be as small as possible, in order to select more agents and achieve a better approximation guarantee.

We give a mechanism that realizes this by imposing a threat, say τ_i, to an agent i in the second part of the mechanism. If the declared cost of agent i exceeds this threat, i will not be selected. In this case, the second term of the payment formula is equivalent to $\int_{c_i}^{\tau_i} x_i(u, \mathbf{b}_{-i})du$, and the goal is for τ_i to be as close to the true cost c_i, while remaining larger or equal to avoid a negative effect on the approximation guarantee. In order to use these threats to bound the payments, τ_i must not increase when agent i declares a higher cost. Note that bounding the payments with τ_i can still be done if τ_i decreases when agent

i declares a higher cost. In our mechanism, we use the following threat τ_i for agent i, which is independent of the declared cost of i:

$$\tau_i = v_i \frac{B}{\alpha(1+\beta)\texttt{opt}_{-i}(N,\mathbf{c})}.$$

Note that this threat τ_i imposes an upper bound on the threshold bid of agent i, i.e., the largest bid agent i can declare such that i is still selected to some extent. Let $\alpha \in (0,1]$ and $\beta > 0$ be some parameters which we fix later. Our mechanism is as follows:

Divisible Agents (DA)

Let $N = \{i \in A \mid c_i \leq B\}$, $n = |N|$ and $i^* = \arg\max_{i \in N} \frac{v_i}{\texttt{opt}_{-i}(N,\mathbf{c})}$

if $\frac{v_{i^*}}{\texttt{opt}_{-i^*}(N,\mathbf{c})} \geq \beta$ **then** set $x_{i^*} = 1$ and $x_i = 0$ for $i \in N \setminus \{i^*\}$

else

  Rename agents s.t. $\frac{v_1}{c_1} \geq \frac{v_2}{c_2} \geq \cdots \geq \frac{v_n}{c_n}$

  Compute $\mathbf{x}$ s.t. $v(\mathbf{x}) = \alpha\texttt{opt}(N,\mathbf{c})$ with $x_i = 1$ for $i < k$, $x_k \in (0,1]$ and $x_i = 0$ for $i > k$, $k \in N$

  for $i \leq k$ **do**

    if $c_i > v_i \frac{B}{\alpha(1+\beta)\texttt{opt}_{-i}(N,\mathbf{c})}$ **then** set $x_i = 0$

For $i \in N$ compute payments p_i according to (1)

return $(\mathbf{x}, \mathbf{p})$

Lemma 1. *Mechanism* **DA** *is truthful and individually rational.*

Proof. We start by showing that the allocation rule is monotone. For convenience, define $\rho_i(N,\mathbf{c}) = \frac{v_i}{\texttt{opt}_{-i}(N,\mathbf{c})}$, $i \in N$. Suppose that, given declared costs $\mathbf{c} = (c_i)_{i \in N}$, $\mathbf{x}$ is computed by the mechanism.

Suppose the mechanism selected agent i^* and suppose i^* decreases their cost to $c' < c_{i^*}$ and let $\mathbf{c}' = (c', \mathbf{c}_{-i^*})$. Then $\rho_i(N,\mathbf{c}')$ of agents $i \neq i^*$ can only decrease, $\rho_{i^*}(N,\mathbf{c}')$ stays the same and therefore i^* remains fully selected. Now suppose i^* increases their cost to $c' > c_{i^*}$. Then $\rho_i(N,\mathbf{c}')$ of agents $i \neq i^*$ can only increase and again $\rho_{i^*}(N,\mathbf{c}')$ stays the same. Therefore i^* remains fully selected if $\forall i \in N$: $\rho_{i^*}(N,\mathbf{c}') \geq \rho_i(N,\mathbf{c}')$, otherwise i^* is not selected. When i^* increases their cost to $c' > B$, i^* is definitely not selected.

Otherwise the mechanism computed $\mathbf{x}$ such that (initially) $v(\mathbf{x}) = \alpha\texttt{opt}(N,\mathbf{c})$. Suppose some winning agent i decreases their cost to $c' < c_i$ and let $\mathbf{c}' = (c', \mathbf{c}_{-i})$. Then $\rho_j(N,\mathbf{c}')$ of agents $j \neq i$ can only decrease and $\rho_i(N,\mathbf{c}')$ stays the same, so the allocation vector is still computed in the second part of the mechanism. The ratio $\frac{v_i}{c'}$ increases, so i can only move further to the front of the ordering. In addition, $\texttt{opt}(N,\mathbf{c}')$ can only increase and therefore i will be selected to the same extent or more. Note that the threat will not deselect agent i, as the value of the threat does not change and i decreases their cost. Now suppose i increases their cost to $c' > c_i$. Then $\rho_j(N,\mathbf{c}')$ of agents $j \neq i$ can only increase and $\rho_i(N,\mathbf{c}')$ stays the same. If the mechanism now selects agent i^*, agent i is not selected. Otherwise, the ratio $\frac{v_i}{c'}$ decreases, so i can only move further to the back of the ordering. In addition, $\texttt{opt}(N, \mathbf{c}')$ can only decrease and therefore i will be selected to the same extent or less. When i increases their cost to $c' > \min\{B, t_i\}$, i is definitely not selected. Therefore, the allocation rule is monotone.

Note that a winning agent i definitely loses when bidding $c' > B$ and as $x_i \in [0,1]$, we have $\int_0^\infty x_i(u, \mathbf{b}_{-i})du \leq 1 \cdot B < \infty$. Therefore, the mechanism is truthful and individually rational by Theorem 1. □

Lemma 2. *Mechanism* **DA** *is budget feasible.*

Proof. If, given declared costs $\mathbf{c} = (c_i)_{i \in N}$, $\mathbf{x}$ is computed by selecting agent i^* we have $\sum_{i \in N} p_i = p_{i^*} \leq 1 \cdot B$, as $x_{i^*} \in [0,1]$ and i^* definitely loses when bidding $c' > B$.

Otherwise, the mechanism computed $\mathbf{x}$ such that (initially) $v(\mathbf{x}) = \alpha \texttt{opt}(N, \mathbf{c})$. By construction we know that the threshold bid for agent i is smaller than or equal to our threat τ_i. Therefore, and because the allocation rule is monotone, the payment of agent i can be bounded by $x_i \tau_i$ and

$$\sum_{i \in N} p_i = \sum_{i=1}^{k} p_i \leq \sum_{i=1}^{k} x_i v_i \frac{B}{\alpha(1+\beta)\texttt{opt}_{-i}(N, \mathbf{c})} \leq \sum_{i=1}^{k} x_i v_i \frac{B}{\alpha \texttt{opt}(N, \mathbf{c})} = B,$$

were the last inequality follows from $\texttt{opt}(N, \mathbf{c}) \leq \texttt{opt}_{-i}(N, \mathbf{c}) + v_i < (1 + \beta)\texttt{opt}_{-i}(N, \mathbf{c})$, as the mechanism did not select agent i^*. Hence, the mechanism is budget feasible. □

In order to prove the approximation guarantee, we need the following lemma.

Lemma 3. *Let* $\mathbf{x}^*$ *be the corresponding solution of* $\texttt{opt}(A, \mathbf{c})$ *and assume agents* $i \in A$*,* $n = |A|$*, are ordered such that* $\frac{v_1}{c_1} \geq \cdots \geq \frac{v_n}{c_n}$*. Let* $k \leq n$ *be an integer such that* $\sum_{i=1}^{k-1} v_i x_i^* < \alpha \texttt{opt}(A, \mathbf{c}) \leq \sum_{i=1}^{k} v_i x_i^*$ *with* $\alpha \in (0,1]$*. It follows that* $\frac{c_k}{v_k}(1-\alpha)\texttt{opt}(A, \mathbf{c}) \leq B$*.*

Proof. For convenience let $\texttt{opt} = \texttt{opt}(A, \mathbf{c})$. Note that we can split $\mathbf{x}^*$ into $\mathbf{x} = (x_1^*, \ldots, x_{k-1}^*, x_k, 0, \ldots, 0)$ and $\mathbf{y} = (0, \ldots, 0, x_k^* - x_k, x_{k+1}^*, \ldots, x_n^*)$ with $x_k \leq x_k^*$ such that $\mathbf{x}^* = \mathbf{x} + \mathbf{y}$, $v(\mathbf{x}) = \alpha \texttt{opt}$ and $v(\mathbf{y}) = (1-\alpha)\texttt{opt}$. By feasibility of $\mathbf{x}^*$, and thus $\mathbf{y}$, and the ordering of the agents it follows that

$$B \geq c(\mathbf{y}) = \sum_{i=k}^{n} c_i y_i = \sum_{i=k}^{n} \frac{c_i}{v_i} v_i y_i \geq \frac{c_k}{v_k} \sum_{i=k}^{n} v_i y_i = \frac{c_k}{v_k}(1-\alpha)\texttt{opt}.$$

□

Lemma 3 can be interpreted in the following way. If agents k up to some $j \in A$, $j \geq k$, together contribute a fraction of $(1-\alpha)$ of the value of the optimal solution, then 'paying' these agents a cost per unit value of $\frac{c_k}{v_k}$ is budget feasible, as these agents have a greater or equal cost per unit value.

Lemma 4. *Mechanism* **DA** *has an approximation guarantee of* $\frac{\sqrt{5}+1}{\sqrt{5}-1}$ *if* $\alpha = \frac{\sqrt{5}-1}{\sqrt{5}+1}$ *and* $\beta = \frac{1}{2}\sqrt{5} - \frac{1}{2}$.

Proof. Note that $N = A$ by Lemma 1. Suppose given declared costs $\mathbf{c} = (c_i)_{i \in A}$, $\mathbf{x}$ is computed by the mechanism. If the mechanism selected agent i^*, we have

$$v_{i^*} \geq \beta \texttt{opt}_{-i^*}(N, \mathbf{c}) \geq \beta \texttt{opt}(N, \mathbf{c}) - \beta v_{i^*} \quad \Leftrightarrow \quad v(\mathbf{x}) \geq \frac{\beta}{1+\beta} \texttt{opt}(N, \mathbf{c}).$$

Otherwise the mechanism computed $\mathbf{x}$ such that $v(\mathbf{x}) = \alpha \texttt{opt}(N, \mathbf{c})$. If no agent $i \leq k$ is deselected by our threat ($c_i \leq \tau_i$), then by construction $v(\mathbf{x}) \geq \alpha \texttt{opt}(N, \mathbf{c})$. So we want

$$c_i \leq v_i \frac{B}{(1-\alpha)\texttt{opt}(N, \mathbf{c})} \leq t_i \quad \Leftrightarrow \quad \texttt{opt}(N, \mathbf{c}) \geq \underbrace{\frac{\alpha(1+\beta)}{(1-\alpha)}}_{=:R} \texttt{opt}_{-i}(N, \mathbf{c}),$$

where the upper bound on c_i follows from Lemma 3. As $\texttt{opt}(N, \mathbf{c}) \geq \texttt{opt}_{-i}(N, \mathbf{c})$, then surely the approximation guarantee holds if $R \leq 1$. Balancing the approximation ratios $\frac{\beta+1}{\beta} = \frac{1}{\alpha}$ subject to $R \leq 1$ leads to $\alpha = \frac{\sqrt{5}-1}{\sqrt{5}+1} \approx 0.38$, $\beta = \frac{1}{2}\sqrt{5} - \frac{1}{2} \approx 0.61$ and an approximation guarantee of $\frac{\sqrt{5}+1}{\sqrt{5}-1} \approx 2.62$. □

For the explanation why mechanism **DA** is computationally efficient, we refer to Appendix A. From Lemmas 1, 2 and 4, we arrive at the following theorem.

Theorem 2. *Mechanism* **DA** *with* $\alpha = \frac{\sqrt{5}-1}{\sqrt{5}+1}$ *and* $\beta = \frac{1}{2}\sqrt{5} - \frac{1}{2}$ *is truthful, individually rational, budget feasible and has an approximation guarantee of* $\frac{\sqrt{5}+1}{\sqrt{5}-1}$.

Lower Bound. Next we show that no deterministic, truthful, individually rational, budget feasible and deterministic mechanism exists with an approximation guarantee of $(\frac{5}{4} - \epsilon_1)$, for some $\epsilon_1 > 0$. For contradiction, assume such a mechanism does exist. Consider the following two instances, both with budget B and two agents with equal valuations: $\mathcal{I}_1 = (B, \mathbf{v} = \mathbf{1}, \mathbf{c} = (B, B))$ and $\mathcal{I}_2 = (B, \mathbf{v} = \mathbf{1}, \mathbf{c} = (\epsilon_2, B))$. In the first instance $\texttt{opt} = 1$, so for the approximation guarantee to hold, an allocation vector must satisfy $v(\mathbf{x}) > 4/(5 - 4\epsilon_1)$. Therefore, any such mechanism must have some agent i for which $x_i > 2/(5-4\epsilon_1)$ and assume w.l.o.g. that this is agent 1. In the second instance $\texttt{opt} = 2 - \epsilon_2/B$, so for the approximation guarantee to hold, an allocation vector must satisfy $v(\mathbf{x}) > (2 - \epsilon_2/B)(4/(5 - 4\epsilon_1))$. By the previous instance and individual rationality, it follows that agent 1 can guarantee himself a utility greater than $u = (B - \epsilon_2)2/(5 - 4\epsilon_1)$ by deviating to B. As agent 1 must be somewhat selected to achieve the approximation guarantee, $p_1 > \epsilon_2 x_1 + u$ by truthfulness. Therefore in the best case, if agent 1 is entirely selected, this leads to a budget left smaller than $B' = B - \epsilon_2 - u$. By spending this all on agent 2, this leads to an allocation vector with value smaller than $2 - \epsilon_2/B - u/B$. With elementary calculations, one can show that this value is smaller than $\alpha\texttt{opt}$ if $\epsilon_2 < (8B\epsilon_1)/(4\epsilon_1 + 1)$, resulting in a contradiction.

4 Divisible Agents in a θ-Competitive Market

It is common for lower bound proofs to use multiple instances in order to show that a mechanism cannot satisfy all properties in each instance. In our proof, and also in the lower bound proofs by Singer [1] and Gravin et al. [3], this leads to very specific instances. One instance has agents with equal efficiency while the other instance has agents for which the difference in efficiency cannot be bounded by a finite constant. Both instances are plausible in, say, a mature market were the efficiencies of the agents are close, or a premature market were the efficiencies of the agents differ a lot.

However, it is reasonable to assume that after some time the efficiencies of the agents are somewhat bounded, as agents will cease to exist if they cannot compete with the other agents in terms of efficiency. We therefore introduce a setting in which some bound on the efficiencies of the agents is known and seek a tighter approximation guarantee for this setting. We formalize this with the following definition.

Definition 1. *An instance $I = (B, (v_i)_{i\in A}, (c_i)_{i\in A})$ of the procurement auction is θ-competitive with $\theta \geq 1$ if*

$$\max_{i\in A: c_i \leq B} \frac{v_i}{c_i} \leq \theta \min_{i\in A: c_i\leq B} \frac{v_i}{c_i}. \tag{2}$$

In this setting, we will also say that the agents are *θ-competitive*. Note that if $\theta = 1$ then all agents are equally competitive. If $\theta \to \infty$ the competitiveness of the agents is unbounded, which corresponds to the original setting.

The Mechanism. Under the assumption that an instance is θ-competitive, an agent can only increase their declared costs up to some c' before becoming the agent with worst efficiency. Again, we use the payment rule of Theorem 1, and want the second term of (1) to be as small as possible. We give a mechanism that realizes this, not by directly imposing a threat, but by setting the parameters of the mechanism to specific values. This will ensure that if, in the second part of the mechanism, the declared cost of an agent exceeds some c', this agent will not be selected. Let $\alpha \in (0, 1]$ and $\beta > 0$ be some parameters which we fix later. Our mechanism is as follows:

Divisible θ-competitive Agents (DA-θ)

Let $N = \{i \in A \mid c_i \leq B\}$, $n = |N|$ and $i^* = \arg\max_{i\in N} \frac{v_i}{\mathtt{opt}_{-i}(N,\mathbf{c})}$
if $\frac{v_{i^*}}{\mathtt{opt}_{-i^*}(N,\mathbf{c})} \geq \beta$ **then** set $x_{i^*} = 1$ and $x_i = 0$ for $i \in N \setminus \{i^*\}$
else
  Rename agents such that $\frac{v_1}{c_1} \geq \frac{v_2}{c_2} \geq \cdots \geq \frac{v_n}{c_n}$
  Compute $\mathbf{x}$ s.t. $v(\mathbf{x}) = \alpha\mathtt{opt}(N, \mathbf{c})$ with $x_i = 1$ for $i < k$, $x_k \in (0, 1]$ and $x_i = 0$ for $i > k$, $k \in N$
For $i \in N$ compute payments p_i according to (1)
return $(\mathbf{x}, \mathbf{p})$

Lemma 5. *Mechanism* **DA-**θ *is truthful and individually rational.*

The proof of Lemma 5 is identical to the proof of Lemma 1 with one exception: In the second part of the mechanism, an agent i that is selected will definitely not be selected if i would have declared a cost $c' > \min\{B, \theta c_i\}$, were the reasoning of the second argument is stated in the proof of Lemma 7. In order to prove budget feasibility, we will need the following lemma which states that if agents 1 up to some $k \in A$ together contribute a fraction of α of the value of the optimal solution, then the corresponding total cost of these agents cannot exceed a fraction of α of the budget.

Lemma 6. *Let* $\mathbf{x}^*$ *be the corresponding solution of* $\texttt{opt}(A, \mathbf{c})$ *and assume agents* $i \in A$*,* $n = |A|$*, are ordered such that* $\frac{v_1}{c_1} \geq \cdots \geq \frac{v_n}{c_n}$*. Let* $v(\mathbf{x}) = \alpha v(\mathbf{x}^*)$ *with* $\alpha \in (0, 1]$*, integer* $k \leq n$*,* $x_i = x_i^*$ *for* $i < k$*,* $x_k \in (0, 1]$*,* $x_i = 0$ *for* $i > k$*. Then* $\sum_{i=1}^{k} c_i x_i \leq \alpha B$*.*

Proof. For convenience let $\texttt{opt} = \texttt{opt}(A, \mathbf{c})$. Similar to the proof of Lemma 3, split $\mathbf{x}^*$ into $\mathbf{x}$ and $\mathbf{y}$ such that $\mathbf{x}^* = \mathbf{x} + \mathbf{y}$, $v(\mathbf{x}) = \alpha\texttt{opt}$ and $v(\mathbf{y}) = (1-\alpha)\texttt{opt}$. For contradiction, suppose $c(\mathbf{x}) = \sum_{i=1}^{k} c_i x_i > \alpha B$. Then it must be that $c(\mathbf{y}) < (1-\alpha)B$, as otherwise $c(\mathbf{x}^*) = c(\mathbf{x}) + c(\mathbf{y}) > B$. We have

$$\alpha B < \sum_{i=1}^{k} c_i x_i = \sum_{i=1}^{k} \frac{c_i}{v_i} v_i x_i \leq \frac{c_k}{v_k} \sum_{i=1}^{k} v_i x_i = \frac{c_k}{v_k} \alpha\texttt{opt} \quad \Leftrightarrow \quad B < \frac{c_k}{v_k}\texttt{opt},$$

and

$$(1-\alpha)B > \sum_{i=k}^{n} c_i y_i = \sum_{i=k}^{n} \frac{c_i}{v_i} v_i y_i \geq \frac{c_k}{v_k} \sum_{i=k}^{n} v_i y_i = \frac{c_k}{v_k}(1-\alpha)\texttt{opt} \quad \Leftrightarrow \quad B > \frac{c_k}{v_k}\texttt{opt},$$

resulting in a contradiction. □

Lemma 7. *Mechanism* **DA-**θ *is budget feasible if* $\alpha \leq \min\left\{\frac{1}{\theta}, \frac{1}{1+\beta}\right\}$*.*

Proof. If, given declared costs $\mathbf{c} = (c_i)_{i \in A}$, $\mathbf{x}$ is computed by selecting agent i^*, the proof is identical to Lemma 2. Otherwise the mechanism computed $\mathbf{x}$ such that $v(\mathbf{x}) = \alpha\texttt{opt}(N, \mathbf{c})$. Suppose a winning agent i increases their cost to $c' = \theta c_i$ and let $\mathbf{c}' = (c', \mathbf{c}_{-i})$. We proof that i loses in this case if $\alpha \leq \frac{1}{1+\beta}$. As the allocation vector was initially computed in the second part of the mechanism, we have $v_i \leq \beta\texttt{opt}_{-i}(N, \mathbf{c})$ leading to

$$\texttt{opt}(N, \mathbf{c}') \leq \texttt{opt}(N, \mathbf{c}) \leq \texttt{opt}_{-i}(N, \mathbf{c}) + v_i \leq (1+\beta)\texttt{opt}_{-i}(N, \mathbf{c}) \quad \Leftrightarrow$$

$$\alpha\texttt{opt}(N, \mathbf{c}') \leq \frac{1}{1+\beta}\texttt{opt}(N, \mathbf{c}') \leq \texttt{opt}_{-i}(N, \mathbf{c}) \leq \sum_{j \in N, j \neq i} v_j.$$

The above inequality points out that the total valuation of agents $j \neq i$ is greater than or equal to $\alpha\texttt{opt}(N, \mathbf{c}')$. Therefore, as agent i will be last in the ordering when increasing their cost to c', i loses if the allocation vector is computed in the second part of the mechanism. If the allocation vector is computed by selecting i^*, agent i loses as i cannot be i^* in this case. Therefore, and by monotonicity of the allocation rule, we can bound the sum of the payments with

$$\sum_{i \in N} p_i = \sum_{i=1}^{k} p_i \leq \sum_{i=1}^{k} \theta c_i x_i \leq \theta \alpha B,$$

were the last inequality follows from Lemma 6 as it also holds for an optimal solution $\mathbf{x}^*$ of the form $x_i^* = 1$ for $i < k$, $x_k^* \in [0,1]$ and $x_i^* = 0$ for $k < i \leq n$. So if $\alpha \leq \frac{1}{\theta}$ and $\alpha \leq \frac{1}{1+\beta}$, the payments are budget feasible. □

For the explanation why the mechanism **DA-**θ is computationally efficient, we again refer to Appendix A. As the mechanism is truthful, $N = A$ and the approximation guarantee is equal to $\max\left\{\frac{\beta+1}{\beta}, \frac{1}{\alpha}\right\}$, where the second argument is by construction and the first argument follows from exactly the same reasoning as in Lemma 4. Minimizing this max-expression subject to $\alpha \leq \min\left\{\frac{1}{\theta}, \frac{1}{1+\beta}\right\}$ (budget feasibility) leads to $\alpha = \frac{1}{2}$ and $\beta = 1$. Therefore, the following theorem follows from Lemmas 5 and 7.

Theorem 3. *If $\theta \in [1,2]$, mechanism **DA-**θ with $\alpha = \frac{1}{2}$ and $\beta = 1$ is truthful, individually rational, budget feasible and has an approximation guarantee of 2.*

It can be seen in the left plot of Figure 1 that if $\theta > 2$, it is optimal to set $\alpha = \frac{1}{\theta}$ and β to any number in $\left[\frac{\alpha}{1-\alpha}, \frac{1-\alpha}{\alpha}\right]$, as then α is the limiting factor in the approximation guarantee ($\frac{1}{\alpha} \geq (1+\beta)/\beta$) and the value of α is feasible ($\alpha \leq 1/(1+\beta)$). So if it is known that agents are θ-competitive with $\theta < 2.62$, it is beneficial to use mechanism **DA-**θ instead of mechanism **DA**.

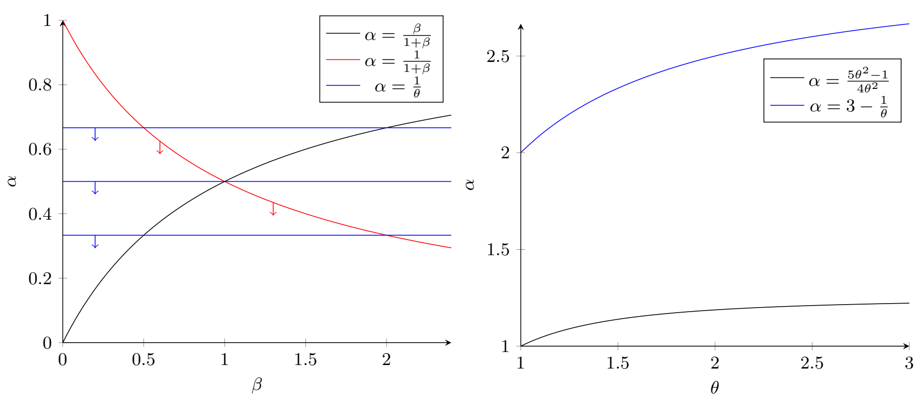

Fig. 1. Left: Plot of when the approximation guarantees are equal (black) and parameter constraints for α (red and blue). Right: Lower bound on the approximation guarantee for different values of θ for the divisible (black) and indivisible case (blue).

Lower Bound. Similar to Section 3, one can show that if $\theta \geq 2$ no deterministic, truthful, individually rational and budget feasible mechanism exists with an approximation guarantee of $(\frac{19}{16} - \epsilon)$, for some $\epsilon > 0$. Additionally, this instance can be adjusted to a specific value of θ to prove that no mechanism exists with an approximation guarantee of $((5\theta^2 - 1)/(4\theta^2) - \epsilon)$, for some $\epsilon > 0$. For details we refer to Appendix B. The general lower bound is plotted in Figure 1 and, as in Section 3, converges to $\frac{5}{4}$ if $\theta \to \infty$.

For the indivisible case, one can show that if $\theta \geq 2$ no deterministic, truthful, individually rational and budget feasible mechanism exists with an approximation guarantee of $(\frac{5}{2} - \epsilon)$ compared to the fractional knapsack optimum, for some $\epsilon > 0$. This lower bound can also be adjusted to a specific value of θ and, as in Gravin et al. [3], converges to 3 if $\theta \to \infty$. For details we also refer to Appendix B.

5 Concave Additive Valuation Function

In this section, we extend the original setting by two extra elements: namely types of agents and a concave non-decreasing valuation function for each type. In the Introduction, we described how this relates to the example in which an auctioneer wants to organize a comedy show.

An instance $I = (B, (v_i)_{i \in A}, (c_i)_{i \in A}, (t_i)_{i \in A}, (l_j)_{j \in T})$ of this type of procurement auction is defined as follows: Each agent $i \in A$ has a type $t_i \in \{1, 2, \ldots, t\}$, with $t \leq n$. If agents have the same type, they are substitutable, meaning that they are offering a similar good or service. For every type $j \in T$, the auctioneer has a valuation function $l_j : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, which maps the accumulative selected value of type j, to the actual value obtained by the auctioneer. We assume that the auctioneer becomes more saturated when selecting more value of a certain type, but selecting more value will never have a negative impact. More formally, for each type $j \in T$: l_j is concave and non-decreasing, $l_j(x) \leq x$ and $l_j(0) = 0$.

In this section, $\texttt{opt}(A, \mathbf{c})$ corresponds to the value of the optimal solution of the concave program stated below, regarding the set of agents A, costs $\mathbf{c} = (c_i)_{i \in A}$, values $\mathbf{v} = (v_i)_{i \in A}$, types $\mathbf{t} = (t_i)_{i \in A}$ and functions $(l_j)_{j \in T}$. Similarly, $\mathbf{v}$, $\mathbf{t}$ and $(l_j)_{j \in T}$ are omitted as arguments for readability and $\texttt{opt}_{-i}(A, \mathbf{c})$ regards the set of agents $A \setminus \{i\}$.

$$\begin{array}{rll} \max & \sum_{j \in T} l_j(V_j) & \\ \text{s.t.} & \sum_{i \in A} c_i x_i \leq B & \\ & \sum_{i \in A : t_i = j} v_i x_i = V_j & \forall j \in T \\ & x_i \in [0, 1] & \forall i \in A \end{array}$$

As the objective function is concave (sum of concave functions), there exists an optimal solution that we can find in polynomial time[1]. Additionally, there exists an optimal solution $\mathbf{x}^*$ to this program with the following structure. Let $S_t = \{i \in A : t_i = t\}$ be the set of agents with type $t \in T$. Then when ordering

[1] Note that the feasibility region is compact.

and renaming the agents in this set according to decreasing efficiency, there exists an integer $k \leq |S_t|$ such that the fractions selected are $x_i^* = 1$ for $i < k$, $x_k^* \in [0,1]$ and $x_i^* = 0$ for $k < i \leq |S_t|$. For agents $i \in \{1, \dots, |S_t|\}$ we define

$$v_i^* = l_t(x) - l_t(y) \quad \text{with} \quad x = \sum_{j=1}^{i} v_j x_j^* \quad \text{and} \quad y = \sum_{j=1}^{i-1} v_j x_j^*,$$

so $v_i^* = 0$ for agents $i \in \{k+1, \dots, |S_t|\}$. When we refer to an optimal solution $\mathbf{x}^*$ of $\texttt{opt}(A, \mathbf{c})$ in this section, we refer to the $\mathbf{x}^*$ with this structure.

Additionally, given a vector of declared costs $\mathbf{c} = (c_i)_{i \in A}$, we define $\hat{v}_i$ for all $i \in A$. Again let S_t be the set of agents with type $t \in T$ and reorder and rename the agents in this set according to decreasing efficiency. For agents $i \in \{1, \dots, |S_t|\}$ we define

$$\hat{v}_i = l_t(x) - l_t(y) \quad \text{with} \quad x = \sum_{j=1}^{i} v_j \quad \text{and} \quad y = \sum_{j=1}^{i-1} v_j.$$

We will give two mechanisms for this setting: one for the case when the efficiencies of the agents are unbounded and one for the case when the agents are θ-competitive.

The Mechanism. Mechanism **DA** can be altered to fit the setting of this section. We therefore consider a different ordering of the agents in the second part of the mechanism and alter the threats and $\rho_i(N, \mathbf{c})$'s to fit the valuation function. Note that in this mechanism, the threats τ_i are not independent of the declared cost of agent i, but as τ_i might decrease when agent i declares a higher cost, this imposes no problem. Again, let $\alpha \in (0,1]$ and $\beta > 0$ be some parameters which we fix later.

Divisible Agents & Concave Valuation (DA-CV)

Let $N = \{i \in A \mid c_i \leq B\}$, $n = |N|$ and $i^* = \arg\max_{i \in N} \frac{l_{t_i}(v_i)}{\texttt{opt}_{-i}(N, \mathbf{c})}$
if $\frac{l_{t_{i^*}}(v_i^*)}{\texttt{opt}_{-i^*}(N, \mathbf{c})} \geq \beta$ **then** set $x_{i^*} = 1$ and $x_i = 0$ for $i \in N \setminus \{i^*\}$
else
  Compute the optimal solution $\mathbf{x}^*$ of $\texttt{opt}(N, \mathbf{c})$, compute $(v_i^*)_{i \in N}$ and set $\mathbf{x} = \mathbf{x}^*$ and $j = n$
  Rename agents s.t. $\frac{v_1^*}{c_1} \geq \frac{v_2^*}{c_2} \geq \dots \geq \frac{v_n^*}{c_n}$
  while $v(\mathbf{x}) > \alpha v(\mathbf{x}^*)$ **do**
    Set $x_j = 0$ and $j = j - 1$
    if $v(\mathbf{x}) < \alpha v(\mathbf{x}^*)$ **then** set $x_{j+1} \in [0,1]$ s.t. $v(\mathbf{x}) = \alpha v(\mathbf{x}^*)$
  for $i \leq j + 1$ **do**
    if $c_i > \hat{v}_i \frac{B}{\alpha(1+\beta)\texttt{opt}_{-i}(N, \mathbf{c})}$ **then** set $x_i = 0$
For $i \in N$ compute payments p_i according to (1)
return $(\mathbf{x}, \mathbf{p})$

Lemma 8. *Mechanism* **DA-CV** *is truthful and individually rational.*

Proof. Proving that the mechanism is truthful and individually rational is identical to the proof of Lemma 1 with one exception, namely $\rho_i(N, \mathbf{c}) = \frac{l_{t_i}(v_i)}{\texttt{opt}_{-i}(N,\mathbf{c})}$ for $i \in N$. Note that by the way we define the construction of $\mathbf{x}^*$, if in the second part of the mechanism an agent increases (decreases) their declared bid, their fraction selected can only decrease (increase) or remain the same. Also, note that the threat will not deselect agent i if i decreases their cost, as the value of the threat then might increase or remain the same. □

Lemma 9. *Mechanism* **DA-CV** *is budget feasible.*

Proof. If, given declared costs $\mathbf{c} = (c_i)_{i \in A}$, $\mathbf{x}$ is computed by selecting agent i^*, the proof is identical to Lemma 2. Otherwise the allocation vector is computed in the second part of the mechanism. By construction we know that the threshold bid for agent i is smaller than or equal to our threat τ_i, as τ_i can only decrease or remain the same when agent i increases their bid. Therefore, and because the allocation rule is monotone, the payment of agent i can be bounded by $x_i \tau_i$ and

$$\sum_{i \in N} p_i = \sum_{i=1}^{k} p_i \leq \sum_{i=1}^{k} x_i \hat{v}_i \frac{B}{\alpha(1+\beta)\texttt{opt}_{-i}(N,\mathbf{c})} \leq \sum_{i=1}^{k} x_i \hat{v}_i \frac{B}{\alpha \texttt{opt}(N,\mathbf{c})} \leq B,$$

were the second to last inequality follows from $\texttt{opt}(N, \mathbf{c}) \leq \texttt{opt}_{-i}(N, \mathbf{c}) + l_{t_i}(v_i) < (1+\beta)\texttt{opt}_{-i}(N, \mathbf{c})$, as the mechanism did not select agent i^*. The last inequality follows as the functions l_t are concave, so $x_i \hat{v}_i \leq v_i^*$. Hence, the mechanism is budget feasible. □

Lemma 10. *Mechanism* **DA-CV** *has an approximation guarantee of* $\frac{\sqrt{5}+1}{\sqrt{5}-1}$ *if* $\alpha = \frac{\sqrt{5}-1}{\sqrt{5}+1}$ *and* $\beta = \frac{1}{2}\sqrt{5} - \frac{1}{2}$.

Proof. Note that $N = A$ by Lemma 8. Given declared costs $\mathbf{c} = (c_i)_{i \in A}$, let $\mathbf{x}$ be computed by the mechanism. If the mechanism selected agent i^*, we have

$$l_{t_i^*}(v_{i^*}) \geq \beta \texttt{opt}_{-i^*}(N, \mathbf{c}) \geq \beta \texttt{opt}(N, \mathbf{c}) - \beta l_{t_i^*}(v_{i^*}) \Leftrightarrow v(\mathbf{x}) \geq \frac{\beta}{1+\beta} \texttt{opt}(N, \mathbf{c}).$$

Otherwise, the mechanism computed $\mathbf{x}$ such that $v(\mathbf{x}) = \alpha \texttt{opt}(N, \mathbf{c})$. If no agent $i \leq k$ is deselected by our threat ($c_i \leq \tau_i$), then by construction $v(\mathbf{x}) \geq \alpha \texttt{opt}(N, \mathbf{c})$. So we want

$$c_i \leq v_i^* \frac{B}{(1-\alpha)\texttt{opt}(N,\mathbf{c})} \leq t_i \quad \Leftrightarrow \quad \texttt{opt}(N, \mathbf{c}) \geq \frac{\alpha(1+\beta)}{(1-\alpha)} \frac{v_i^*}{\hat{v}_i} \texttt{opt}_{-i}(N, \mathbf{c}),$$

where the first inequality follows from Lemma 3 by replacing v_i with v_i^*. By construction $\frac{v_i^*}{\hat{v}_i} \leq 1$, and the proof now follows from similar reasoning as in the proof of Lemma 4. □

Theorem 4. *Mechanism* **DA-CV** *with* $\alpha = \frac{\sqrt{5}-1}{\sqrt{5}+1}$ *and* $\beta = \frac{1}{2}\sqrt{5} - \frac{1}{2}$ *is truthful, individually rational, budget feasible and has an approximation guarantee of* $\frac{\sqrt{5}+1}{\sqrt{5}-1}$.

Theorem 4 follows from Lemmas 8, 9 and 10. If the θ-competitive market assumption is added to the setting where the auctioneer has a concave additive valuation function, we can construct a mechanism **Divisible θ-competitive Agents & Concave Valuation (DA-θ-CV)**, by removing lines 9 and 10 of the mechanism **DA-CV**. For the proof of Theorem 5, we refer to Appendix C as it is similar to earlier proofs.

Theorem 5. *If* $\theta \in [1, 2]$, *mechanism* **DA-θ-CV** *with* $\alpha = \frac{1}{2}$ *and* $\beta = 1$ *is truthful, individually rational, budget feasible and has an approximation guarantee of* 2.

If there is only one type of agent, i.e., there is only one function l_j, then both mechanisms are computationally efficient by a similar reasoning as in Appendix A. If, however, there are multiple types of agents, there are (rather artificial) cases where the functions l_j may intersect infinitely many times and the number of breakpoints of the payment function is not polynomially bounded. Our mechanism remains computationally efficient if the number of such breakpoints is polynomially bounded. Note that for both settings with concave additive valuation function, the lower bounds of Section 3 and 4 still hold. In both lower bound examples one can simply define that each agent has the same type t and let $l_t(x) = x$.

6 Conclusion and Future Work

In this paper, we considered budget feasible mechanisms for procurement auctions in the divisible case. Our mechanism has an approximation guarantee of 2.62, which leaves room for improvement on the upper or lower bound (1.25) of the approximation guarantee. We introduced a notion of competitiveness between agents and gave a mechanism with an approximation guarantee of 2 for this setting. In this setting the gap between the upper and lower bound (1.18) of the approximation guarantee also remains. It would be interesting to see if both these gaps can be tightened or even closed. Both our results can be extended to the setting with concave non-decreasing valuation functions.

It would be natural to continue the study of the divisible case in settings with different valuation functions. Additionally, one could study the problem with an underlying structure, such as matroids in the indivisible case, to capture a relation between goods.

Acknowledgements. We thank Georgios Amanatidis for proposing to study budget feasible mechanisms for divisible agents when he was a postdoc at CWI.

References

1. Singer, Y.: Budget feasible mechanisms. In: Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science, pp. 765–774. IEEE, (2010)
2. Chen, N., Gravin, N., Lu, P.: On the approximability of budget feasible mechanisms. In: Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms, pp 685–699. SIAM, (2011)
3. Gravin, N., Jin, Y., Lu, P., Zhang, C.: Optimal budget-feasible mechanisms for additive valuations. ACM Transactions on Economics and Computation (TEAC), **8**(4), 1-15 (2020)
4. Anari, N., Goel, G., Nikzad, A.: Mechanism design for crowdsourcing: An optimal 1-1/e competitive budget-feasible mechanism for large markets. In: IEEE 55th Annual Symposium on Foundations of Computer Science, pp. 266-275. IEEE, (2014)
5. Jalaly Khalilabadi, P., Tardos, É.: 2018. Simple and efficient budget feasible mechanisms for monotone submodular valuations. In: International Conference on Web and Internet Economics (WINE), vol. 11316, pp 246–263. Springer, (2018)
6. Goel, G., Nikzad, A., Singla, A.: Mechanism design for crowdsourcing markets with heterogeneous tasks. In: Proceedings of the 2nd AAAI Conference on Human Computation and Crowdsourcing. AAAI, (2014)
7. Leonardi, S., Monaco, G., Sankowski, P., Zhang, Q.: Budget Feasible Mechanisms on Matroids. In: Proceedings of the 19th International Conference on Integer Programming and Combinatorial Optimization (IPCO), vol. 10328, pp. 368-379. Springer, (2017)
8. Chan, H., Chen, J.: Truthful multi-unit procurements with budgets. In: International Conference on Web and Internet Economics (WINE), pp. 89–105. Springer, (2014)
9. Archer, A., Tardos, É.: Truthful mechanisms for one-parameter agents. In: Proceedings 42nd IEEE Symposium on Foundations of Computer Science, pp. 482-491. IEEE, (2001)
10. Dobzinski, S., Papadimitriou, C. H., Singer, Y.: Mechanisms for complement-free procurement. In: Proceedings of the 12th ACM conference on Electronic commerce, pp. 273-282. (2011)
11. Bei, X., Chen, N., Gravin, N., Lu, P.: Budget feasible mechanism design: from prior-free to bayesian. In: Proceedings of the 44th annual ACM symposium on Theory of computing, pp. 449-458. (2012)
12. Amanatidis, G., Birmpas, G., Markakis, E.: Coverage, matching, and beyond: new results on budgeted mechanism design. In: International Conference on Web and Internet Economics (WINE), pp. 414-428. Springer, (2016)

A Computationally efficient

We explain why the mechanism **DA** is computationally efficient. This is trivial for the majority of the mechanism, but may not be for the computation of the payment vector. Computing the payments according to (1) can also be done in polynomial time, as these payment functions are piecewise and have a polynomially bounded number of breakpoints. Note that $p_i = 0$ if $x_i = 0$.

If the mechanism selects agent i^*, the payment function of i^* only has one breakpoint, namely at B or at some $c' < B$ for which another agent becomes i^*. Otherwise agents are selected such that $v(\mathbf{x}) = \alpha\texttt{opt}$. We will consider the maximum number of breakpoints the payment function can have in this case for some agent i. Suppose the agents are ordered according to decreasing efficiency, and let $j \leq n$ be the highest index of an agent selected in $\texttt{opt}$. If initially $x_i = 1$, agent i can increase their bid until i becomes agent k and is selected fractionally. This is the first breakpoint of the payment function. If i then continues to increase their bid, i will at some point move a place up in the ordering and could then still be fractionally selected. This is another breakpoint of the payment function, as i's fraction selected decreases by a substantial amount at this point. When i increases their bid, the index j also decreases at some point. This can be another breakpoint of the payment function, as the slope might then change. When i continues to increase their cost, at some point i's place in the ordering will be equal to the index j of that moment. In this case i will not be selected as the fractional agent as then we would have $v_i > (1-\alpha)\texttt{opt} \geq \beta\texttt{opt}_{-i}$, by the values of α and β. The number of breakpoints of the payment function of agent i is therefore bounded by $j - i \leq n - 1$. Note that the number of breakpoints can be smaller than $j - i$ by the threat we impose or by the mechanism selecting i^*. Additionally, the sub-functions of the payment function are well-defined and can be found efficiently.

Note that for the mechanism **DA-**θ, the number of breakpoints of the payment function of an agent i is also bounded by $n - 1$. If agent $i = 1$ would move all the way to the back of the ordering, i will not be selected as then $v_i > (1-\alpha)\texttt{opt} \geq \beta\texttt{opt}_{-i}$, by the values of α and β.

B Lower Bounds Section 4

Divisible Case: Next we show that if $\theta \geq 2$ no deterministic, truthful, individually rational and budget feasible mechanism exists with an approximation guarantee of $(\frac{19}{16} - \epsilon)$, for some $\epsilon > 0$. For contradiction, assume such a mechanism does exist. Consider the following two instances, both with budget B and two agents with equal valuations:

$$\mathcal{I}_1 = \Big(B, \mathbf{v} = (1,1), \mathbf{c} = \Big(\frac{2}{3}B, \frac{2}{3}B\Big)\Big) \text{ and } \mathcal{I}_2 = \Big(B, \mathbf{v} = (1,1), \mathbf{c} = \Big(\frac{1}{3}B, \frac{2}{3}B\Big)\Big).$$

Note that both instances satisfy (2) for $\theta = 2$. In the first instance $\texttt{opt} = 1.5$, so for the approximation guarantee to hold, an allocation vector must satisfy

$v(\mathbf{x}) > \frac{24}{19}$. Therefore, any such mechanism must have some agent i for which $x_i > \frac{12}{19}$ and assume w.l.o.g. that this is agent 1. In the second instance $\texttt{opt} = 2$, so for the approximation guarantee to hold, an allocation vector must now satisfy $v(\mathbf{x}) > \frac{32}{19}$. By the previous instance and individual rationality, it follows that agent 1 can guarantee itself a utility of at least $\frac{4}{19}B$ by deviating to $\frac{2}{3}B$. As agent 1 must be somewhat selected to achieve the approximation guarantee, we have $p_1 > \frac{1}{3}Bx_1 + \frac{4}{19}B$ by truthfulness. Therefore in the best case, if agent 1 is entirely selected, this leads to a budget left smaller than $\frac{26}{57}B$. By spending this on agents 2, the value that can be acquired is smaller than $\frac{39}{57}$, leading to an allocation vector with value smaller than $\frac{96}{57}$ and resulting in a contradiction.

Note that this instance can be adjusted to a specific value of θ as follows:

$$\mathcal{I}_1 = \left(B, \mathbf{v}, \mathbf{c} = \left(\frac{\theta}{\theta+1}B, \frac{\theta}{\theta+1}B\right)\right) \text{ and } \mathcal{I}_2 = \left(B, \mathbf{v}, \mathbf{c} = \left(\frac{1}{\theta+1}B, \frac{\theta}{\theta+1}B\right)\right),$$

to prove that no mechanism exists with an approximation guarantee of $((5\theta^2 - 1)/(4\theta^2) - \epsilon)$, for some $\epsilon > 0$.

Indivisible Case: Using two almost similar instances as in Gravin et al. [3], we show that if $\theta \geq 2$ in the indivisible case no deterministic, truthful, individually rational and budget feasible mechanism exists with an approximation guarantee of $(\frac{5}{2} - \epsilon)$ compared to the fractional knapsack optimum, for some $\epsilon > 0$. For contradiction, assume such a mechanism does exist. Consider the following two instances, both with budget B and 3 agents with equal valuations:

$$\mathcal{I}_1 = (B, \mathbf{v} = \mathbf{1}, \mathbf{c} = (c^*, c^*, c^*)) \quad \text{and} \quad \mathcal{I}_2 = (B, \mathbf{v} = \mathbf{1}, \mathbf{c} = (c^*/2, c^*, c^*)),$$

with $c^* = B/(2 - \frac{\epsilon}{2})$. Note that both instances satisfy (2) for $\theta = 2$. In the first instance $\texttt{opt} = 2 - \frac{\epsilon}{2}$. A mechanism can only satisfy all the properties if one agent is (fully) selected and this agent is paid at least c^*. Assume w.l.o.g. that this is agent 1. In the second instance $\texttt{opt} = 2.5 - \frac{\epsilon}{2}$, so two agents must be (fully) selected to achieve the approximation guarantee. By truthfulness, the payment of agent 1 must be at least c^*, as otherwise agent 1 could have been better off declaring a cost of c^*. As the payment to the other agent (2 or 3) that is (fully) selected must also be at least c^* by individual rationality, this results in a contradiction.

Note that this instance can be adjusted to a specific value of θ by setting $c_1 = c^*/\theta$ in the second instance to prove that no deterministic, truthful, individually rational and budget feasible mechanism exists with an approximation guarantee of $(3 - \frac{1}{\theta} - \epsilon)$, for some $\epsilon > 0$.

C Proof of Theorem 5

Proof. Proving that the mechanism is truthful and individually rational is identical to the proof in Lemma 8, but again with one exception. In the second part of the mechanism, an agent i that is selected will definitely not be selected if i would have declared a cost $c' > \min\{B, \theta c_i\}$.

By using Lemma 6 where we replace v_i by v_i^*, the mechanism is budget feasible if $\alpha \leq \min\left\{\frac{1}{\theta}, \frac{1}{1+\beta}\right\}$ by the proof of Lemma 7 and modifying the second part of the proof accordingly: Otherwise the mechanism computed $\mathbf{x}$ such that $v(\mathbf{x}) = \alpha v(\mathbf{x}^*) \ldots$ As the allocation vector was initially computed in the second part of the mechanism, we have $l_{t_i}(v_i) \leq \beta \texttt{opt}_{-i}(N, \mathbf{c})$ leading to

$$\texttt{opt}(N, \mathbf{c}') \leq \texttt{opt}(N, \mathbf{c}) \leq \texttt{opt}_{-i}(N, \mathbf{c}) + l_{t_i}(v_i) \leq (1+\beta)\texttt{opt}_{-i}(N, \mathbf{c}) \quad \Leftrightarrow$$

$$\alpha \texttt{opt}(N, \mathbf{c}') \leq \frac{1}{1+\beta}\texttt{opt}(N, \mathbf{c}') \leq \texttt{opt}_{-i}(N, \mathbf{c}) \leq \sum_{j \in N, j \neq i} \hat{v}_i.$$

As the mechanism is truthful, $N = A$ and the approximation guarantee is equal to $\max\left\{\frac{\beta+1}{\beta}, \frac{1}{\alpha}\right\}$, were the second argument is by construction and the first argument follows from exactly the same reasoning as in Lemma 10. Again, minimizing this maximizing subject to $\alpha \leq \min\left\{\frac{1}{\theta}, \frac{1}{1+\beta}\right\}$ (budget feasibility) leads to $\alpha = \frac{1}{2}$ and $\beta = 1$ and an approximation guarantee of 2. □