



UvA-DARE (Digital Academic Repository)

Partial Allocations in Budget-Feasible Mechanism Design: Bridging Multiple Levels of Service and Divisible Agents

Amanatidis, G.; Klumper, S.; Markakis, E.; Schäfer, G.; Tsikiridis, A.

DOI

[10.1007/978-3-031-48974-7_3](https://doi.org/10.1007/978-3-031-48974-7_3)

Publication date

2024

Document Version

Final published version

Published in

Web and Internet Economics

License

Article 25fa Dutch Copyright Act (<https://www.openaccess.nl/en/in-the-netherlands/you-share-we-take-care>)

[Link to publication](#)

Citation for published version (APA):

Amanatidis, G., Klumper, S., Markakis, E., Schäfer, G., & Tsikiridis, A. (2024). Partial Allocations in Budget-Feasible Mechanism Design: Bridging Multiple Levels of Service and Divisible Agents. In J. Garg, M. Klimm, & Y. Kong (Eds.), *Web and Internet Economics: 19th International Conference, WINE 2023, Shanghai, China, December 4–8, 2023 : proceedings* (pp. 41-58). (Lecture Notes in Computer Science; Vol. 14413), (Advanced Research in Computing and Software Science). Springer. https://doi.org/10.1007/978-3-031-48974-7_3

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)



Partial Allocations in Budget-Feasible Mechanism Design: Bridging Multiple Levels of Service and Divisible Agents

Georgios Amanatidis^{1,4}, Sophie Klumper^{2,5}, Evangelos Markakis^{3,4,6},
Guido Schäfer^{2,7}, and Artem Tsikiris²✉

¹ University of Essex, Colchester, UK

² Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
artem.tsikiris@cwi.nl

³ Athens University of Economics and Business, Athens, Greece

⁴ Archimedes/Athena RC, Athens, Greece

⁵ Vrije Universiteit, Amsterdam, The Netherlands

⁶ Input Output Global (IOG), Athens, Greece

⁷ University of Amsterdam, Amsterdam, The Netherlands

Abstract. Budget-feasible procurement has been a major paradigm in mechanism design since its introduction by Singer [24]. An auctioneer (buyer) with a strict budget constraint is interested in buying goods or services from a group of strategic agents (sellers). In many scenarios it makes sense to allow the auctioneer to only partially buy what an agent offers, e.g., an agent might have multiple copies of an item to sell, they might offer multiple levels of a service, or they may be available to perform a task for any fraction of a specified time interval. Nevertheless, the focus of the related literature has been on settings where each agent's services are either fully acquired or not at all. A reason for this is that in settings with partial allocations, without any assumptions on the costs, there are strong inapproximability results (see, e.g., Chan and Chen [10], Anari et al. [5]). Under the mild assumption of being able to afford each agent entirely, we are able to circumvent such results. We design a polynomial-time, deterministic, truthful, budget-feasible, $(2 + \sqrt{3})$ -approximation mechanism for the setting where each agent offers multiple levels of service and the auctioneer has a discrete separable concave valuation function. We then use this result to design a deterministic, truthful and budget-feasible mechanism for the setting where any fraction of a service can be acquired and the auctioneer's valuation function is separable concave (i.e., the sum of concave functions). The approximation ratio of this mechanism depends on how "nice" the concave functions are, and is $O(1)$ for valuation functions that are sums of $O(1)$ -regular functions (e.g., functions like $\log(1+x)$). For the special case of a linear valuation function, we improve the best known approximation ratio from $1 + \phi$ (by Klumper and Schäfer [17]) to 2. This establishes a separation result between this setting and its indivisible counterpart.

Keywords: Procurement Auctions · Budget-Feasible Mechanism Design · Multiple Levels of Service

1 Introduction

Consider a procurement auction where the agents have *private* costs on the services that they offer, and the auctioneer associates a value for each possible set of selected agents. This forms a single parameter auction environment, where the agents may strategically misreport their cost to their advantage: obtaining higher payments. Imagine now that the auctioneer additionally has a strict budget constraint that they cannot violate. Under these considerations, a natural goal for the auctioneer is to come up with a truthful mechanism for hiring a subset of the agents, that maximizes her procured value, such that the total payments to the agents respect the budget limitation. This is precisely the model that was originally proposed by Singer [24] for *indivisible* agents, i.e., with a binary decision to be made for each agent (to be hired or not). Given also that even the non-strategic version of such budget-constrained problems tend to be NP-hard, the main focus is on providing budget-feasible mechanisms that achieve approximation guarantees on the auctioneer’s optimal potential value.

Ever since the work of Singer [24], a large body of works has emerged, devoted to obtaining improved results on the original model, as well as proposing a number of extensions. These extensions include, among others, additional feasibility constraints, richer objectives, more general valuation functions and additional assumptions, such as Bayesian modeling. Undoubtedly, all these results have significantly enhanced our understanding for the indivisible scenario. In this paper, we move away from the case of indivisible agents and concentrate on two settings that have received much less attention in the literature. In both of the models that we study, instead of hiring agents entirely or not at all, the auctioneer has more flexibility and is allowed to partially procure the services offered by each agent. We assume that the auctioneer’s valuation function is the sum of individual valuation functions, each associated with a particular agent.

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 [10] 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. Chan and Chen [10] proposed randomized, truthful, and budget-feasible mechanisms for this setting, with approximation guarantees that depend on the number of agents. 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 [24], when it is assumed that the auctioneer can afford to hire any individual agent entirely. Anari et al. [5] were the first to study the divisible setting. In their work they employed a *large market* assumption, which, in the context of budget-feasible mechanism design, roughly means that the cost of each agent for their entire service is insignificant, compared to the budget of the auctioneer. Additionally, they notice that in the divisible setting, no truthful mechanism with a finite approximation guarantee exists without any restriction on the costs. Very recently, Klumper and Schäfer [17] revisited this problem without the large market assumption but under the much milder assumption that the auctioneer can afford to hire any individual agent entirely (which is standard in the literature for the indivisible setting, but here it does restrict the bidding space). They present a deterministic, truthful and budget-feasible mechanism that achieves an approximation ratio of $1 + \phi \approx 2.62$ for linear valuation functions and extend it to the setting in which all agents are associated with the same concave valuation function.

The two discussed settings of procurement auctions have a number of practical applications in various domains. As previously mentioned, the divisible setting would for example be useful to model the time availability of a worker in the context of crowdsourcing. Moreover, these types of auctions can also be applied to other industries, such as transportation and logistics, where the delivery of goods and services can be broken down into multiple levels of service. For instance, in the transportation industry, the first level of service can represent the basic delivery service, while the higher levels can represent more premium and specialized services, such as express delivery or temperature-controlled shipping. The auctioneer can then choose to hire each agent up to an available level of service, not necessarily the best offered, based on the budget constraint and the value of the services provided.

Our Contributions. In this work, we propose deterministic, truthful and budget-feasible mechanisms for settings with partial allocations. Specifically:

- We present a mechanism, SORT-&-REJECT(k) (Mechanism 1), with an approximation ratio of $2 + \sqrt{3} \approx 3.73$ for the indivisible agent setting with multiple levels of service and concave valuation functions (Sect. 3, Theorem 2). The main idea behind our novel mechanism is to apply a backwards greedy approach, in which we start from an optimal fractional solution and we discard single levels of service one by one, until a carefully chosen stopping condition is met. For this setting, no constant-factor approximation mechanism was previously known.
- We use SORT-&-REJECT(k) as a subroutine in order to design the mechanism for the setting with divisible agents, CHUNK-&-SOLVE (Mechanism 2), that

achieves an approximation ratio of $L(1 + \phi + o(k^{-1}))$ for L -regular concave valuation functions (Sect. 4.1, Theorem 3), where k is a discretization parameter. This is the first result for the problem that is independent of the number of agents n . Note that L -regularity is a Lipschitz-like condition and for $L = 1$ the problem reduces to the setting with linear valuation functions. In this case, our ratio retrieves the best known guarantee of Klumper and Schäfer [17] as k grows. On a technical level, we exploit the correspondence between the discrete and the continuous settings; as the number of service levels grows large, the former converges to the latter.

- We improve on the aforementioned best known result for $L = 1$, by suggesting a 2-approximation mechanism, PRUNE-&-ASSIGN (Mechanism 4), for the divisible setting with linear valuation functions (Sect. 4.2, Theorem 4). This mechanism is inspired by the randomized 2-approximate mechanism proposed by Gravin et al. [14] for the indivisible setting.

As we mentioned above, all our results are under the mild assumption that we can afford each agent entirely. For the setting with divisible agents this is necessary in order to achieve any non-trivial factor [5], and it was also assumed by Klumper and Schäfer [17]. Even for the discrete setting with multiple levels of service this assumption circumvents a strong lower bound of Chan and Chen [10] (see also Remark 1). In both settings our assumptions are much weaker than the large market assumptions often made in the literature (see, e.g., [5, 16]).

Further Related Work. The design of truthful budget-feasible mechanisms for indivisible agents was introduced by Singer [24], who gave a deterministic mechanism for additive valuation functions with an approximation guarantee of 5, along with a lower bound of 2 for deterministic mechanisms. This guarantee was subsequently improved to $2 + \sqrt{2} \approx 3.41$ by Chen et al. [11], who also provided a lower bound of 2 for randomized mechanisms and a lower bound of $1 + \sqrt{2} \approx 2.41$ for deterministic mechanisms. Gravin et al. [14] gave a 3-approximate deterministic mechanism, which is the best known guarantee for deterministic mechanisms to this day, along with a lower bound of 3 when the guarantee is with respect to the optimal non-strategic fractional solution. Regarding randomized mechanisms, Gravin et al. [14] settled the question by providing a 2-approximate randomized mechanism, matching the lower bound of Chen et al. [11]. Finally, the question has also been settled under the large market assumption by Anari et al. [5], who extended their $\frac{e}{e-1} \approx 1.58$ mechanism for the setting with divisible agents to the indivisible setting. As mentioned earlier, Klumper and Schäfer [17] study the divisible setting without the large market assumption, but under the assumption that the private cost of each agent is bounded by the budget and give, among other results, a deterministic $(1 + \phi)$ -approximate mechanism for linear valuation functions, i.e., non-identical valuations.

For indivisible agents, the problem has also been extended to richer valuation functions. This line of inquiry also started by Singer [24], who gave a randomized algorithm with an approximation guarantee of 112 for a monotone submodular objective. Once again, this result was improved by Chen et al. [11]

to a 7.91 guarantee, and the same authors devised a deterministic mechanism with a 8.34 approximation. Subsequently, the bound for randomized mechanisms was improved by Jalaly and Tardos [16] to 5. Very recently, Balkanski et al. [8] proposed a new method of designing mechanisms that goes beyond the sealed-bid auction paradigm. Instead, Balkanski et al. [8] presented mechanisms in the form of deterministic clock auctions and, for the monotone submodular case, present a deterministic clock auction which achieves a 4.75 guarantee.

Beyond monotone submodular valuations, it becomes much harder to obtain truthful mechanisms with small constants as approximation guarantees. Namely, for non-monotone submodular objectives, the first randomized mechanism that runs in polynomial time is due to Amanatidis et al. [3] and its approximation guarantee is 505. This guarantee was improved to 64 by Balkanski et al. [8] who provided a deterministic mechanism for the problem and Huang et al. [15] who gave a further improvement of $(3 + \sqrt{5})^2$ for randomized mechanisms. In both [8] and [15] the mechanisms take the form of clock auctions, procedures in which bidders are offered prices in multiple rounds, see also [21].

Richer valuations that have been studied are XOS valuation functions (see Bei et al. [9], Amanatidis et al. [2]) and subadditive valuation functions (see Dobzinski et al. [12], Bei et al. [9], Balkanski et al. [8]). For subadditive valuation functions, no mechanism achieving a constant approximation is known. However, Bei et al. [9] have proved that such a mechanism should exist, using a non-constructive argument. Finding such a mechanism is an intriguing open question.

Other settings that have been studied include environments with underlying feasibility constraints, such as downward-closed environments (Amanatidis et al. [1], Huang et al. [15]) and matroid constraints (Leonardi et al. [18]). Other environments in which the auctioneer wants to get a set of heterogeneous tasks done and each task requires that the hired agent has a certain skill, have been studied as well, see Goel et al. [13], Jalaly and Tardos [16]. Recently, Li et al. [19] studied facility location problems under the lens of budget-feasibility, in which facilities have private facility-opening costs. Finally, the problem has been studied in a beyond worst-case analysis setting by Rubinstein and Zhao [23].

2 Model and Preliminaries

We first define the standard budget-feasible mechanism design model below which constitutes the basis of the more general models considered in this paper. The multiple service level model is introduced in Sect. 2.2 and the divisible agent model in Sect. 2.3.

2.1 Basic Model

We consider a procurement auction consisting of a set of agents $N = \{1, \dots, n\}$ and an auctioneer who has an available budget $B \in \mathbb{R}_{>0}$. Each agent $i \in N$ offers a service and has a private cost parameter $c_i \in \mathbb{R}_{>0}$, representing their

true cost for providing this service. The auctioneer derives some value $v_i \in \mathbb{R}_{\geq 0}$ from the service of agent i which is assumed to be public information.

A deterministic mechanism \mathcal{M} in this setting consists of an allocation rule $\mathbf{x} : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}^n$ and a payment rule $\mathbf{p} : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}^n$. To begin with, the auctioneer collects a profile $\mathbf{b} = (b_i)_{i \in N} \in \mathbb{R}_{\geq 0}^n$ of declared costs from the agents. Here, b_i denotes the cost declared by agent $i \in N$, which may differ from their true cost c_i . Given the declarations, the auctioneer determines an allocation (hiring scheme) $\mathbf{x}(\mathbf{b}) = (x_1(\mathbf{b}), \dots, x_n(\mathbf{b}))$, where $x_i(\mathbf{b}) \in \mathbb{R}_{\geq 0}$ is the allocation decision for agent i , i.e., to what extent agent i is hired. Generally, we distinguish between the *divisible* and *indivisible* agent setting by means of the corresponding allocation rule. In the divisible setting, each agent i can be allocated fractionally, i.e., $x_i(\mathbf{b}) \in \mathbb{R}_{\geq 0}$. In the indivisible setting, each agent i can only be allocated integrally, i.e., $x_i(\mathbf{b}) \in \mathbb{N}_{\geq 0}$. Given an allocation \mathbf{x} , we define $W(\mathbf{x}) = \{i \in N \mid x_i > 0\}$ as the set of agents who are positively allocated under \mathbf{x} . The auctioneer also determines a vector of payments $\mathbf{p}(\mathbf{b}) = (p_1(\mathbf{b}), \dots, p_n(\mathbf{b}))$, where $p_i(\mathbf{b})$ is the payment agent i will receive for their service.

We assume that agents have quasi-linear utilities, i.e., for a deterministic mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$, the utility of agent $i \in N$ for a profile \mathbf{b} is $u_i(\mathbf{b}) = p_i(\mathbf{b}) - c_i \cdot x_i(\mathbf{b})$. We are interested in mechanisms that satisfy three properties for any true profile \mathbf{c} and any declared profile \mathbf{b} :

- *Individual rationality*: Each agent $i \in N$ receives non-negative utility, i.e., $u_i(\mathbf{b}) \geq 0$.
- *Budget-feasibility*: The sum of all payments made by the auctioneer does not exceed the budget, i.e., $\sum_{i \in N} p_i(\mathbf{b}) \leq B$.
- *Truthfulness*: Each agent $i \in N$ does not have an incentive to misreport their true cost, regardless of the declarations of the other agents, i.e., $u_i(c_i, \mathbf{b}_{-i}) \geq u_i(\mathbf{b})$ for any b_i and \mathbf{b}_{-i} .

Given an allocation \mathbf{x} , the total value that the auctioneer obtains is denoted by $v(\mathbf{x})$. The exact form of this function depends on the respective model we are studying and will be defined in the subsections below.

All the models that are studied in this paper are single-parameter settings and so the characterization of Myerson [22] applies.¹ It is therefore sufficient to focus on the class of mechanisms with *monotone non-increasing* (called *monotone* for short) allocation rules. An allocation rule is monotone non-increasing if for every agent $i \in N$, every profile \mathbf{b} , and all $b'_i \leq b_i$, it holds that $x_i(\mathbf{b}) \leq x_i(b'_i, \mathbf{b}_{-i})$. We will use this together with Theorem 1 below to design truthful mechanisms.

Theorem 1 ([7, 22]). *A monotone non-increasing allocation rule $\mathbf{x}(\mathbf{b})$ admits a payment rule that is truthful and individually rational if and only if for all agents $i \in N$ and all bid profiles \mathbf{b}_{-i} , we have $\int_0^\infty x_i(z, \mathbf{b}_{-i}) dz < \infty$. In this case, we can take the payment rule $p(\mathbf{b})$ to be*

$$p_i(\mathbf{b}) = b_i x_i(\mathbf{b}) + \int_{b_i}^\infty x_i(z, \mathbf{b}_{-i}) dz. \quad (1)$$

¹ We refer the reader to [6] for a rigorous treatment of the uniqueness property of Myerson's characterization result.

In this paper, we will exclusively derive monotone allocation rules that are implemented with the payment rule as defined in (1). Therefore, in the remainder of this paper, we adopt the convention of referring to the true cost profile \mathbf{c} of the agents as input (rather than distinguishing it from the declared cost profile \mathbf{b}). Also, throughout the paper we will omit the explicit reference to the respective cost profile \mathbf{c} whenever it is clear from the context.

2.2 k -Level Model

We consider the following multiple service level model as a natural extension of the standard model introduced above (see also [10]). Throughout the paper, 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.² Here, $v_i(j)$ denotes the auctioneer's valuation of the first j levels of service of agent i for the auctioneer. Observe that in this setting each agent $i \in N$ is indivisible and the range of the allocation rule is constrained to $\{0, \dots, k\}$, i.e., $x_i : \mathbb{R}_{\geq 0}^n \rightarrow \{0, \dots, k\}$. Note also that the total cost of agent i is linear (as defined above), 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 $v(\mathbf{x}(\mathbf{c})) = \sum_{i \in N} v_i(x_i(\mathbf{c}))$.

Cost Restrictions: We consider different assumptions with respect to the ability of the auctioneer to hire multiple service levels. In the *all-in setting*, we assume that the auctioneer can afford to hire all levels of each single agent, i.e., given a cost profile \mathbf{c} , for every agent $i \in N$ it holds that $k \cdot c_i \leq B$. In contrast, in the *best-in setting*, which is equivalent to the setting of Chan and Chen [10], the auctioneer is guaranteed only to be able to afford the first service level, i.e., given a cost profile \mathbf{c} , for every agent $i \in N$ it holds that $c_i \leq B$.³ We focus on the all-in setting throughout this extended abstract; see the full version of our work [4] for an almost tight result on the best-in setting.

Remark 1. For the best-in setting, Chan and Chen [10] show a lower bound of k for the approximation guarantee of any deterministic, truthful, budget-feasible mechanism and a lower bound of $\ln(k)$ for the approximation guarantee of any

² Our results very easily extend to the setting where there is a different (public) k_i associated with each agent i . We use a common k for the sake of presentation.

³ Whenever we use one of these assumptions, we implicitly constrain the space of the (declared) cost profiles. That is, we assume that any agent who violates the respective condition is discarded up front from further considerations, e.g., by running a pre-processing step that removes such agents.

randomized, universally truthful, budget-feasible mechanism. For these bounds, a single agent is used and then it is claimed that they generalize to nk and $\ln(nk)$, respectively, for n agents. The former is not correct, as we show in the full version of this paper [4], where we present a $(k + 2 + o(1))$ -approximate mechanism, named GREEDY-BEST-IN(k), almost settling the deterministic case. Note that the mechanism suggested by Chan and Chen [10] is $4(1 + \ln(nk))$ -approximate.

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*⁴ (see, e.g., [20] 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 a mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ 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 relaxation of the above problem as a proxy for $\text{OPT}_I^k(\mathbf{c}, B)$. The definition and further details about this are deferred to Sect. 2.4 below.

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}).

2.3 Divisible Agent Model

Next, we introduce the fractional model that we study in this paper. Throughout the paper, we refer to this model as the *divisible agent model*: Here, the auctioneer is allowed to hire each agent for an arbitrary fraction of the full service. More precisely, each agent $i \in N$ is divisible and the range of the allocation rule is constrained to $[0, 1]$, i.e., $x_i : \mathbb{R}_{\geq 0}^n \rightarrow [0, 1]$. Each agent $i \in N$ has an associated valuation function $\bar{v}_i : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ (which is public information), where $\bar{v}_i(x)$ represents how valuable a fraction $x \in [0, 1]$ of the service of agent i is to the auctioneer.

Valuation Functions: Also here, we assume without loss of generality that each v_i is normalized such that $\bar{v}_i(0) = 0$. We focus on the general class of non-decreasing and concave valuation functions. The total value that the auctioneer derives from an allocation $\mathbf{x}(\mathbf{c})$ is defined as $v(\mathbf{x}(\mathbf{c})) = \sum_{i \in N} \bar{v}_i(x_i(\mathbf{c}))$.

L-Regularity Condition: We introduce the following regularity condition for the valuation functions which will be crucial in our analysis of the divisible agent model below. Given a function $f : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$, we say that f is *L-regular Lipschitz* (or just *L-regular* for short) for $L \geq 1$ if

$$f(x) \leq xL f(1) \quad \forall x \in [0, 1]. \quad (3)$$

⁴ Note that for $k = 1$, the problem reduces to the well-known *0-1 Knapsack Problem*.

Note that if we remove $f(1)$ from the above definition, then this definition coincides with the standard Lipschitz definition. We say that an instance of the divisible agent model is L -regular for some $L \geq 1$, if for each agent $i \in N$ the valuation function \bar{v}_i is L -regular as defined in (3).

Cost Restrictions: We assume that the auctioneer can afford each agent to the full extent. More formally, given a cost profile \mathbf{c} it must hold that for each agent $i \in N$, $c_i \leq B$. With respect to this assumption, the same remarks as given above apply (see Footnote 3).

Benchmark: As above, the performance of a mechanism is measured by comparing $v(\mathbf{x}(\mathbf{c}))$ with the underlying (non-strategic) combinatorial optimization problem, which we refer to as the *Fractional Concave Knapsack Problem*:

$$\text{OPT}_F(\mathbf{c}, B) := \max \sum_{i=1}^n \bar{v}_i(x_i) \quad \text{s.t.} \quad \sum_{i=1}^n c_i x_i \leq B, \quad x_i \in [0, 1] \quad \forall i \in N. \quad (4)$$

In the divisible agent model, a mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ is α -approximate with $\alpha \geq 1$ if $v(\mathbf{x}(\mathbf{c})) \geq \frac{1}{\alpha} \text{OPT}_F(\mathbf{c})$.

An instance I of the divisible agent model will be denoted by a tuple $I = (N, \mathbf{c}, B, (\bar{v}_i)_{i \in N})$. As mentioned before, we will omit the explicit reference of certain input parameters if they are clear from the context.

2.4 Fractional k -Bounded Knapsack Problem

We also consider the *Fractional k -Bounded Knapsack Problem* that follows from the k -Bounded Knapsack Problem in (2) by relaxing the integrality constraint:

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

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.

We state the following as a fact and refer the reader to [4] for details. More generally, all the proofs that are omitted here can be found in [4].

Fact 1. *Given an instance $I = (N, \mathbf{c}, B, k, (v_i)_{i \in N})$ of the Fractional k -Bounded Knapsack Problem, a simple greedy algorithm computes in time $O(kn \log(kn))$ an optimal solution \mathbf{x}^* that has at most one coordinate with a non-integral value.*

The next fact relates the values of instances which only differ with respect to their budget.

Fact 2. Let $I = (N, \mathbf{c}, B, k, (v_i)_{i \in N})$ and $I' = (N, \mathbf{c}, B', k, (v_i)_{i \in N})$ with $B < B'$ be two instances of the Fractional k -Bounded Knapsack Problem. Then, $\text{OPT}_F^k(\mathbf{c}, B)/B \geq \text{OPT}_F^k(\mathbf{c}, B')/B'$.

In most cases the budget is going to be clear from the context, so usually we are going to omit B from $\text{OPT}_F^k(\mathbf{c}, B)$ and $\text{OPT}_I^k(\mathbf{c}, B)$, and simply write $\text{OPT}_F^k(\mathbf{c})$ and $\text{OPT}_I^k(\mathbf{c})$, respectively.

3 Budget-Feasible Mechanism for Multiple Service Levels

We derive a natural truthful and budget-feasible greedy mechanism for the k -level model. This mechanism will also be used in our CHUNK-&-SOLVE mechanism for the divisible agent model (see Sect. 4.1).

3.1 A Truthful Greedy Mechanism

The main idea underlying our mechanism 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 service levels of 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 (Fact 1). We then repeatedly discard the worst service level (in terms of marginal value-per-cost) of an agent from this allocation, 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.

The main result of this section is the following theorem:

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

The polynomial running time for the allocation is straightforward. In the remainder of this section, we prove several lemmata to establish the properties stated in Theorem 2. Technically, the most challenging part is to prove that the mechanism is budget-feasible (see Sect. 3.2).

We first show that the allocation rule of SORT-&-REJECT(k) is monotone.

Lemma 1. The allocation rule of SORT-&-REJECT(k) is monotone.

Mechanism 1: SORT-&-REJECT(k)

\triangleright **Input:** A profile \mathbf{c} and a parameter $\alpha \in (0, 1)$;

- 1 Set $i^* = \arg \max_{i \in N} v_i(k) / \text{OPT}_F^k(\mathbf{c}_{-i})$
- 2 **if** $v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \cdot \text{OPT}_F^k(\mathbf{c}_{-i^*})$ **then**
- 3 \lfloor set $x_{i^*} = k$ and $x_i = 0$ for all $i \in N \setminus \{i^*\}$
- 4 **else**
- 5 Compute an optimal allocation $\mathbf{x}^*(\mathbf{c})$ of $\text{OPT}_F^k(\mathbf{c})$. ;
- 6 Initialize $\mathbf{x} = (\lfloor x_1^*(\mathbf{c}) \rfloor, \dots, \lfloor x_n^*(\mathbf{c}) \rfloor)$. ;
- 7 **for** $i \in N$ and $j = 1, \dots, x_i$ **do**
- 8 \lfloor add the marginal value-per-cost ratio $m_i(j)/c_i$ to a list \mathcal{L} .
- 9 Sort \mathcal{L} in non-increasing order and let ℓ be the index of the last agent of $W(\mathbf{x})$ in \mathcal{L} .
- 10 **while** $v(\mathbf{x}) - m_\ell(x_\ell) \geq \alpha \text{OPT}_F^k(\mathbf{c})$ **do**
- 11 Set $x_\ell = x_\ell - 1$. ;
- 12 \lfloor Remove the last element from \mathcal{L} and update ℓ . ;
- 13 Allocate \mathbf{x} and determine the payments \mathbf{p} according to (1).

Since the payments are computed according to (1), we conclude that the mechanism is truthful and individually rational. We continue by showing that SORT-&-REJECT(k) achieves the claimed approximation guarantee.

Lemma 2. *Let $\mathbf{x}(\mathbf{c})$ be the allocation computed by SORT-&-REJECT(k) for a cost profile \mathbf{c} . It holds that $v(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}_F^k(\mathbf{c})$.*

3.2 Making Sort-&-Reject(k) Budget-Feasible

It remains to prove that SORT-&-REJECT(k) is budget-feasible. We introduce some auxiliary notation: Consider a cost profile \mathbf{c} and an agent $i \in W(\mathbf{x}(\mathbf{c}))$. Let $j \in \{1, \dots, x_i(\mathbf{c})\}$ be an arbitrary service level allocated to i . Intuitively, we refer to the *critical payment* $p_{ij}(\mathbf{c}_{-i})$ for service level j of i as the largest cost q that i can declare and still obtain service level j (see Fig. 1 for an illustration). More formally, we define $Q_{ij}(\mathbf{c}_{-i})$ as the set of all points q satisfying $\lim_{z \rightarrow q^-} x_i(z, \mathbf{c}_{-i}) \geq j$ and $\lim_{z \rightarrow q^+} x_i(z, \mathbf{c}_{-i}) \leq j$ and let $p_{ij}(\mathbf{c}_{-i}) = \sup(Q_{ij}(\mathbf{c}_{-i}))$. Note that such a point q must always exist and $c_i \leq q \leq \frac{B}{k}$. To see this, note that $x_i(c_i, \mathbf{c}_{-i}) \geq j$ which implies that $c_i \leq q$ and that $x_i(z, \mathbf{c}_{-i}) = 0 < j$ for all $z > \frac{B}{k}$ (by our assumption that $c_i \leq \frac{B}{k}$) which implies that $q \leq \frac{B}{k}$.⁵

It is easy to see that the payment of an agent i can be written as the sum over these critical payments for the levels of service i was hired for.

Fact 3. *Let \mathbf{c} be a cost profile and let $i \in W(\mathbf{x}(\mathbf{c}))$. It holds that $p_i(\mathbf{c}) = \sum_{j=1}^{x_i(\mathbf{c})} p_{ij}(\mathbf{c}_{-i})$.*

⁵ It is not hard to see that the set $Q_{ij}(\mathbf{c}_{-i})$ is also closed and thus the supremum always exists.




Fig. 1. Illustration of the critical payments of agent i .

Lemma 3 is the main technical tool needed to establish budget-feasibility for the **else** part of SORT-&-REJECT(k). It is also used in the proof of Theorem 3 in the divisible agent setting.

Lemma 3 (Budget Feasibility Lemma). *Let \mathbf{c} be a cost profile such that $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \text{OPT}_F^k(\mathbf{c}_{-i^*})$. Then,*

$$\sum_{i=1}^n p_i(\mathbf{c}) \leq \frac{B}{1-\alpha} \left(\frac{m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} + \frac{\alpha}{1-\alpha} \right).$$

Observe that using Lemma 3, we can determine a range for values of α , for which SORT-&-REJECT(k) is budget-feasible:

Lemma 4. SORT-&-REJECT(k) is budget-feasible for $\alpha \leq \frac{1}{2+\sqrt{3}}$.

It is clear that Lemma 4 along with Lemmata 1 and 2 conclude the proof of Theorem 2.

4 Two Budget-Feasible Mechanisms for Divisible Agents

We consider the divisible agent model and derive two truthful and budget-feasible mechanisms. The first one is obtained by discretizing the valuation functions and reducing the problem to the k -level model (Sect. 4.1). The second one is an improved 2-approximate mechanisms for the divisible agent model with linear valuation functions (Sect. 4.2).

4.1 Using Sort-&-Reject(k) for Divisible Agents

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 $i \in N$.

Mechanism 2: CHUNK-&-SOLVE

- ▷ **Input:** A profile \mathbf{c} and a positive integer k
- 1 Initialize for each $i \in N$, $v_i : \{0, \dots, k\} \rightarrow \mathbb{R}_+$ with $v_i(0) = 0$. ;
 - 2 **for** $i \in N$ **and** $j \in \{1, \dots, k\}$ **do** set $v_i(j) := \bar{v}_i(j/k)$.
 - 3 Set $\alpha(L, k) = (3k + L - \sqrt{5k^2 + 6kL + L^2})/2k$.
 - 4 Let $\tilde{I} = (N, \mathbf{c}, kB, k, (v_i)_{i \in N})$ denote the resulting discretized instance of the k -level model. ;
 - 5 Compute \mathbf{x} by running SORT-&-REJECT(k)($\mathbf{c}, \alpha(L, k)$) on \tilde{I} . ;
 - 6 **for** $i \in N$ **do** set $\bar{x}_i = x_i/k$.
 - 7 Allocate $\bar{\mathbf{x}}$ and determine the payments $\bar{\mathbf{p}}$ according to (1).
-

Throughout this section, we assume that all valuation functions are L -regular for some $L \geq 1$ as defined in (3).

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 Sect. 3 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:

Theorem 3. *CHUNK-&-SOLVE is truthful, individually rational, 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 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.

4.2 A Mechanism for Divisible Agents and Linear Valuations

CHUNK-&-SOLVE retrieves the best known approximation guarantee of $1 + \phi$ for $L = 1$ and as $k \rightarrow \infty$ [17] (i.e., for divisible agents with linear valuations). Below, we improve upon this and give a simple 2-approximate budget-feasible mechanism for this setting. Our mechanism is inspired by the randomized 2-approximate budget-feasible mechanism by Gravin et al. [14] for indivisible agents. We also prove a lower bound of 2 for deterministic, truthful, individually rational and budget feasible mechanisms with independent allocation rules (as defined below).

Mechanism 3: PRUNING by Gravin et al. [14]

▷ **Input:** A profile \mathbf{c} with the agents relabeled so that $\frac{v_1}{c_1} \geq \dots \geq \frac{v_n}{c_n}$.

- 1 Let $r := \frac{1}{B} \max\{v_i \mid i \in N\}$.
- 2 **foreach** $i \in N$ **do** set $\bar{x}_i = 1$ if $\frac{v_i}{c_i} \geq r$ and $\bar{x}_i = 0$ otherwise.
- 3 Let $\ell := \arg \max\{i \mid \bar{x}_i = 1\}$.
- 4 **while** $rB < \sum_{i=1}^{\ell} v_i - \max\{v_i \mid i = 1, \dots, \ell\}$ **do**
- 5 Continuously increase rate r .
- 6 **if** $\frac{v_{\ell}}{c_{\ell}} \leq r$ **then** set $\bar{x}_{\ell} = 0$ and $\ell = \ell - 1$.
- 7 **return** $(r, \bar{\mathbf{x}})$

Phase 1: Pruning Mechanism for Divisible Agents. We first extend the PRUNING mechanism of Gravin et al. [14] to the divisible setting. This mechanism constitutes a crucial building block for both their deterministic 3-approximate mechanism and their randomized 2-approximate mechanism for indivisible agents [14]. As we show below, it serves as a useful starting point for the divisible setting as well.

Given a profile \mathbf{c} , this mechanism computes an allocation $\bar{\mathbf{x}}(\mathbf{c})$, which we refer to as the *provisional allocation*, and a positive quantity $r(\mathbf{c})$, which we refer to as the *rate*. We assume that the agents are initially relabeled by their decreasing value-per-cost ratio, i.e., $\frac{v_1}{c_1} \geq \frac{v_2}{c_2} \geq \dots \geq \frac{v_n}{c_n}$. The mechanism proceeds as described in Mechanism 3.

Gravin et al. [14] showed that PRUNING is monotone. In fact, an even stronger *robustness* property holds (and is implicit in the proof of Lemma 3.1 in [14]):

each agent i that is a winner in the provisional allocation cannot alter the outcome of PRUNING unilaterally while remaining a winner in the provisional allocation.

Lemma 5 (implied by Lemma 3.1 of [14]). *Let \mathbf{c} be a profile. Consider an agent $i \in N$ with $\bar{x}_i(\mathbf{c}) = 1$. Then, for all c'_i such that $\bar{x}_i(c'_i, \mathbf{c}_{-i}) = 1$, it holds that $\bar{\mathbf{x}}(c'_i, \mathbf{c}_{-i}) = \bar{\mathbf{x}}(\mathbf{c})$ and $r(c'_i, \mathbf{c}_{-i}) = r(\mathbf{c})$.*

Given this robustness property, PRUNING can be used as a first filtering step to discard inefficient agents, followed by a subsequent allocation scheme which takes $(r(\mathbf{c}), \bar{\mathbf{x}}(\mathbf{c}))$ as input. If the subsequent allocation scheme is monotone, then the sequential composition of PRUNING with this allocation scheme is monotone as well. This *composability property* is proven in Lemma 3.1 of [14].

Let $(r, \bar{\mathbf{x}})$ be the output of PRUNING for a cost profile \mathbf{c} . Given $\bar{\mathbf{x}}$, we define S as the set of agents that are provisionally allocated, i^* as the highest value agent in S , and T as the set of remaining agents. More formally, we define

$$S = \{i \in N \mid \bar{x}_i = 1\}, \quad i^* = \arg \max\{v_i \mid i \in S\} \quad \text{and} \quad T = S \setminus \{i^*\}. \quad (5)$$

Note that the definitions of S , i^* and T depend on $\bar{\mathbf{x}}$ (and thus the cost profile \mathbf{c}). For notational convenience, we do not state this reference explicitly if it is clear from the context.

The following properties were proved in [14] and are useful in our analysis.

Lemma 6 (Lemma 3.2 of [14]). *Given a profile \mathbf{c} , let $(r, \bar{\mathbf{x}})$ be the output of PRUNING. Let $S = T \cup \{i^*\}$ be defined as in (5) with respect to $\bar{\mathbf{x}}$. We have*

1. $c_i \leq \frac{v_i}{r} \leq B$ for all $i \in S$.
2. $v(T) \leq rB < v(S)$.
3. $\text{OPT}_F \leq v(S) + r \cdot (B - c(S))$.

Phase 2: Independent Allocation Schemes. Our mechanism combines the PRUNING mechanism above with the allocation schemes defined in (6) below. We refer to the resulting mechanism as PRUNE-&-ASSIGN (see Mechanism 4).

First, we need to define the following constants:

$$q_{i^*} = \begin{cases} \frac{1}{2} - q & \text{if } v_{i^*} \leq v(T) \\ \frac{1}{2} & \text{otherwise} \end{cases}, \quad q_i = 1 - q_{i^*} - q, \forall i \in T, \quad \text{where } q = \frac{1}{2} \frac{v(S) - rB}{\min\{v_{i^*}, v(T)\}}.$$

Note that the constant q_i for all agents $i \in T$ is the same. It is not hard to prove that $q \in [0, \frac{1}{2}]$ (see [14, Lemma 5.1]). The constants above are chosen so that $rB/2 = q_{i^*}v_{i^*} + (1 - q_{i^*} - q)v(T)$.

We can now define our (fractional) allocation function $x_i(\mathbf{c}) = x_i(c_i)$ for each agent $i \in T \cup \{i^*\}$:

$$x_i(z) = q_i + \frac{v_i - rz}{2v_i} \quad \text{for } z \in \left[0, \frac{v_i}{r}\right]. \quad (6)$$

Note that $x_i(c_i, \mathbf{c}_{-i}) = x_i(c_i)$ only depends on agent i 's cost c_i . We call such allocation rules *independent*. Further, note that by property (1) of Lemma 6, the cost c_i of each agent $i \in S$ after pruning is at most $\frac{v_i}{r}$, i.e., $x_i(z)$ will be determined by some value $z \in [0, \frac{v_i}{r}]$. It is not hard to verify that x_i is well-defined (given the chosen parameters q_{i^*} , q_T and q above).

Theorem 4. *PRUNE-&-ASSIGN is truthful, individually rational, budget-feasible and 2-approximate for instances of the divisible agent model with linear valuations.*

Given the lower bound of $1 + \sqrt{2} \approx 2.41$ by Chen et al. [11] in this setting with indivisible agents, Theorem 4 establishes a separation between the indivisible agent model and the divisible agent model with linear valuations.

Finally, we give a lower bound of 2 for deterministic, truthful, individually rational and budget feasible mechanisms with independent allocation rules. Note that PRUNE-&-ASSIGN does not belong to this class of mechanisms (due to PRUNING). However, our analysis of PRUNE-&-ASSIGN is tight (see [4]).

Theorem 5. *Let \mathcal{M} be a deterministic, truthful, individually rational and budget feasible mechanism with an approximation guarantee of α . If \mathcal{M} has independent allocation rules, then $\alpha \geq 2$.*

Mechanism 4: PRUNE-&-ASSIGN for Divisible Agents

- ▷ **Input:** A profile \mathbf{c} with the agents relabeled so that $\frac{v_1}{c_1} \geq \frac{v_2}{c_2} \geq \dots \geq \frac{v_n}{c_n}$
- 1 Obtain $(r, \bar{\mathbf{x}})$ by running PRUNING for profile \mathbf{c} .
 - 2 Let $S = T \cup \{i^*\}$ be as defined in (5).
 - 3 Determine the fractional allocation: $x_{i^*}(c_{i^*})$ and $x_i(c_i), \forall i \in T$ as in (6).
 - 4 Allocate \mathbf{x} and determine the payments \mathbf{p} as in (1).
-

5 Conclusion and Future Work

In this work we revisited two settings where partial allocations are allowed and draw clear connections between them. Under mild assumptions like being able to afford each agent entirely and having “nice” concave valuation functions (i.e., $O(1)$ -regular), we give deterministic, truthful and budget-feasible mechanisms with constant approximation guarantees. We believe these are settings that are both interesting and relevant to applications and there are several open questions we do not settle here. A natural direction, not considered at all in this work, is to deal with additional combinatorial constraints, like matching, matroid, or even polymatroid (for the k -level setting) constraints. For the k -level setting, it would be interesting to understand whether we can obtain mechanisms with approximation guarantees closer to those possible for single-level settings, or alternatively, determine whether allowing multiple levels of service is an inherently harder problem. As far as simple settings are concerned, the most important open problem is still the indivisible agents case with additive valuations, for which the best-possible approximation ratio is in $[1 + \sqrt{2}, 3]$ (due to [11, 14]). The corresponding range for the divisible agent setting is $[e/(e-1), 2]$ (due to [5] and our Theorem 4). Any progress on these fronts may give rise to novel techniques, which may be also used for problems in richer environments.

Acknowledgment. This work was supported by the Dutch Research Council (NWO) through its Open Technology Program, proj. no. 18938, and the Gravitation Project NETWORKS, grant no. 024.002.003. It has also been partially supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0, funded by the European Union under the NextGenerationEU Program and the EU Horizon 2020 Research and Innovation Program under the Marie Skłodowska-Curie Grant Agreement, grant no. 101034253.

Moreover, it was supported by the “1st Call for HFRI Research Projects to support faculty members and researchers and the procurement of high-cost research equipment” (proj. no. HFRI-FM17-3512). Finally, part of this work was done during AT’s visit to the University of Essex that was supported by COST Action CA16228 (European Network for Game Theory).

References

1. Amanatidis, G., Birmpas, G., Markakis, E.: Coverage, matching, and beyond: new results on budgeted mechanism design. In: Cai, Y., Vetta, A. (eds.) WINE 2016.

- LNCS, vol. 10123, pp. 414–428. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-54110-4_29
2. Amanatidis, G., Birmpas, G., Markakis, E.: On budget-feasible mechanism design for symmetric submodular objectives. In: Devanur, N.R., Lu, P. (eds.) WINE 2017. LNCS, vol. 10660, pp. 1–15. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71924-5_1
 3. Amanatidis, G., Kleer, P., Schäfer, G.: Budget-feasible mechanism design for non-monotone submodular objectives: offline and online. In: Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, pp. 901–919 (2019)
 4. Amanatidis, G., Klumper, S., Markakis, E., Schäfer, G., Tsikiris, A.: Partial allocations in budget-feasible mechanism design: bridging multiple levels of service and divisible agents. CoRR/arXiv abs/2307.07385 (2023)
 5. Anari, N., Goel, G., Nikzad, A.: Budget feasible procurement auctions. Oper. Res. **66**(3), 637–652 (2018)
 6. Apt, K.R., Heering, J.: Characterization of incentive compatible single-parameter mechanisms revisited. J. Mech. Inst. Design **7**(1), 113–129 (2022)
 7. Archer, A., Tardos, É.: Truthful mechanisms for one-parameter agents. In: Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science, pp. 482–491 (2001)
 8. Balkanski, E., Garimidi, P., Gkatzelis, V., Schoepflin, D., Tan, X.: Deterministic budget-feasible clock auctions. In: Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, pp. 2940–2963 (2022)
 9. Bei, X., Chen, N., Gravin, N., Lu, P.: Worst-case mechanism design via Bayesian analysis. SIAM J. Comput. **46**(4), 1428–1448 (2017)
 10. Chan, H., Chen, J.: Truthful multi-unit procurements with budgets. In: Liu, T.-Y., Qi, Q., Ye, Y. (eds.) WINE 2014. LNCS, vol. 8877, pp. 89–105. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-13129-0_7
 11. Chen, N., Gravin, N., Lu, P.: On the approximability of budget feasible mechanisms. In: Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, pp. 685–699 (2011)
 12. Dobzinski, S., Papadimitriou, C.H., Singer, Y.: Mechanisms for complement-free procurement. In: Proceedings of the 12th ACM Conference on Electronic Commerce, EC 2011, pp. 273–282 (2011)
 13. Goel, G., Nikzad, A., Singla, A.: Allocating tasks to workers with matching constraints: truthful mechanisms for crowdsourcing markets. In: Proceedings of the 23rd International Conference on World Wide Web, WWW 2014, pp. 279–280 (2014)
 14. Gravin, N., Jin, Y., Lu, P., Zhang, C.: Optimal budget-feasible mechanisms for additive valuations. ACM Trans. Econ. Comput. (TEAC) **8**(4), 1–15 (2020)
 15. Huang, H., Han, K., Cui, S., Tang, J.: Randomized pricing with deferred acceptance for revenue maximization with submodular objectives. In: Proceedings of the ACM Web Conference 2023, pp. 3530–3540 (2023)
 16. Jalaly Khalilabadi, P., Tardos, É.: Simple and efficient budget feasible mechanisms for monotone submodular valuations. In: Christodoulou, G., Harks, T. (eds.) WINE 2018. LNCS, vol. 11316, pp. 246–263. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-04612-5_17
 17. Klumper, S., Schäfer, G.: Budget feasible mechanisms for procurement auctions with divisible agents. In: Kanellopoulos, P., Kyropoulou, M., Voudouris, A. (eds.) SAGT 2022. LNCS, vol. 13584, pp. 78–93. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-15714-1_5

18. Leonardi, S., Monaco, G., Sankowski, P., Zhang, Q.: Budget feasible mechanisms on matroids. In: Eisenbrand, F., Koenemann, J. (eds.) IPCO 2017. LNCS, vol. 10328, pp. 368–379. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59250-3_30
19. Li, M., Wang, C., Zhang, M.: Budget feasible mechanisms for facility location games with strategic facilities. *Auton. Agent. Multi-Agent Syst.* **36**(2), 35 (2022)
20. Martello, S., Toth, P.: *Knapsack Problems: Algorithms and Computer Implementations*. Wiley, Hoboken (1990)
21. Milgrom, P., Segal, I.: Clock auctions and radio spectrum reallocation. *J. Polit. Econ.* **128**(1), 1–31 (2020)
22. Myerson, R.: Optimal auction design. *Math. Oper. Res.* **6**(1), 58–73 (1981)
23. Rubinstein, A., Zhao, J.: Beyond worst-case budget-feasible mechanism design. In: 14th Innovations in Theoretical Computer Science Conference (ITCS 2023). Schloss Dagstuhl-Leibniz-Zentrum für Informatik (2023)
24. Singer, Y.: Budget feasible mechanisms. In: In Proceedings of the 51st Annual Symposium on Foundations of Computer Science, FOCS 2010, pp. 765–774 (2010)