

Optimal Budget-Feasible Mechanisms for Additive Valuations

NICK GRAVIN, Shanghai University of Finance and Economics, China

YAONAN JIN, Columbia University

PINYAN LU, Shanghai University of Finance and Economics, China

CHENHAO ZHANG, Northwestern University, Illinois

In this article, we show a tight approximation guarantee for budget-feasible mechanisms with an additive buyer. We propose a new simple randomized mechanism with approximation ratio of 2, improving the previous best known result of 3. Our bound is tight with respect to either the optimal offline benchmark or its fractional relaxation. We also present a simple deterministic mechanism with the tight approximation guarantee of 3 against the fractional optimum, improving the best known result of $(2 + \sqrt{2})$ for the weaker integral benchmark.

CCS Concepts: • **Theory of computation** → **Computational pricing and auctions; Algorithmic mechanism design;**

Additional Key Words and Phrases: Budget-feasible auction, mechanism design, approximation

ACM Reference format:

Nick Gravin, Yaonan Jin, Pinyan Lu, and Chenhao Zhang. 2020. Optimal Budget-Feasible Mechanisms for Additive Valuations. *ACM Trans. Econ. Comput.* 8, 4, Article 21 (October 2020), 15 pages.

<https://doi.org/10.1145/3417746>

1 INTRODUCTION

In a typical procurement setting, a buyer wants to purchase items from a set A of agents. Each agent $i \in A$ can supply an item (or provide a service) at an incurred cost of c_i to himself, and the buyer wants to optimize his valuation for the set of acquired items taking into account the costs of items. Because the agents may strategically report their costs, this setting is usually considered as a truthful mechanism design problem.

A preliminary version of this article appeared at the 2019 ACM EC conference.

Nick Gravin and Pinyan Lu are supported by Science and Technology Innovation 2030—“New Generation of Artificial Intelligence” Major Project No. (2018AAA0100903), NSFC grant 61922052 and 61932002, Innovation Program of Shanghai Municipal Education Commission, Program for Innovative Research Team of Shanghai University of Finance and Economics (IRTSHUFE), and the Fundamental Research Funds for the Central Universities.

Yaonan Jin is supported by NSF IIS-1838154, NSF CCF-1703925, NSF CCF-1814873, and NSF CCF-1563155.

Authors’ addresses: N. Gravin and P. Lu, Shanghai University of Finance and Economics, 318 Wuchuan Rd, Wujiaochang, Yangpu District, Shanghai, China, 200083; emails: {nikolai, lu.pinyan}@mail.shufe.edu.cn; Y. Jin, Columbia University, New York City, New York, 10025; email: jin.yaonan@columbia.edu; C. Zhang, Northwestern University, 633 Clark St, Evanston, Illinois, 60201; email: chenhao.zhang.rea@u.northwestern.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

2167-8375/2020/10-ART21 \$15.00

<https://doi.org/10.1145/3417746>

These problems have been extensively studied by the AGT community. The earlier work analyzed the case where the buyer's valuation takes 0–1 values (see, e.g., Ref. [5]) in the frugality framework, with the objective of payment minimization. A more recent line of work on the budget-feasible mechanism design (see, e.g., Ref. [24]) studies more general valuation functions with a budget constraint of B on the buyer's total payment. Our work belongs to the latter category.

Research in the budget-feasible framework focuses on different classes of complement-free valuations (ranging from the class of additive valuations to the most general class of subadditive valuations), and has many applications such as procurement in crowdsourcing markets [26], experimental design [18], and advertising in social networks [25]. The central problem for these online labor markets is to properly price each task. The budget feasibility mechanism design model is a very reasonable model that naturally captures the budget limitation on the buyer and also uncertainty about workers costs.

This setting corresponds to the most basic additive valuation of the buyer, which is the topic of our paper. That is, we assume that every hired worker $i \in W$ generates a value of $v_i \geq 0$ to the buyer, whose total valuation from all the hired workers W is equal to $v(W) = \sum_{i \in W} v_i$. Without any incentive constraints, this naturally defines the Knapsack optimization problem:

$$\text{Find workers } S \subseteq A: \quad \max_{S \subseteq A} v(S) = \sum_{i \in S} v_i, \text{ subject to } \sum_{i \in S} c_i \leq B.$$

In the budget-feasible framework, the goal is to design truthful direct-revelation mechanisms¹ that decide (1) which workers $W \subseteq A$ to select and (2) how much to pay them under the budget constraint. A mechanism is evaluated against the benchmark of the optimal solution to the Knapsack problem. Over all possible choices of the value v_i 's and the cost c_i 's, the worst-case multiplicative gap between the outcome $v(W)$ and the optimal Knapsack solution is called the *approximation ratio* of this mechanism.

For the above problem with an additive buyer, Singer [24] gave the first 5-approximation mechanism. Later, the result was improved by Chen et al. [11] via a $(2 + \sqrt{2})$ -approximation deterministic mechanism and a 3-approximation randomized mechanism, which still remain the best known upper bounds for the problem for nearly a decade. Further, the best known lower bounds are $(\sqrt{2} + 1)$ for the deterministic mechanisms and 2 for the randomized ones [11]. Thus, there are gaps for both the deterministic mechanisms $[\sqrt{2} + 1, 2 + \sqrt{2}]$ and the randomized ones [2, 3]. Since these two intervals intersect, it is even unclear whether the best randomized mechanism is indeed better than the best deterministic one.

Also, for the above problem with an additive buyer, Anari et al. [4] studied an important special case of large markets (i.e., the setting where each worker has vanishingly small cost compared to the buyer's budget) and acquired the tight bound of $\frac{e}{e-1}$.

Fractional Knapsack. Interestingly, all previous work on budget-feasible mechanisms for an additive buyer actually obtained results against the stronger benchmark of the optimal solution to Fractional Knapsack, i.e., the fractional relaxation of the Knapsack problem. (Nonetheless, the lower bounds apply to the Knapsack benchmark instead of the Fractional Knapsack benchmark.) Indeed, although Knapsack is a well-known NP-hard problem, its fractional relaxation admits an efficient solution by a simple greedy algorithm, and generally has much better behavior than the integral optimum. We also compare the performance of our mechanisms to the Fractional Knapsack benchmark.

¹Typically, there are no assumptions in the literature about the prior distribution of the agents' costs. The truthfulness condition means that the strategy of reporting the true cost is ex post a dominant and individually rational strategy for every single agent.

Our Results. We propose two natural mechanisms that both achieve tight guarantees against the Fractional Knapsack benchmark. Namely, we prove a 3-approximation guarantee for a deterministic mechanism and a 2-approximation guarantee for a randomized one. Given the matching lower bound of 2 even against the weaker Knapsack benchmark, the guarantee from our randomized mechanism is also tight against the standard benchmark. Our results establish a clear separation between the respective power of randomized and deterministic mechanisms: no deterministic mechanism has an approximation guarantee better than $(\sqrt{2} + 1)$, whereas our randomized mechanism already achieves a 2-approximation.

Concretely, we propose a new natural design principle of two-stage mechanisms. In the first stage, we greedily exclude the items with low *value-per-cost* ratios.² Then, in the second stage, we leverage the simple *posted-price* schemes based on the values of the remaining items. Both of our randomized and deterministic mechanisms share the first stage, which stops earlier than its analogues from the previous work. A remarkable property of the first stage, which we call *pruning* (similar to the pruning approach in the frugality literature [10, 20]) is that it can be *composed* (in the sense of Ref. [1]) with any truthful follow-up mechanism that runs on the items left to the second stage. The difference between our randomized and deterministic mechanisms lies in the follow-up posted-price schemes—the randomized mechanism uses non-adaptive posted prices with the total sum below the budget, whereas our deterministic mechanism employs adaptive pricing that depends on whether the previous agents accepted or rejected their posted-price offers.

Intuition behind our mechanism. The pruning stage of both mechanisms allows the buyer to reduce the choice complexity, and gives a reasonable upper bound on the payment to each remaining agent. The value of the fractional optimum never decreases too much, especially when the individual true cost c_i of each remaining agent is a non-negligible fraction of the budget B . We prove that the fractional optimum drops at most by a factor of two after the pruning stage for an arbitrary set of values and costs.

The idea behind the pruning stage is that the removed agents can be safely ignored by the mechanism, since the remaining items suffice to get the desired approximation to the fractional optimum. Moreover, the mechanism should naturally prefer the items with higher value-per-cost ratios. Our pruning process is based on the value-per-cost ratio, and works specifically for an additive-valuation buyer. That is, it is still unknown how to extend such a pruning stage to more general classes of valuation functions.

The second stage of our randomized mechanism draws a random vector of budget-feasible posted prices. This is the same type of the mechanism as was used by Bei et al. [9] to establish the tight approximation ratio of 2 for a subadditive buyer in the promise version of the problem (i.e., where the buyer is ensured to have a budget higher than the total cost of all items). Their result holds in the Bayesian setting, which by the *minimax principle* implies the existence of a randomized posted-price mechanism with the same approximation ratio in the worst-case setting. In our problem with an additive buyer, we explicitly construct a desired distribution over the posted-price vectors. Such posted-price schemes seem to be useful and easily adaptable to more general classes of valuation functions.

1.1 Related Work

A complementary concept of budget-feasible mechanism design is *frugality*, for which the objective is payment minimization under the feasibility constraint on the set of winning agents. In that

²This is essentially the main approach used in the previous work, had we continued until the remaining items (as a whole) become budget-feasible.

framework, there is a rich literature studying different systems of feasible sets, including matroid set systems [19], path and k -paths auctions [5, 10, 12, 15, 27], vertex cover, and k -vertex cover [14, 17, 20].

The framework of budget-feasible mechanism design was proposed by Singer [24]. Beyond additive valuations, other more general classes of complement-free valuations also have been considered in the literature:

$$\text{submodular} \subset \text{fractionally subadditive} \subset \text{subadditive}.$$

Singer gave an 112-approximation mechanism for submodular valuations [24]. This bound was improved to 7.91 and 8.34, respectively, for the randomized and deterministic mechanisms by Chen et al. [11], and then to 4 and 5 by Jalaly and Tardos in Ref. [21]. For fractionally subadditive valuations, Bei et al. [9] gave a 768-approximation randomized mechanism. For subadditive valuations, Dobzinski et al. [13] first gave an $O(\log^2 n)$ -approximation randomized mechanism and an $O(\log^3 n)$ -approximation deterministic mechanism. Later, Bei et al. [9] showed the existence of an $O(1)$ -approximation mechanism in this most general setting. Nonetheless, an explicit description of such a mechanism is still unknown.

There also have been many interesting and practically motivated adjustments to the original budget feasibility model. In particular, Anari et al. [4] investigated the variant with the additional large market assumption (namely, every agent has a negligible cost compared to the whole budget) and attained the tight result of $\frac{e}{e-1}$ for an additive buyer. Leonardi et al. [22] explored an additive-valuation model where the winning agents must form an independent set from a matroid. Amanatidis et al. [2, 3] investigated the variants with several important subclasses of submodular and fractionally subadditive valuations. Badanidiyuru et al. [6] studied the family of online pricing mechanisms in the budget feasibility model, motivated by practical restrictions given by the existing platforms. Balkanski and Hartline [7] obtained improved guarantees in the Bayesian framework. Goel et al. [16] concerned more complex scenarios on a crowdsourcing platform, where the buyer hires the workers to complete more than one task. Balkanski and Singer [8] considered fair mechanisms (instead of truthful mechanisms) in the budget feasibility model.

2 PRELIMINARIES

In the procurement auction, there are n items for sale, each held by a single agent $i \in [n]$ with a privately known cost $c_i \geq 0$ and a publicly known value $v_i > 0$ for the buyer. The buyer has an *additive* valuation function $v(A) = \sum_{i \in A} v_i$ for purchasing a subset $A \subseteq [n]$ of items. Due to the revelation principle, we only consider direct-revelation mechanisms. Upon receiving *bids* $\mathbf{b} = (b_i)_{i=1}^n$ of the claimed costs from the agents, a mechanism determines a set $W \subseteq [n]$ of *winning* agents and the *payments* $\mathbf{p} = (p_i)_{i=1}^n$ to the agents.

In the budget feasibility model, a *deterministic* mechanism \mathcal{M} is specified by an allocation function $\mathbf{x}(\mathbf{b}) : \mathbb{R}_+^n \rightarrow \{0, 1\}^n$ (thus, the winning set $W \stackrel{\text{def}}{=} \{i \in [n] \mid x_i(\mathbf{b}) = 1\}$) and a payment function $\mathbf{p}(\mathbf{b}) : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$. We use the notation b_i to denote the i -th entry of the bid vector \mathbf{b} , and the notation \mathbf{b}_{-i} the bid vector without bidder $i \in [n]$. We are interested in those truthful mechanisms that satisfy the following properties for any $\mathbf{b} = (b_i)_{i=1}^n$ and any $\mathbf{c} = (c_i)_{i=1}^n$.

- *Individual rationality*: $p_i(\mathbf{b}) \geq c_i$ and, thus, $u_i(\mathbf{b}) = p_i(\mathbf{b}) - c_i \geq 0$ for every $i \in W$, while $p_i(\mathbf{b}) \geq 0$ and, thus, $u_i(\mathbf{b}) = p_i(\mathbf{b}) \geq 0$ for every $i \notin W$. Namely, every agent $i \in [n]$ gets a non-negative utility.
- *Budget feasibility*: the total payment $\sum_{i \in W} p_i(\mathbf{b})$ is capped with a given budget $B \in \mathbb{R}_+$.
- *Truthfulness*: every agent $i \in [n]$ maximizes his utility when he bids the true cost $b_i = c_i$; namely, $u_i(c_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i})$ for any c_i and any $\mathbf{b} = (b_i, \mathbf{b}_{-i})$.

It is well known (see Ref. [23]) that truthfulness holds if and only if: (1) the allocation function $x_i(b_i, \mathbf{b}_{-i})$ is monotone in bid b_i , i.e., each winning agent $i \in W$ keeps winning when he unilaterally claims a lower bid $b_i \leq c_i$; and (2) the payment $p_i(\mathbf{b})$ to each winning agent $i \in W$ is the threshold/maximum bid for him to keep winning, i.e., $p_i(\mathbf{b}) = \sup\{b_i \in \mathbb{R}_+ \mid x_i(b_i, \mathbf{b}_{-i}) = 1\}$.

In general, a mechanism can have *randomized* allocation and payment rules. We restrict our attention to the mechanisms that can be described as a probability distribution over truthful deterministic mechanisms. Namely, any realization of such a randomized mechanism is some deterministic truthful mechanism that satisfies the above properties. A randomized mechanism of this type is called a *universally truthful* mechanism. We notice that most of the previous work on budget feasible mechanism only studies universally truthful mechanisms.

We denote by alg the value $\sum_{i \in W} v_i$ derived from a deterministic mechanism, or the expected value $\mathbb{E}\left[\sum_{i \in W} v_i\right]$ in case of a randomized mechanism. W.l.o.g., we assume $c_i \leq B$ for each agent $i \in [n]$, since this agent cannot win when $c_i > B$ (due to the individual rationality and the budget feasibility constraints). If the buyer knows the private costs $\mathbf{c} = (c_i)_{i=1}^n$ of the items, he would simply select the subset of items with the maximum total value, under the budget constraint. Let opt denote the optimal solution to this Knapsack problem:

$$\text{opt} \stackrel{\text{def}}{=} \max_{(x_i)_{i=1}^n \in \{0,1\}^n} \sum_{i=1}^n x_i \cdot v_i, \quad \text{subject to} \quad \sum_{i=1}^n x_i \cdot c_i \leq B. \quad (\text{Knapsack})$$

We also consider the fractional relaxation of the problem, and define its optimum as

$$\text{fopt} \stackrel{\text{def}}{=} \max_{(x_i)_{i=1}^n \in [0,1]^n} \sum_{i=1}^n x_i \cdot v_i, \quad \text{subject to} \quad \sum_{i=1}^n x_i \cdot c_i \leq B. \quad (\text{Fractional Knapsack})$$

Although opt is NP-hard to calculate, finding fopt is easy: one greedily and divisibly takes the items in the decreasing order of their value-per-cost ratios,³ until the budget is exhausted or no item is left. Under our assumption that $c_i \leq B$ for all $i \in [n]$, we have $1 \leq \frac{\text{fopt}}{\text{opt}} \leq 2$.⁴

We say that a mechanism achieves an α -approximation against the benchmark opt , if under whatever values $\mathbf{v} = (v_i)_{i=1}^n$ and costs $\mathbf{c} = (c_i)_{i=1}^n$, the outcome value alg is at least an $\frac{1}{\alpha}$ -fraction of the Knapsack solution opt . In what follows, we usually evaluate a mechanism against the stronger benchmark fopt , i.e., the solution to the Fractional Knapsack problem.

$$\alpha \leq \max_{\mathbf{v}, \mathbf{c}, B} \frac{\text{opt}}{\text{alg}} \quad \Leftrightarrow \quad \alpha \leq \max_{\mathbf{v}, \mathbf{c}, B} \frac{\text{fopt}}{\text{alg}}.$$

3 COMPOSITION OF MECHANISMS: PRUNING

Every mechanism presented in this work can be described as a composition of two stages. In particular, all of our mechanisms share the same first stage, called PRUNING-MECHANISM (see Figure 1), which serves to exclude the items with low value-per-cost ratios.

In Step (2b), if there are multiple items such that $\frac{v_k}{c_k} \leq r$, we discard them *one by one* in lexicographical order, and stop discarding items once the While-Loop meets the Stop-Condition. The output set $S(r)$ is always nonempty, since the Stop-Condition of the While-Loop is violated when $S(r)$ contains only one item.

PRUNING-MECHANISM possesses a remarkable *composability property*: the combination of it with any truthful follow-up mechanism \mathcal{M} running on the remaining items $i \in S(r)$ is still a truthful

³Namely, the decreasing order $(\sigma_i)_{i=1}^n$ is a permutation of $[n]$ such that $\frac{v_{\sigma_1}}{c_{\sigma_1}} \geq \frac{v_{\sigma_2}}{c_{\sigma_2}} \geq \dots \geq \frac{v_{\sigma_n}}{c_{\sigma_n}}$.

⁴Without this assumption, the gap between the two optima $\frac{\text{fopt}}{\text{opt}}$ can be arbitrarily large.

PRUNING-MECHANISM

- (1) Let $r \stackrel{\text{def}}{=} \frac{1}{B} \cdot \max\{v_i \mid i \in [n]\}$ and $S(r) \stackrel{\text{def}}{=} \{i \in [n] \mid \frac{v_i}{c_i} \geq r\}$
- (2) While $rB < v(S(r)) - \max\{v_i \mid i \in S(r)\}$ do:
 - (a) Continuously increase ratio r
 - (b) If $\frac{v_k}{c_k} \leq r$, then discard item k : $S(r) \leftarrow S(r) \setminus \{k\}$
- (3) Return pair $(r, S(r))$

Fig. 1. The first stage, PRUNING-MECHANISM, shared by all of our mechanisms.

MECHANISM-TEMPLATE

- (0) Receive the pair $(r, S(r))$ from PRUNING-MECHANISM
- (1) Run mechanism \mathcal{M} on the set $S(r)$:
 - (a) Select the winning set \overline{W} from $S(r)$ according to \mathcal{M}
 - (b) Cap the payment with $\frac{v_i}{r}$, i.e., $\overline{p}_i \stackrel{\text{def}}{=} \min\{p_i, \frac{v_i}{r}\}$, for each $i \in \overline{W}$

Fig. 2. The template of a composition mechanism $\overline{\mathcal{M}}$.

mechanism. More concretely (see Figure 2), the composition mechanism $\overline{\mathcal{M}} = (\overline{x}, \overline{p})$ of PRUNING-MECHANISM with a follow-up mechanism $\mathcal{M} = (x, p)$ works as follows:

LEMMA 3.1 (COMPOSABILITY). *If a follow-up mechanism \mathcal{M} is individually rational, budget-feasible, and truthful, then so is the composition mechanism $\overline{\mathcal{M}}$.*

PROOF. By Step (1b) of PRUNING-MECHANISM, every item $i \in S(r)$ has a value-per-cost ratio at least r , which means $c_i \leq \frac{v_i}{r}$. Thus, capping the payment with $\frac{v_i}{r}$ does not break the individual rationality. The follow-up mechanism \mathcal{M} itself is budget-feasible, and the composition mechanism $\overline{\mathcal{M}}$ can only reduce the payment for a winning item. Given these, we are left to show the truthfulness of $\overline{\mathcal{M}}$.

We claim that no winning item $i \in S(r)$ may change the output of PRUNING-MECHANISM by manipulating its bid to c'_i , unless this item gets excluded from $S(r)$ because of a too high bid c'_i . Indeed, suppose item i is still winning with the bid c'_i ; then, item i was never removed from the set $S(r)$, i.e., $\frac{v_i}{c'_i} \geq r$ at all times in the While-Loop of the PRUNING-MECHANISM. Given that item i stays in the set $S(r)$, the Stop-Condition of the While-Loop and the order in which we discard other items do not depend on the exact bid c'_i of item i .

Since the follow-up mechanism \mathcal{M} has a *monotone* allocation rule, so does the composition mechanism $\overline{\mathcal{M}}$. Regarding a losing item $i \notin \overline{W}$ (i.e., item i loses in $\overline{\mathcal{M}}$ when it bids truthfully), reporting a higher bid $c'_i > c_i$ does not help this item to pass the PRUNING-MECHANISM stage. As we discussed above, suppose that item i passes the PRUNING-MECHANISM stage by bidding $c'_i > c_i$, namely, $i \in S'(r')$; then, the two outcomes of PRUNING-MECHANISM under the two bids c'_i and c_i must be the same, namely $(r', S'(r')) = (r, S(r))$. In other words, when item i reports the true cost c_i , it passes the PRUNING-MECHANISM stage as well, but then loses in the follow-up mechanism \mathcal{M} . Given that the follow-up mechanism \mathcal{M} is truthful and runs on the same pair $(r', S'(r')) = (r, S(r))$ in both scenarios, item i will lose again in the follow-up mechanism \mathcal{M} , when it reports the higher bid $c'_i > c_i$.

The payment $\overline{p}_i = \min\{p_i, \frac{v_i}{r}\}$ of the composition mechanism $\overline{\mathcal{M}}$ is exactly the *threshold bid* for an item $i \in \overline{W}$ to keep winning: (1) passing the PRUNING-MECHANISM stage requires a bid of at least $\frac{v_i}{r}$; and (2) winning in the follow-up mechanism \mathcal{M} (after passing the PRUNING-MECHANISM stage) requires a bid of at least p_i .

In addition, a winning item $i \in \overline{W}$ cannot improve its utility by reporting a lower bid $c'_i < c_i$. As mentioned, when this winning item bids a lower $c'_i < c_i$, the PRUNING-MECHANISM returns the same pair $(r', S'(r')) = (r, S(r))$. Since the follow-up mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ is truthful (i.e., a monotone allocation rule and a threshold-based payment rule), item i gets the same payment $p'_i = p_i$ under either bid c'_i or c_i . The composition mechanism $\overline{\mathcal{M}}$ thus has the same payment $\overline{p}'_i = \min\{p'_i, \frac{v_i}{r'}\} = \min\{p_i, \frac{v_i}{r}\} = \overline{p}_i$ in both scenarios.

This completes the proof of Lemma 3.1. \square

We show now several useful properties of the output $(r, S(r))$ of PRUNING-MECHANISM.

LEMMA 3.2 (PRUNING MECHANISM). *Let $i^* \in \arg \max\{v_i \mid i \in S(r)\}$ denote the highest-value item or one of the highest-value items,⁵ and let $T \stackrel{\text{def}}{=} S(r) \setminus \{i^*\}$. Then, the following hold:*

- (a) $c_i \leq \frac{v_i}{r} \leq B$ for each item $i \in S(r)$.
- (b) $v(T) \leq rB < v(S(r))$.
- (c) $\text{fopt} \leq v(S(r)) + r \cdot (B - c(S(r))) < 2 \cdot v(S(r))$.

PROOF. **Property (a).** The first inequality follows from Step (1b) of PRUNING-MECHANISM; the second inequality holds, since the ratio r is initialized to be $\frac{1}{B} \cdot \max\{v_i \mid i \in [n]\}$ and keeps increasing during the While-Loop.

Property (b). We observe that the first inequality is a reformulation of the Stop-Condition of the While-Loop. To prove the second inequality, we note that there are two possibilities that can lead to the termination of the While-Loop, and $rB < v(S(r))$ holds in both cases.

- [Increase of ratio r]. Continuous increase of r implies $rB = v(T) < v(S(r))$.
- [Discard of an item k]. Value-per-cost ratio r is fixed before and after the discard. Before the discard, in that Stop-Condition has not been invoked,

$$rB < v(S(r)) + v_k - \max\{v_{i^*}, v_k\} \leq v(S(r)).$$

Property (c). The second inequality follows from Property (b). We show the first inequality based on case analysis. Let $\mathbf{x} = (x_i)_{i=1}^n$ denote the solution to the Fractional Knapsack problem. We have either $S(r) \subseteq \{i \in [n] \mid x_i = 1\}$ or $S(r) \supseteq \{i \in [n] \mid x_i > 0\}$. This claim holds since: (1) PRUNING-MECHANISM discards the items in increasing order of the value-per-cost ratios; but (2) the greedy algorithm takes the items in decreasing order of the value-per-cost ratios; and (3) in both processes, we break ties lexicographically.

- [When $S(r) \subseteq \{i \in [n] \mid x_i = 1\}$]. We notice that $c(S(r)) \leq \sum_{i \in [n]} x_i \cdot c_i \leq B$. Namely, regarding the Fractional Knapsack optimum, the total cost $\sum_{i \in [n]} x_i \cdot c_i$ is at least the cost on the items in $S(r)$, and is at most the budget B . In addition, every item $i \notin S(r)$ has a value-per-cost ratio $\frac{v_i}{c_i} \leq r$. Consequently, the total value of the items beyond set $S(r)$ is $\sum_{i \notin S(r)} x_i \cdot v_i \leq r \cdot \sum_{i \notin S(r)} x_i \cdot c_i \leq r \cdot (B - c(S(r)))$.
- [When $S(r) \supseteq \{i \in [n] \mid x_i > 0\}$]. We have $\sum_{i \in [n]} x_i \cdot c_i \leq c(S(r))$ and $\sum_{i \in [n]} x_i \cdot c_i \leq B$, and every item $i \in S(r)$ has a value-per-cost ratio $\frac{v_i}{c_i} \geq r$. As a result, $v(S(r)) - \text{fopt} = \sum_{i \in S(r)} (1 - x_i) \cdot v_i \geq r \cdot \sum_{i \in S(r)} (1 - x_i) \cdot c_i \geq r \cdot (c(S(r)) - B)$.

This completes the proof of properties (a), (b), and (c). \square

Mechanisms in the Second Stage. Given Lemma 3.1, PRUNING-MECHANISM can be composed with any follow-up truthful mechanism. Actually, we focus on the class of *posted-price*

⁵When there are multiple highest-value items, we break ties lexicographically.




Fig. 3. A 4-approximation deterministic budget-feasible mechanism.




Fig. 4. A new $(2 + \sqrt{2})$ -approximation deterministic budget-feasible mechanism.

mechanisms.⁶ Such a mechanism is determined by a set of prices $(B_i)_{i \in S(r)}$ subject to the budget constraint $\sum_{i \in S(r)} B_i \leq B$, and naturally meets the individual rationality, the budget feasibility, and the truthfulness.⁷

To illustrate how to analyze the approximability of a two-stage posted-price mechanism, and as a warm-up exercise, below, we discuss two simple mechanisms.

Warm-Up. Our first mechanism (see Figure 3) chooses the higher-value subset between $\{i^*\}$ and T as the winning set W , where $i^* \in \arg \max\{v_i \mid i \in S(r)\}$ is the highest-value item and $T = S(r) \setminus \{i^*\}$ (see Lemma 3.2), by offering price $\frac{v_i}{r}$ to each $i \in \{i^*\}$ or to each $i \in T$. Hence, we deduce from Lemma 3.2(c) that $\text{fopt} \leq 2 \cdot v(S(r)) \leq 4 \cdot \max\{v_{i^*}, v(T)\} = 4 \cdot \text{alg}$.

Our second posted-price mechanism (see Figure 4) recovers the best known result of $(2 + \sqrt{2})$ by Chen et al. [11]. This statement is formalized as the following theorem.

THEOREM 3.3. *SECOND-WARM-UP-MECHANISM is a $(2 + \sqrt{2})$ -approximation mechanism (individually rational, budget-feasible, and truthful) against the Fractional Knapsack benchmark.*

PROOF. We only show the approximability via case analysis; the other properties are obvious.

– [Case 1 that $v_{i^*} \geq \sqrt{2} \cdot v(T)$]. The highest-value item i^* is the only winner, and, thus, the outcome value $\text{alg} = v_{i^*}$. Then, according to Lemma 3.2(c), we have

$$\text{fopt} \leq 2 \cdot v(S(r)) = 2 \cdot (v_{i^*} + v(T)) \leq (2 + \sqrt{2}) \cdot v_{i^*} = (2 + \sqrt{2}) \cdot \text{alg}.$$

– [Case 2 that $v_{i^*} < \sqrt{2} \cdot v(T)$]. There are two possibilities. First, when $c_{i^*} \leq B - \frac{v(T)}{r}$, all items $i \in S(r)$ together form the winning set W , i.e., $\text{alg} = v(S(r))$. Due to Lemma 3.2(c), $\text{fopt} \leq 2 \cdot v(S(r)) = 2 \cdot \text{alg}$. Second, when $c_{i^*} > B - \frac{v(T)}{r}$, only the items $i \in T$ are chosen as

⁶To obtain our 3-approximation deterministic mechanism in Section 4, we actually use an adaptive posted-price scheme. Namely, the take-it-or-leave price offered to a specific item $i \in S(r)$ can change, depending on whether the items that have already made decisions accepted or rejected their posted-price offers.

⁷In the case of a randomized mechanism, any realization is given by a particular set of budget-feasible posted prices $(B_i)_{i \in S(r)}$, i.e., a truthful deterministic mechanism. Thus, this randomized mechanism is universally truthful.

DETERMINISTIC-MECHANISM

- (0) Receive the pair $(r, S(r))$ from PRUNING-MECHANISM
- (1) If $v_{i^*} \leq \frac{1}{2} \cdot v(T)$, get items T by offering price $\frac{v_i}{r}$ to each item $i \in T$
- (2) Else if $v_{i^*} \geq 2 \cdot v(T)$, get item i^* by offering price $\frac{v_{i^*}}{r}$ to i^*
- (3) Else, i.e., when $\frac{1}{2} \cdot v(T) < v_{i^*} < 2 \cdot v(T)$:
- (a) Offer price $B_{i^*} \stackrel{\text{def}}{=} \min\left\{\frac{v_{i^*}}{r}, \frac{2 \cdot v_{i^*} - v(T)}{v(S(r))} \cdot B\right\}$ to item i^*
- (b) If $c_{i^*} \leq B_{i^*}$,^a offer $B_i \stackrel{\text{def}}{=} \min\left\{\frac{v_i}{r}, \frac{v_i}{v(T)} \cdot (B - B_{i^*})\right\}$ to each item $i \in T$
- (c) Else, get items T by offering price $\frac{v_i}{r}$ to each item $i \in T$

^aIf $c_{i^*} \leq B_{i^*}$, item i^* will accept offer B_{i^*} . Otherwise, $c_{i^*} > B_{i^*}$ and item i^* will reject offer B_{i^*} , and then each item $i \in T$ will accept offer B_i .

Fig. 5. The 3-approximation deterministic budget-feasible mechanism.

the winners, i.e., $\text{alg} = v(T)$. Consequently,

$$\begin{aligned}
 \text{fopt} &\leq v(S(r)) + r \cdot (B - c(S(r))) && \text{(Lemma 3.2(c))} \\
 &\leq v(S(r)) + v(T) && \left(\text{as } c(S(r)) \geq c_{i^*} > B - \frac{v(T)}{r}\right) \\
 &= v(i^*) + 2 \cdot v(T) && \text{(as } v(S(r)) = v(i^*) + v(T)) \\
 &< (2 + \sqrt{2}) \cdot \text{alg}. && \text{(as } v_{i^*} < \sqrt{2} \cdot v(T) = \sqrt{2} \cdot \text{alg}).
 \end{aligned}$$

This completes the proof of Theorem 3.3. \square

We emphasize that our SECOND-WARM-UP-MECHANISM achieves a 2-approximation, when $v_{i^*} < \sqrt{2} \cdot v(T)$ and $c_{i^*} \leq B - \frac{v(T)}{r}$. One might ask a natural question: is it possible to achieve a better trade-off between this 2-approximation case and the $(2 + \sqrt{2})$ -approximation cases? In the next section, we will confirm this guess by presenting a slightly more complicated *adaptive* posted-price scheme, resulting in a 3-approximation deterministic mechanism.

4 DETERMINISTIC MECHANISM

The warm-up mechanisms have merely a few possible outcomes and do not adapt to the decisions of the items: either the highest-value item i^* , or the remaining items T , or, rarely, both item i^* and items T win; all the posted prices $(B_i)_{i \in S(r)}$ are almost equal to the maximum possible values $(\frac{v_i}{r})_{i \in S(r)}$. Such rigid structure hinders both warm-up mechanisms from achieving better performance guarantees than a $(2 + \sqrt{2})$ -approximation.

Now, we give a mechanism (called DETERMINISTIC-MECHANISM, see Figure 5) that achieves a better approximation. This mechanism (first stage) gets the pair $(r, S(r))$ via the PRUNING-MECHANISM given in Section 3, and then (second stage) applies an adaptive posted-price scheme.

THEOREM 4.1. *DETERMINISTIC-MECHANISM is a 3-approximation mechanism (individually rational, budget-feasible, and truthful) against the Fractional Knapsack benchmark.*

PROOF. The individual rationality and the truthfulness are easy to see, regarding the pricing nature of DETERMINISTIC-MECHANISM, Lemma 3.1, and Lemma 3.2(a). To show the budget feasibility, we consider either Case (3b) or Case (3c) in the mechanism:

- [Case (3b)]. $\sum_{i \in W} B_i \leq B_{i^*} + \sum_{i \in T} \frac{v_i}{v(T)} \cdot (B - B_{i^*}) = B$.
- [Case (3c)]. Since $W = T$, we know from Lemma 3.2(b) that $\sum_{i \in W} \frac{v_i}{r} = \frac{v(T)}{r} \leq B$.

We now show the approximation guarantee. Both Case (1) and Case (2), where either $v_{i^*} \leq \frac{1}{2} \cdot v(T)$ or $v_{i^*} \geq 2 \cdot v(T)$, are easy to analyze. Since $\text{alg} = \max\{v_{i^*}, v(T)\}$ in either case,

$$\text{fopt} < 2 \cdot v(S(r)) = 2 \cdot (v_{i^*} + v(T)) \leq 3 \cdot \max\{v_{i^*}, v(T)\} = 3 \cdot \text{alg},$$

where the first step applies Lemma 3.2(c), and the third step holds since we have $2 \cdot v_{i^*} \leq v(T)$ or $v_{i^*} \geq 2 \cdot v(T)$ in both cases.

From now on, we safely assume $\frac{1}{2} \cdot v(T) < v_{i^*} < 2 \cdot v(T)$. Conditioned on either $c_{i^*} \leq B_{i^*}$ or $c_{i^*} > B_{i^*}$, we are only left to deal with Case (3b) and Case (3c).

[Case (3b) that $c_{i^*} \leq B_{i^*}$]. We denote by $U \stackrel{\text{def}}{=} \{i \in T \mid c_i \leq B_i\}$ the set of winners in T , so the outcome value $\text{alg} = v_{i^*} + v(U)$. Of course, a losing item $i \in (T \setminus U)$ rejects the offered price $B_i = \min\{\frac{v_i}{r}, \frac{v_i}{v(T)} \cdot (B - B_{i^*})\}$ (by definition), since it has a too large cost $c_i > B_i$. But this losing item was not discarded during PRUNING-MECHANISM, so it has a high enough value-per-cost ratio $\frac{v_i}{c_i} \geq r$ (see Lemma 3.2(a)) and thus a cost $c_i \leq \frac{v_i}{r}$. For these reasons, the price offered to this losing item is exactly $B_i = \frac{v_i}{v(T)} \cdot (B - B_{i^*})$. We deduce that

$$c(S(r)) \geq \sum_{i \in (T \setminus U)} c_i > \sum_{i \in (T \setminus U)} B_i = \frac{v(T \setminus U)}{v(T)} \cdot (B - B_{i^*}). \quad (1)$$

By Lemma 3.2(c), $\text{fopt} \leq v(S(r)) + r \cdot (B - c(S(r)))$. We plug Inequality (1) into it and get

$$\begin{aligned} \text{fopt} &\stackrel{(1)}{<} v(S(r)) + r \cdot \left(\frac{v(U)}{v(T)} \cdot B + \frac{v(T \setminus U)}{v(T)} \cdot B_{i^*} \right) \\ &< v(S(r)) \cdot \left(1 + \frac{v(U)}{v(T)} + \frac{v(T \setminus U)}{v(T)} \cdot \frac{B_{i^*}}{B} \right) && \text{(Lemma 3.2(b): } rB < v(S(r))\text{)} \\ &\leq v(S(r)) \cdot \left(1 + \frac{v(U)}{v(T)} \right) + \frac{v(T \setminus U)}{v(T)} \cdot (2 \cdot v_{i^*} - v(T)) && \left(\text{as } B_{i^*} \leq \frac{2 \cdot v_{i^*} - v(T)}{v(S(r))} \cdot B \right) \\ &= 3 \cdot v_{i^*} + v(U) \cdot \left(2 - \frac{v_{i^*}}{v(T)} \right) && \left(\text{as } v(S(r)) = v_{i^*} + v(T) \right) \\ &\leq 3 \cdot v_{i^*} + 3 \cdot v(U) = 3 \cdot \text{alg}. \end{aligned}$$

[Case (3c) that $c_{i^*} > B_{i^*}$]. According to Lemma 3.2(a), $c_{i^*} \leq \frac{v_{i^*}}{r}$, and $c_i \leq \frac{v_i}{r}$ for any $i \in T$. Since $B_{i^*} < c_{i^*} \leq \frac{v_{i^*}}{r}$, we have $B_{i^*} = \min\{\frac{v_{i^*}}{r}, \frac{2 \cdot v_{i^*} - v(T)}{v(S(r))} \cdot B\} = \frac{2 \cdot v_{i^*} - v(T)}{v(S(r))} \cdot B$.

In this case, the highest-value item i^* rejects its offer, but all the remaining items $i \in T$ accept their offers. Thus, the winning set is $W = T$, and the outcome value is $\text{alg} = v(T)$. We then deduce that

$$\begin{aligned} \text{fopt} &\leq v(S(r)) + r \cdot (B - c(S(r))) && \text{(Lemma 3.2(c))} \\ &\leq v(S(r)) + r \cdot (B - B_{i^*}) && \left(\text{as } c(S(r)) \geq c_{i^*} > B_{i^*} \right) \\ &\leq v(S(r)) \cdot \left(2 - \frac{B_{i^*}}{B} \right) && \text{(Lemma 3.2(b): } rB < v(S(r))\text{)} \\ &= 3 \cdot v(T) = 3 \cdot \text{alg}. && \left(\text{as } B_{i^*} = \frac{2 \cdot v_{i^*} - v(T)}{v(S(r))} \cdot B \right) \end{aligned}$$

To conclude, we have $3 \cdot \text{alg} \geq \text{fopt}$ in all cases, which completes the proof of Theorem 4.1. \square

RANDOMIZED-MECHANISM

- (0) Receive the pair $(r, S(r))$ from PRUNING-MECHANISM
- (1) Let $q \stackrel{\text{def}}{=} \frac{1}{2} \cdot \frac{v(S(r)) - rB}{\min\{v_{i^*}, v(T)\}}$
- (2) If $v_{i^*} \leq v(T)$, let $q_{i^*} \stackrel{\text{def}}{=} (\frac{1}{2} - q)$ and $q_T \stackrel{\text{def}}{=} \frac{1}{2}$
- (3) Else, let $q_{i^*} \stackrel{\text{def}}{=} \frac{1}{2}$ and $q_T \stackrel{\text{def}}{=} (\frac{1}{2} - q)$
- (4) Offer price^a B_{i^*} to item i^* , where B_{i^*} is defined as follows:
- (a) With probability q_{i^*} , let $B_{i^*} \stackrel{\text{def}}{=} \frac{v_{i^*}}{r}$
 - (b) With probability q_T , let $B_{i^*} \stackrel{\text{def}}{=} B - \frac{v(T)}{r}$
 - (c) With probability q , draw $B_{i^*} \sim \text{Uniform}[B - \frac{v(T)}{r}, \frac{v_{i^*}}{r}]$
- (5) Offer price $B_i \stackrel{\text{def}}{=} \frac{v_i}{v(T)} \cdot (B - B_{i^*})$ to each item $i \in T$

^aFor every item $i \in S(r)$, price B_i is well defined in range $[0, \frac{v_i}{r}] \subseteq [0, B]$, by Lemma 3.2 (a).

Fig. 6. The 2-approximation randomized budget-feasible mechanism.

4.1 Matching Lower Bound

Against the Fractional Knapsack benchmark, our DETERMINISTIC-MECHANISM turns out to have the best possible approximation ratio among all deterministic mechanisms. To see so, we now construct a matching lower-bound instance, which is similar to Ref. [24, Proposition 5.2].

THEOREM 4.2. *No deterministic mechanism (truthful, individually rational, and budget-feasible) has an approximation ratio less than 3 against the Fractional Knapsack benchmark, even if there are only three items.*

PROOF. For the sake of contradiction, assume that there is a $(3 - \varepsilon)$ -approximation deterministic mechanism for some constant $\varepsilon > 0$. Consider the following two scenarios with three items having values $v_1 = v_2 = v_3 = 1$. Let $c^* \stackrel{\text{def}}{=} \frac{B}{2 - \varepsilon/2}$; notice that $2c^* > B$.

- [With costs (c^*, c^*, c^*)]. Due to the individual rationality, each winning item gains a payment of at least c^* . To guarantee the promised approximation ratio of $(3 - \varepsilon)$ under budget feasibility, there is exactly one winning item. W.l.o.g., we assume that the winner is the first item.
- [With costs $(0, c^*, c^*)$]. By the truthfulness, item 1 wins once again, getting the same payment of at least c^* . As a result, the budget left is at most $(B - c^*) < c^*$. Regarding the budget feasibility and individual rationality, neither item 2 nor item 3 can win.

In the later scenario, the mechanism generates value $\text{alg} = 1$, yet the Fractional Knapsack benchmark achieves value $\text{fopt} = 1 + \frac{B}{c^*} = 3 - \frac{\varepsilon}{2} > 3 - \varepsilon$. This contradicts our assumption that the mechanism is $(3 - \varepsilon)$ -approximation, concluding the proof of the theorem. \square

5 MAIN RESULT: RANDOMIZED MECHANISM

We now present the main result of our work, a randomized mechanism (called RANDOMIZED-MECHANISM, see Figure 6) that achieves a 2-approximation to the Fractional Knapsack benchmark. Regarding the matching lower bound by Chen et al. [11, Theorem 4.2] against the weaker Knapsack benchmark, this approximation guarantee is tight for both benchmarks. Our mechanism (first stage) gets the pair $(r, S(r))$ from the PRUNING-MECHANISM given in Section 3, and then (second stage) applies a randomized non-adaptive posted-price scheme.

We first verify that all quantities in RANDOMIZED-MECHANISM are well defined.

LEMMA 5.1. $0 \leq q = \frac{1}{2} \cdot \frac{v(S(r)) - rB}{\min\{v_{i^*}, v(T)\}} \leq \frac{1}{2}$ and $0 \leq B - \frac{v(T)}{r} < \frac{v_{i^*}}{r}$.

PROOF. The first inequality is due to Lemma 3.2(b) that $v(S(r)) > rB$. Lemma 3.2 further implies $rB \geq v_{i^*}$ and $rB \geq v(T)$, i.e., $rB \geq \max\{v_{i^*}, v(T)\}$. Now, the second inequality in Lemma 5.1 follows, as $q = \frac{1}{2} \cdot \frac{v_{i^*} + v(T) - rB}{\min\{v_{i^*}, v(T)\}} \leq \frac{1}{2} \cdot \frac{v_{i^*} + v(T) - \max\{v_{i^*}, v(T)\}}{\min\{v_{i^*}, v(T)\}} = \frac{1}{2}$. Finally, rearranging $v(T) \leq rB < v(S(r)) = v_{i^*} + v(T)$ leads to the last two inequalities. \square

Similar to DETERMINISTIC-MECHANISM in Section 4, we also slightly abuse notations and also refer to RANDOMIZED-MECHANISM as the composition of two mechanisms: PRUNING-MECHANISM with RANDOMIZED-MECHANISM.

THEOREM 5.2. RANDOMIZED-MECHANISM is a 2-approximation mechanism (individually rational, budget-feasible, and universally truthful) against the Fractional Knapsack benchmark.

PROOF. Since RANDOMIZED-MECHANISM is a posted-price scheme, it is individually rational. Since each random realization of the prices $(B_i)_{i \in S(r)}$ is budget-feasible, i.e., $\sum_{i \in S(r)} B_i = B$ by construction, the mechanism is also budget-feasible. Note that (1) all random choices in RANDOMIZED-MECHANISM, i.e., the prices $(B_i)_{i \in S(r)}$, can be made before execution of the mechanism; and (2) for each such choice, the resulting posted-price mechanism is obviously truthful. Due to Lemma 3.1, all desired properties extend to the composition mechanism, hence, being individually rational, budget-feasible, and universally truthful.

In the rest of the proof, we show that RANDOMIZED-MECHANISM is a 2-approximation to fopt . Let $(x_i)_{i \in S(r)}$ denote the allocation probabilities; then, the mechanism generates an expected value of $\text{alg} = \sum_{i \in S(r)} v_i \cdot x_i$. In order to prove the approximation guarantee, we need the following: Equation (2), Inequality (3), and Inequality (4), which will be proved later.

$$rB = 2q_{i^*} \cdot v_{i^*} + 2q_T \cdot v(T), \quad (2)$$

$$v_{i^*} \cdot x_{i^*} \geq q_{i^*} \cdot v_{i^*} + \frac{1}{2} \cdot (v_{i^*} - r \cdot c_{i^*}), \quad (3)$$

$$v_i \cdot x_i \geq q_T \cdot v_i + \frac{1}{2} \cdot (v_i - r \cdot c_i), \quad \forall i \in T. \quad (4)$$

Indeed, these mathematical facts together with Lemma 3.2(c) imply that $2 \cdot \text{alg} \geq \text{fopt}$.

$$\begin{aligned} \text{fopt} &\leq v(S(r)) + r \cdot (B - c(S(r))) \\ &\stackrel{(2)}{=} (v_{i^*} + v(T)) + 2 \cdot (q_{i^*} \cdot v_{i^*} + q_T \cdot v(T)) - r \cdot (c_{i^*} + c(T)) \\ &\stackrel{(3,4)}{\leq} 2v_{i^*} \cdot x_{i^*} + 2 \cdot \sum_{i \in T} v_i \cdot x_i \\ &= 2 \cdot \text{alg}. \end{aligned}$$

Now, we are only left to prove Equation (2), Inequality (3) and Inequality (4).

[Equation (2)]. By the definitions of q_{i^*} and q_T , in either case of Step (2) or Step (3),

$$\begin{aligned} q_{i^*} \cdot v_{i^*} + q_T \cdot v(T) &= \frac{1}{2} \cdot (v_{i^*} + v(T)) - q \cdot \min\{v_{i^*}, v(T)\} \\ &= \frac{1}{2} \cdot (v_{i^*} + v(T)) - \frac{1}{2} \cdot (v(S(r)) - rB) \quad (\text{definition of } q) \\ &= \frac{1}{2} \cdot rB. \end{aligned}$$

[Inequality (3)]. It is equivalent to showing that $\Pr[B_{i^*} \geq c_{i^*}] = x_{i^*} \geq q_{i^*} + \frac{v_{i^*} - r \cdot c_{i^*}}{2v_{i^*}}$.

– [When $c_{i^*} \leq B - \frac{v(T)}{r}$]. Item i^* always accepts price B_{i^*} , i.e., $\Pr[B_{i^*} \geq c_{i^*}] = 1$, which gives us the desired bound of $1 \geq q_{i^*} + \frac{v_{i^*} - r \cdot c_{i^*}}{2v_{i^*}}$, because $q_{i^*} \leq \frac{1}{2}$.

– [When $c_{i^*} > B - \frac{v(T)}{r}$]. Due to Lemma 3.2(a), $\frac{v_{i^*}}{r} \geq c_{i^*}$. We consider the random events in Step (4a) that $B_{i^*} = \frac{v_{i^*}}{r}$ and in Step (4c) that $B_{i^*} \sim \text{Uniform}[B - \frac{v(T)}{r}, \frac{v_{i^*}}{r}]$. Since $\Pr[\text{Step (4a)}] = q_{i^*}$ and $\Pr[\text{Step (4c)}] = q$, putting everything together gives

$$\begin{aligned} \Pr[B_{i^*} \geq c_{i^*}] &= q_{i^*} + q \cdot \frac{v_{i^*}/r - c_{i^*}}{v_{i^*}/r - (B - v(T)/r)} && \left(\text{Lemma 3.2(a)} : \frac{v_i}{r} \geq c_i \right) \\ &= q_{i^*} + \frac{1}{2} \cdot \frac{v_{i^*} - r \cdot c_{i^*}}{\min\{v_{i^*}, v(T)\}} && \text{(definition of } q) \\ &\geq q_{i^*} + \frac{v_{i^*} - r \cdot c_{i^*}}{2v_{i^*}}. && \left(\text{as } v_{i^*} \geq \min\{v_{i^*}, v(T)\} \text{ and } \frac{v_i}{r} \geq c_i \right) \end{aligned}$$

[Inequality (4)]. The argument is similar to the above. For each item $i \in T$, we claim that $\Pr[B_i \geq c_i] = x_i \geq q_T + \frac{v_i - r \cdot c_i}{2v_i}$.

– [When $c_i \leq \frac{v_i}{v(T)} \cdot (B - \frac{v_i^*}{r})$]. Item i always accepts price B_i , i.e., $\Pr[B_i \geq c_i] = 1$, which gives us the desired bound of $1 \geq q_T + \frac{v_i - r \cdot c_i}{2v_i}$, in that $q_T \leq \frac{1}{2}$.

– [When $c_i > \frac{v_i}{v(T)} \cdot (B - \frac{v_i^*}{r})$]. By Step (5), $B_i \geq c_i$ if and only if $B_{i^*} \leq B - v(T) \cdot \frac{c_i}{v_i}$. We consider the random events in Step (4b) that $B_{i^*} = B - \frac{v(T)}{r}$ and in Step (4c) that $B_{i^*} \sim \text{Uniform}[B - \frac{v(T)}{r}, \frac{v_{i^*}}{r}]$. Because $\Pr[\text{Step (4b)}] = q_T$ and $\Pr[\text{Step (4c)}] = q$,

$$\begin{aligned} \Pr[B_i \geq c_i] &= q_T + q \cdot \frac{(B - v(T) \cdot c_i/v_i) - (B - v(T)/r)}{v_{i^*}/r - (B - v(T)/r)} \\ &= q_T + q \cdot \frac{v(T)}{v(S(r)) - rB} \cdot \frac{v_i - r \cdot c_i}{v_i} && \text{(as } v(S(r)) = v_{i^*} + v(T)) \\ &= q_T + \frac{1}{2} \cdot \frac{v(T)}{\min\{v_{i^*}, v(T)\}} \cdot \frac{v_i - r \cdot c_i}{v_i} && \text{(definition of } q) \\ &\geq q_T + \frac{v_i - r \cdot c_i}{2v_i}, && \left(\text{Lemma 3.2(a)} : \frac{v_i}{r} \geq c_i \right) \end{aligned}$$

This completes the proof of Theorem 5.2. \square

6 CONCLUSION AND OPEN QUESTION

In this work, we proposed a budget-feasible randomized mechanism with the best possible approximation guarantee for an additive buyer. In addition, our deterministic mechanism still leaves some room for improvement: the best possible approximation guarantee is somewhere between $[\sqrt{2} + 1, 3]$. However, our instance from Theorem 4.2 clearly demonstrates that quite a different approach that is specifically tailored to the real Knapsack optimum (rather than the fractional relaxation solution) is needed.

The class of additive valuations is the most basic class of valuations in the research agenda for budget-feasible mechanisms. We hope that our results may lead to new mechanisms and improved analysis for broader valuation classes. Indeed, given the same factor 2-approximation result of Ref. [9] for the promise version of the problem for a subadditive buyer, we are even so bold as to conjecture that the true approximation guarantee for a subadditive buyer is still 2 (leaving all computational considerations aside).

Our composition approach has a lot of resemblance to the pruning ideas from the frugality literature. This demonstrates that ideas and approaches from one area of reverse auction design might be beneficial to another. We believe that there could be more interesting connections between these two complementary agendas.

Finally, our mechanisms use posted prices in the second stage. Besides the practical interest and motivation of posted-price mechanisms in the prior work, our work gives additional support to study this family of mechanisms in budget-feasible framework from a theoretical viewpoint.

ACKNOWLEDGMENTS

We are grateful to the anonymous reviewers for their dedication in carefully reading through this article; they offered many invaluable comments and suggestions. Work done in part when the second and the last authors were research assistants at ITCS, Shanghai University of Finance and Economics.

REFERENCES

- [1] Gagan Aggarwal and Jason D. Hartline. 2006. Knapsack auctions. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, January 22–26, 2006, Miami, Florida*. ACM Press, 1083–1092. <http://dl.acm.org/citation.cfm?id=1109557.1109677>.
- [2] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2016. Coverage, matching, and beyond: New results on budgeted mechanism design. In *Proceedings of the 12th International Conference on Web and Internet Economics, WINE 2016, December 11–14, 2016, Montreal, Canada*. Yang Cai and Adrian Vetta (Eds.), Vol. 10123. Springer, 414–428. DOI: https://doi.org/10.1007/978-3-662-54110-4_29
- [3] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2017. On budget-feasible mechanism design for symmetric submodular objectives. In *Proceedings of the 13th International Conference on Web and Internet Economics, WINE 2017, December 17–20, 2017, Bangalore, India*. Nikhil R. Devanur and Pinyan Lu (Eds.), Vol. 10660. Springer, 1–15. DOI: https://doi.org/10.1007/978-3-319-71924-5_1
- [4] Nima Anari, Gagan Goel, and Afshin Nikzad. 2014. Mechanism design for crowdsourcing: An optimal $1-1/e$ competitive budget-feasible mechanism for large markets. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, October 18–21, 2014, Philadelphia, PA*. IEEE Computer Society, 266–275. DOI: <https://doi.org/10.1109/FOCS.2014.36>
- [5] Aaron Archer and Éva Tardos. 2007. Frugal path mechanisms. *ACM Trans. Algorithms* 3, 1 (2007), 3:1–3:22. DOI: <https://doi.org/10.1145/1219944.1219948>
- [6] Ashwinkumar Badanidiyuru, Robert Kleinberg, and Yaron Singer. 2012. Learning on a budget: Posted price mechanisms for online procurement. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC 2012, June 4–8, 2012, Valencia, Spain*. Boi Faltings, Kevin Leyton-Brown, and Panos Ipeirotis (Eds.). ACM, 128–145. DOI: <https://doi.org/10.1145/2229012.2229026>
- [7] Eric Balkanski and Jason D. Hartline. 2016. Bayesian budget feasibility with posted pricing. In *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11–15, 2016*, Jacqueline Bourdeau, Jim Hendler, Roger Nkambou, Ian Horrocks, and Ben Y. Zhao (Eds.). ACM, 189–203. DOI: <https://doi.org/10.1145/2872427.2883032>
- [8] Eric Balkanski and Yaron Singer. 2015. Mechanisms for fair attribution. In *Proceedings of the 16th ACM Conference on Economics and Computation, EC 2015, June 15–19, 2015, Portland, OR*. Tim Roughgarden, Michal Feldman, and Michael Schwarz (Eds.). ACM, 529–546. DOI: <https://doi.org/10.1145/2764468.2764505>
- [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. DOI: <https://doi.org/10.1137/16M1067275>
- [10] Ning Chen, Edith Elkind, Nick Gravin, and Fedor Petrov. 2010. Frugal mechanism design via spectral techniques. In *51st Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23–26, 2010, Las Vegas, NV*. IEEE Computer Society, 755–764. DOI: <https://doi.org/10.1109/FOCS.2010.77>
- [11] Ning Chen, Nick Gravin, and Pinyan Lu. 2011. On the approximability of budget feasible mechanisms. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, January 23–25, 2011, San Francisco, CA*. Dana Randall (Ed.). SIAM, 685–699. DOI: <https://doi.org/10.1137/1.9781611973082.54>
- [12] Ning Chen and Anna R. Karlin. 2007. Cheap labor can be expensive. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, January 7–9, 2007, New Orleans, LA*. Nikhil Bansal, Kirk Pruhs, and Clifford Stein (Eds.). SIAM, 707–715. <http://dl.acm.org/citation.cfm?id=1283383.1283459>.
- [13] Shahar Dobzinski, Christos H. Papadimitriou, and Yaron Singer. 2011. Mechanisms for complement-free procurement. In *Proceedings 12th ACM Conference on Electronic Commerce, EC 2011, June 5–9, 2011, San Jose, CA*. Yoav Shoham, Yan Chen, and Tim Roughgarden (Eds.). ACM, 273–282. DOI: <https://doi.org/10.1145/1993574.1993615>
- [14] Edith Elkind, Leslie Ann Goldberg, and Paul W. Goldberg. 2007. Frugality ratios and improved truthful mechanisms for vertex cover. In *Proceedings 8th ACM Conference on Electronic Commerce (EC-2007), June 11–15, 2007, San Diego*,

- CA. Jeffrey K. MacKie-Mason, David C. Parkes, and Paul Resnick (Eds.). ACM, 336–345. DOI : <https://doi.org/10.1145/1250910.1250959>
- [15] Edith Elkind, Amit Sahai, and Kenneth Steiglitz. 2004. Frugality in path auctions. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004, January 11–14, 2004, New Orleans, LA*. J. Ian Munro (Ed.). SIAM, 701–709. <http://dl.acm.org/citation.cfm?id=982792.982900>.
- [16] Gagan Goel, Afshin Nikzad, and Adish Singla. 2014. Mechanism design for crowdsourcing markets with heterogeneous tasks. In *Proceedings of the 2nd AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2014, November 2–4, 2014, Pittsburgh, PA*. Jeffrey P. Bigham and David C. Parkes (Eds.). AAAI. <http://www.aaai.org/ocs/index.php/HCOMP/HCOMP14/paper/view/8968>.
- [17] Mohammad Taghi Hajiaghayi, Mohammad Reza Khani, and Saeed Seddighin. 2018. Frugal auction design for set systems: Vertex cover and knapsack. In *Proceedings of the 2018 ACM Conference on Economics and Computation, June 18–22, 2018, Ithaca, NY*. Éva Tardos, Edith Elkind, and Rakesh Vohra (Eds.). ACM, 645–662. DOI : <https://doi.org/10.1145/3219166.3219229>
- [18] Thibaut Horel, Stratis Ioannidis, and S. Muthukrishnan. 2014. Budget feasible mechanisms for experimental design. In *Proceedings of the 11th Latin American Symposium, LATIN 2014, March 31– April 4, 2014, Montevideo, Uruguay*. Alberto Pardo and Alfredo Viola (Eds.), Vol. 8392. Springer, 719–730. DOI : https://doi.org/10.1007/978-3-642-54423-1_62
- [19] Anna R. Karlin, David Kempe, and Tami Tamir. 2005. Beyond VCG: Frugality of truthful mechanisms. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2005, October 23–25, 2005, Pittsburgh, PA*. IEEE Computer Society, 615–626. DOI : <https://doi.org/10.1109/SFCS.2005.25>
- [20] David Kempe, Mahyar Salek, and Christopher Moore. 2010. Frugal and truthful auctions for vertex covers, flows and cuts. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23–26, 2010, Las Vegas, NV*. IEEE Computer Society, 745–754. DOI : <https://doi.org/10.1109/FOCS.2010.76>
- [21] Pooya Jalaly Khalilabadi and Éva Tardos. 2018. Simple and efficient budget feasible mechanisms for monotone submodular valuations. In *Proceedings of the 14th International Conference, WINE 2018, December 15–17, 2018, Oxford, UK*, George Christodoulou and Tobias Harks (Eds.), Vol. 11316. Springer, 246–263. DOI : https://doi.org/10.1007/978-3-030-04612-5_17
- [22] Stefano Leonardi, Gianpiero Monaco, Piotr Sankowski, and Qiang Zhang. 2017. Budget feasible mechanisms on matroids. In *Proceedings of the 19th International Conference on Integer Programming and Combinatorial Optimization, IPCO 2017, June 26–28, 2017, Waterloo, ON, Canada*, Friedrich Eisenbrand and Jochen Könnemann (Eds.), Vol. 10328. Springer, 368–379. DOI : https://doi.org/10.1007/978-3-319-59250-3_30
- [23] Roger B. Myerson. 1981. Optimal auction design. *Math. Oper. Res.* 6, 1 (1981), 58–73. DOI : <https://doi.org/10.1287/moor.6.1.58>
- [24] Yaron Singer. 2010. Budget feasible mechanisms. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23–26, 2010, Las Vegas, NV*. IEEE Computer Society, 765–774. DOI : <https://doi.org/10.1109/FOCS.2010.78>
- [25] Yaron Singer. 2012. How to win friends and influence people, truthfully: Influence maximization mechanisms for social networks. In *Proceedings of the 5th International Conference on Web Search and Web Data Mining, WSDM 2012, February 8–12, 2012, Seattle, WA*. Eytan Adar, Jaime Teevan, Eugene Agichtein, and Yoelle Maarek (Eds.). ACM, 733–742. DOI : <https://doi.org/10.1145/2124295.2124381>
- [26] Yaron Singer and Manas Mittal. 2013. Pricing mechanisms for crowdsourcing markets. In *Proceedings of the 22nd International World Wide Web Conference, WWW 2013, May 13–17, 2013, Rio de Janeiro, Brazil*. Daniel Schwabe, Virgílio A. F. Almeida, Hartmut Glaser, Ricardo Baeza-Yates, and Sue B. Moon (Eds.). International World Wide Web Conferences Steering Committee/ACM, 1157–1166. DOI : <https://doi.org/10.1145/2488388.2488489>
- [27] Kunal Talwar. 2003. The price of truth: Frugality in truthful mechanisms. In *Proceedings of the 20th Annual Symposium on Theoretical Aspects of Computer Science, STACS 2003, February 27– March 1, 2003, Seattle, WA*. Helmut Alt and Michel Habib (Eds.), Vol. 2607. Springer, 608–619. DOI : https://doi.org/10.1007/3-540-36494-3_53

Received September 2019; revised July 2020; accepted July 2020