



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

GEORGIOS AMANATIDIS, Athens University of Economics and Business, Athens, Greece and Archimedes / Athena RC, Athens, Greece

SOPHIE KLUMPER, Centrum Wiskunde en Informatica, Amsterdam, Netherlands and University of Amsterdam, Amsterdam, Netherlands

EVANGELOS MARKAKIS, Athens University of Economics and Business, Athens, Greece, Archimedes / Athena RC, Athens, Greece, and Input Output (IOG), Athens, Greece

GUIDO SCHÄFER, Centrum Wiskunde en Informatica, Amsterdam, Netherlands and University of Amsterdam, Amsterdam, Netherlands

ARTEM TSIKIRIDIS, Centrum Wiskunde en Informatica, Amsterdam, Netherlands

Budget-feasible procurement has been a major paradigm in mechanism design since its introduction by Singer [28]. 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, might offer multiple levels of a service, or 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 in which each agent's services are either fully acquired or not at all. A reason for this is that in settings with partial allocations, such as the ones mentioned, there are strong inapproximability results (see, e.g., Anari et al. [5], Chan and Chen [10]). 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 in which each agent offers multiple levels of service and the auctioneer has a valuation function that is separable concave, i.e., it is the sum of concave functions. We then use this result to design a deterministic, truthful, and budget-feasible

This work was supported by the Dutch Research Council (NWO) through its Open Technology Program (project 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 Hellenic Foundation for Research and Innovation (HFRI), partially by the "1st Call for HFRI Research Projects to support faculty members and researchers and the procurement of high-cost research equipment" (project no. HFRI-FM17-3512) and partially by the HFRI call "Basic Research Financing (Horizontal Support of all Sciences)" under the National Recovery and Resilience Plan "Greece 2.0" funded by the European Union – NextGenerationEU (project no. 15877). Finally, part of this work was done during AT's visit to the University of Essex, which was supported by COST Action CA16228 (European Network for Game Theory).

Authors' Contact Information: Georgios Amanatidis, Athens University of Economics and Business, Athens, Greece and Archimedes / Athena RC, Athens, Greece; e-mail: g.d.amanatidis@gmail.com; Sophie Klumper, Centrum Wiskunde en Informatica, Amsterdam, Netherlands and University of Amsterdam, Amsterdam, Netherlands; e-mail: Sophie.Klumper@cw.nl; Evangelos Markakis, Athens University of Economics and Business, Athens, Greece, Archimedes / Athena RC, Athens, Greece, and Input Output (IOG), Athens, Greece; e-mail: markakis@gmail.com; Guido Schäfer, Centrum Wiskunde en Informatica, Amsterdam, Netherlands and University of Amsterdam, Amsterdam, Netherlands; e-mail: g.schaefer@cw.nl; Artem Tsikiridis, Centrum Wiskunde en Informatica, Amsterdam, Netherlands; e-mail: artem.tsikiridis@gmail.com.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2167-8375/2025/04-ART9

<https://doi.org/10.1145/3723883>

$O(1)$ -approximation mechanism for the setting in which any fraction of a service can be acquired, again for separable concave objectives. For the special case in which the objective is the sum of linear valuation functions, we improve the best known approximation ratio for the problem from $(3 + \sqrt{5})/2$ (by Klumper and Schäfer [19]) to 2. This establishes a separation between this setting and its indivisible counterpart.

CCS Concepts: • **Theory of computation** → **Algorithmic mechanism design**; **Packing and covering problems**; **Discrete optimization**; • **Mathematics of computing** → **Combinatorial optimization**; *Combinatorial algorithms*;

Additional Key Words and Phrases: Procurement auctions, budget-feasible mechanism design, multiple levels of service

ACM Reference Format:

Georgios Amanatidis, Sophie Klumper, Evangelos Markakis, Guido Schäfer, and Artem Tsikiris. 2025. Partial Allocations in Budget-Feasible Mechanism Design: Bridging Multiple Levels of Service and Divisible Agents. *ACM Trans. Econ. Comput.* 13, 2, Article 9 (April 2025), 28 pages. <https://doi.org/10.1145/3723883>

1 Introduction

Consider a procurement auction in which the agents have *private* costs on the services that they can offer and the auctioneer associates a value for each possible set of selected agents. This forms a single parameter auction environment in which the agents may strategically misreport their cost to their advantage for obtaining higher payments. Imagine now that the auctioneer additionally has a strict budget constraint that cannot be violated. 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 one's procured value and such that the total payments to the agents respect the budget limitations. This is precisely the model that was originally proposed by Singer [28] for *indivisible* agents, i.e., with a binary decision to be made for each agent (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.

Since the work of Singer [28], a large body of works has emerged devoted to obtaining improved results on the original model as well as to 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. All of these results have significantly enhanced our understanding for the indivisible scenario. In this article, 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, which directly generalizes Singer's setting [28], 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 of the levels that the auctioneer 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 was 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 total number of levels/copies. 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 single agent.

Divisible Agents: A related setting is the one 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 the agent is 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 [28], when it is assumed that the auctioneer can afford to hire any individual divisible 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 the agent's entire service is insignificant compared with the budget of the auctioneer. Additionally, they noticed 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 [19] revisited this problem without the large market assumption but rather 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 presented a deterministic, truthful and budget-feasible mechanism that achieves an approximation ratio of $\frac{3+\sqrt{5}}{2} \approx 2.62$ for linear valuation functions and extended it to the setting in which all agents are associated with the same concave valuation function.

The two aforementioned settings of procurement auctions have a number of practical applications in various domains. As previously mentioned, the divisible setting would be useful, for example, to model the time availability of a worker in the context of crowdsourcing. These types of auctions can also be applied to other industries, such as transportation and logistics, in which 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, whereas 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.

1.1 Our Contributions

In this work, we propose deterministic, truthful, and budget-feasible mechanisms for settings with partial allocations, as follows.

- We present a mechanism, SORT- ϕ -REJECT $_{\alpha}^k$ (Mechanism 2), 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 (Section 3.1, Theorem 3.1). The main idea behind our novel mechanism is to apply a backwards greedy approach, in which we start from an optimal fractional solution and 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. Note that the work of Chen et al. [11] implies that no deterministic mechanism can achieve an approximation ratio better than $1 + \sqrt{2} \approx 2.41$ for this setting (even when $k = 1$). We further explore how this guarantee improves as one moves towards a *large market* (Section 3.3, Theorem 3.2).

- We use $\text{SORT-}\hat{\phi}\text{-REJECT}_\alpha^k$ as a subroutine in order to design a mechanism for the setting with divisible agents, $\text{CHUNK-}\hat{\phi}\text{-SOLVE}$ (Mechanism 3), which achieves an approximation ratio of $4 + 2\sqrt{3} \approx 7.46$ for non-decreasing concave valuation functions (Section 4.1, Theorem 4.1). This is the first result for the problem that is independent of the number of agents n and the derivative of the valuation functions close to 0, and is a significant improvement over the conference version of this work [4]. Note that Anari et al. [5] have shown that no mechanism can achieve an approximation ratio better than $e/(e-1) \approx 1.58$ for this setting. On a technical level, we exploit the correspondence between the discrete and the continuous settings; as the number of levels of service grows arbitrarily large, the former converges to the latter. This behavior, however, does not imply a polynomial reduction. For a modest number of levels, one could say for each discrete valuation function that it either approximates well its continuous counterpart or that its value is concentrated in its first level.
- We improve on the state-of-the-art for the divisible setting with linear valuation functions [19] by suggesting a deterministic 2-approximation mechanism, $\text{PRUNE-}\hat{\phi}\text{-ASSIGN}$ (Mechanism 5; see Section 4.2, Theorem 4.2). Note that the $e/(e-1)$ lower bound of Anari et al. [5] holds even in this case, as the instance they construct has a linear valuation function. By proving this result, we establish a separation of the divisible agents model and its indivisible counterpart, for which a lower bound of $1 + \sqrt{2}$ is known owing to Chen et al. [11]. Our mechanism is inspired by the *randomized 2-approximation* 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 [19]. Even for the discrete setting with multiple levels of service, this assumption circumvents a strong lower bound of Chan and Chen [10], which is linear to the maximum number of levels offered by any agent (see also Remark 2.1 and Appendix A). In both settings, our assumptions are much weaker than the large market assumptions often made in the literature (see, e.g., [5, 18]).

1.2 Further Related Work

The design of truthful budget-feasible mechanisms for indivisible agents was introduced by Singer [28], 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 deterministic 3-approximation 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 randomized 2-approximation 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 [19] studied the divisible setting without the large market assumption, but rather under the assumption that the private cost of each agent is bounded by the budget and gave, among other results, a deterministic $\frac{3+\sqrt{5}}{2}$ -approximation mechanism for linear valuation functions.

For indivisible agents, the problem has also been extended to richer valuation functions. This line of inquiry was also started by Singer [28], 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; the same authors devised a deterministic mechanism with an 8.34 approximation. Subsequently, the bound for randomized mechanisms was improved by Jalaly and Tardos [18] to 5. More 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, presented a 4.75-approximation mechanism. Very recently, these guarantees were further improved by Han et al. [15], who devised a randomized mechanism and a deterministic mechanism with approximation ratios 4.3 and 4.45, respectively, under the same paradigm.

Beyond monotone submodular valuation functions, it becomes significantly 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 was by Amanatidis et al. [3] with an approximation guarantee of 505. This guarantee was improved to 64 by Balkanski et al. [8] who provided the first deterministic mechanism for the problem and Huang et al. [17] who gave a further improvement of $(3 + \sqrt{5})^2$ for randomized mechanisms. The state-of-the-art randomized mechanism for non-monotone submodular valuation functions is by Han et al. [15], achieving an approximation ratio of 12. Richer valuation functions, such as XOS valuation functions (see Amanatidis et al. [2], Bei et al. [9], Neogi et al. [26]) and subadditive valuation functions (see Balkanski et al. [8], Bei et al. [9], Dobzinski et al. [12], Neogi et al. [26]), have been extensively studied. Until very recently, no explicit construction of a mechanism achieving a constant approximation for subadditive valuation functions was known. However, Bei et al. [9] had shown that such a mechanism should exist assuming a *demand query oracle*, using a non-constructive argument. Interestingly, Neogi et al. [26] recently constructed such a mechanism, marking a major breakthrough.

Other settings that have been studied include environments with underlying feasibility constraints, such as downward-closed environments (Amanatidis et al. [1], Huang et al. [17]) and matroid constraints (Leonardi et al. [20]). 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 [18]). Recently, Li et al. [21] 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 [27].

For further details, we refer the interested reader to the very recent survey of Liu et al. [22].

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 article. The multiple-levels-of-service model is introduced in Section 2.2 and the divisible-agent model is in Section 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 the true cost for providing this service in full. 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 the 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 a (possibly fractional) 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 that agent i will receive for the 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 the 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 models that are studied in this article are single-parameter settings; thus, the characterization of Myerson [25] 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 2.1 to design truthful mechanisms.

THEOREM 2.1 ([7, 25]). *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 that $\int_0^\infty x_i(z, \mathbf{b}_{-i}) dz < \infty$. In this case, we can take the payment rule $\mathbf{p}(\mathbf{b})$ to be*

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

In this article, we will exclusively derive monotone allocation rules that are implemented with the payment rule as defined in Equation (1). Therefore, in the remainder of this article, 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}), yet we will omit the explicit reference to \mathbf{c} if it is clear from the context. Finally, whenever tie breaking is needed for any of our mechanisms, we assume a lexicographic precedence.

2.2 k -Level Model

We consider the following multiple levels of service model as a natural extension of the standard model introduced above (see also [10]). Throughout the article, we refer to this model as the *k -level model* for short: Suppose that 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 value for the first j levels of service of agent i . 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.,

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

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

$x_i : \mathbb{R}_{\geq 0}^n \rightarrow \{0, \dots, k\}$. Note also that the total cost of agent i is linear (as defined above), i.e., the cost of using $x_i = j$ levels of service 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 $j \in \{1, \dots, k\}$, with the convention that $m_i(0) = 0$. 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 levels of service. 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 $k \cdot c_i \leq B$. Note that, under this assumption, we implicitly constrain the space of the (declared) cost profiles.³ In contrast, in the *best-in setting*, which is equivalent to the setting of Chan and Chen [10], the auctioneer is guaranteed to be able to afford only the first level of service, i.e., given a cost profile \mathbf{c} , for every agent $i \in N$, it holds that $c_i \leq B$. We focus on the all-in setting throughout this work, with the only exception being Appendix A, in which we derive an almost tight result on the best-in setting.

REMARK 2.1. 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 randomized, universally truthful, budget-feasible mechanism. In Appendix A, we present a $(k+2+o(1))$ -approximation mechanism, named *GREEDY-BEST-IN^k* (Mechanism 6), almost settling the deterministic case. Although the best-in setting is not central in our work, this might be of independent interest. Note that the randomized mechanism suggested in Chan and Chen [10] has an approximation ratio of $4(1 + \ln(nk))$.

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* (see, e.g., [23] for a classification of knapsack problems):

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

The *k-Bounded Knapsack Problem* is NP-hard in general, since for $k = 1$ it reduces to the well-known *0-1 Knapsack Problem*. We say that a mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ is an α -approximation mechanism with $\alpha \geq 1$ if $\alpha \cdot v(\mathbf{x}(\mathbf{c})) \geq \text{OPT}^k(\mathbf{c})$. We also consider the relaxation of the above problem as a proxy for $\text{OPT}^k(\mathbf{c})$. The definition and further details about this are deferred to Section 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., we often refer to an 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 work. Throughout the paper, we refer to it 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

³That is, we assume that any agent who violates the respective condition is discarded up front from further consideration, e.g., by running a pre-processing step that removes such agents.

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 $v(\mathbf{x}(\mathbf{c})) = \sum_{i \in N} \bar{v}_i(x_i(\mathbf{c}))$.

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 $i \in N$, $c_i \leq B$. The observation about the cost restrictions in Section 2.2 applies to this assumption as well, i.e., we do constrain the bidding space.

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*:

$$\text{OPT}_F(\mathbf{c}) := \max \sum_{i=1}^n \bar{v}_i(x_i) \quad \text{s.t.} \quad \sum_{i=1}^n c_i x_i \leq B, \quad x_i \in [0, 1] \quad \forall i \in N. \quad (3)$$

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

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 defined in (2) by relaxing the integrality constraint:

$$\text{OPT}_F^k(\mathbf{c}) := \max \sum_{i=1}^n v_i(\lfloor x_i \rfloor) + m_i(\lceil x_i \rceil)(x_i - \lfloor x_i \rfloor) \quad \text{s.t.} \quad \sum_{i=1}^n c_i x_i \leq B, \quad x_i \in [0, k] \quad \forall i \in N. \quad (4)$$

Naturally, it holds that $\text{OPT}_F^k(\mathbf{c}) \geq \text{OPT}^k(\mathbf{c})$. Note that $\text{OPT}_F^1(\mathbf{c}) = \text{OPT}_F(\mathbf{c})$.

It is not hard to see that, due to the fact that the objective is separable concave, $\text{OPT}_F^k(\mathbf{c})$ inherits the well-known properties of its one-dimensional analogue. This includes the fact that an optimal solution can be computed by an adaptation of the standard greedy algorithm that sorts the elements in decreasing marginal density (marginal value per cost) and selects as many as possible [16]. To the best of our knowledge, separable concave functions are the *widest* class of objectives for which the Fractional k -Bounded Knapsack Problem admits optimal solutions with this structure, which we crucially utilize here. For completeness, we state the algorithm as Algorithm 1 below, along with an easy fact that we will repeatedly use.

FACT 2.1. *Given an instance $I = (N, \mathbf{c}, B, k, (v_i)_{i \in N})$ of the Fractional k -Bounded Knapsack Problem, Algorithm 1 computes in time $O(kn \log(kn))$ an optimal solution \mathbf{x}^* that has at most one coordinate with a non-integral value.*

3 A Budget-Feasible Mechanism for Multiple Levels of Service

We derive a natural truthful and budget-feasible greedy mechanism for the k -level model. This mechanism will also be used in our **CHUNK- \mathcal{C} -SOLVE** mechanism for the divisible-agent model (see

ALGORITHM 1: A Greedy Algorithm for Fractional k -Bounded Knapsack

▷ **Input:** An instance $I = (N, c, B, k, (v_i)_{i \in N})$ of the k -level model.

- 1 Initialize an empty list \mathcal{L} of $k \cdot n$ elements.
- 2 **for** $i \in N$ **do**
- 3 **for** $j \in \{1, \dots, k\}$ **do**
- 4 Add $\frac{m_i(j)}{c_i}$ to the list \mathcal{L} .
- 5 Sort \mathcal{L} in decreasing order. // we break ties lexicographically with respect to the agents' indices
- 6 Let $\mathbf{x}^* = \mathbf{0}$ and $j = 1$.
- 7 **while** $\sum_{i \in N} c_i x_i^* < B$ **do**
- 8 Let $\ell \in [n]$ be the agent corresponding to the j -th marginal value-per-cost in \mathcal{L} .
- 9 Set $x_\ell^* = x_\ell^* + \min\left(\frac{B - \sum_{i \in N} c_i x_i^*}{c_\ell}, 1\right)$.
- 10 Set $j = j + 1$.
- 11 Return \mathbf{x}^* .

Section 4.1). Our mechanism falls into the family of mechanisms that truncate a greedy optimal solution, as is often the case in this literature (see Further Related Work).

3.1 A Truthful Greedy Mechanism

The main idea underlying our mechanism is as follows. If there is an agent i^* whose maximum value $v_{i^*}(k)$ is high enough (in a certain sense), then we simply pick all levels of service of this agent. Otherwise, we compute an allocation using the following greedy procedure. We first compute an optimal allocation $\mathbf{x}^*(c)$ 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. The intuition is that this allocation is close to the optimal fractional solution because $\mathbf{x}^*(c)$ has at most one fractional component (Fact 2.1). We then repeatedly discard the worst level of service (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. A similar approach was employed by Klumper and Schäfer [19], albeit without the subtleties of our discrete setting, as they were working on the divisible agent model.

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. Note 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 2.

The main result of this section is the following theorem.

THEOREM 3.1. *For $\alpha = \frac{1}{2+\sqrt{3}}$, SORT- \hat{c} -REJECT $_\alpha^k$ is a truthful, individually rational, budget-feasible $(2 + \sqrt{3})$ -approximation mechanism for instances of the k -level model, and runs in time polynomial in n and k .*

The polynomial run time for computing the final allocation \mathbf{x} is straightforward given Fact 2.1. After determining \mathbf{x} , computing the payment p_i of each agent $i \in N$ can be done efficiently, as we describe in Appendix B. In the remainder of this section, we prove several lemmata to establish the properties stated in Theorem 3.1. Technically, the most challenging part is to prove that the mechanism is budget feasible (see Section 3.2).

The following property directly follows by construction of the mechanism.

Mechanism 2: SORT- $\hat{\mathcal{C}}$ -REJECT $_{\alpha}^k$

▷ **Input:** An instance $I = (N, \mathbf{c}, B, k, (v_i)_{i \in N})$ of the k -level model.

- 1 Let $i^* \in \arg \max_{i \in N} v_i(k) / \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i})$. // recall that ties are broken lexicographically
- 2 **if** $v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \cdot \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})$ **then**
- 3 Set $x_{i^*} = k$ and $x_i = 0$ for all $i \in N \setminus \{i^*\}$.
- 4 **else**
- 5 Compute an optimal fractional allocation $\mathbf{x}^*(\mathbf{c})$ using Algorithm 1. // i.e., $\mathbf{x}^*(\mathbf{c})$ is almost integral
- 6 Initialize $\mathbf{x} = (\lfloor x_1^*(\mathbf{c}) \rfloor, \dots, \lfloor x_n^*(\mathbf{c}) \rfloor)$. // by the construction of $\mathbf{x}^*(\mathbf{c})$ and the all-in assumption, $\mathbf{x} \neq \mathbf{0}$
- 7 **for** $i \in W(\mathbf{x})$ **do** // recall that $W(\mathbf{x}) = \{i \in N \mid x_i > 0\}$
- 8 **for** $j \in \{1, \dots, x_i\}$ **do**
- 9 Add the marginal value-per-cost ratio $m_i(j)/c_i$ to a list \mathcal{L} .
- 10 Sort \mathcal{L} in non-increasing order and let ℓ be the agent corresponding to the last marginal value-per-cost in \mathcal{L} .
- 11 **while** $v(\mathbf{x}) - m_{\ell}(x_{\ell}) \geq \alpha \text{OPT}_{\text{F}}^k(\mathbf{c})$ **do**
- 12 Set $x_{\ell} = x_{\ell} - 1$.
- 13 Remove the last element from \mathcal{L} and update ℓ .
- 14 Allocate \mathbf{x} and determine the payments \mathbf{p} according to (1).

FACT 3.1. *The allocation \mathbf{x} returned by SORT- $\hat{\mathcal{C}}$ -REJECT $_{\alpha}^k$ satisfies $x_i \leq x_i^*$ for every $i \in N$.*

We now show that the allocation rule of SORT- $\hat{\mathcal{C}}$ -REJECT $_{\alpha}^k$ is monotone.

LEMMA 3.1. *For any $\alpha \in (0, 1)$, the allocation rule of SORT- $\hat{\mathcal{C}}$ -REJECT $_{\alpha}^k$ is monotone.*

PROOF. Let \mathbf{c} be a cost profile. We distinguish the following two cases.

Case 1: $v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \cdot \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})$. In this case, i^* is hired for k levels of service. Suppose that i^* unilaterally deviates and decreases the cost to $c'_{i^*} < c_{i^*}$. Such a deviation has no influence on which agent is chosen in line 1 and does not alter the condition of this case. To see that, note that the quantity $\text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})$ does not depend on the bid of i^* , whereas $\text{OPT}_{\text{F}}^k((c'_{i^*}, \mathbf{c}_{-i^*})_{-i}) \geq \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i})$ for any $i \neq i^*$. Therefore, for any such deviation, i^* will remain the sole winner and will be hired for k levels of service, i.e., $x_{i^*}(c'_{i^*}, \mathbf{c}_{-i^*}) = x_{i^*}(\mathbf{c}) = k$ (in a slight abuse of notation, we introduce arguments to the components of \mathbf{x} to distinguish between the two runs). No other agent was winning in this case, hence, there is no need to examine deviations by other agents.

Case 2: $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \cdot \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})$. Here, the allocation rule of SORT- $\hat{\mathcal{C}}$ -REJECT $_{\alpha}^k$ is allocating to a set $W(\mathbf{x}(\mathbf{c}))$. Note that because we are in the all-in setting, $W(\mathbf{x}(\mathbf{c}))$ starts as a non-empty set, and our **while** condition guarantees that it remains so. Fix an agent i in $W(\mathbf{x}(\mathbf{c}))$ and suppose that i unilaterally deviates and declares $c'_i < c_i$. First, note that for every $j \in N \setminus \{i\}$, it holds that

$$\frac{v_j(k)}{\text{OPT}_{\text{F}}^k(c'_i, \mathbf{c}_{-\{j,i\}})} \leq \frac{v_j(k)}{\text{OPT}_{\text{F}}^k(\mathbf{c}_{-j})} \leq \frac{v_{i^*}(k)}{\text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})},$$

where the first inequality follows from the monotonicity of $\text{OPT}_{\text{F}}^k(\cdot)$ with respect to each cost, whereas for i we have that

$$\frac{v_i(k)}{\text{OPT}_{\text{F}}^k(\mathbf{c}_{-i})} \leq \frac{v_{i^*}(k)}{\text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})},$$

from before i deviated.

Therefore, even if some agent $\hat{i} \neq i^*$ is chosen in line 1 under the profile (c'_i, \mathbf{c}_{-i}) , the mechanism will always execute the **else** case. Moreover, by the way that $\mathbf{x}^*(c'_i, \mathbf{c}_{-i})$ is constructed, the corresponding initial integral solution \mathbf{x} in line 7 will have an i -th coordinate at least as large as before. Now, consider what happens to the **while** condition of the mechanism. On the one hand, it can only be that $\text{OPT}_F^k(c'_i, \mathbf{c}_{-i}) \geq \text{OPT}_F^k(\mathbf{c})$. On the other hand, the marginal value-per-cost ratios of agent i under the profile (c'_i, \mathbf{c}_{-i}) have a better position in the ordering constructed by $\text{SORT-}\hat{\sigma}\text{-REJECT}_\alpha^k$. Therefore, agent i will be hired to at least the same extent or more, i.e., for the final allocations we have that $x_i(c'_i, \mathbf{c}_i) \geq x_i(\mathbf{c})$, proving monotonicity. \square

Since the payments are computed according to Equation (1), we conclude that the mechanism is truthful and individually rational. We continue by showing that $\text{SORT-}\hat{\sigma}\text{-REJECT}_\alpha^k$ achieves the claimed approximation guarantee.

LEMMA 3.2. *Fix any $\alpha \in (0, \frac{3-\sqrt{5}}{2})$. For the allocation $\mathbf{x}(\mathbf{c})$ computed by $\text{SORT-}\hat{\sigma}\text{-REJECT}_\alpha^k$ with input the cost profile \mathbf{c} , it holds that $v(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}^k(\mathbf{c})$.*

PROOF. For a cost profile \mathbf{c} , we will prove the claimed guarantee against the optimal value of the fractional relaxation of the bounded knapsack instance, i.e., we will show that $v(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}_F^k(\mathbf{c})$. As mentioned in Section 2.2, this establishes our guarantee since $\text{OPT}_F^k(\mathbf{c}) \geq \text{OPT}^k(\mathbf{c})$. Again, we distinguish the following two cases.

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

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

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

Case 2: $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \cdot \text{OPT}_F^k(\mathbf{c}_{-i^*})$. In this case, whenever the **while** loop of $\text{SORT-}\hat{\sigma}\text{-REJECT}_\alpha^k$ runs at least once, we immediately obtain that $v(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}_F^k(\mathbf{c})$ for any $\alpha < 1$. Alternatively, consider an instance for which the **while** condition evaluates to *False* the first time it is checked. In this case, the initial version of \mathbf{x} cannot be equal to $\mathbf{x}^*(\mathbf{c})$. As a consequence, under $\mathbf{x}^*(\mathbf{c})$, the last level of service of some agent—say, agent f —is fractionally included. In particular, this means that $B - \sum_{i \in N} \lfloor x_i^* \rfloor c_i < c_f$. We argue that, for $\alpha \leq \frac{3-\sqrt{5}}{2}$, we still obtain $v(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}_F^k(\mathbf{c})$ and the lemma follows. Toward a contradiction, suppose that this is not the case. This implies that

$$\alpha \text{OPT}_F^k(\mathbf{c}) > v(\mathbf{x}(\mathbf{c})) = \text{OPT}_F^k(\mathbf{c}) - m_f \left(\lceil x_f^* \rceil \right) \frac{B - \sum_{i \in N} \lfloor x_i^* \rfloor c_i}{c_f} > \text{OPT}_F^k(\mathbf{c}) - m_f \left(\lceil x_f^* \rceil \right),$$

which, by rearranging terms, yields

$$m_f \left(\lceil x_f^* \rceil \right) > (1 - \alpha) \text{OPT}_F^k(\mathbf{c}). \quad (5)$$

At the same time, we have that

$$\frac{m_f(\lceil x_f^* \rceil)}{\text{OPT}_F^k(\mathbf{c})} \leq \frac{v_f(k)}{\text{OPT}_F^k(\mathbf{c}_{-f})} \leq \frac{v_{i^*}(k)}{\text{OPT}_F^k(\mathbf{c}_{-i^*})} < \frac{\alpha}{1-\alpha}, \quad (6)$$

where the first inequality follows from straightforward bounds on both the numerator and the denominator, the second inequality holds by the choice of i^* , and the third one is the very definition of Case 2.

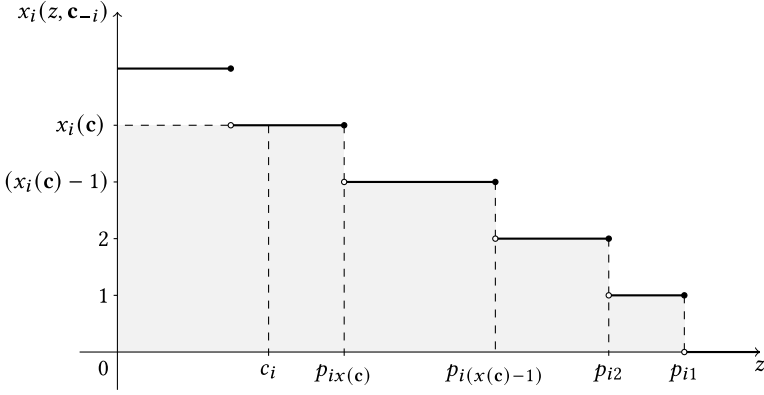


Fig. 1. Illustration of the critical payments of agent i .

Now, by combining inequalities (5) and (6), we obtain that α must be such that $\frac{\alpha}{1-\alpha} > 1 - \alpha$ or, equivalently, that $\alpha > (3 - \sqrt{5})/2$, a contradiction. \square

3.2 Ensuring the Budget Feasibility of SORT- \mathcal{C} -REJECT $_{\alpha}^k$

It remains to prove that SORT- \mathcal{C} -REJECT $_{\alpha}^k$ 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 any level of service among the ones for which agent i is hired. Intuitively, we refer to the *critical payment* $p_{ij}(\mathbf{c}_{-i})$ for level of service j of i as the largest cost p that i can declare and still have i 's level of service j picked by the mechanism (see Figure 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 $p_{ij}(\mathbf{c}_{-i}) = \sup(Q_{ij}(\mathbf{c}_{-i}))$. Such a number q must always exist (e.g., $c_i \in Q_{ij}(\mathbf{c}_{-i})$ and $c_i \leq q \leq \frac{B}{k}$). To see the latter, 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 all-in assumption that would prune any agent declaring cost more than $\frac{B}{k}$), which implies that $q \leq \frac{B}{k}$. Given that $Q_{ij}(\mathbf{c}_{-i})$ is nonempty and bounded from above, $p_{ij}(\mathbf{c}_{-i})$ always exists. 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.

LEMMA 3.3. *Let \mathbf{c} be a cost profile and assume that $i \in W(\mathbf{x}(\mathbf{c}))$. It holds that $p_i(\mathbf{c}) = \sum_{j=1}^{x_i(\mathbf{c})} p_{ij}(\mathbf{c}_{-i})$.*

PROOF. Starting with Equation (1) from Theorem 2.1, we obtain that

$$p_i(\mathbf{c}) = c_i x_i(\mathbf{c}) + \int_{c_i}^{\infty} x_i(z, \mathbf{c}_{-i}) dz = c_i x_i(\mathbf{c}) + \sum_{j=1}^{x_i(\mathbf{c})-1} \int_{p_{i(j+1)}(\mathbf{c}_{-i})}^{p_{ij}(\mathbf{c}_{-i})} j dz + \int_{c_i}^{p_{ix_i(\mathbf{c})}(\mathbf{c}_{-i})} x_i(\mathbf{c}) dz = \sum_{j=1}^{x_i(\mathbf{c})} p_{ij}(\mathbf{c}_{-i}). \quad \square$$

Lemma 3.4 is the main technical tool needed to establish budget feasibility for the **else** part of the SORT- \mathcal{C} -REJECT $_{\alpha}^k$ mechanism. It is also used in the proofs of Theorems 3.2 and 4.1 (implicitly in the latter) for the large market and the divisible agent setting, respectively.

LEMMA 3.4. *Fix any $\alpha \in (0, 1)$. Let \mathbf{c} be a cost profile such that $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})$. Then,*

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

Observe that, using Lemma 3.4, we can determine a range for values of α , for which SORT- \mathcal{C} -REJECT $_{\alpha}^k$ is budget feasible.

LEMMA 3.5. For $\alpha \in (0, \frac{1}{2+\sqrt{3}}]$, SORT- \mathcal{C} -REJECT $_{\alpha}^k$ is budget feasible.

PROOF. Let \mathbf{c} be a cost profile. If $v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})$, agent i^* is the only agent hired and for k levels of services. By Lemma 3.3, we obtain that

$$p_{i^*}(\mathbf{c}) = \sum_{j=1}^k p_{i^*j}(\mathbf{c}_{-i^*}) = k \cdot p_{i^*k}(\mathbf{c}_{-i^*}) \leq k \frac{B}{k} = B.$$

The second equality follows from the definition of the critical payments and from the fact that the allocation of i^* will only be smaller than k when another agent becomes i^* or when $c'_i > \frac{B}{k}$. The inequality follows from the fact that $p_{i^*k}(\mathbf{c}_{-i^*}) \leq \frac{B}{k}$, as argued above.

Consider now the other case, i.e., a profile \mathbf{c} such that $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})$. In this case, we may invoke Lemma 3.4 and, thus, we get that

$$\begin{aligned} \sum_{i=1}^n p_i(\mathbf{c}) &< \frac{B}{1-\alpha} \left(\frac{m_{\ell}(x_{\ell}(\mathbf{c}))}{\text{OPT}_{\text{F}}^k(\mathbf{c}_{-\ell})} + \frac{\alpha}{1-\alpha} \right) \leq \frac{B}{1-\alpha} \left(\frac{v_{\ell}(k)}{\text{OPT}_{\text{F}}^k(\mathbf{c}_{-\ell})} + \frac{\alpha}{1-\alpha} \right) \\ &\leq \frac{B}{1-\alpha} \left(\frac{v_{i^*}(k)}{\text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})} + \frac{\alpha}{1-\alpha} \right) < \frac{B}{1-\alpha} \left(\frac{\alpha}{1-\alpha} + \frac{\alpha}{1-\alpha} \right) \leq B \frac{2\alpha}{(1-\alpha)^2}. \end{aligned}$$

The second inequality follows by observing that $m_{\ell}(x_{\ell}(\mathbf{c})) \leq v_{\ell}(k)$ since $v_{\ell}(\cdot)$ is non-decreasing. The next two inequalities are due to the definition of i^* and the fact that, in this case, $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})$.

To obtain the budget feasibility SORT- \mathcal{C} -REJECT $_{\alpha}^k$, we must ensure that $\frac{2\alpha}{(1-\alpha)^2} \leq 1$ or, equivalently, that $\alpha \leq \frac{1}{2+\sqrt{3}}$. The proof follows. \square

It is clear that Lemma 3.5 along with Lemmata 3.1 and 3.2 conclude the proof of Theorem 3.1.

The rest of this section is devoted to proving Lemma 3.4. We start by presenting a series of auxiliary statements that will prove to be useful in our analysis. The purpose of these statements is to characterize and give upper bounds on the individual payments of winning agents, whenever the **else** part of the mechanism is executed. We begin with Lemma 3.6, in which we derive an upper bound on the costs of winning agents. Observe that the final value of ℓ is the index of the agent with the smallest value-per-cost ratio in the allocation computed by the mechanism.

LEMMA 3.6. Fix any $\alpha \in (0, 1)$ and let \mathbf{c} be a cost profile for an instance with $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})$. At the end of a run of SORT- \mathcal{C} -REJECT $_{\alpha}^k$ it holds that

$$c_{\ell} < \frac{B}{1-\alpha} \cdot \frac{m_{\ell}(x_{\ell}(\mathbf{c}))}{\text{OPT}_{\text{F}}^k(\mathbf{c})}.$$

PROOF. Let \mathbf{x} be the allocation vector at the end of a run of SORT- \mathcal{C} -REJECT $_{\alpha}^k$ with input \mathbf{c} and use $\mathbf{x}^* := \mathbf{x}(\mathbf{c})$ for brevity. Observe that, since $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i^*})$, it must be that $v(\mathbf{x}) - m_{\ell}(x_{\ell}) < \alpha \text{OPT}_{\text{F}}^k(\mathbf{c})$, which implies that

$$\text{OPT}_{\text{F}}^k(\mathbf{c}) - v(\mathbf{x}) + m_{\ell}(x_{\ell}) > (1-\alpha) \text{OPT}_{\text{F}}^k(\mathbf{c}). \quad (7)$$

By Fact 2.1, there exists at most one agent in $W(\mathbf{x}^*)$ with a non-integer allocation. We denote that agent by f (if no such agent exists, let f be an arbitrary agent in $W(\mathbf{x}^*)$). Note that it is possible that $f = \ell$. We have that

$$\begin{aligned}
B &\geq \sum_{i \in W(\mathbf{x}^*)} c_i x_i^* \geq c_\ell (x_\ell^* - x_\ell + 1) + \sum_{i \in W(\mathbf{x}^*) \setminus \{\ell\}} c_i (x_i^* - x_i) \\
&= \sum_{j=x_\ell}^{\lfloor x_\ell^* \rfloor} \frac{c_\ell}{m_\ell(j)} m_\ell(j) + \sum_{i \in W(\mathbf{x}^*) \setminus \{\ell\}} \sum_{j=x_i+1}^{\lfloor x_i^* \rfloor} \frac{c_i}{m_i(j)} m_i(j) + \frac{c_f}{m_f(\lceil x_f^* \rceil)} (x_f^* - \lfloor x_f^* \rfloor) m_f(\lceil x_f^* \rceil) \\
&\geq \frac{c_\ell}{m_\ell(x_\ell)} \left(\sum_{j=x_\ell}^{\lfloor x_\ell^* \rfloor} m_\ell(j) + \sum_{i \in W(\mathbf{x}^*) \setminus \{\ell\}} \sum_{j=x_i+1}^{\lfloor x_i^* \rfloor} m_i(j) + (x_f^* - \lfloor x_f^* \rfloor) m_f(\lceil x_f^* \rceil) \right) \\
&= \frac{c_\ell}{m_\ell(x_\ell)} \left(m_\ell(x_\ell) + \sum_{i \in W(\mathbf{x}^*)} \sum_{j=x_i+1}^{\lfloor x_i^* \rfloor} m_i(j) + (x_f^* - \lfloor x_f^* \rfloor) m_f(\lceil x_f^* \rceil) \right) \\
&= \frac{c_\ell}{m_\ell(x_\ell)} \left(\text{OPT}_F^k(\mathbf{c}) - v(\mathbf{x}) + m_\ell(x_\ell) \right) > \frac{c_\ell}{m_\ell(x_\ell)} (1 - \alpha) \text{OPT}_F^k(\mathbf{c}).
\end{aligned}$$

The first inequality follows by the feasibility of \mathbf{x}^* , whereas the second inequality follows by the fact that $x_\ell \geq 1$ and, thus, every cost now has a smaller or equal coefficient than before. The next inequality is due to the marginal value-per-cost ordering that Algorithm 1 uses to build \mathbf{x}^* and SORT- $\hat{\varphi}$ -REJECT $_\alpha^k$ uses (in the reverse order) to obtain \mathbf{x} . Finally, the last inequality is due to Equation (7). The lemma follows by rearranging terms. \square

We now proceed to obtaining an upper bound on the payments that each agent receives for each level of service.

LEMMA 3.7. Fix any $\alpha \in (0, 1)$ and let \mathbf{c} be a cost profile such that $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \text{OPT}_F^k(\mathbf{c}_{-i^*})$. Moreover, let $i \in W(\mathbf{x}(\mathbf{c}))$. For $1 \leq j \leq x_i(\mathbf{c})$, it holds that

$$p_{ij}(\mathbf{c}_{-i}) \leq \frac{B}{1-\alpha} \cdot \frac{m_i(j)}{\text{OPT}_F^k(\mathbf{c}_{-i})}.$$

PROOF. Let $p \in Q_{ij}(\mathbf{c}_{-i})$, i.e., p is a cost declaration that guarantees agent i at least j levels of service, and let λ be the index of the agent with the smallest value-per-cost ratio in the allocation computed by SORT- $\hat{\varphi}$ -REJECT $_\alpha^k$ with input (p, \mathbf{c}_{-i}) (as opposed to ℓ when the input is \mathbf{c}). By the definition of p and λ , we have that $m_i(j)/p \geq m_\lambda(x_\lambda(p, \mathbf{c}_{-i}))/c_\lambda$. By rearranging and applying Lemma 3.6 for the profile (p, \mathbf{c}_{-i}) and λ , we get that

$$p \leq \frac{c_\lambda m_i(j)}{m_\lambda(x_\lambda(p, \mathbf{c}_{-i}))} < \frac{B}{1-\alpha} \cdot \frac{m_i(j)}{\text{OPT}_F^k(p, \mathbf{c}_{-i})} \leq \frac{B}{1-\alpha} \cdot \frac{m_i(j)}{\text{OPT}_F^k(\mathbf{c}_{-i})},$$

where the last inequality follows by the monotonicity of $\text{OPT}_F^k(\cdot)$ with respect to subinstances. This implies that

$$p_{ij}(\mathbf{c}_{-i}) = \sup(Q_{ij}(\mathbf{c}_{-i})) \leq \frac{B}{1-\alpha} \cdot \frac{m_i(j)}{\text{OPT}_F^k(\mathbf{c}_{-i})},$$

as claimed. \square

The final component needed for the proof of Lemma 3.4 is a lower bound on the optimal fractional objective when one agent is excluded.

LEMMA 3.8. Fix any $\alpha \in (0, 1)$ and let \mathbf{c} be a cost profile such that $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \text{OPT}_F^k(\mathbf{c}_{-i^*})$. For every agent $i \in N$, it holds that

$$\text{OPT}_F^k(\mathbf{c}_{-i}) > \frac{1-\alpha}{\alpha} (v(\mathbf{x}(\mathbf{c})) - m_\ell(x_\ell)).$$

PROOF. By the stopping condition of the **while** loop, we have that

$$\begin{aligned} v(\mathbf{x}(\mathbf{c})) - m_\ell(x_\ell) &< \alpha \text{OPT}_F^k(\mathbf{c}) \leq \alpha \left(\text{OPT}_F^k(\mathbf{c}_{-i}) + v_i(k) \right) \leq \alpha \text{OPT}_F^k(\mathbf{c}_{-i}) \left(1 + \frac{v_{i^*}(k)}{\text{OPT}_F^k(\mathbf{c}_{-i^*})} \right) \\ &< \alpha \text{OPT}_F^k(\mathbf{c}_{-i}) \left(1 + \frac{\alpha}{1-\alpha} \right) = \frac{\alpha}{1-\alpha} \text{OPT}_F^k(\mathbf{c}_{-i}). \end{aligned}$$

The third inequality follows by the definition of i^* , whereas the last inequality follows directly by the assumption of the lemma. \square

We finally present the proof of Lemma 3.4.

PROOF OF LEMMA 3.4. We can upper bound the total payments as follows:

$$\begin{aligned} \sum_{i=1}^n p_i(\mathbf{c}) &= \sum_{i \in W(\mathbf{x}(\mathbf{c}))} p_i(\mathbf{c}) = \sum_{i \in W(\mathbf{x}(\mathbf{c}))} \sum_{j=1}^{x_i(\mathbf{c})} p_{ij}(\mathbf{c}_{-i}) \\ &\leq \frac{B}{1-\alpha} \sum_{i \in W(\mathbf{x}(\mathbf{c}))} \sum_{j=1}^{x_i(\mathbf{c})} \frac{m_i(j)}{\text{OPT}_F^k(\mathbf{c}_{-i})} = \frac{B}{1-\alpha} \sum_{i \in W(\mathbf{x}(\mathbf{c}))} \frac{v_i(x_i(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-i})} \\ &= \frac{B}{1-\alpha} \left(\sum_{i \in W(\mathbf{x}(\mathbf{c})) \setminus \{\ell\}} \frac{v_i(x_i(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-i})} + \frac{v_\ell(x_\ell(\mathbf{c})) - m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} + \frac{m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} \right) \\ &< \frac{B}{1-\alpha} \left(\frac{\alpha}{1-\alpha} \cdot \frac{\sum_{i \in W(\mathbf{x}(\mathbf{c}))} v_i(x_i(\mathbf{c})) - m_\ell(x_\ell(\mathbf{c}))}{v(\mathbf{x}(\mathbf{c})) - m_\ell(x_\ell(\mathbf{c}))} + \frac{m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} \right) \\ &= \frac{B}{1-\alpha} \left(\frac{\alpha}{1-\alpha} + \frac{m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} \right). \end{aligned}$$

The first equality reflects the fact that $p_{ij}(\mathbf{c}_{-i}) = 0$ for $i \notin W(\mathbf{x}(\mathbf{c}))$, whereas the second equality is due to Lemma 3.3. The first inequality follows by applying Lemma 3.7 for every agent $i \in W(\mathbf{x}(\mathbf{c}))$ and every $j \in \{1 \dots, x_i(\mathbf{c})\}$. Finally, the second inequality follows by applying Lemma 3.8 to the denominators of all of the terms except from $m_\ell(x_\ell(\mathbf{c}))/\text{OPT}_F^k(\mathbf{c}_{-\ell})$. \square

3.3 SORT- \mathcal{S} -REJECT $_\alpha^k$ for Large Markets

We conclude this section with a note on the performance of SORT- \mathcal{S} -REJECT $_\alpha^k$ for instances in which no single agent can significantly impact the total value attainable by the entire market. This property is frequently observed in applications such as internet marketplaces and crowdsourcing environments. In the economics and computation literature, such markets are often referred to as *large* (see, e.g., [24] for one of the first works). The notion has also been studied in the context of budget-feasible mechanism design under multiple definitions (see, e.g., the work of Anari et al. [5] and Jalaly and Tardos [18]).

Given an instance $I = (N, \mathbf{c}, B, k, (v_i)_{i \in N})$ of the k -level model, let $\theta = \theta(I) := \max_{j \in N} m_j(1)/\text{OPT}^k(\mathbf{c})$ be the *largeness* of its underlying market. Intuitively, an instance I with a small θ implies that there is no agent in N whose first offered level of service is significantly valuable for the auctioneer when comparing with the highest value the auctioneer can attain under the optimal hiring scheme (under truthful declarations).

DEFINITION 3.1. A k -level instance $I = (N, \mathbf{c}, B, k, (v_i)_{i \in N})$ models a large market when $\theta \rightarrow 0$.

Definition 3.1 is a natural generalization to the k -level model of the value-based alternative large market assumption given by Anari et al. [5]. The main result here is Theorem 3.2, which implies that for large markets the approximation ratio of SORT- $\tilde{\mathcal{C}}$ -REJECT $_{\alpha}^k$ becomes 2.62 while the mechanism retains all of its good properties.

THEOREM 3.2. *SORT- $\tilde{\mathcal{C}}$ -REJECT $_{\alpha}^k$ with $\alpha = \frac{2(1-\theta)}{3+\sqrt{5+4\theta}}$ is a truthful, individually rational, budget-feasible $\frac{3+\sqrt{5+4\theta}}{2(1-\theta)}$ -approximation mechanism for instances of the k -level model with largeness θ , and runs in time polynomial in n and k .*

Before proving Theorem 3.2, we show a natural lower bound on the largeness θ of an instance I for which SORT- $\tilde{\mathcal{C}}$ -REJECT $_{\alpha}^k$ may allocate to more than one bidder. Lemma 3.9 will be useful to prove the budget feasibility of SORT- $\tilde{\mathcal{C}}$ -REJECT $_{\alpha}^k$ for θ -large instances.

LEMMA 3.9. *For every cost profile \mathbf{c} such that $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \text{OPT}_{\mathbf{F}}^k(\mathbf{c}_{-i^*})$ and every bidder $\ell \in W(\mathbf{x}(\mathbf{c}))$, it holds that $\theta \geq (1-\alpha)m_{\ell}(x_{\ell})/\text{OPT}_{\mathbf{F}}^k(\mathbf{c}_{-\ell})$.*

PROOF. Since $\text{OPT}_{\mathbf{F}}^k(\mathbf{c}) \leq \text{OPT}_{\mathbf{F}}^k(\mathbf{c}_{-\ell}) + v_{\ell}(k) \leq (1 + \frac{\alpha}{1-\alpha})\text{OPT}_{\mathbf{F}}^k(\mathbf{c}_{-\ell}) = \frac{1}{1-\alpha} \text{OPT}_{\mathbf{F}}^k(\mathbf{c}_{-\ell})$, we have that

$$\frac{m_{\ell}(x_{\ell})}{\text{OPT}_{\mathbf{F}}^k(\mathbf{c}_{-\ell})} \leq \frac{m_{\ell}(x_{\ell})}{(1-\alpha)\text{OPT}_{\mathbf{F}}^k(\mathbf{c})} \leq \frac{m_{\ell}(x_{\ell})}{(1-\alpha)\text{OPT}^k(\mathbf{c})} \leq \frac{m_{\ell}(1)}{(1-\alpha)\text{OPT}^k(\mathbf{c})} \leq \frac{\theta}{1-\alpha}.$$

The second inequality follows by the fact that the fractional problem is a relaxation of the integral problem and the third inequality is due to the concavity of v_{ℓ} , whereas the last inequality follows by the definition of θ . \square

PROOF OF THEOREM 3.2. As in Section 3, truthfulness and individual rationality follows by Lemma 3.1 and the fact that we determine the payments according to Theorem 2.1. Furthermore, the approximation guarantee follows by Lemma 3.2 since, here, $\alpha = \frac{2(1-\theta)}{3+\sqrt{5+4\theta}} = \frac{3-\sqrt{5+4\theta}}{2} < \frac{3-\sqrt{5}}{2}$.

What remains to be shown here is that SORT- $\tilde{\mathcal{C}}$ -REJECT $_{\alpha}^k$ is budget feasible for this configuration. Let \mathbf{c} be a cost profile for an instance with largeness θ . If $v_{i^*}(k) \geq \frac{\alpha}{1-\alpha} \text{OPT}_{\mathbf{F}}^k(\mathbf{c}_{-i^*})$, we work exactly as in the proof of Lemma 3.5. When $v_{i^*}(k) < \frac{\alpha}{1-\alpha} \text{OPT}_{\mathbf{F}}^k(\mathbf{c}_{-i^*})$, we resort to Lemma 3.4 and Lemma 3.9 and have that

$$\sum_{i \in W(\mathbf{x}(\mathbf{c}))} p_i(\mathbf{c}) \leq \frac{B}{1-\alpha} \left(\frac{m_{\ell}(x_{\ell}(\mathbf{c}))}{\text{OPT}_{\mathbf{F}}^k(\mathbf{c}_{-\ell})} + \frac{\alpha}{1-\alpha} \right) \leq \frac{B}{1-\alpha} \left(\frac{\theta}{1-\alpha} + \frac{\alpha}{1-\alpha} \right),$$

where the first inequality is obtained as in the proof of Lemma 3.5 and the second follows from Lemma 3.9. Budget feasibility follows by observing that for $\alpha = \frac{3-\sqrt{5+4\theta}}{2}$, it holds that $\theta + \alpha = (1-\alpha)^2$. \square

Note that the approximation guarantee of Theorem 3.2 improves almost linearly for small θ . For instance, for $0 \leq \theta \leq 0.04$ (i.e., in the case in which no single level of service is more than 4% of the optimal value), the approximation guarantee is roughly $2.618 + 3.17\theta$ (in the sense that, for this range of θ , we have that $|\frac{3+\sqrt{5+4\theta}}{2(1-\theta)} - (2.618 + 3.17\theta)| < 10^{-3}$). Moreover, the theorem bridges the gap between the threshold for budget feasibility of Lemma 3.5 (which forces the approximation ratio to be at least $2 + \sqrt{3}$) and the approximation guarantee bottleneck of Lemma 3.2.

4 Mechanisms for Divisible Agents

We consider the divisible agent model and derive two truthful, budget-feasible mechanisms with constant approximation guarantees. The first mechanism is obtained by discretizing the valuation

Mechanism 3: CHUNK- $\tilde{\mathcal{C}}$ -SOLVE

▷ **Input:** An instance $I = (N, \tilde{\mathbf{c}}, B, (\bar{v}_i)_{i \in N})$ of the divisible-agent model.

- 1 Initialize for each $i \in N$, $v_i : \{0, \dots, n\} \rightarrow \mathbb{R}_+$ with $v_i(0) = 0$.
- 2 **for** $i \in N$ and $j \in \{1, \dots, n\}$ **do**
- 3 Set $v_i(j) = \bar{v}_i(j/n)$.
- 4 **for** $i \in N$ **do**
- 5 Set $c_i = \tilde{c}_i/n$.
- 6 Let $J = (N, \mathbf{c}, B, n, (v_i)_{i \in N})$ be the resulting discretized instance of the k -level model.
 // Note that $k = n$, i.e., we construct an instance with n levels of service.
- 7 Set $\alpha = (2 + \sqrt{3})^{-1}$ and $k = n$, and compute \mathbf{x} by running SORT- $\tilde{\mathcal{C}}$ -REJECT $_{\alpha}^k$ on J .
- 8 **for** $i \in N$ **do**
- 9 Set $\bar{x}_i = x_i/n$.
- 10 Allocate $\bar{\mathbf{x}}$ and determine the payments $\bar{\mathbf{p}}$ according to (1).

functions and reducing the general case of the problem to the k -level model (Section 4.1). The second mechanism is an improved 2-approximation for the special case in which all agents have linear valuation functions (Section 4.2). Note that the $e/(e-1)$ lower bound of Anari et al. [5] holds even in the latter restricted setting, as the instance they construct has a linear valuation function.

4.1 A Constant Factor Approximation Using SORT- $\tilde{\mathcal{C}}$ -REJECT $_{\alpha}^k$

Recall that in the divisible-agent model, we have that $\mathbf{x} \in [0, 1]^n$ and concave non-decreasing valuation functions $\bar{v}_i : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ with $\bar{v}_i(0) = 0$ for all $i \in N$.

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 a constant approximation for the setting with divisible agents by applying this discretization for $k = n$ and then using SORT- $\tilde{\mathcal{C}}$ -REJECT $_{\alpha}^k$ from Section 3. Our CHUNK- $\tilde{\mathcal{C}}$ -SOLVE mechanism exploits this idea; a detailed description is given in Mechanism 3.

THEOREM 4.1. *CHUNK- $\tilde{\mathcal{C}}$ -SOLVE is a truthful, individually rational, budget-feasible $(4 + 2\sqrt{3})$ -approximation mechanism for instances of the divisible-agent model and runs in time polynomial in n .*

For notational convenience, throughout Section 4.1, we use $\tilde{\mathbf{c}}$ to denote cost profiles in the divisible-agent model and \mathbf{c} to denote cost profiles in the k -level model. The following lemma will be key in proving the approximation guarantee of the theorem.

LEMMA 4.1. *Let $I = (N, \tilde{\mathbf{c}}, B, (\bar{v}_i)_{i \in N})$ be an instance of the divisible agent model. Consider that the instance $J = (N, \mathbf{c}, B, n, (v_i)_{i \in N})$ as defined in line 6 of the mechanism CHUNK- $\tilde{\mathcal{C}}$ -SOLVE. It holds that $\text{OPT}_F(\tilde{\mathbf{c}}) \leq 2 \cdot \text{OPT}_F^n(\mathbf{c})$.*

PROOF. Let $\mathbf{x}^* = (x_1^*, \dots, x_n^*) \in [0, 1]^n$ be an optimal solution to the respective Fractional Concave Knapsack Problem for I (see (3) in Section 2.3). Let $z_i := \lfloor x_i^* n \rfloor$ for all $i \in N$, and $\mathbf{z} := (z_1, \dots, z_n)$. Define $\mathbf{1} = (1, 1, \dots, 1)$. We have that

$$\begin{aligned}
 \text{OPT}_F(\tilde{\mathbf{c}}) &= \sum_{i \in N} \bar{v}_i(x_i^*) \leq \sum_{i \in N} \bar{v}_i\left(\min\left\{\frac{z_i + 1}{n}, 1\right\}\right) = \sum_{i \in N} v_i(\min\{z_i + 1, n\}) \\
 &= \sum_{i \in N} v_i(\min\{z_i + 1, n\} - 1) + m_i(\min\{z_i + 1, n\}) \\
 &\leq \sum_{i \in N} v_i(z_i) + v_i(1) = v(\mathbf{z}) + v(\mathbf{1}).
 \end{aligned} \tag{8}$$

For the first inequality, we use the facts that $x_i^* \leq 1$ and $x_i^* n \leq \lfloor x_i^* n \rfloor + 1 = z_i + 1$ for all $i \in N$. The second equality follows by expressing \bar{v}_i in terms of v_i for each $i \in N$ (see line 3 of CHUNK- $\bar{\phi}$ -SOLVE). The final inequality is due to the monotonicity of the functions \bar{v}_i and v_i for each $i \in N$, as well as their concavity, which implies that $m_i(j) \leq m_i(1) \leq v_i(1)$ for any $j \in [n]$.

Next, we show that both \mathbf{z} and $\mathbf{1}$ are feasible solutions for J . For \mathbf{z} , it holds by construction that

$$\sum_{i \in N} c_i z_i = \sum_{i \in N} \frac{\bar{c}_i}{n} z_i = \sum_{i \in N} \frac{\bar{c}_i}{n} \lfloor x_i^* n \rfloor \leq \sum_{i \in N} \bar{c}_i x_i^* \leq B.$$

The last inequality follows from the fact that \mathbf{x}^* is a feasible solution for I as it is, by assumption, optimal.

Similarly, we can verify that $\mathbf{1}$ is a feasible solution for J by observing that

$$\sum_{i \in N} c_i = \sum_{i \in N} \frac{\bar{c}_i}{n} = \frac{\sum_{i \in N} \bar{c}_i}{n} \leq \frac{n \cdot \max_{i \in N} \bar{c}_i}{n} = \max_{i \in N} \bar{c}_i \leq B,$$

where the final inequality holds by the cost restriction assumption of the divisible agent model (see Section 2.3).

Since both \mathbf{z} and $\mathbf{1}$ are feasible solutions for J , we obtain that

$$v(\mathbf{z}) \leq \text{OPT}_F^n(\mathbf{c}) \quad \text{and} \quad v(\mathbf{1}) \leq \text{OPT}_F^n(\mathbf{c}).$$

By combining this with Equation (8), the lemma follows. \square

PROOF OF THEOREM 4.1. CHUNK- $\bar{\phi}$ -SOLVE reduces the divisible instance to an instance with $k = n$ multiple levels of service and uses SORT- $\bar{\phi}$ -REJECT $_{\alpha}^k$ parameterized with $\alpha = (2 + \sqrt{3})^{-1}$ to compute an allocation. Thus, the mechanism inherits the truthfulness and individual rationality properties through Lemma 3.1. Moreover, for the approximation guarantee, we have that

$$\bar{v}(\bar{\mathbf{x}}(\bar{\mathbf{c}})) = v(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}_F^n(\mathbf{c}) \geq \frac{\alpha}{2} \text{OPT}_F(\bar{\mathbf{c}}),$$

where the first inequality follows from Lemma 3.2 and the second inequality follows from Lemma 4.1.

What remains to be shown is that CHUNK- $\bar{\phi}$ -SOLVE is budget feasible. We will show that, for an instance $I = (N, \bar{\mathbf{c}}, B, (\bar{v}_i)_{i \in N})$ of the divisible-agent model, the payments prescribed by line 10 of CHUNK- $\bar{\phi}$ -SOLVE coincide with the payments made by SORT- $\bar{\phi}$ -REJECT $_{\alpha}^k$ for the instance $J = (N, \mathbf{c}, B, n, (v_i)_{i \in N})$, which is constructed in line 6 of CHUNK- $\bar{\phi}$ -SOLVE. Specifically, we will show that for each agent $i \in N$, it holds that $\bar{p}_i(\bar{\mathbf{c}}) = p_i(\mathbf{c})$. Let $w = f(x) = x/n$. Indeed, for each $i \in N$, we have that

$$\begin{aligned} \bar{p}_i(\bar{\mathbf{c}}) &= \bar{c}_i \cdot \bar{x}_i(\bar{\mathbf{c}}) + \int_{\bar{c}_i}^B \bar{x}_i(z, \bar{\mathbf{c}}_{-i}) dz = c_i n \cdot \frac{x_i(\mathbf{c})}{n} + \int_{c_i n}^B \frac{x_i(z/n, \mathbf{c}_{-i})}{n} dz \\ &= c_i x_i(\mathbf{c}) + \int_{c_i n}^B \frac{1}{n} \cdot x_i\left(\frac{z}{n}, \mathbf{c}_{-i}\right) dz = c_i x_i(\mathbf{c}) + \int_{c_i n}^B f'(z) \cdot x_i(f(z), \mathbf{c}_{-i}) dz \\ &= c_i x_i(\mathbf{c}) + \int_{f(c_i n)}^{f(B)} x_i(w, \mathbf{c}_{-i}) dw = c_i x_i(\mathbf{c}) + \int_{c_i}^{B/n} x_i(w, \mathbf{c}_{-i}) dw = p_i(\mathbf{c}). \end{aligned}$$

The first equality follows from applying Theorem 2.1 (Myerson's payment identity) to the allocation $\bar{\mathbf{x}}(\bar{\mathbf{c}})$ for I in CHUNK- $\bar{\phi}$ -SOLVE. The second equality follows from lines 5 and 9 since, by definition, $\bar{x}_i(\bar{\mathbf{b}}) = x_i(\frac{1}{n}\bar{\mathbf{b}})/n$ holds for all $\bar{\mathbf{b}} \in [0, B]^n$; here, $\bar{\mathbf{b}} = (z, \bar{\mathbf{c}}_{-i})$ and $\frac{1}{n}(z, \bar{\mathbf{c}}_{-i}) = (z/n, \mathbf{c}_{-i})$. The subsequent four equalities are obtained through simple calculus, involving a change of the

Mechanism 4: PRUNING (by Gravin et al. [14])

▷ **Input:** $I = (N, \bar{\mathbf{c}}, B, (v_i)_{i \in N})$ with the agents relabeled so that $\frac{v_1}{c_1} \geq \dots \geq \frac{v_n}{c_n}$

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

variable of the integral term. Finally, the last equality holds because SORT- $\hat{\mathcal{C}}$ -REJECT $_{\alpha}^k$ also prescribes payments according to Theorem 2.1 for instance J . Thus, as SORT- $\hat{\mathcal{C}}$ -REJECT $_{\alpha}^k$ is budget feasible by Lemma 3.5, budget feasibility of CHUNK- $\hat{\mathcal{C}}$ -SOLVE follows.

We conclude the proof of the theorem by noting that CHUNK- $\hat{\mathcal{C}}$ -SOLVE runs in time polynomial in n . Recall that, by Theorem 3.1, Mechanism SORT- $\hat{\mathcal{C}}$ -REJECT $_{\alpha}^k$ runs in time polynomial in n and k . Since Mechanism CHUNK- $\hat{\mathcal{C}}$ -SOLVE, given an instance I with n agents, constructs an instance of the n -level model, i.e., $k = n$, its run time is polynomial in n . \square

4.2 An Improved Mechanism for Linear Valuation Functions

The best known approximation guarantee for the special case in which all divisible agents have *linear* valuation functions is $(3 + \sqrt{5})/2 \approx 2.62$, by Klumper and Schäfer [19]. Below, we improve upon this by presenting a simple budget-feasible 2-approximation mechanism for this setting, thereby establishing that the optimal approximation ratio for this problem lies in the interval $[e/(e-1), 2]$. Our mechanism is inspired by the randomized 2-approximation mechanism by Gravin et al. [14] for indivisible agents. To avoid unnecessarily heavy notation in this special case, let $v_i := \bar{v}_i(1)$ and, thus, $\bar{v}_i(x) = v_i x$, for all $i \in N$.

4.2.1 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-approximation mechanism and their randomized 2-approximation 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 whole-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 4.

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 4.2 (IMPLIED BY LEMMA 3.1 OF [14]). *Let \mathbf{c} be a cost profile. Consider an agent $i \in N$ with $\bar{x}_i(\mathbf{c}) = 1$. Then, for all \mathbf{c}'_i such that $\bar{x}_i(\mathbf{c}'_i, \mathbf{c}_{-i}) = 1$, it holds that $\bar{\mathbf{x}}(\mathbf{c}'_i, \mathbf{c}_{-i}) = \bar{\mathbf{x}}(\mathbf{c})$ and $r(\mathbf{c}'_i, \mathbf{c}_{-i}) = r(\mathbf{c})$.*

Given this robustness property, PRUNING can be used as a first filtering step to discard inefficient agents, followed by a subsequent allocation scheme that takes $(r(\mathbf{c}), \bar{\mathbf{x}}(\mathbf{c}))$ as input. If the subsequent

Mechanism 5: PRUNE- \hat{c} -ASSIGN

-
- ▷ **Input:** $I = (N, \bar{c}, B, (v_i)_{i \in N})$ 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{x}) by running PRUNING for profile c .
 - 2 Let $S = \{i \in N \mid \bar{x}_i = 1\}$, $i^* \in \arg \max\{v_i \mid i \in S\}$, $T = S \setminus \{i^*\}$.
 - 3 Determine the fractional allocations $x_i(c_i)$, $\forall i \in S$, as in Equation (11), and set $x_i(c_i) = 0$, $\forall i \in N \setminus S$.
 - 4 Allocate x and determine the payments p according to Equation (1).
-

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{x}) be the output of PRUNING for a cost profile c . Given \bar{x} , we define S as the set of agents that are provisionally allocated, i^* as the highest value agent in S (where ties are broken lexicographically), and T as the set of remaining agents. Formally,

$$S = \{i \in N \mid \bar{x}_i = 1\}, \quad i^* \in \arg \max\{v_i \mid i \in S\}, \quad \text{and} \quad T = S \setminus \{i^*\}. \quad (9)$$

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

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

LEMMA 4.3 (LEMMA 3.2 OF [14]). *Given a profile c , let (r, \bar{x}) be the output of PRUNING. Let $S = T \cup \{i^*\}$ be defined as in Equation (9) with respect to \bar{x} . We have that*

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

4.2.2 Phase 2: The Final Allocation Scheme. Our mechanism combines the PRUNING mechanism above with the allocation scheme defined in Equation (11) below. We refer to the resulting mechanism as PRUNE- \hat{c} -ASSIGN (see Mechanism 5).

First, we need to define the following constants:

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

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

$$\frac{rB}{2} = q_{i^*} v_{i^*} + (1 - q_{i^*} - q) v(T). \quad (10)$$

For given r and S , as is the case in line 3 of PRUNE- \hat{c} -ASSIGN, we are going to define our (fractional) allocation function $x_i(c)$, for $i \in S$, so that it only depends on agent i 's cost c_i . Thus, we will slightly abuse notation and write $x_i(c_i)$ in what follows. For each agent $i \in S$, let

$$x_i(z) = q_i + \frac{v_i - rz}{2v_i} \quad \text{for } z \in \left[0, \frac{v_i}{r}\right]. \quad (11)$$

It is not hard to verify that $0 \leq x_i(z) \leq 1$ given the chosen parameters q_{i^*} , q_T , and q above. Further, by property (i) of Lemma 4.3, the cost c_i of each agent $i \in S$ is at most v_i/r and, therefore, $x_i(c_i)$ is well defined.

THEOREM 4.2. *PRUNE- \hat{c} -ASSIGN is a truthful, individually rational, budget-feasible 2-approximation mechanism for instances of the divisible-agent model with linear valuation functions, and runs in time polynomial in n .*

Given the lower bound of $1 + \sqrt{2} \approx 2.41$ by Chen et al. [11] in this setting with indivisible agents, Theorem 4.2 establishes a separation between the indivisible agent model and the divisible agent model with linear valuation functions. The theorem follows by the three lemmata below.

LEMMA 4.4. *PRUNE- \hat{c} -ASSIGN is monotone.*

PROOF. As argued above, PRUNING is monotone and it suffices to show that the allocation scheme in Equation (11) is monotone. Fix an arbitrary agent i and a cost profile \mathbf{c} . Let $c'_i < c_i$. We need to prove that $x_i(c'_i, \mathbf{c}_{-i}) \geq x_i(\mathbf{c})$.

Let $\bar{\mathbf{x}}(\mathbf{c})$ be the provisional allocation obtained from PRUNING. If $\bar{x}_i(\mathbf{c}) = 0$, the claim follows trivially. Assume that $\bar{x}_i(\mathbf{c}) = 1$. Because PRUNING is monotone, we have that $\bar{x}_i(c'_i, \mathbf{c}_{-i}) = 1$. Also, because PRUNING is robust, the output remains the same, i.e., $(\bar{\mathbf{x}}(\mathbf{c}), r(\mathbf{c})) = (\bar{\mathbf{x}}(c'_i, \mathbf{c}_{-i}), r(c'_i, \mathbf{c}_{-i}))$. This implies that the respective rates that determine the fractional allocation in Equation (11) are the same. The claim now follows by observing that the allocation functions in (11) are monotone non-increasing in z . \square

LEMMA 4.5. *PRUNE- \hat{c} -ASSIGN has an approximation ratio of 2.*

PROOF. By property (iii) of Lemma 4.3, we have that

$$\text{OPT}_F \leq v(S) + r \cdot (B - c(S)) = \sum_{i \in S} (v_i - rc_i) + rB.$$

By the definition in Equation (11), we have for every $i \in T \cup \{i^*\}$ that

$$x_i(z) = q_i + \frac{v_i - rz}{2v_i} \quad \text{or, equivalently,} \quad v_i - rz = 2v_i(x_i(z) - q_i).$$

Combining these two inequalities, we obtain that

$$\text{OPT}_F \leq 2 \sum_{i \in S} v_i x_i(c_i) - 2 \sum_{i \in S} v_i q_i + rB = 2v(\mathbf{x}) - 2(q_{i^*} v_{i^*} + (1 - q_{i^*} - q)v(T)) + rB = 2v(\mathbf{x}(\mathbf{c})),$$

where, for the last equality, we exploit the identity in Equation (10). This concludes the proof. \square

LEMMA 4.6. *PRUNE- \hat{c} -ASSIGN is budget feasible.*

PROOF. Given a cost profile \mathbf{c} , the payment of agent $i \in S$ is defined as

$$p_i(\mathbf{c}) = c_i x_i(\mathbf{c}) + \int_{c_i}^B x_i(z, \mathbf{c}_{-i}) dz = c_i x_i(c_i) + \int_{c_i}^{v_i/r} x_i(z) dz \leq \int_0^{v_i/r} x_i(z) dz.$$

Here, the second equality holds because PRUNING ensures that $x_i(z, \mathbf{c}_{-i}) = 0$ for $z > v_i/r$. The inequality holds because $x_i(\cdot)$ is strictly decreasing in $[0, c_i]$.

Consider an agent $i \in T \cup \{i^*\}$. We have that

$$\int_0^{v_i/r} x_i(z) dz = \int_0^{v_i/r} \left(q_i + \frac{v_i - rz}{2v_i} \right) dz = \frac{v_i}{r} \left(q_i + \frac{1}{2} \right) - \frac{r}{2v_i} \int_0^{v_i/r} z dz = \frac{v_i}{r} \left(q_i + \frac{1}{4} \right).$$

Summing over all agents, we obtain that the total payment is at most

$$\sum_{i=1}^n p_i(\mathbf{c}) = \sum_{i \in S} p_i(\mathbf{c}) \leq \sum_{i \in S} \frac{v_i}{r} \left(q_i + \frac{1}{4} \right) = \frac{1}{r} (q_{i^*} v_{i^*} + (1 - q_{i^*} - q)v(T)) + \frac{v(S)}{4r} = \frac{B}{2} + \frac{v(S)}{4r},$$

where the last equality uses the identity in Equation (10).

We conclude the proof by showing that $v(S)/4r \leq B/2$ or, equivalently, $rB \geq v(S)/2$. Starting with Equation (10) and applying the definition of q_{i^*} , we have that

$$\frac{rB}{2} = (q_{i^*}v_{i^*} + (1 - q_{i^*} - q)v(T)) = \frac{1}{2}(v_{i^*} + v(T)) - q \cdot \min\{v_{i^*}, v(T)\} \geq \frac{1}{2}v(S) - qv(T).$$

The claim follows by observing that $v(T) \leq rB$ by property (ii) of Lemma 4.3 and $q \leq \frac{1}{2}$ by definition. \square

We conclude the section by showing that our analysis of PRUNE- $\hat{\mathcal{C}}$ -ASSIGN is tight. Indeed, consider an instance with 2 agents with $v_1 = v_2 = 1$, $c_1 = \epsilon$, $c_2 = 1 - \epsilon$, with $\epsilon \in (0, 1)$, and the budget of the auctioneer being $B = 1$. The optimal hiring scheme is $\mathbf{x}^* = (1, 1)$ and, therefore, $v(\mathbf{x}^*(\mathbf{c})) = 2$. Consider now the outcome of PRUNE- $\hat{\mathcal{C}}$ -ASSIGN for this instance. Initially, PRUNING does not discard either bidder since the **while** condition for $r = 1/B$ evaluates to *False*. Thus, $S = \{1, 2\}$ and $r = 1/B$. Then, it is easy to observe that the final hiring scheme of PRUNE- $\hat{\mathcal{C}}$ -ASSIGN prescribes an allocation of $x_1(0) = x_2(1) = 1/2$ (by viewing bidder 1 as i^*). Therefore, $v(\mathbf{x}(\mathbf{c})) = 1$, which matches the guarantee of Theorem 4.2.

5 Conclusion and Future Work

In this work, we revisited two budget-feasible mechanism design settings in which partial allocations are allowed and draw clear connections between them. Under the mild assumption of being able to afford each agent entirely, 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, such as matching, matroid, or even polymatroid (for the k -level setting) constraints. For the k -level setting, it would also 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. At the same time, it would be interesting to obtain an improved approximation guarantee for concave valuation functions in the divisible-agents model. A possible avenue here could be to come up with a more nuanced discretization procedure than that of CHUNK- $\hat{\mathcal{C}}$ -SOLVE. Finally, as far as simple settings are concerned, the most important open problem is still the indivisible agents case with additive valuation functions, 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 4.2). Any progress on these fronts may give rise to novel techniques, which could further be used for problems in richer environments.

References

- [1] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2016. Coverage, matching, and beyond: New results on budgeted mechanism design. In *Proceedings of Web and Internet Economics: 12th International Conference (WINE'16)*. 414–428.
- [2] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2017. On budget-feasible mechanism design for symmetric submodular objectives. In *Proceedings of Web and Internet Economics: 13th International Conference (WINE'17)*. Springer, 1–15.
- [3] Georgios Amanatidis, Pieter Kleer, and Guido Schäfer. 2019. Budget-feasible mechanism design for non-monotone submodular objectives: Offline and online. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC'19)*. 901–919.
- [4] Georgios Amanatidis, Sophie Klumper, Evangelos Markakis, Guido Schäfer, and Artem Tsikirisidis. 2023. Partial allocations in budget-feasible mechanism design: Bridging multiple levels of service and divisible agents. In *Proceedings of the 9th International Conference on Web and Internet Economics, WINE'23*, Jugal Garg, Max Klimm, and Yuqing Kong (Eds.). *Lecture Notes in Computer Science*, Vol. 14,413. Springer-Verlag, Berlin, 41–58.

- [5] Nima Anari, Gagan Goel, and Afshin Nikzad. 2018. Budget feasible procurement auctions. *Operations Research* 66, 3 (2018), 637–652.
- [6] Krzysztof R. Apt and Jan Heering. 2022. Characterization of incentive compatible single-parameter mechanisms revisited. *The Journal of Mechanism and Institution Design* 7, 1 (2022), 113–129.
- [7] Aaron Archer and Éva Tardos. 2001. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*. IEEE, 482–491.
- [8] Eric Balkanski, Pranav Garimidi, Vasilis Gkatzelis, Daniel Schoepflin, and Xizhi Tan. 2022. Deterministic budget-feasible clock auctions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2022*. 2940–2963.
- [9] Xiaohui Bei, Ning Chen, Nick Gravin, and Pinyan Lu. 2017. Worst-case mechanism design via bayesian analysis. *SIAM J. Comput.* 46, 4 (2017), 1428–1448.
- [10] Hau Chan and Jing Chen. 2014. Truthful multi-unit procurements with budgets. In *Web and Internet Economics: 10th International Conference, WINE 2014, Proceedings*. 89–105.
- [11] Ning Chen, Nick Gravin, and Pinyan Lu. 2011. On the approximability of budget feasible mechanisms. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA'11)*. 685–699.
- [12] Shahar Dobzinski, Christos H. Papadimitriou, and Yaron Singer. 2011. Mechanisms for complement-free procurement. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC'11)*. 273–282.
- [13] Gagan Goel, Afshin Nikzad, and Adish Singla. 2014. Allocating tasks to workers with matching constraints: Truthful mechanisms for crowdsourcing markets. In *Proceedings of the 23rd International Conference on World Wide Web (WWW'14)*. 279–280.
- [14] Nick Gravin, Yaonan Jin, Pinyan Lu, and Chenhao Zhang. 2020. Optimal budget-feasible mechanisms for additive valuations. *ACM Transactions on Economics and Computation (TEAC)* 8, 4 (2020), 1–15.
- [15] Kai Han, You Wu, He Huang, and Shuang Cui. 2023. Triple eagle: Simple, fast and practical budget-feasible mechanisms. In *37th Conference on Neural Information Processing Systems*. 33894–33911.
- [16] Dorit S. Hochbaum. 1995. A nonlinear knapsack problem. *Operations Research Letters* 17, 3 (1995), 103–110.
- [17] He Huang, Kai Han, Shuang Cui, and Jing Tang. 2023. Randomized pricing with deferred acceptance for revenue maximization with submodular objectives. In *Proceedings of the ACM Web Conference 2023*. 3530–3540.
- [18] Pooya Jalaly and Éva Tardos. 2018. Simple and efficient budget feasible mechanisms for monotone submodular valuations. In *Proceedings of Web and Internet Economics: 14th International Conference (WINE'18)*. 246–263.
- [19] Sophie Klumper and Guido Schäfer. 2022. Budget feasible mechanisms for procurement auctions with divisible agents. In *Proceedings of Algorithmic Game Theory: 15th International Symposium (SAGT'22)*. 78–93.
- [20] Stefano Leonardi, Gianpiero Monaco, Piotr Sankowski, and Qiang Zhang. 2017. Budget feasible mechanisms on matroids. In *Proceedings of Integer Programming and Combinatorial Optimization: 19th International Conference (IPCO'17)*. 368–379.
- [21] Minming Li, Chenhao Wang, and Mengqi Zhang. 2022. Budget feasible mechanisms for facility location games with strategic facilities. *Autonomous Agents and Multi-Agent Systems* 36, 2 (2022), 35.
- [22] Xiang Liu, Hau Chan, Minming Li, and Weiwei Wu. 2024. Budget feasible mechanisms: A survey. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI'24)*, Kate Larson (Ed.). International Joint Conferences on Artificial Intelligence Organization, 8132–8141. Survey Track.
- [23] Silvano Martello and Paolo Toth. 1990. *Knapsack Problems: Algorithms and Computer Implementations*. John Wiley & Sons, Inc.
- [24] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54, 5 (2007), 22–es.
- [25] Roger Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- [26] Rian Neogi, Kanstantsin Pashkovich, and Chaitanya Swamy. 2024. Budget-feasible mechanism design: Simpler, better mechanisms and general payment constraints. In *15th Innovations in Theoretical Computer Science Conference (ITCS'24)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik. 84:1–84:22.
- [27] Aviad Rubinfeld and Junyao Zhao. 2023. Beyond worst-case budget-feasible mechanism design. In *14th Innovations in Theoretical Computer Science Conference (ITCS'23)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [28] Yaron Singer. 2010. Budget feasible mechanisms. In *Proceedings of the 51st Annual Symposium on Foundations of Computer Science (FOCS'10)*. 765–774.

Appendices

A GREEDY-BEST-IN^k: Almost Tight for the Best-in Setting

As described in Section 2.2, throughout this work we consider the k -level model under the *all-in setting*, i.e., given a cost profile \mathbf{c} , for every agent $i \in N$ it holds that $k \cdot c_i \leq B$. However, the k -level model has also been studied in the *best-in setting* (see [10]), where the weaker assumption that $c_i \leq B$ for all $i \in N$ is assumed instead. In this appendix, we show that the truthful and budget-feasible mechanism SORT- $\tilde{\mathcal{C}}$ -REJECT $_{\alpha}^k$ (Mechanism 2 defined in Section 3.2), requires just a few modifications to achieve an approximation guarantee that is almost tight for this setting. More precisely, for k levels of service and concave valuation functions, we re-parameterize our mechanism SORT- $\tilde{\mathcal{C}}$ -REJECT $_{\alpha}^k$ and obtain a guarantee of $k + 2 + o(1)$, almost matching the known lower bound of k of Chan and Chen [10] (see also Remark 2.1 in Section 2.2). Note that Theorem A.1 directly shows that the lower bounds of Chan and Chen [10] cannot be generalized for instances with multiple agents and remain linear or logarithmic to the total number of levels of service across all agents. In that sense, their upper bound is close to being tight for a constant number of agents, but its gap with the lower bound grows as an unbounded function of n . Instead, the guarantee of our Theorem A.1 is almost best possible for any n .

For presentation purposes, we restate this parameterized version of the mechanism separately in this section, “rebrand” it as GREEDY-BEST-IN^k and analyze it in a mostly self-contained manner.

Mechanism 6: GREEDY-BEST-IN^k

```

▷ Input: A profile  $\mathbf{c}$ .
1 Set  $\alpha = (3 + k - \sqrt{9 + 2k + k^2}) / 2k$  and  $\beta = (1 - 2\alpha) / (\alpha k + 1)$ .
2 Set  $i^* \in \arg \max_{i \in N} v_i(1) / \text{OPT}_F^k(\mathbf{c}_{-i})$ . // we break ties lexicographically with respect to the agents' indices
3 if  $v_{i^*}(1) \geq \beta \cdot \text{OPT}_F^k(\mathbf{c}_{-i^*})$  then
4   | Set  $x_{i^*} = 1$  and  $x_i = 0$  for all  $i \in N \setminus \{i^*\}$ .
5 else
6   | Compute an optimal fractional allocation  $\mathbf{x}^*(\mathbf{c})$  using Algorithm 1. // i.e.,  $\mathbf{x}^*(\mathbf{c})$  is almost integral
7   | Initialize  $\mathbf{x} = (\lfloor x_1^*(\mathbf{c}) \rfloor, \dots, \lfloor x_n^*(\mathbf{c}) \rfloor)$ . // by the construction of  $\mathbf{x}^*(\mathbf{c})$  and the best-in assumption,  $\mathbf{x} \neq \mathbf{0}$ 
8   | for  $i \in W(\mathbf{x})$  do // recall that  $W(\mathbf{x}) = \{i \in N \mid x_i > 0\}$ 
9     |   | for  $j \in \{1, \dots, x_i\}$  do
10      |   |   | Add the marginal value-per-cost ratio  $m_i(j)/c_i$  to a list  $\mathcal{L}$ .
11   | Sort  $\mathcal{L}$  in non-increasing order and let  $\ell$  be the index of the last agent of  $W(\mathbf{x})$  in  $\mathcal{L}$ .
12   | while  $v(\mathbf{x}) - m_{\ell}(x_{\ell}) \geq \alpha \text{OPT}_F^k(\mathbf{c})$  do
13     |   | Set  $x_{\ell} = x_{\ell} - 1$ .
14     |   | Remove the last element from  $\mathcal{L}$  and update  $\ell$ .
15 Allocate  $\mathbf{x}$  and determine the payments  $\mathbf{p}$  according to (1).
```

THEOREM A.1. *Mechanism GREEDY-BEST-IN^k is a truthful, individually rational, budget-feasible $(k + 2 + o(1))$ -approximation mechanism for instances of the k -level model and the best-in setting, and runs in time polynomial in n and k .*

Intuitively, the fact that the mechanism may allocate a single level of service of an agent i^* while excluding everyone else—even if more levels of i^* could be afforded at a first glance—makes the **if** part of GREEDY-BEST-IN^k significantly weaker than that of SORT- $\tilde{\mathcal{C}}$ -REJECT $_{\alpha}^k$. However, this is needed to ensure budget-feasibility.

To prove Theorem A.1, similar arguments to those of Section 3.2 are used, albeit slightly adapted for the best-in environment. We first present two auxiliary lemmata, one lemma for the proof of the approximation guarantee and one lemma that will prove to be useful for showing budget feasibility for the best-in setting.

LEMMA A.1. *For the allocation $\mathbf{x}(\mathbf{c})$ computed by GREEDY-BEST-IN^k with input the cost profile \mathbf{c} , it holds that $v(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}_F^k(\mathbf{c})$, with $\alpha = (3 + k - \sqrt{9 + 2k + k^2}) / 2k$.*

PROOF. Similar to Lemma 3.2, for a cost profile \mathbf{c} we will prove the claimed guarantee against the optimal solution to the fractional relaxation of the bounded knapsack instance, i.e., we will show that $v(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}_F^k(\mathbf{c})$. We distinguish the following two cases.

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

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

where the last inequality follows by the concavity of $v_{i^*}(\cdot)$. Rearranging terms leads to

$$v_{i^*}(1) \geq \frac{\beta}{1 + \beta k} \text{OPT}_F^k(\mathbf{c}) = \alpha \text{OPT}_F^k(\mathbf{c}),$$

which concludes the case since $v_{i^*}(1) = v(\mathbf{x}(\mathbf{c}))$.

Case 2: $v_{i^*}(1) < \beta \cdot \text{OPT}_F^k(\mathbf{c}_{-i^*})$. In this case, we immediately obtain that $v(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}_F^k(\mathbf{c})$ whenever the **while** loop runs at least once. We argue that $v(\mathbf{x}(\mathbf{c})) \geq \alpha \text{OPT}_F^k(\mathbf{c})$ still holds for our choice of α and β when the **while** condition on line 12 always evaluates to *False*. Toward a contradiction, suppose that this is not the case. For the best-in setting, it is true that, similar to the proof of Lemma 3.2, for the fractionally included agent f it holds that

$$\frac{m_f(\lceil x_f^* \rceil)}{\text{OPT}_F^k(\mathbf{c})} \leq \frac{m_f(1)}{\text{OPT}_F^k(\mathbf{c})} \leq \frac{v_f(1)}{\text{OPT}_F^k(\mathbf{c}_{-f})} \leq \frac{v_{i^*}(1)}{\text{OPT}_F^k(\mathbf{c}_{-i^*})} < \beta, \quad (12)$$

where the first inequality follows from the concavity of each valuation function, the third inequality holds by the choice of i^* , and the last one is the very definition of Case 2 for this setting.

At the same time, Equation (5) holds regardless of whether we are in the best-in or the all-in setting.

By combining inequalities (5) and (12), we obtain that α must be such that $\beta > 1 - \alpha$. However, it is a matter of simple calculations to verify that this is not possible for any $k > 0$ for our choice of α and β , leading to a contradiction. \square

LEMMA A.2. *Let \mathbf{c} be a cost profile so that $v_{i^*}(1) < \beta \text{OPT}_F^k(\mathbf{c}_{-i^*})$ in line 3 of GREEDY-BEST-IN^k. Then, for every $i \in W(\mathbf{x}(\mathbf{c}))$ it holds that*

$$\text{OPT}_F^k(\mathbf{c}) \leq (1 + \beta k) \text{OPT}_F^k(\mathbf{c}_{-i}).$$

PROOF. Indeed, by concavity of $v_i(\cdot)$, the definition of i^* in GREEDY-BEST-IN^k and the fact that $v_{i^*}(1) < \beta \text{OPT}_F^k(\mathbf{c}_{-i^*})$, we obtain that

$$\begin{aligned} \text{OPT}_F^k(\mathbf{c}) &\leq \text{OPT}_F^k(\mathbf{c}_{-i}) + v_i(k) \leq \text{OPT}_F^k(\mathbf{c}_{-i}) + kv_{i^*}(1) \\ &\leq \text{OPT}_F^k(\mathbf{c}_{-i}) \left(1 + \frac{kv_{i^*}(1)}{\text{OPT}_F^k(\mathbf{c}_{-i^*})} \right) \leq (1 + \beta k) \text{OPT}_F^k(\mathbf{c}_{-i}). \end{aligned} \quad \square$$

We are now ready to prove Theorem A.1.

PROOF OF THEOREM A.1. Observe that truthfulness and individual rationality of GREEDY-BEST-IN^k follow by the monotonicity of the allocation rule, as proven by Lemma 3.1, and the fact that we determine the payments according to Equation (1). To prove that GREEDY-BEST-IN^k achieves the claimed guarantee, it is sufficient to observe that the guarantee of $1/\alpha$ we proved in Lemma A.1 is upper bounded by $k + 2 + o(k)$. Indeed, $\lim_{k \rightarrow \infty} 1/\alpha = k + 2$.

What remains to be proved is the budget feasibility of GREEDY-BEST-IN^k. To do that, we distinguish the following cases.

Case 1: $v_{i^*}(1) \geq \beta \text{OPT}_F^k(\mathbf{c}_{-i^*})$. In this case, agent i^* is the only agent hired and for 1 level of service. By Theorem 2.1, we obtain that $p_{i^*}(\mathbf{c}) = p_{i^*1}(\mathbf{c}_{-i^*}) \leq B$ and budget feasibility follows by directly invoking the best-in assumption, i.e., $c_{i^*} \leq B$, for all c_{i^*} .

Case 2: $v_{i^*}(1) < \beta \text{OPT}_F^k(\mathbf{c}_{-i^*})$. As in the proof of Lemma 3.5, we have that

$$\begin{aligned}
 \sum_{i=1}^n p_i(\mathbf{c}) &= \sum_{i \in W(\mathbf{x}(\mathbf{c}))} p_i(\mathbf{c}) = \sum_{i \in W(\mathbf{x}(\mathbf{c}))} \sum_{j=1}^{x_i(\mathbf{c})} p_{ij}(\mathbf{c}_{-i}) \\
 &\leq \frac{B}{1-\alpha} \sum_{i \in W(\mathbf{x}(\mathbf{c}))} \sum_{j=1}^{x_i(\mathbf{c})} \frac{m_i(j)}{\text{OPT}_F^k(\mathbf{c}_{-i})} = \frac{B}{1-\alpha} \sum_{i \in W(\mathbf{x}(\mathbf{c}))} \frac{v_i(x_i(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-i})} \\
 &= \frac{B}{1-\alpha} \left(\sum_{i \in W(\mathbf{x}(\mathbf{c})) \setminus \{\ell\}} \frac{v_i(x_i(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-i})} + \frac{v_\ell(x_\ell(\mathbf{c})) - m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} + \frac{m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} \right) \\
 &\leq \frac{B}{1-\alpha} \left((1 + \beta k) \frac{\sum_{i \in W(\mathbf{x}(\mathbf{c}))} v_i(x_i(\mathbf{c})) - m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c})} + \frac{m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} \right) \\
 &\leq \frac{B}{1-\alpha} \left(\alpha(1 + \beta k) \frac{\sum_{i \in W(\mathbf{x}(\mathbf{c}))} v_i(x_i(\mathbf{c})) - m_\ell(x_\ell(\mathbf{c}))}{v(\mathbf{x}(\mathbf{c})) - m_\ell(x_\ell(\mathbf{c}))} + \frac{m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} \right) \\
 &= \frac{B}{1-\alpha} \left(\alpha(1 + \beta k) + \frac{m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} \right) \leq \frac{B}{1-\alpha} (\alpha(1 + \beta k) + \beta) .
 \end{aligned}$$

The first equality reflects the fact that $p_{ij}(\mathbf{c}_{-i}) = 0$ for $i \notin W(\mathbf{x}(\mathbf{c}))$, whereas the second equality is due to an adaptation of Lemma 3.3 for the best-in setting.⁴ The first inequality follows by applying Lemma 3.7 for every agent $i \in W(\mathbf{x}(\mathbf{c}))$ and every $j \in \{1, \dots, x_i(\mathbf{c})\}$. Then, the second inequality is due to Lemma A.2, whereas the third inequality is a consequence of the **while** condition of GREEDY-BEST-IN^k on line 12. Finally, the last inequality is due to the fact that

$$\frac{m_\ell(x_\ell(\mathbf{c}))}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} \leq \frac{v_\ell(1)}{\text{OPT}_F^k(\mathbf{c}_{-\ell})} \leq \frac{v_{i^*}(1)}{\text{OPT}_F^k(\mathbf{c}_{-i^*})} < \beta .$$

Budget feasibility follows by observing that since $\beta = \frac{1-2\alpha}{\alpha k+1}$, it holds that

$$\frac{\alpha(1 + \beta k) + \beta}{1-\alpha} = \frac{\alpha}{1-\alpha} + \beta \cdot \frac{\alpha k + 1}{1-\alpha} = \frac{\alpha}{1-\alpha} + \frac{1-2\alpha}{1-\alpha} = 1.$$

This concludes the proof. □

⁴The definition of the critical payments for the best-in setting is completely analogous to the one for SORT- ϕ -REJECT $_{\alpha}^k$ in Section 3.2 for the all-in setting.

B Efficient Computation of Payments

In this section, we show how the payments returned by $\text{SORT-}\hat{\varphi}\text{-REJECT}_\alpha^k$ can be computed in polynomial time. For ease of reference, we repeat the necessary notation here. Given an instance $I = (N, c, B, k, (v_i)_{i \in N})$ of the k -level model, $\mathbf{x}(c) \in \{0, \dots, k\}^n$ is the allocation returned by $\text{SORT-}\hat{\varphi}\text{-REJECT}_\alpha^k$. According to line 14 of the mechanism, the payments $\mathbf{p}(c) = (p_1(c), \dots, p_n(c))$ are computed as defined in Equation (1) of Theorem 2.1.

Recall from Section 3.2 that for each $i \in W(\mathbf{x}(c))$ and $j = 1, \dots, x_i(c)$, $Q_{ij}(\mathbf{c}_{-i})$ is the set of all points $q \in [c_i, B/k]$ satisfying $\lim_{z \rightarrow q^-} x_i(z, \mathbf{c}_{-i}) \geq j$ and $\lim_{z \rightarrow q^+} x_i(z, \mathbf{c}_{-i}) \leq j$. Furthermore, recall that $p_{ij}(\mathbf{c}_{-i}) = \sup(Q_{ij}(\mathbf{c}_{-i}))$, the critical payment of the j -th level of service of agent i . By Lemma 3.3, we can express the total payment of each agent $i \in W(\mathbf{x}(c))$ as the sum of these per-level critical payments, i.e.,

$$p_i(c) = \sum_{r=1}^{x_i(c)} p_{ir}(\mathbf{c}_{-i}).$$

Now, for every $i, j \in N$, let

$$\rho_i(c, j) = v_j(k) \cdot \frac{\text{OPT}_F^k(\mathbf{c}_{-i})}{\text{OPT}_F^k(\mathbf{c}_{-j})}.$$

For a profile c , we use $i^*(c)$ to denote i^* as computed in line 1 of $\text{SORT-}\hat{\varphi}\text{-REJECT}_\alpha^k$. We distinguish the following two cases for any given agent $i \in W(\mathbf{x})$.

Case 1: $i = i^*(c)$ and $v_i(k) \geq \frac{\alpha}{1-\alpha} \text{OPT}_F^k(\mathbf{c}_{-i})$. By line 2, the **if** part of $\text{SORT-}\hat{\varphi}\text{-REJECT}_\alpha^k$ is executed and we have that $x_i(c) = k$. Moreover, it is not hard to see that $p_{i1}(\mathbf{c}_{-i}) = \dots = p_{ik}(\mathbf{c}_{-i})$, as if agent i were to unilaterally declare a cost $c'_i \in [c_i, B/k]$, they would still be hired for k levels of service, as long as $i = i^*(c'_i, \mathbf{c}_{-i})$. Otherwise, whenever there exists a $c'_i \in [c_i, B/k]$ for which $i \neq i^*(c'_i, \mathbf{c}_{-i})$, we have that the **if** part of $\text{SORT-}\hat{\varphi}\text{-REJECT}_\alpha^k$ is still executed as $v_i(k) \leq \rho_i((c'_i, \mathbf{c}_{-i}), i^*(c'_i, \mathbf{c}_{-i}))$, and $x_i(c'_i, \mathbf{c}_{-i}) = 0$.⁵ Therefore, for $r = 1, \dots, x_i(c)$, we can write

$$p_{ir}(\mathbf{c}_{-i}) = \inf \left(\left\{ \frac{B}{k} \right\} \cup \left\{ c'_i \in [c_i, \frac{B}{k}] \mid i \neq i^*(c'_i, \mathbf{c}_{-i}) \right\} \right).$$

If there is a $c'_i \in [c_i, B/k]$ such that $i \neq i^*(c'_i, \mathbf{c}_{-i}) = j$, we can compute it in polynomial time. To compute this c'_i , consider the non-increasing ordering of marginal value-per-cost ratios for the agents $N \setminus \{j\}$ with respect to $(\mathbf{c}_{-\{j,i\}}, c_i)$. We start by computing the *costs of interest*. These are all the costs $b \in [c_i, B/k]$ for which there is a change in this ordering due to the unilateral deviation of agent i to b . There are at most $(n-1)k^2$ such points in the interval $[c_i, B/k]$. For notational convenience, define

$$\theta_i(j, \mathbf{c}_{-i}) := v_i(j) \frac{\text{OPT}_F^k(\mathbf{c}_{-i})}{v_i(k)}.$$

Starting with the smallest cost of interest, denoted by c_i^1 , in this interval, one can check whether $\text{OPT}_F^k(\mathbf{c}_{-\{j,i\}}, c_i^1) \leq \theta_i(j, \mathbf{c}_{-i})$. If this is true, one can find the point c'_i where equality is reached in polynomial time. This is either at $c_i^1 = c'_i$, or in $[c_i, c_i^1)$. For the latter, this point can be found by ordering the marginal values for $N \setminus \{j\}$ with respect to $(\mathbf{c}_{-\{j,i\}}, c_i)$ and finding the point where the first s marginal values and a fraction of the $(s+1)$ -th marginal value add up to $\theta_i(j, \mathbf{c}_{-i})$. One can then find c'_i such that this allocation is exactly budget feasible. If $\text{OPT}_F^k(\mathbf{c}_{-\{j,i\}}, c_i^1) > \theta_i(j, \mathbf{c}_{-i})$, continue the same procedure with c_i^2 and c_i^1 instead of c_i^1 and c_i , respectively.

Case 2: $v_{i^*(c)}(k) < \frac{\alpha}{1-\alpha} \text{OPT}_F^k(\mathbf{c}_{-i^*(c)})$. In this case, the **else** part of $\text{SORT-}\hat{\varphi}\text{-REJECT}_\alpha^k$ is executed. Recall the definition of $\ell(\mathbf{x}(c))$ as the index in $W(\mathbf{x}(c))$ of the agent hired for the least efficient level

⁵Note that this is the only possible alternative outcome. This is implied by Lemma 3.1.

of service in $\mathbf{x}(\mathbf{c})$. For notational convenience, we define

$$\sigma(\mathbf{c}) := \frac{m_{\ell(\mathbf{x}(\mathbf{c}))}(x_i(\mathbf{c}))}{c_{\ell(\mathbf{x}(\mathbf{c}))}}.$$

Fix a level of service $r \in [x_i(\mathbf{c})]$ for agent i . By the definition of $\ell(\mathbf{x}(\mathbf{c}))$, we have that $\frac{m_i(r)}{c_i} \geq \sigma(\mathbf{c})$.

Intuitively, there are two events under which agent i will no longer be hired for the agent's r -th level of service when declaring a unilateral deviation $c'_i \in [c_i, B/k]$. The first event occurs when, under (c'_i, \mathbf{c}_{-i}) , the r -th level of service is no longer efficient enough to be selected by SORT- $\hat{\mathcal{C}}$ -REJECT $_{\alpha}^k$. Formally, c'_i is such that $\frac{m_i(r)}{c'_i} < \sigma(c'_i, \mathbf{c}_{-i})$. The second event occurs when, under (c'_i, \mathbf{c}_{-i}) , the **if** condition in line 2 becomes **true**, meaning:

$$v_{i^*(c'_i, \mathbf{c}_{-i})}(k) \geq \frac{\alpha}{1 - \alpha} \text{OPT}_{\text{F}}^k(c'_i, \mathbf{c}_{-\{i, i^*(\mathbf{c})\}})$$

and agent i 's allocation under (c'_i, \mathbf{c}_{-i}) becomes 0.⁶

Thus, the critical payment for agent i for the r -th level of service is the smallest value in $[c_i, B/k]$ at which at least one of these two events is triggered. If this is not the case, then the critical payment for this level is simply B/k . Formally:

$$\begin{aligned} p_{ir}(\mathbf{c}_{-i}) &= \inf \left(\left(\frac{B}{k} \right) \cup \left\{ c'_i \in \left[c_i, \frac{B}{k} \right] \mid \frac{m_i(r)}{c'_i} < \sigma(c'_i, \mathbf{c}_{-i}) \vee v_{i^*(c'_i, \mathbf{c}_{-i})}(k) \right. \right. \\ &\quad \left. \left. \geq \frac{\alpha}{1 - \alpha} \text{OPT}_{\text{F}}^k(c'_i, \mathbf{c}_{-\{i, i^*(\mathbf{c})\}}) \right\} \right). \end{aligned}$$

We now describe how the critical payment $p_{ir}(\mathbf{c}_{-i})$ is computed by focusing on each event separately.

To compute the smallest cost such that $m_i(r)/c'_i < \sigma(c'_i, \mathbf{c}_{-i})$ if such a cost exists, it is sufficient to consider the points $c'_i \in [c_i, B/k]$ for which the ratio $m_i(r)/c'_i$ equals the marginal value-per-cost ratio of less efficient levels of service according to the ordering of Algorithm 1, which are at most nk . These will be the *costs of interest* for this event. Consider these costs of interest in increasing order, and check whether for any of these costs the **while** condition in line 11 evaluates to **true**. Otherwise, suppose that this happens for a cost \hat{c}_i and let \bar{c}_i be the cost considered before \hat{c}_i (if this happens for the first cost of interest, then $\bar{c}_i = c_i$). Then,

$$p_{ir}(\mathbf{c}_{-i}) = \inf \left(\{ \hat{c}_i \} \cup \left\{ c'_i \in [\bar{c}_i, \hat{c}_i] \mid (v(\mathbf{x}(\mathbf{c}_{-i}, \bar{c}_i)) - m_i(r)) \frac{1}{\alpha} = \text{OPT}_{\text{F}}^k(\mathbf{c}_{-i}, c'_i) \right\} \right),$$

which can be computed in polynomial time. Note that in order to compute a $c'_i \in [\bar{c}_i, \hat{c}_i]$ in the set above (if it exists), one should again look into the costs for which the ordering of marginal value-per-cost ratios of Algorithm 1 changes.

For the second event, we need to find the smallest $c'_i \in [c_i, B/K]$ such that the **if** condition in line 2 is **true** for $i^*(c'_i, \mathbf{c}_{-i})$ if such a cost exists. This can be done similarly to Case 1. We can check for each agent $j \in N \setminus \{i\}$ whether there exists such a c'_i for which the **if** condition in line 2 is **true** and then take the smallest c'_i . The difference is that instead of comparing with $\theta_i(j, \mathbf{c}_{-i})$, we compare with $v_j(k)(1 - \alpha)/\alpha$.

Received 22 January 2024; revised 17 February 2025; accepted 25 February 2025

⁶Note that by the definition of this case, it cannot be that $i = i^*(c'_i, \mathbf{c}_{-i})$.