On the Approximability of Budget Feasible Mechanisms

Ning Chen* Nick Gravin*† Pinyan Lu‡

Abstract

Budget feasible mechanisms, recently initiated by Singer (FOCS 2010), extend algorithmic mechanism design problems to a realistic setting with a budget constraint. We consider the problem of designing truthful budget feasible mechanisms for monotone submodular functions: We give a randomized mechanism with an approximation ratio of 7.91 (improving on the previous best-known result 233.83), and a deterministic mechanism with an approximation ratio of 8.34. We also study the knapsack problem, which is a special submodular function, give a $2 + \sqrt{2}$ approximation deterministic mechanism (improving on the previous best-known result 5), and a 3 approximation randomized mechanism. We provide similar results for an extended knapsack problem with heterogeneous items, where items are divided into groups and one can pick at most one item from each group.

Finally we show a lower bound of $1 + \sqrt{2}$ for the approximation ratio of deterministic mechanisms and 2 for randomized mechanisms for knapsack, as well as the general monotone submodular functions. Our lower bounds are unconditional, and do not rely on any computational or complexity assumptions.

1 Introduction

It is well-known that a mechanism may have to pay a large amount to enforce incentive compatibility (i.e., truthfulness). For example, the seminal VCG mechanism may have unbounded payment (compared to the shortest path) in path auctions [1]. The negative effect of truthfulness on payments leads to a broad study of frugal mechanism design, i.e., how should one minimize his payment to get a desired output with incentive agents? While a class of results have been established [1, 23, 10, 11, 4], in practice, one cannot expect a negative overhead for a few perspectives, e.g., budget or resource limit.

Recently, Singer [21] considered mechanism design problems from a reverse angle and initiated a study on truthful mechanism design with a sharp budget constraint: The total payment of a mechanism is upper bounded by a given value B. Formally, in a marketplace each agent/item has a *privately* known incurred cost c_i. For any given subset S of agents, there is a *publicly* known valuation $v(S)$, meaning the social welfare derived from S. A mechanism selects a subset S of agents and decides a payment p_i to each $i \in S$. Agents bid strategically on their costs and would like to maximize their utility $p_i - c_i$. The objective is to design truthful budget feasible mechanisms with outputs approximately close to a socially optimal solution. In other words, it studies the "price of being truthful" in a budget constraint framework[1].

Although budget is a realistic condition that appears almost everywhere in daily life, it has not received much attention until very recently [7, 2, 3, 21]. In the framework of worst case analysis, most results are negative [7]. The introduction of budget adds another dimension to mechanism design; it further limits the searching space, especially given the (already) strong restriction of truthfulness. Designing budget feasible mechanisms even requires us to bound the threshold payment of each individual, which, not surprisingly, is tricky to analyze.

While the problem in general does not admit any budget feasible mechanism[2], Singer [21] studied an important class of valuation functions, i.e., monotone submodular functions. He gave a randomized truthful

*Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore. Email: ningc@ntu.edu.sg, ngravin@pmail.ntu.edu.sg. Work done while visiting Microsoft Research Asia.

†St. Petersburg Department of Steklov Mathematical Institute RAS, Russia.

‡Microsoft Research Asia. Email: pinyanl@microsoft.com.

[1]Note that if we do not consider truthful mechanism design, the problem is purely an optimization question with an extra capacity (i.e., budget) constraint, which has been well-studied in, e.g., [17, 22, 13, 8, 14], in the framework of submodularity with different conditions. It is well-known that a simple greedy algorithm gives the best possible approximation ratio $1 - 1/e$ [17] for maximizing a monotone submodular function with a capacity constraint. When agents are weighted (corresponding to costs in our setting), the simple greedy algorithm may have an unbounded approximation ratio [9]; a variant of the greedy algorithm which picks the maximum of the original greedy solution and the agent with the largest value yields the same approximation ratio $1 - 1/e$ [13].

[2]For example, one with budget $B = 1$ would like to purchase a path from s to t in a network $\{(s, v), (v, t)\}$ where each edge has incurred cost 0. In any truthful mechanism that guarantees to buy the path (i.e., outputs the socially optimum solution), one has to pay each edge the threshold value B, leading to a total payment $2B$ which exceeds the given budget.

	Submodular functions				Knapsack			
	deterministic		randomized		deterministic		randomized	
	upper	lower	upper	lower	upper	lower	upper	lower
Singer [21]	–	2	233.83	–	5	2	–	–
Our results	8.34*	$1+\sqrt{2}$	7.91	2	$2+\sqrt{2}$	$1+\sqrt{2}$	3	2

*It may require exponential running time for general monotone submodular functions.

mechanism with a constant approximation ratio of 233.83 for any monotone submodular functions, and deterministic mechanisms for special cases including knapsack (ratio 5) and coverage. Further, he showed that no deterministic truthful mechanism can obtain an approximation ratio better than 2, even for knapsack.

1.1 Our Results. In this paper, we improve upper and lower bounds of budget feasible mechanisms for monotone submodular functions and knapsack, summarized in the above table.

In truthful mechanism design, if there is no restriction on total payment, it is sufficient to focus on designing monotone allocations — the payment to each individual winner is the unique threshold to maintain the winning status [16]. With a sharp budget constraint, in addition to designing monotone allocations, we have to upper bound the sum of threshold payments. For submodular functions, the natural greedy algorithm is a good candidate for designing budget feasible mechanisms due to its nice monotonicity and small approximation ratio. However, the threshold payment to each winner can be very complicated because an agent can manipulate its ranking position in the greedy algorithm, which results in different computations of the marginal contributions for the rest agents, and therefore unpredictably change the set of winners. Singer [21] bounded the threshold of each winner by considering all possible ranking positions for his bid and taking the maximum of the thresholds of all these positions. In Section 3, we give a clean and tight analysis for the upper bound on threshold payment by applying the combinatorial structure of submodular functions (Lemma 3.2). These upper bounds on payments suggest appropriate parameters in our randomized mechanism, which, roughly speaking, selects the greedy algorithm or the agent with the largest value at a certain probability.

A difficulty of deriving deterministic mechanisms is related to the agent i^* with the largest value $v(i^*)$. The greedy algorithm may not take it due to its (possibly large) cost, which could result in a solution with an arbitrarily bad ratio. However, we cannot simply compare the solution of greedy algorithm with $v(i^*)$ because this breaks monotonicity as the agent i^* is able to manipulate the greedy solution by his bid (this is exactly where randomization helps). To get around of this issue, we drop i^* out of the market and compare $v(i^*)$ with the remaining agents in an appropriate way — now i^* is completely independent of the rest of the market and cannot affect its output — this gives our deterministic mechanisms for monotone submodular functions and knapsack small approximation ratios (note that we still need to be careful about the agents in the remaining market as they are still able to manipulate their bids to beat $v(i^*)$).

On the other hand, it is interesting to explore limitations of budget feasible mechanisms. Singer gave a simple lower bound of 2 on the approximation ratio and proposed that exploring the lower bounds that are dictated by budget feasibility is "perhaps the most interesting question" [21]. In Section 4, we prove a stronger lower bound of $1+\sqrt{2}$ for deterministic mechanisms. In most lower bounds proofs for truthful mechanisms, a number of related instances are constructed and one shows that a truthful mechanism cannot do well for all of them [5, 12, 18, 20]. (For example, in Singer's proof, three instances are constructed.) Our lower bound proof uses a slightly different approach: We first establish a property of a truthful mechanism for all instances provided that the mechanism has a good approximation ratio (Lemma 4.1), then we conclude that this property is inconsistent with the budget feasibility condition for a carefully constructed instance. Furthermore, we show a lower bound of 2 for universally randomized budget feasible mechanisms. Both our lower bounds are independent of computational assumptions and hold for instances with a small number of agents.

While submodular functions admit good approximation budget feasible mechanisms, extending them to more general functions seems to be a very difficult task. It was proved that we do not have any good approximation mechanisms for instances like the path and spanning tree [21]. In Section 5, we take a first step of this generalization by considering an extended knapsack problem with heterogeneous items (i.e., a group constraint), where items are of different types and we are only allowed to pick one item from each type. Here we cannot apply the same greedy mechanism for the original knapsack as it may not even generate a feasible solution; its approximation ratio can be arbitrarily bad if we only take the first item from each type. To construct a truthful mechanism with a good approximation, we employ a greedy strategy with *deletions* — in the process of the greedy algorithm, we either add a new item whose type has not been considered or replace an existing item with the new one of the same type. Although there are deletions, the greedy algorithm is still monotone (but its proof is much more involved), based on which we have similar approximation mechanisms for heterogeneous knapsack. We believe that the greedy strategy with deletions can be extended to a number of interesting non-submodular settings to derive budget feasible mechanisms with good approximations.

2 Preliminaries

In a marketplace there are n agents (or items), denoted by $A = \{1, \dots, n\}$. Each agent i has a privately known incurred *cost* c_i (or denoted by $c(i)$). For any given subset $S \subseteq A$ of agents, there is a publicly known valuation $v(S)$, meaning the social welfare derived from S. We assume $v(\emptyset) = 0$ and $v(S) \leq v(T)$ for any $S \subset T \subseteq A$ throughout this paper. We say the valuation function is *submodular* if $v(S)+v(T) \geq v(S\cap T)+v(S\cup T)$ for any $S, T \subseteq A$.

Upon receiving a *bid* cost b_i from each agent, a mechanism decides an *allocation* $S \subseteq A$ as winners and a *payment* p_i to each $i \in A$. We assume that the mechanism has no positive transfer (i.e., $p_i = 0$ if $i \notin S$) and is individually rational (i.e., $p_i \geq b_i$ if $i \in S$). Agents bid strategically on their costs and would like to maximize their utilities, which is $p_i - c_i$ if i is a winner and 0 otherwise. We say a mechanism is *truthful* if it is of the best interests for each agent to report his true cost. For randomized mechanisms, we consider universal truthfulness in this paper (i.e., a randomized mechanism takes a distribution over deterministic truthful mechanisms).

Our setting is in single parameter domain, as each agent has one private cost. It is well-known [16] that a mechanism is truthful if and only if its allocation rule is monotone (i.e., a winner keeps winning if he unilaterally decreases his bid) and the payment to each winner is his threshold bid (i.e., the maximal bid for which the agent still wins). Therefore, we will only focus on designing monotone allocations and do not specify the payment to each winner explicitly.

A mechanism is said to be *budget feasible* if $\sum_i p_i \leq B$, where B is a given sharp budget constraint. Assume without loss of generality that $c_i \leq B$ for any agent $i \in A$, since otherwise he will never win in any (randomized) budget feasible truthful mechanism. Our objective is to design truthful budget feasible mechanisms with outputs approximately close to the social optimum. That is, we want to minimize the *approximation ratio* of a mechanism, which is defined as $\max_I \frac{opt(I)}{\mathcal{M}(I)}$, where $\mathcal{M}(I)$ is the (expected) value of mechanism $\mathcal{M}$ on instance I and $opt(I)$ is the optimal value of the integer program: $\max_{S\subseteq A} v(S)$ subjected to $c(S) \leq B$, where $c(S) = \sum_{i\in S} c_i$.

3 Budget Feasible Mechanisms

For any given monotone submodular function, we denote the marginal contribution of an item i with respect to set S by $m_S(i) = v(S \cup \{i\}) - v(S)$. We assume that agents are sorted according to their non-increasing marginal contributions relative to their costs, recursively defined by: $i + 1 = \arg\max_{j\in A\setminus S_i} \frac{m_{S_i}(j)}{c_j}$, where $S_i = \{1, \dots, i\}$ and $S_0 = \emptyset$. To simplify notations we will denote this order by $[n]$ and write m_i instead of $m_{S_{i-1}}(i)$. This sorting, in the presence of submodularity, implies that

$$\frac{m_1}{c_1} \geq \frac{m_2}{c_2} \geq \cdots \geq \frac{m_n}{c_n}.$$

Notice that $v(S_k) = \sum_{i\leq k} m_i$ for all $k \in [n]$.

The following greedy scheme is the core of our mechanism (where the parameters denote the set of agents A and available budget $B/2$).

GREEDY-SM$(A, B/2)$

1. Let $k = 1$ and $S = \emptyset$
2. While $k \leq |A|$ and $c_k \leq \frac{B}{2} \cdot \frac{m_k}{\sum_{i\in S\cup\{k\}} m_i}$
   - $S \leftarrow S \cup \{k\}$
   - $k \leftarrow k + 1$
3. Return winning set S

Our mechanism for general monotone submodular functions is as follows.[3]

RANDOM-SM

1. Let $A = \{i \mid c_i \leq B\}$ and $i^* \in \arg\max_{i\in A} v(i)$
2. with probability 0.4, return i^*
3. with probability 0.6, return GREEDY-SM$(A, B/2)$

[3]Our mechanism has a similar flavor to Singer's mechanism [21] for the greedy scheme and randomness between the greedy and the item with the largest value. Indeed, both are due to the algorithm that maximizes monotone submodular functions with weighted items [13]. Our mechanism, however, treats the greedy scheme and random selection in a slightly different way, which yields a much better approximation ratio.

In the above mechanism, if it returns i^*, the payment to i^* is B; if it returns Greedy-SM$(A, B/2)$, the payment is more complicated and is given in [21]. Actually, we do not need this explicit payment formula to prove our result.

Theorem 3.1. Random-SM *is a budget feasible universally truthful mechanism for a submodular valuation function with an approximation ratio of* $\frac{5e}{e-1}(\approx 7.91)$.

3.1 Analysis of Random-SM**.** In this subsection we analyze Random-SM in terms of three respects: Truthfulness, budget feasibility and approximation. They together yield the proof for Theorem 3.1.

3.1.1 Universal Truthfulness. Our mechanism is a simple random combination of two mechanisms. To prove that the Random-SM is universally truthful, it suffices to prove that these two mechanisms are truthful respectively, i.e., the allocation rule is monotone.

The scheme where it simply returns i^* is obviously truthful. Also it is easy to see in the prior step that throwing away the agents having costs greater than B does not affect truthfulness. The greedy scheme Greedy-SM$(A, B/2)$ is monotone as well, since any item out of a winning set cannot increase its bid to become a winner.

3.1.2 Budget Feasibility. While truthfulness is quite straightforward, the budget feasibility analysis turns out to be quite tricky. The difficulties arise when we compute the payment to each item. Indeed, it can happen that an item changes its bid (while still remaining in the winning set) to force the mechanism to change its output. In other words, an item can control the output of the mechanism. Fortunately, in such a case no item can reduce the valuation of the output too much. That enables us to write an upper bound on the bid of each item in case of submodularity; summing up these bounds yields budget feasibility.

If the mechanism returns i^*, his payment is B and it is clearly budget feasible. It still remains to prove budget feasibility for Greedy-SM$(A, B/2)$. A similar but weaker result has been proven in [21], using the characterization of payments and arguing that the total payment is not larger than B. Here we directly show that the payment to any item i in the winning set S is bounded above by $\frac{m_i}{v(S)} \cdot B$; then the total payment will be bounded by B since $\frac{\sum_{i\in S} m_i}{v(S)} \cdot B = B$. Before doing that, we first prove a useful lemma.

Lemma 3.1. *Consider any* $S \subset T \subseteq [n]$ *and* $t_0 = \arg\max_{t\in T\setminus S} \frac{m_S(t)}{c(t)}$. *Then*

$$\frac{v(T) - v(S)}{c(T) - c(S)} \le \frac{m_S(t_0)}{c(t_0)}.$$

Proof. Assume for contradiction that the lemma does not hold, then for all $t \in T \setminus S$, we have

$$\frac{v(T) - v(S)}{c(T) - c(S)} > \frac{m_S(t)}{c(t)}.$$

Then add all inequalities each multiplied by $\frac{c(t)}{\sum_{t\in T\setminus S} c(t)}$ together, we have

$$\frac{v(T) - v(S)}{c(T) - c(S)} > \frac{\sum_{t\in T\setminus S} m_S(t)}{\sum_{t\in T\setminus S} c(t)} = \frac{\sum_{t\in T\setminus S} m_S(t)}{c(T) - c(S)}.$$

This implies that $v(T) - v(S) > \sum_{t\in T\setminus S} m_S(t)$, which contradicts the submodularity. □

Let $1, \ldots, k$ be the order of items in which we add them to the winning set. Let $\emptyset = S_0 \subset S_1 \subset \ldots \subset S_k \subseteq [n]$ be the sequence of winning sets that we pick at each step by applying our mechanism. Thus we have $S_j = [j]$. Now, since v is sumbodular, we can write the following chain of inequalities (note that marginal contribution is smaller for larger sets).

$$\frac{m_{S_0}(1)}{c_1} \ge \frac{m_{S_1}(2)}{c_2} \ge \ldots \ge \frac{m_{S_{k-1}}(k)}{c_k} \ge \frac{2v(S_k)}{B}.$$

The following is our main lemma.

Lemma 3.2. *No item* $j \in S =$ Greedy-SM$(A, B/2)$ *can bid more than* $m_{S_{j-1}}(j)\frac{B}{v(S)}$ *and still get into the winning set. Thus the payment to* j *is upper bounded by* $m_{S_{j-1}}(j)\frac{B}{v(S)}$.

Proof. Assume that $S = S_k$ is the winning set and there is $j \in S_k$ such that it can bid $b_j > m_{S_{j-1}}(j)\frac{B}{v(S_k)}$ and still win (given fixed bids of others). We will use notation b instead of c to emphasize that we consider a new scenario where j has increased its bid to b_j and others remain the same.

Note that

$$\frac{m_{S_0}(1)}{c_1} \ge \frac{m_{S_1}(2)}{c_2} \ge \ldots \ge \frac{m_{S_{j-1}}(j)}{c_j} \ge \frac{m_{S_{j-1}}(j)}{b_j}.$$

Thus the agents in S_{j-1} still get into the winning set.

For bid vector b, the set we have chosen right before j (denoted by T) is included into the winning set. Thus, by the rule of the greedy mechanism, we have

$$j = \arg\max_{i\in[n]\setminus T} \frac{m_T(i)}{b_i}, \tag{3.1}$$

$$\frac{m_T(j)}{b_j} \ge \frac{2v(T\cup\{j\})}{B}. \tag{3.2}$$

We may assume $S_k \cup T \supset T \cup \{j\}$. Indeed, otherwise $T \cup \{j\} = S_k \cup T$ and

$$\frac{m_{S_{j-1}}(j)}{b_j} \geq \frac{m_T(j)}{b_j} \geq \frac{2v(T \cup \{j\})}{B} \geq \frac{2v(S_k)}{B} \geq \frac{v(S_k)}{B}.$$

Thus $b_j \leq m_{S_{j-1}} \frac{B}{v(S_k)}$ and we get a contradiction.

Let $R = S_k \setminus T$. Applying equation (3.1) and Lemma 3.1 to $S_k \cup T$ and $T \cup \{j\}$, we know that for some $r_0 \in R \setminus \{j\}$,

$$\frac{v(S_k \cup T) - v(T \cup \{j\})}{b(S_k \cup T) - b(T \cup \{j\})} \leq \frac{m_{T\cup\{j\}}(r_0)}{b(r_0)} \leq \frac{m_T(j)}{b_j}.$$

On the other hand, since $b_j > m_{S_{j-1}}(j)\frac{B}{v(S_k)}$, we have

$$\frac{m_T(j)}{b_j} < \frac{m_T(j)}{m_{S_{j-1}}(j)} \frac{v(S_k)}{B} < \frac{v(S_k)}{B}.$$

Combining these inequalities, we get

$$\frac{v(S_k \cup T) - v(T \cup \{j\})}{b(S_k \cup T) - b(T \cup \{j\})} < \frac{v(S_k)}{B}.$$

We have

$$b(S_k \cup T) - b(T \cup \{j\}) = b(R \setminus \{j\}) = c(R \setminus \{j\}) \leq c(S_k).$$

Recall that $\frac{m_{S_{i-1}}(i)}{c_i} \geq \frac{2v(S_k)}{B}$ for $i \in [k]$. Thus $c_i \leq m_{S_{i-1}}(i)\frac{B}{2v(S_k)}$ and $c(S_k) = \sum_{i=1}^{k} c(i) \leq \frac{B}{2}$. We get

$$\begin{aligned}\frac{v(S_k) - v(T \cup \{j\})}{B/2} &\leq \frac{v(S_k) - v(T \cup \{j\})}{c(S_k)} \\ &\leq \frac{v(S_k \cup T) - v(T \cup \{j\})}{b(S_k \cup T) - b(T \cup \{j\})} \\ &< \frac{v(S_k)}{B}\end{aligned}$$

Thus, $v(S_k) < 2v(T \cup \{j\})$.

Applying inequality (3.2) we derive

$$\frac{m_{S_{j-1}}(j)}{b_j} \geq \frac{m_T(j)}{b_j} \geq \frac{2v(T \cup \{j\})}{B} > \frac{v(S_k)}{B},$$

which is contradictory to the fact that $b_j > m_{S_{j-1}}(j)\frac{B}{v(S_k)}$. □

3.1.3 Approximation Ratio.

3.1.3 Approximation Ratio. Before analyzing the performance of our mechanism, we consider the following simple greedy algorithm (without considering bidding strategies): Order items according to their marginal contributions divided by costs and add as many items as possible (i.e., it stops when we cannot add the next item as the sum of c_i otherwise will be bigger than B). Moreover we can consider the fractional variant of that, i.e., for the remaining budget we take a portion of the item at which we have stopped. Let ℓ be the maximal index for which $\sum_{i=1,\ldots,\ell} c_i \leq B$. Let $c'_{\ell+1} = B - \sum_{i=1,\ldots,\ell} c_i$ and $m'_{\ell+1} = m_{\ell+1} \cdot \frac{c'_{\ell+1}}{c_{\ell+1}}$. Hence, the fractional greedy solution is defined as

$$fgre(A) \triangleq \sum_{i=1}^{\ell} m_i + m'_{\ell+1}.$$

It is well-known that the greedy algorithm is a $1-1/e$ approximation of maximizing monotone submodular functions with a cardinality constraint [17]. Also it was shown that the simple greedy algorithm has an unbounded approximation ratio in case of weighted items with a capacity constraint. Nevertheless, a variant of greedy algorithm was suggested in [13] which gives the same $1 - 1/e$ approximation to the weighted case. The following lemma, which is fundamental to our analysis, establishes the same approximation ratio for the fractional greedy algorithm described above. (The proof is deferred to Appendix A.)

Lemma 3.3. *The fractional greedy solution has an approximation ratio of $1-1/e$ for the weighted submodular maximization problem. That is,*

$$fgre(A) \geq (1 - 1/e) \cdot opt(A),$$

where $opt(A)$ is the value of the optimal integral solution for the given instance A.

Now we are ready to analyze the approximation ratio of the mechanism Random-SM. Let $S = \{1, \ldots, k\}$ be the subset returned by Greedy-SM$(A, \frac{B}{2})$. For any $j = k+1, \ldots, \ell$, we have

$$\frac{c_j}{m_j} \geq \frac{c_{k+1}}{m_{k+1}} > \frac{B}{2\sum_{i=1}^{k+1} m_i},$$

where the last inequality follows from the fact that the greedy strategy stops at item $k+1$. Hence, we have $c_j > B \cdot \frac{m_j}{2\sum_{i=1}^{k+1} m_i}$. The same analysis shows that $c'_{\ell+1} > B \cdot \frac{m'_{\ell+1}}{2\sum_{i=1}^{k+1} m_i}$. Therefore,

$$B \cdot \frac{\sum_{j=k+1}^{\ell} m_j + m'_{\ell+1}}{2\sum_{i=1}^{k+1} m_i} < \sum_{j=k+1}^{\ell} c_j + c'_{\ell+1} \leq B.$$

This implies that $2\sum_{i=1}^{k+1} m_i > \sum_{j=k+1}^{\ell} m_j + m'_{\ell+1}$ and

$m_{k+1} + 2\sum_{i=1}^{k} m_i > \sum_{j=k+2}^{\ell} m_j + m'_{\ell+1}$. Hence,

$$\begin{aligned} fgre(A) &= \sum_{i=1}^{\ell} m_i + m'_{\ell+1} \\ &= \sum_{i=1}^{k+1} m_i + \sum_{j=k+2}^{\ell} m_j + m'_{\ell+1} \\ &< 3\sum_{i\in S} m_i + 2m_{k+1} \\ &\leq 3\sum_{i\in S} m_i + 2v(i^*) \end{aligned}$$

Together with Lemma 3.3, we can bound the optimal solution as

(3.3)
$$opt(A) \leq \frac{e}{e-1}\Big(3\text{Greedy-SM}(A, B/2) + 2v(i^*)\Big).$$

Therefore, the expected value of our randomized mechanism is $\frac{3}{5}\text{Greedy-SM}(A, B/2) + \frac{2}{5}v(i^*) \geq \frac{e-1}{5e}opt$.

3.2 Deterministic Mechanism. In this section, we provide a deterministic truthful mechanism which is budget feasible and has a constant approximation ratio. In the following description, $opt(A\setminus\{i^*\}, B)$ denotes the value of the optimal solution for the weighted submodular maximization problem for the given instance $A\setminus\{i^*\}$ with budget B.

Deterministic-SM

1. Let $A = \{i \mid c_i \leq B\}$ and $i^* \in \arg\max_{i\in A} v(i)$
2. If $\frac{1+4e+\sqrt{1+24e^2}}{2(e-1)} \cdot v(i^*) \geq opt(A \setminus \{i^*\}, B)$,[4] return i^*
3. Otherwise, return Greedy-SM$(A, B/2)$

[4] Our deterministic mechanism in general is not in polynomial time because of the hardness of computing an optimal solution for submodular maximization problems. However, we may substitute it by the optimum of the fractional problem; therefore for special problems like knapsack (discussed in the following subsection), we can get a polynomial time deterministic mechanism. Note however that we cannot replace it by the simple greedy solution as it breaks monotonicity.

Indeed, even if one is given unbounded computational power, we are still unable to solve the budget feasible mechanism design problem optimally (in particular, our lower bounds in the subsequent section still apply). Our mechanism suggests a natural question on the power of computation in (budget feasible) mechanism design at the price of being truthful [19, 6]. In particular, can an (exponential runtime) mechanism beat the lower bound of all polynomial time mechanisms? We leave this as future work.

Theorem 3.2. Deterministic-SM *is a deterministic budget feasible truthful mechanism for monotone submodular functions with an approximation ratio of* $\frac{6e-1+\sqrt{1+24e^2}}{2(e-1)} (\approx 8.34)$.

Proof. Note that the bid of i^* is independent to the value of $opt(A \setminus \{i^*\}, B)$. Therefore, the mechanism is truthful (a detailed similar argument is given in the proof of Theorem B.1 in Appendix B). Budget feasibility follows from Lemma 3.2 and the observation that Step 2 only gives additional upper bounds on the thresholds of winners from Greedy-SM$(A, B/2)$.

If the following, we prove the approximate ratio. Let

$$x = \frac{1+4e+\sqrt{1+24e^2}}{2(e-1)} (\approx 7.34).$$

We observe that

$$opt(A, B) - v(i^*) \leq opt(A \setminus \{i^*\}, B) \leq opt(A, B).$$

If the condition in Step 2 holds and the mechanism outputs i^*, then

$$opt(A, B) \leq opt(A \setminus \{i^*\}, B) + v(i^*) \leq (x + 1) \cdot v(i^*).$$

Otherwise, the condition in Step 2 fails and the mechanism outputs Greedy-SM$(A, B/2)$ in Step 3. Applying formula (3.3), we have

$$\begin{aligned} x \cdot v(i^*) &< opt(A \setminus \{i^*\}, B) \\ &\leq opt(A, B) \\ &\leq \frac{e}{e-1}\Big(3\text{Greedy-SM}(A, B/2) + 2v(i^*)\Big). \end{aligned}$$

This implies that

$$v(i^*) \leq \frac{3e}{x(e-1) - 2e}\text{Greedy-SM}(A, B/2).$$

Hence,

$$\begin{aligned} opt &\leq \frac{e}{e-1}\Big(3\text{Greedy-SM}(A, B/2) + 2v(i^*)\Big) \\ &\leq \frac{e}{e-1}\left(3 + \frac{6e}{x(e-1) - 2e}\right) \cdot \text{Greedy-SM}(A, B/2). \end{aligned}$$

Simple calculations show that

$$\begin{aligned} 1 + x &= \frac{6e - 1 + \sqrt{1+24e^2}}{2(e-1)} \\ &= \frac{e}{e-1}\left(3 + \frac{6e}{x(e-1) - 2e}\right). \end{aligned}$$

Therefore, we have $opt \leq (x+1) \cdot \text{Greedy-SM}(A, B/2)$ in the both cases, which concludes the proof of the claim with an approximation ratio of $\frac{e}{e-1}\left(3 + \frac{6e}{x(e-1)-2e}\right) (\approx 8.34)$. □

3.3 Improved Mechanisms for Knapsack. In this subsection, we consider a special model of submodular functions where the valuations of agents are additive, i.e., $v(S) = \sum_{i \in S} v_i$ for $S \subseteq [n]$. This leads to an instance of the Knapsack problem, where items correspond to agents and the size of the knapsack corresponds to budget B. Singer [21] gave a 5-approximation deterministic mechanism. By applying approaches from the previous subsections, we have the following results (proofs are deferred to Appendix B).

THEOREM 3.3. *There are $2+\sqrt{2}$ approximation deterministic and 3 approximation randomized polynomial truthful budget feasible mechanisms for knapsack.*

4 Lower Bounds

In this section we focus on lower bounds for the approximation ratio of truthful budget feasible mechanisms for knapsack. Note that the same lower bounds can be applied to the general monotone submodular functions as well. In [21], a lower bound of 2 is obtained by the following argument: Consider the case with two items, both of unit value (the value of two items together is 2). If their costs are $(B-\epsilon, B-\epsilon)$, at least one item should win, otherwise the approximation ratio is infinite. Without loss of generality, we can assume that the first item wins, and as a result its payment is at least $B-\epsilon$. Now consider another profile $(\epsilon, B-\epsilon)$, the first item should also win (due to monotonicity) and get payment of at least $B-\epsilon$ by truthfulness. The second item then could not win because of the budget constraint and individual rationality. Therefore, the mechanism can only achieve a value of 1 for that instance while the optimal solution is 2. This gives us the lower bound of 2.

We improve the deterministic lower bound to $1+\sqrt{2}$ by a more involved proof. We also prove a lower bound of 2 for universally randomized truthful mechanisms. All our lower bounds are unconditional, which implies that we do not impose any complexity assumptions and constraints of the running time on the mechanism. Our lower bounds rely only on truthfulness and budget feasibility.

4.1 Deterministic Lower Bound

THEOREM 4.1. *No deterministic truthful budget feasible mechanism can achieve an approximation ratio better than $1+\sqrt{2}$, even if there are only three items.*

Assume otherwise that there is a budget feasible truthful mechanism that can achieve a ratio better than $1+\sqrt{2}$. We consider the following scenario: Budget $B=1$, and values $v_1=\sqrt{2}$, $v_2=v_3=1$. Then the mechanism on a scenario has the following two properties: (i) If all items are winners in the optimal solution, the mechanism must output at least two items; and (ii) if $\{1,2\}$ or $\{1,3\}$ is the optimal solution, the mechanism cannot output either $\{2\}$ or $\{3\}$ (i.e., a single item with unit value). For any item i, let function $p_i(c_j, c_k)$ be the payment offered to item i given that the bids of the other two items are c_j and c_k. That is, $p_i(c_j, c_k)$ is the threshold bid of i to be a winner.

LEMMA 4.1. *For any $c_3 > 0.5$ and any domain $(a,b) \subset (0, 1-c_3)$, there is $c_2 \in (a,b)$ such that $p_1(c_2, c_3) < 1-c_2$.*

Proof. Assume otherwise that there are $c_3 > 0.5$ and domain $(a,b) \subset (0,1-c_3)$ such that for any $c_2 \in (a,b)$, $p_1(c_2,c_3) \geq 1-c_2$. Let $c_1 = 1-c_3-b$, then $c_1+c_2+c_3 < 1 = B$, which implies that the mechanism has to output at least two items. Since $0 < c_1 = 1-c_3-b < 1-c_2 \leq p_1(c_2,c_3)$, item 1 is a winner. Further, $p_1(c_2,c_3) \geq 1-c_2 > 0.5$, which together with budget feasibility implies that item 3 cannot be a winner. Therefore, item 2 must be a winner with payment $p_2(c_1,c_3) = c_2$ due to individual rationality and budget feasibility. The same analysis still holds if the true cost of item 2 becomes $c_2' = \frac{c_2+b}{2}$, i.e., item 2 is still a winner with payment c_2'. Thus for the sample (c_1,c_2,c_3) the payment satisfies $p_2(c_1,c_3) \geq c_2' > c_2$, a contradiction. □

Since item 2 and 3 are identical, the above lemma still holds if we switch item 2 and 3 in the claim. We are now ready to prove Theorem 4.1.

Proof of Theorem 4.1. Define $c_3 = 0.7$ and $(a,b) = (0.2, 0.3)$. Note that c_3 and (a,b) satisfy the condition of Lemma 4.1. Hence, there is $c \in (0.2, 0.3)$ such that $p_1(c, 0.7) < 1-c$. Define $p_1(c,0.7) = 1-c-x$, where $x > 0$. Symmetrically, define $c_2 = 0.7$ and $(a',b') = (c, \min\{0.3, c+x\})$. Again by Lemma 4.1, there is $d \in (a',b')$ such that $p_1(0.7,d) < 1-d$. Define $p_1(0.7,d) = 1-d-y$, where $y > 0$. Pick $c_1 = 1-d-\epsilon$, where $\epsilon > 0$ is sufficiently small so that $c_1 \in (1-c-x, 1-c) \cap (1-d-y, 1-d)$. Note that since $d \in (c, c+x)$, c_1 is well-defined.

Consider a true cost vector $(c_1, c, 0.7)$. Since $p_1(c,0.7) = 1-c-x < c_1$, item 1 cannot be a winner. Since $c_1 + c = 1-d-\epsilon+c < 1$, the optimal solution has a value of at least $v_1+v_2 = 1+\sqrt{2}$; therefore the mechanism has to output both items 2 and 3. Hence, $p_3(c_1,c) \geq c_3 = 0.7$.

Similarly, consider true cost vector $(c_1, 0.7, d)$; we have $p_2(c_1,d) \geq c_2 = 0.7$. Finally, consider cost vector (c_1,c,d). By the above two inequalities, both items 2 and 3 are the winners; this contradicts the budget feasibility.

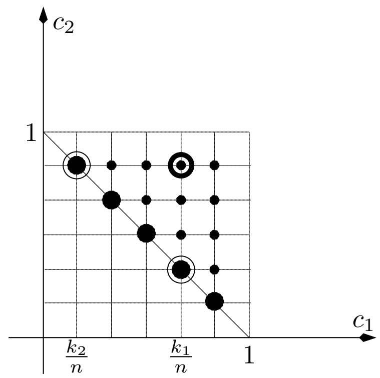

Figure 1: Distribution for $n = 6$. Width of a point emphasizes its probability.

4.2 Randomized Lower Bound

Theorem 4.2. *No randomized (universally) truthful budget feasible mechanism can achieve an approximation ratio better than* 2*, even in the case of two items.*

Proof. We use Yao's min-max principle, which is a typical tool used to prove lower bounds, where we need to design a distribution of instances and argue that any deterministic budget feasible mechanism cannot get an expected approximation ratio which is better than 2.

All the instances contain two items both with value 1. Their costs (c_1, c_2) are drawn from the following distribution (see Fig. 1 for an example):

1. $(\frac{kB}{n}, \frac{(n-k)B}{n})$ with probability $\frac{1-\epsilon}{n-1}$, where $k = 1, 2, \ldots, n-1$,
2. $(\frac{iB}{n}, \frac{jB}{n})$ with probability $\frac{2\epsilon}{(n-1)(n-2)}$, where $i, j \in \{1, \ldots, n-1\}$ and $i + j > n$,

where $1 > \epsilon > 0$ and n is a large integer.

We first claim that for any deterministic truthful budget feasible mechanism with finite expected approximation ratio, there is at most one instance, for which both items win in the mechanism. Assume for contradiction that there are at least two such instances. Note that for the second distribution $(\frac{iB}{n}, \frac{jB}{n})$, where $i + j > n$, it cannot be the case that both items win due to the budget constraint. Hence, the two instances must be of the first type; denote them as $(\frac{k_1B}{n}, \frac{(n-k_1)B}{n})$ and $(\frac{k_2B}{n}, \frac{(n-k_2)B}{n})$, where $k_1 > k_2$. Consider then the instance $(\frac{k_1B}{n}, \frac{(n-k_2)B}{n})$. Since $k_1 + n - k_2 > n$, this is the instance of the second type in our distribution. Therefore it has nonzero probability (see Fig. 1). The mechanism has a finite approximation ratio, thus it should have a finite approximation ratio on the instance $(\frac{k_1B}{n}, \frac{(n-k_2)B}{n})$ as well. As a result, it cannot be the case that both items lose. We assume that item 1 wins (the proof for the other case is similar); the payment to him is at least $\frac{k_1B}{n}$ due to individual rationality. Then consider the original instance $(\frac{k_2B}{n}, \frac{(n-k_2)B}{n})$; item 1 should also win and get a threshold payment, which is equal to or greater than $\frac{k_1B}{n}$. Therefore the payment to the second item is at most $B - \frac{k_1B}{n} = \frac{(n-k_1)B}{n}$ because of the budget constraint. Since $\frac{(n-k_1)B}{n} < \frac{(n-k_2)B}{n}$, we arrive at a contradiction with either individual rationality or the assumption that both items won in the instance $(\frac{k_2B}{n}, \frac{(n-k_2)B}{n})$.

On the other hand, for all instances $(\frac{kB}{n}, \frac{(n-k)B}{n})$, both items win in the optimal solution with value 2. Hence, the expected approximation ratio of any deterministic truthful budget feasible mechanism is at least $\frac{1-\epsilon}{n-1} \cdot 1 + (n-2) \cdot \frac{1-\epsilon}{n-1} \cdot 2 + \epsilon \cdot 1 = 2 - \epsilon - \frac{1-\epsilon}{n-1}$. The ratio approaches 2 when $\epsilon \to 0$ and $n \to \infty$. □

5 Beyond Submodularity

A natural generalization of knapsack is to consider heterogeneous items, i.e., items are partitioned into groups and we can select at most one item from each group. Formally, we are given m different types of items and each item has a (private) cost c_i and a (public) value v_i, as well as an indicator $t_i \in [m]$ standing for the type of item i. The goal is to pick items of different types[5] to maximize the total value given a budget constraint B. The knapsack problem studied in the last section is therefore a special case of the heterogeneous problem when all items are of different types. However, we cannot simply apply the mechanisms for knapsack here because of heterogeneity. (Notice however that the lower bounds established in the last section still work.)

The main difference of this problem with knapsack or general monotone submodular functions is that here not every subset is a feasible solution[6]. A straightfor-

[5]One may consider a relaxed version of heterogeneous knapsack, where any subset is feasible and its value is defined to be the sum of the maximum values of all types. That relaxed version is also known as a OXS function, a subclass of submodular functions defined in [15]; hence, our mechanisms for submodular functions can be applied here.

[6]For example, in some advertising markets, it is required that competitors' ads cannot be listed together due to negative externalities. This extra constraint that one cannot pick more than one item from the same type makes our problem different from the relaxed problem. In particular, we cannot treat our heterogeneous knapsack as a submodular function problem. Moreover, it does not even belong to XOS, a quite general class of valuation functions defined in [15] containing OXS and submodular functions.

ward greedy algorithm could end up with a very poor solution: Consider a situation where every type contains one very small item (both v_i and c_i are very small) but with a large value cost ratio of $\frac{v_i}{c_i}$; greedy algorithm will take all these small items first and therefore not be able to take more since each type already has one item. The overall value of this greedy solution can be arbitrarily bad compared to the optimal solution.

To construct a truthful mechanism for heterogeneous knapsack, we employ a greedy strategy with *deletions*. The main idea is that at every time the algorithm making a greedy move, we consider two possible changes: (i) Add a new item whose type has not been considered before, or (ii) replace an existing item with a new one of same type. Among all the possible choices (of two cases), we greedily select items with the highest value to cost ratio: In the case of adding a new item, its value cost to ratio is defined as usually $\frac{v_i}{c_i}$. For the replacement case where we replace i with j, its marginal value is $v_j - v_i$ and marginal cost is $c_j - c_i$, and hence its value to cost ratio is defined as $\frac{v_j - v_i}{c_j - c_i}$.

As before, now we assume that all the items are ordered according to their appearances in the greedy algorithm (note that some items never appear in the algorithm and we simply ignore them). The following greedy strategy is similar to what we did for the knapsack problem. In Appendix C, we prove that it is monotone (therefore truthful) and budget feasible. (Here for notational simplicity, assume that we already take an item with $c = 0$ and $v = 0$ for each type, thus every greedy step can be viewed as a replacement.)

GREEDY-HK

1. Let $k = 1$, $S = \emptyset$, and $last[j] = 0$ for $j \in [m]$
2. While $k \leq |A|$ and $c(k) - c(last[t_k]) \leq B \cdot \frac{v(k) - v(last[t_k])}{v(k) - v(last[t_k]) + \sum_{i \in S} v(i)}$
   - let $S \leftarrow (S \setminus \{last[t_k]\}) \cup \{k\}$
   - let $last[t_k] = k$
   - let $k \leftarrow k + 1$
3. Return winning set S

By applying the above GREEDY-HK, we have the following claim for heterogeneous knapsack. (Details can be found in Appendix C.)

THEOREM 5.1. *There are $2 + \sqrt{2}$ approximation deterministic and 3 approximation randomized polynomial truthful budget feasible mechanisms for knapsack with heterogeneous items.*

Finally, we comment that greedy approach is typically the first choice when one considers designing truthful mechanisms because it usually has a nice monotone property. However, when we allow cancelations in the greedy process, its monotonicity may fail. In the heterogeneous knapsack problem, fortunately GREEDY-HK is still monotone (although its proof is much more involved) and therefore we are able to apply it to design truthful mechanisms with good approximation ratios. Our idea sheds light on the possibility of exploring budget feasible mechanisms in larger domains beyond submodularity.

Acknowledgements

We thank Yaron Singer for helpful discussions.

References

[1] A. Archer, É. Tardos, *Frugal Path Mechanisms*, SODA 2002, 991-999.

[2] S. Bhattacharya, G. Goel, S. Gollapudi K. Munagala, *Budget Constrained Auctions with Heterogeneous Items*, STOC 2010.

[3] S. Chawla, J. Hartline, D. Malec, B. Sivan, *Sequential Posted Pricing and Multi-Parameter Mechanism Design*, STOC 2010.

[4] N. Chen, E. Elkind, N. Gravin, F. Petrov, *Frugal Mechanism Design via Spectral Techniques*, FOCS 2010.

[5] G. Christodoulou, E. Koutsoupias, A. Vidali, *A lower Bound for Scheduling Mechanisms*, SODA 2007, 1163-1170.

[6] S. Dobzinski, *A Note on the Power of Truthful Approximation Mechanisms*, arXiv:0907.5219v2.

[7] S. Dobzinski, R. Lavi, N. Nisan, *Multi-Unit Auctions with Budget Limits*, FOCS 2008, 260-269.

[8] U. Feige, V. Mirrokni, J. Vondrak, *Maximizing Non-Monotone Submodular Functions*, FOCS 2007, 461-471.

[9] S. Khuller, A. Moss, J. S. Naor, *The Budgeted Maximum Coverage Problem*, Information Processing Letter, V.70(1), 39-45, 1999.

[10] A. R. Karlin, D. Kempe, T. Tamir, *Beyond VCG: Frugality of Truthful Mechanisms*, FOCS 2005, 615-626.

[11] D. Kempe, M. Salek, C. Moore, *Frugal and Truthful Auctions for Vertex Covers, Flows, and Cuts*, FOCS 2010.

[12] E. Koutsoupias, A. Vidali, *A Lower Bound of $1+\phi$ for Truthful Scheduling Mechanisms*, MFCS 2007, 454-464.

[13] A. Krause, C. Guestrin, *A Note on the Budgeted Maximization of Submodular Functions*, CMU-CALD-05-103.

[14] J. Lee, V. Mirrokni, V. Nagarajan, M. Sviridenko, *Non-Monotone Submodular Maximization under Matroid and Knapsack Constraints*, STOC 2009, 323-332.

[15] B. Lehmann, D. Lehmann, N. Nisan, *Combinatorial Auctions with Decreasing Marginal Utilities*, EC 2001, 18-28.

[16] R. Myerson, *Optimal Auction Design*, Mathematics of Operations Research, V.6(1), 1981.
[17] G. Nemhauser, L. Wolsey, M. Fisher, *An Analysis of the Approximations for Maximizing Submodular Set Functions*, Mathematical Programming, V.14, 1978, 265-294, 1978.
[18] N. Nisan, A. Ronen, *Algorithmic Mechanism Design*, Games and Economic Behavior, V.35(1-2), 166-196, 2001.
[19] C. H. Papadimitriou, M. Schapira, Y. Singer, *On the Hardness of Being Truthful*, FOCS 2008, 250-259.
[20] A. D. Procaccia, M. Tennenholtz, *Approximate Mechanism Design without Money*, EC 2009, 177-186.
[21] Y. Singer, *Budget Feasible Mechanisms*, FOCS 2010.
[22] M. Sviridenko, *A Note on Maximizing a Submodular Set Function Subject to a Knapsack Constraint*, Operation Research Letter,, V.32(1), 41-43, 2004.
[23] K. Talwar, *The Price of Truth: Frugality in Truthful Mechanisms*, STACS 2003, 608-619.

A Proof of Lemma 3.3

Proof. Let w_i denote the weight of each item $i \in [n]$. Our goal in the weighted problem is to pick a set S with total weight $\sum_{i\in S} w_i$ not exceeding given capacity W of maximal possible utility $u(S)$, where u is the given monotone submodular function. As the utility $u_f(S_f)$ is for fractional problem we consider the expectation of $u(S)$, where each $i \in [n]$ is selected at random independently in S with probability equal to the fraction of item i in S_f.

Assume that all weights w_i are integers. We reduce our weighted problem with monotone submodular function u to the unweighted one as follows.

- For each item $i \in [n]$ we consider w_i new items of unit weight. Denote them as i_j for $j \in [w_i]$ and call i the type of the unit i_j.
- The new valuation function ν only depends on the amounts of unit items of each type.
- Let a set S contain a_i units of each type i. Independently for each type, pick at random in the set $\mathcal{R}$ with probability $\frac{a_i}{w_i}$ weighted item i. Define $\nu(S) = E(u(\mathcal{R}))$.

Therefore

$$\nu(S) = \frac{1}{w_1 \cdot \ldots \cdot w_n} \sum_{\pi} u\left(S \cdot \pi\right)$$

where π is a sampling of units one for each type (there are $w_1 \cdot \ldots \cdot w_n$ variants for π); $S \cdot \pi$ is a vector of types at which π hits S.

Using this formula it is not hard to verify monotonicity and submodularity of ν. Indeed, e.g. to verify submodularity one only needs to check that the marginal contribution of any unit is smaller for a large set, i.e., for $S \subset T$ and $i_j \notin T$ verify inequality $\nu(S \cap \{i_j\}) - \nu(S) \geq \nu(T \cap \{i_j\}) - \nu(T)$, which is pretty straightforward.

For any $T \subseteq [n]$ if we consider a set of units $S = \{i_k | i \in T, 1 \leq k \leq w_i\}$, then according to the definition $\nu(S) = u(T)$. Hence, the optimal solution to the unit weights problem is equal to or larger than the optimal solution to the original problem.

To conclude the proof it is only left to show that our fractional greedy scheme for an integer weighted instance gives us the same result as the greedy scheme for its unit weighted version. Note that once we have taken a unit of type i we will proceed to take units of type i until it is exhausted completely (we brake ties in favor of the last type we have picked). Indeed, let $i_k, i_{k+1} \notin S$ then

$$\begin{aligned}
& \nu(S \cup \{i_k\}) - \nu(S) \\
= \; & \nu(S \cup \{i_{k+1}\}) - \nu(S) \\
= \; & \frac{1}{w_1 \cdot \ldots \cdot w_n} \sum_{\{\pi | i_{k+1} \in \pi\}} u\left(S \cup \{i_{k+1}\} \cdot \pi\right) - u\left(S \cdot \pi\right) \\
= \; & \frac{1}{w_1 \cdot \ldots \cdot w_n} \sum_{\{\pi | i_{k+1} \in \pi\}} u\left(S \cup \{i_k, i_{k+1}\} \cdot \pi\right) \\
& - u\left(S \cup \{i_k\} \cdot \pi\right) \\
= \; & \nu(S \cup \{i_k, i_{k+1}\}) - \nu(S \cup \{i_k\})
\end{aligned}$$

Therefore, the marginal contribution of the type i does not decrease if we include in the solution units of type i. On the other hand, because ν is submodular, the marginal contribution of any other type cannot increase. So we will take unit i_{k+1} right after i_k.

Assume we already have picked set S and now are picking the first unit of a type i. Hence, S comprises all units of a type set T. Then we have

$$\begin{aligned}
& \nu\left(S \cup \{i_1\}\right) - \nu\left(S\right) \\
= \; & \frac{1}{\prod_{k=1}^{n} w_k} \sum_{\{\pi | i_1 \in \pi\}} u\left(S \cup \{i_1\} \cdot \pi\right) - u\left(S \cdot \pi\right) \\
= \; & \frac{\prod_{k \neq i} w_k}{\prod_{k=1}^{n} w_k} m_T(i) = \frac{m_T(i)}{w_i}
\end{aligned}$$

Thus $i = argmax_{i \notin T} \frac{m_T(i)}{w_i}$ which coincides with the rule of our fractional greedy scheme.

In case of w_i being real weights the same approach can be applied but in a more tedious way. □

B Mechanisms for Knapsack

In this section, we describe our deterministic and randomized mechanisms for knapsack, yielding a proof for Theorem 3.3.

B.1 Deterministic Mechanism. We consider the following greedy strategy studied by Singer [21].

> GREEDY-KS(A)
>
> 1. Order all items in A s.t. $\frac{v_1}{c_1} \geq \frac{v_2}{c_2} \geq \cdots \geq \frac{v_{|A|}}{c_{|A|}}$
> 2. Let $k = 1$ and $S = \emptyset$
> 3. While $k \leq |A|$ and $c_k \leq B \cdot \frac{v_k}{\sum_{i \in S \cup \{k\}} v_i}$
>   - $S \leftarrow S \cup \{k\}$
>   - $k \leftarrow k + 1$
> 4. Return winning set S

It is shown that the above greedy strategy is monotone (and therefore truthful). Actually, it has the following remarkable property: Any $i \in S$ cannot control the output set given that i is guaranteed to be a winner. That is, if the winning sets are S and S' when i bids c_i and c'_i, respectively, where $i \in S \cap S'$, then $S = S'$. Otherwise, consider the item $i_0 \notin S \cap S'$ with the smallest index; assume without loss of generality that $i_0 \in S \setminus S'$. Let $T = \{j \in S \cap S' \mid j < i_0, j \neq i\}$ be the winning items in $S \cap S' \setminus \{i\}$ before i_0. Then

$$c_{i_0} \leq B \cdot \frac{v_{i_0}}{\sum_{j \in S} v_j} \leq B \cdot \frac{v_{i_0}}{\sum_{j \in T} v_j + v_i + v_{i_0}},$$

which implies that i_0 should be a winner in S' as well, a contradiction.

Given the greedy strategy described above, our mechanism for knapsack is as follows (where $fopt(A)$ denotes the value of the optimal fractional solution; for knapsack it can be computed in polynomial time).

> DETERMINISTIC-KS
>
> 1. Let $A = \{i \mid c_i \leq B\}$ and $i^* \in \arg\max_{i \in A} v_i$
> 2. If $(1 + \sqrt{2}) \cdot v_{i^*} \geq fopt(A \setminus \{i^*\})$, return i^*
> 3. Otherwise, return $S =$ GREEDY-KS(A)

THEOREM B.1. DETERMINISTIC-KS *is a* $2 + \sqrt{2}$ *approximation deterministic budget feasible truthful mechanism for knapsack.*

Proof. The proof consists of each property stated in the claim.

- *Truthfulness.* We analyze monotonicity of the mechanism according to the condition of Steps 2 and 3, respectively. If i^* wins in Step 2 (note that the fractional optimal value computed in Step 2 is independent of the bid of i^*), then i^* still wins if he decreases his bid.

  If the condition in Step 2 fails and the mechanism runs to Step 3, for any $i \in S$, the subset S remains the same if i decreases his bid. Note that if $i \neq i^*$, when i decreases his bid, the value of the fractional optimal solution computed in Step 2 will not decrease. Hence i is still a winner, which implies that the mechanism is monotone.

- *Individual rationality and budget feasibility.* If i^* wins in Step 2, his payment is the threshold bid B. Otherwise, assume that all buyers in A are ordered by $1, 2, \ldots, n$; let $S = \{1, \ldots, k\}$. Note that it is possible that $i^* \in S$. For any $i \in S$, let q_i be the maximum cost that i can bid such that the fractional optimal value on instance $A \setminus \{i^*\}$ is still larger than $(1 + \sqrt{2})v_{i^*}$. Note that $c_i \leq q_i$ and as opposed to general submodular case the marginal contribution v_i does not depend on the ranking of i.

  Thus, the payment to any winner $i \in S \setminus \{i^*\}$ is

  $$p_i = \min\left\{v_i \cdot \frac{c_{k+1}}{v_{k+1}}, B \cdot \frac{v_i}{\sum_{j \in S} v_j}, q_i\right\},$$

  and

  $$p_{i^*} = \min\left\{v_{i^*} \cdot \frac{c_{k+1}}{v_{k+1}}, B \cdot \frac{v_{i^*}}{\sum_{j \in S} v_j}\right\},$$

  if $i^* \in S$. It can be seen that the mechanism is individually rational. Further, $\sum_{i \in S} p_i \leq \sum_{i \in S} B \cdot \frac{v_i}{\sum_{j \in S} v_j} = B$, which implies that the mechanism is budget feasible.

- *Approximation.* Assume that all buyers in A are ordered by $1, 2, \ldots, n$, and $T = \{1, \ldots, k\}$ is the subset returned by GREEDY-KS(A). Let ℓ be the maximal item for which $\sum_{i=1,\ldots,\ell} c_i \leq B$. Let $c'_{\ell+1} = B - \sum_{i=1,\ldots,\ell} c_i$ and $v'_{\ell+1} = v_{\ell+1} \cdot \frac{c'_{\ell+1}}{c_{\ell+1}}$. Hence, the optimal fractional solution is

  $$fopt(A) = \sum_{i=1}^{\ell} v_i + v'_{\ell+1}$$

  For any $j = k+1, \ldots, \ell$, we have

  $$\frac{c_j}{v_j} \geq \frac{c_{k+1}}{v_{k+1}} > \frac{1}{v_{k+1}} \cdot B \cdot \frac{v_{k+1}}{\sum_{i=1}^{k+1} v_i},$$

where the last inequality follows from the fact that the greedy strategy stops at item $k+1$. Hence, $c_j > B \cdot \frac{v_j}{\sum_{i=1}^{k+1} v_i}$. The same analysis shows $c'_{\ell+1} > B \cdot \frac{v'_{\ell+1}}{\sum_{i=1}^{k+1} v_i}$. Therefore,

$$B \cdot \frac{\sum_{j=k+1}^{\ell} v_j + v'_{\ell+1}}{\sum_{i=1}^{k+1} v_i} < \sum_{j=k+1}^{\ell} c_j + c'_{\ell+1} < B,$$

which implies that $\sum_{i=1}^{k} v_i > \sum_{j=k+2}^{\ell} v_j + v'_{\ell+1}$. Hence,

$$fopt(A) = \sum_{i=1}^{\ell} v_i + v'_{\ell+1} < 2\sum_{i \in S} v_i + v_{i^*}$$

A basic observation of the mechanism is that

$$fopt(A) - v_{i^*} \le fopt(A \setminus \{i^*\}) \le fopt(A)$$

Hence, if the condition in Step 2 holds and the mechanism outputs i^*, we have

$$fopt(A) \le fopt(A \setminus \{i^*\}) + v_{i^*} \le (2+\sqrt{2}) \cdot v_{i^*}$$

If the condition in Step 3 fails and the mechanism outputs S in Step 4, we have

$$\begin{aligned}(1+\sqrt{2}) \cdot v_{i^*} &< fopt(A \setminus \{i^*\}) \\ &\le fopt(A) \\ &< 2\sum_{i \in S} v_i + v_{i^*}\end{aligned}$$

which implies that $v_{i^*} < \sqrt{2} \cdot \sum_{i \in S} v_i$. Hence,

$$\begin{aligned}opt(A) \le fopt(A) &= \sum_{i=1,\ldots,\ell} v_i + v'_{\ell+1} \\ &< 2\sum_{i \in S} v_i + v_{i^*} \\ &\le (2+\sqrt{2}) \cdot \sum_{i \in S} v_i.\end{aligned}$$

Therefore, the mechanism is a $(2+\sqrt{2})$ approximation.

□

B.2 Randomized Mechanism. Our randomized mechanism for knapsack is as follows.

RANDOM-KS

1. Let $A = \{i \mid c_i \le B\}$ and $i^* \in \arg\max_{i \in A} v_i$
2. With probability $\frac{1}{3}$, return i^*
3. With probability $\frac{2}{3}$, return GREEDY-KS(A)

THEOREM B.2. RANDOM-KS *is a 3 approximation universal truthful budget feasible mechanism for knapsack.*

Proof. Since both mechanisms in Steps 2 and 3 are budget feasible and truthful, it is left only to prove the approximation ratio.

Using the same notation and argument in the proof of Theorem B.1, assume that all buyers in A are ordered by $1, 2, \ldots, n$, and $T = \{1, \ldots, k\}$ is the subset returned by GREEDY-KS(A). Let ℓ be the maximal item for which $\sum_{i=1,\ldots,\ell} c_i \le B$. Let $c'_{\ell+1} = B - \sum_{i=1,\ldots,\ell} c_i$ and $v'_{\ell+1} = c'_{\ell+1} \cdot \frac{v_{\ell+1}}{c_{\ell+1}}$. Hence, the optimal fractional solution is

$$fopt(A) = \sum_{i=1}^{\ell} v_i + v'_{\ell+1}$$

and

$$fopt(A) = \sum_{i=1}^{\ell} v_i + v'_{\ell+1} < v_{i^*} + 2\sum_{i \in S} v_i.$$

The excepted value of RANDOM-KS is therefore

$$\frac{1}{3} v_{i^*} + \frac{2}{3}\sum_{i \in S} v_i = \frac{1}{3}\Big(v_{i^*} + 2\sum_{i \in S} v_i\Big) > \frac{1}{3} opt$$

which completes the proof. □

C Knapsack with Heterogeneous Items

In this section we analyze the heterogeneous knapsack problem and GREEDY-HK, which leads to a proof of Theorem 5.1.

C.1 Optimal Fractional Solution. We start our study again on fractional solutions to the optimization problem. First we have to define what is a fractional relaxation for heterogeneous knapsack or more precisely what is a feasible fractional solution.

A feasible solution for heterogeneous knapsack is an n-tuple of real numbers $(\alpha_1, \ldots, \alpha_n) \in [0,1]^n$ satisfying $\sum_{i=1}^{n} \alpha_i c_i \le B$ and $\sum_{i \in t_j^-} \alpha_i \le 1$ for any $j \in [m]$. An optimal fractional solution is a feasible solution that maximizes $\sum_{i=1}^{n} \alpha_i v_i$.

We have the following observation on the optimal solution.

LEMMA C.1. *For a given budget B we can pick an optimal fractional solution f_{OPT} such that*

- *there are at most two nonzero amounts of items of any type in f_{OPT}.*
- *there is exactly one item of any type in f_{OPT} except maybe only for one type.*

Proof. Consider any optimal solution f'_{OPT}. Fix the price p_j spent on the particular type j in it. We can use only two items of type j in order to provide the maximum value for the price p_j. Indeed, if one draws all items of type j in the plain with the x-coordinate corresponding to the cost and the y-coordinate corresponding to the value of an item together with the point $(0, 0)$, then the condition $\sum_{i \in t_j^-} \alpha_i \leq 1$ will describe a point in the convex hull of the drawn set.

Thus we can take f_{OPT} with at most two items of a type and derive the first part of the lemma.

One can derive the second part of the lemma by changing p_{j_1} and p_{j_2} in f_{OPT} such that $p_{j_1} + p_{j_2}$ remains constant. Indeed, appealing to the picture again, we consider two convex polygons P_1 and P_2 for the types j_1 and j_2. If both prices p_{j_1} and p_{j_2} get strictly inside the corresponding sides of those polygons, then by stirring p_{j_1} and p_{j_2} in f_{OPT} and keeping $p_{j_1} + p_{j_2}$ constant we can get to a vertex of P_1 or P_2 that does not decrease the total value. □

The following algorithm computes an optimal fractional solution for heterogeneous knapsack. (For convenience we add an item numbered by 0 of a new type with cost 0 and value 0; this does not affect any optimal solution.)

FRACTION-HK

1. For each type $j \in [m]$, (partially) order items of type j as follows:
   - let $last = 0$, $tg = 0$ and $A_j = \emptyset$
   - while $v(last) < \max_{i \in t_j^-} v(i)$
     - let $k = \arg\max_{i \in t_j^-} \frac{v(i) - v(last)}{|c(i) - c(last)|}$ and add k to A_j
     - define $tg_k = \frac{v(k) - v(last)}{|c(k) - c(last)|}$
     - let $last = k$
2. Comprise all A_j into one big set A and order all items s.t. $tg_1 \geq \cdots \geq tg_{|A|}$
3. Let $last[j] = 0$ for each $j \in [m]$, $\alpha_i = 0$ for each $i \in [n]$ and $k = 1$
4. While $k \leq |A|$ and $c_k + \sum_{i=1}^{k-1} \alpha_i \cdot c_i \leq B$
   - let $\alpha_{last[t_k]} \leftarrow 0$
   - let $last[t_k] \leftarrow k$, $\alpha_k \leftarrow 1$
   - let $k \leftarrow k + 1$
5. If $k \leq |A|$, then let $\alpha_k = \frac{B - \sum_{i=1}^{k-1} \alpha_i c_i}{c_k}$ and $\alpha_{last[t_k]} = 1 - \alpha_k$
6. Return vector $(\alpha_i)_{i \in [n]}$

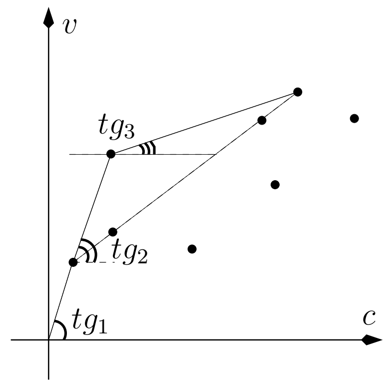

Figure 2: Convex hull

THEOREM C.1. FRACTION-HK *computes an optimal fractional solution for heterogeneous knapsack.*

Proof. If we draw every item $i \in t_j^- \cup \{0\}$ as a point (c_i, v_i) in the plain (see Fig. 2), then all picked items in A_j will correspond to the part of the convex hull's vertices of the drawn set from $(0, 0)$ to the item with maximal value. The computed value of tg will correspond then to the tangent of the side of the convex hull with the right end at the given item.

As in the proof of lemma C.1 one can find the optimal value, that we can get for a type j at the price c, by taking the y-coordinate of the point on a side of the convex hull with c at the x-coordinate. Thus for the optimal fractional solution we only need items from $A = \cup_j A_j$.

Taking everything above into account we can reduce the heterogeneous knapsack to the basic knapsack problem. Fix a type j and construct the instance of the reduced problem $\tilde{K}_j$ as follows. For each item $k \in A_j$ assign the cost $\tilde{c}_k := c_k - c(last[t_k])$ and the value $\tilde{v}_k := v_k - v(last[t_k])$. It is easy to see that the optimal solution to the basic knapsack problem $\tilde{K}_j$ gives the same value as the solution to the original heterogeneous problem restricted to the items of type j for any given budget. Hence the optimal fractional solution to basic knapsack problem $\cup_j \tilde{K}_j$ has the same value as the optimal fractional solution to the original problem.

Now it is easy to check that our algorithm at stages $2 - 5$ computes the optimal fractional solution to the reduced knapsack problem and thus finds the optimal fractional solution to our original problem. □

C.2 Greedy Strategy with Deletions. We consider the following greedy strategy mechanism.

Greedy-HK

1. `Take the same ordered set` A `as in Step 2 of` Fraction-HK
2. `Let` $k = 1$, $S = \emptyset$, `and` $last[j] = 0$ `for` $j \in [m]$
3. `While` $k \le |A|$ `and` $c(k) - c(last[t_k]) \le B \cdot \frac{v(k)-v(last[t_k])}{v(k)-v(last[t_k])+\sum_{i \in S} v(i)}$
   - `let` $S \leftarrow (S \setminus \{last[t_k]\}) \cup \{k\}$
   - `let` $last[t_k] = k$
   - `let` $k \leftarrow k + 1$
4. `Return winning set` S

Recall the notation in the algorithm Fraction-HK, $tg_k = \frac{v(k)-v(last[t_k])}{|c(k)-c(last[t_k])|}$, where $last[t_k]$ is the last item of type t_k in A at the moment when we are about to add k into A. Define $S_k = (S \setminus \{last[t_k]) \cup \{k\}$. Then the second condition in Step 3 of Greedy-HK can be rewritten as

$$tg_k \ge \frac{v(S_k)}{B}$$

We next analyze the mechanism Greedy-HK. Let us denote by $\mathcal{M}_b$ the run of mechanism Greedy-HK on bid b (with the corresponding ordered set A_b, the last item of each type $last_b[t_k]$ and marginal tangent $tg_k(\mathcal{M}_b)$).

Claim C.1. Greedy-HK *is monotone (and therefore truthful).*

Proof. We will show that any losing item cannot bid more and become a winner. Assume otherwise that item j loses with bid c_j but wins with bid $b_j > c_j$, given that all others bid c_i, $i \neq j$.

Note that when j changes his bid, it will only affect the convex hull of items in $t_j^- \cup \{0\}$. The following observations can be verified easily (see Fig. 2):

1. Values $v(S)$ of the set of winners and $v(last[t_k])$ for each type t_k, taken dynamically in the process of the mechanism, keep increasing.
2. Value tg_j decreases when j increases its bid (since point (b_j, v_j) is on the right hand side of point (c_j, v_j)).
3. Ordered set $A_b \setminus t_j^-$ is the same as ordered set $A_c \setminus t_j^-$

By considering the convex hull for t_j^-, one can easily see that if j was not, at any moment, getting into the winning set S in $\mathcal{M}_c$ it also will never get in the winning set in $\mathcal{M}_b$.

Let us explain why when j increases its bid that it cannot help to remain in the winning set if, for the current cost c_j, it has been dropped off.

Note that in the new ordered set A_b, there can be new items of the same type as j (e.g. $last_c[j]$ can be different from $last_b[j]$), but nevertheless $tg_j(\mathcal{M}_b) \le tg_j(\mathcal{M}_c)$. Let $j' \in t_j^-$ be the item that substitutes j in $\mathcal{M}_c$, then $tg_{j'}(\mathcal{M}_c) \le tg_{j'}(\mathcal{M}_b)$ (note that j' necessarily appears in A_b). Let k be an item at which $\mathcal{M}_b$ has stopped, i.e., the first item that we have not taken in the winning set. Assume k stands in A_b not further than j'. Consider two cases.

1. Let $t_k \neq t_j$. Then
   - $tg_k(\mathcal{M}_c) = tg_k(\mathcal{M}_b)$
   - $v(S_{j'}(\mathcal{M}_c)) \ge v(S_k(\mathcal{M}_c))$, as j' stands later than k in A_c
   - $v(S_k(\mathcal{M}_c)) = v(S_k(\mathcal{M}_b))$, since in both $S_k(\mathcal{M}_b)$ and $S_k(\mathcal{M}_c)$, we have taken j for type t_j , and we also have taken the same items for all other types.
2. $t_k = t_j$. Then
   - $tg_{j'}(\mathcal{M}_c) \le tg_{j'}(\mathcal{M}_b) \le tg_k(\mathcal{M}_b)$
   - $v(S_{j'}(\mathcal{M}_c)) \ge v(S_k(\mathcal{M}_b))$. The last equality holds true, because for each type the value of the item in $S_{j'}(\mathcal{M}_c)$ is greater than or equal to the than value of the corresponding item in $S_k(\mathcal{M}_b)$.

In both cases we can write

$$\begin{aligned} tg_k(\mathcal{M}_b) &\ge tg_{j'}(\mathcal{M}_b) \ge tg_{j'}(\mathcal{M}_c) \\ &\ge \frac{v(S_{j'}(\mathcal{M}_c))}{B} \ge \frac{v(S_k(\mathcal{M}_b))}{B} \end{aligned}$$

Thus we have to take k in $\mathcal{M}_b$ to the winning set. Hence we arrive at a contradiction. Hence we have taken j' to the winning set in $\mathcal{M}_b$ and therefore exclude j. □

Unfortunately, in contrast to the knapsack case this scheme does not possess the following property: Any $i \in S$ cannot control the output set given that i is guaranteed to be a winner.

Claim C.2. *Let S be the winning set of* Greedy-HK *on cost vector c. Then no item $j \in S$ can remain a winner with bid b_j satisfying*

$$b_j > (v(j) - v(last_c[t_j])) \cdot \frac{B}{V(S)} + c(last_c[t_j])$$

Proof. Assume to the contrary that there exists such j and bid b_j. We can write

$$\begin{aligned} tg_j(\mathcal{M}_b) &= \frac{v(j) - v(last_b[t_j])}{b_j - c(last_b[t_j])} \\ &\le \frac{v(j) - v(last_c[t_j])}{b_j - c(last_c[t_j])} \\ &< \frac{v(S)}{B} \end{aligned}$$

Consider the ordered set A_c and let k be the last item we have taken in the winning set in $\mathcal{M}_c$. Now consider any item $i \in [1, k]$ where $t_j \neq t_i$. We have $\frac{v(S)}{B} \le tg_k(\mathcal{M}_c) \le tg_i(\mathcal{M}_c) = tg_i(\mathcal{M}_b)$. By the assumption that j is in the winning set in $\mathcal{M}_b$ and $tg_j(\mathcal{M}_b) < \frac{v(S)}{B} \le tg_i(\mathcal{M}_b)$, we get that $S_j(\mathcal{M}_b)$ contains an item i' with $t_i = t_{i'}$ and $v(i') \ge v(i)$. Since j is in S and in $S_j(\mathcal{M}_b)$ we get $v(S_j(\mathcal{M}_b)) \ge v(S)$. Hence

$$\frac{v(S)}{B} > tg_j(\mathcal{M}_b) \ge \frac{v(S_j(\mathcal{M}_b))}{B} \ge \frac{v(S)}{B}$$

which gives a contradiction. □

Claim C.3. *Greedy scheme* Greedy-HK *is budget feasible.*

Proof. Let S be a winning set for $\mathcal{M}$. By Claim C.2, we have an upper bound on the payment p_j to each item $j \in S$, i.e.,

$$p_j \le (v(j) - v(last_c[t_j])) \cdot \frac{B}{V(S)} + c(last_c[t_j])$$

Let $0 = i_0, i_1, \ldots, i_r, i_{r+1} = j$ be the items of type t_j that have appeared in the winning set. We have $tg_{i_\ell} \ge \frac{v(S)}{B}$ for each $\ell = 1, \ldots, r$. Hence

$$c(i_\ell) - c(i_{\ell-1}) \le (v(i_\ell) - v(i_{\ell-1}))\frac{B}{v(S)}$$

Now if we sum up the above inequalities on $c(i_l) - c(i_{l-1})$ for all $\ell = 1, \ldots, r$ and plug it in the bound on p_j, we get

$$p_j \le \frac{B}{v(S)} \sum_{\ell=1}^{r+1} v(i_\ell) - v(i_{\ell-1}) = v(j)\frac{B}{v(S)}$$

Therefore, $\sum_{j \in S} p_j \le B$, which concludes the proof. □

C.3 Mechanisms. Given the greedy strategy described above, our mechanism for heterogeneous knapsack is as follows.

Deterministic-HK

1. Let $A = \{i \mid c_i \le B\}$ and $i^* \in \arg\max_{i \in A} v_i$
2. If $(1 + \sqrt{2}) \cdot v_{i^*} \ge$ Fraction-HK$(A \setminus \{i^*\})$, return i^*
3. Otherwise, return $S =$ Greedy-HK

Theorem C.2. Deterministic-HK *is a* $2 + \sqrt{2}$ *approximation deterministic budget feasible truthful mechanism for heterogeneous knapsack.*

Proof. The proof consists of each property stated in the claim.

- *Truthfulness.* The same proof as for knapsack also works here.
- *Individual rationality and budget feasibility.* If i^* wins in Step 2, his payment is the threshold bid B. Otherwise, payment to each item has an upper bound from the payment rule in Greedy-HK and thus according to the claim C.3 final total payment will be below given budget B.
- *Approximation.* Return back to the algorithm for optimal fractional heterogeneous knapsack. Consider the stage where we add item k to a set A_j, let us define $\tilde{v}(k) = v(k) - v(last[t_k])$ and $\tilde{c}(k) = c(k) - c(last[t_k])$ to be modified value and cost of item k. Let us consider the fractional knapsack $\tilde{FK}$ problem for those modified costs and values for all items in A. It turns out that for any budget this new problem $\tilde{FK}$ has the same answer as initial heterogeneous knapsack HK. Note that our greedy scheme Greedy-KS for modified costs and values and our greedy scheme Greedy-HK for original heterogeneous knapsack also give the same answer. Thus applying the part *approximation* of claim B.1 to the modified problem we obtain the desired bound.

□

We can also have the following randomized mechanism with an approximation ratio of 3 (its proof is similar to Theorem B.2).

Random-HK

1. Let $A = \{i \mid c_i \le B\}$ and $i^* \in \arg\max_{i \in A} v_i$
2. With probability $\frac{1}{3}$, return i^*
3. With probability $\frac{2}{3}$, return $S =$ Greedy-HK

Theorem C.3. Random-HK *is a* 3 *approximation universal truthful budget feasible mechanism for heterogeneous knapsack.*