Check for updates

Frugal Auction Design for Set Systems: Vertex Cover and Knapsack

MOHAMMADTAGHI HAJIAGHAYI, University of Maryland, USA
MOHAMMAD REZA KHANI, Microsoft R&D, USA
SAEED SEDDIGHIN, University of Maryland, USA

We study mechanism design for procurement auctions in which the goal is to buy a subset of items or hire a team of providers. In order to measure the efficiency of a mechanism, one defines an appropriate benchmark which denotes a reasonable expectation of the payments and defines the overpayment of a mechanism based on the benchmark. This ratio is called the *frugality ratio* of the mechanism. Procurement auctions are well-studied and benchmarks proposed for these auctions have evolved over a sequence of papers [2, 5, 8, 12, 13]. In this work, we introduce a newer benchmark, and based on that, study classic procurement auctions. Our benchmark addresses critical issues raised by the unintuitive behavior of the previous benchmarks. We show two attractive properties for our benchmark which have been lacking in the previous proposals: *monotonicity* and *smoothness*.

Based on our benchmark, we provide positive results for vertex cover and knapsack auctions. Prior to this work, Kempe *et al.* [13] propose a constant approximation mechanism for vertex cover auctions. However, their analysis suffers from an error. We give a correct analysis to the mechanism of Kempe *et al.* [13] *with respect to our benchmark*. In particular, we prove their mechanism is optimal up to a constant factor. Our analysis is different from what Kempe *et al.* [13] propose. We also study the knapsack auctions and give a truthful mechanism for such auctions with a bounded frugality ratio. We show that this is almost tight by presenting a lower bound on the frugality ratio of any truthful mechanism for such auctions. *All our results depend on both properties of the benchmark.* [1]

CCS Concepts: • **Theory of computation → Algorithmic game theory and mechanism design**;

Additional Key Words and Phrases: Frugal, Knapsack, Vertex Cover

1 INTRODUCTION

Suppose we want to purchase a number of different items or hire a team to accomplish certain tasks. It is quite common that the items or people are interchangeable. In other words, every task can be performed by different groups of people which naturally leads to a competition between individuals. As an example, suppose a government wishes to connect a city to another one, by building roads between intermediary cities. Constructing each road can be entrusted with a contractor at a certain

Supported in part by NSF CAREER award CCF-1053605, NSF BIGDATA grant IIS-1546108, NSF AF:Medium grant CCF-1161365, DARPA GRAPHS/AFOSR grant FA9550-12-1-0423, and another DARPA SIMPLEX grant

[1] In the interest of space, some of the proofs and results have been omitted in this version and included in the full-version.

Authors' addresses: MohammadTaghi Hajiaghayi, University of Maryland, USA, hajiagha@cd.umd.edu; Mohammad Reza Khani, Microsoft R&D, USA, rezak@microsoft.com; Saeed Seddighin, University of Maryland, USA, saeedreza.seddighin@gmail.com.

ACM EC'18, June 18–22, 2018, Ithaca, NY, USA. ACM ISBN 978-1-4503-5829-3/18/06…$15.00
https://doi.org/10.1145/3219166.3219229

cost. It is not natural to expect the contractors to reveal the true cost of their service since they are selfishly motivated. Therefore one of the key difficulties, especially from the government's perspective, is to find out which subset of roads to build and what the actual cost of each construction is. A very effective and commonly used approach (see *e.g.* [10, 11]) is to design a procurement auction and hope that the competition between contractors keeps the costs low enough. The goal is then to design a mechanism for the auction to (i) make the providers truthful (i.e., leave no incentive for the providers to misrepresent their costs) (ii) do not overpay by much.

Procurement auctions have been extensively studied in the literature [4, 6, 7, 10, 11, 15]. A mechanism for a procurement auction is an algorithm that receives the bids as input and specifies a subset of bidders as the winning set. In addition to the selection of the providers, the payments are also determined by the mechanism. It is shown that if the selection rule is monotonic (decreasing the bid of each provider increases his likelihood of winning), the payments are uniquely determined by Myerson's characterization of truthful mechanisms. Truthfulness makes the market transparent and reveals the actual costs. However, as a downside, the payments of a truthful mechanism are inevitably higher than the announced costs. Therefore, it is crucial to define a measure to evaluate the overpayment and as such, the efficiency of mechanisms. One natural way to define the overpayment is to compare the total payment of the auctioneer to the total cost of the cheapest option. However, this formulation becomes unbounded for all truthful mechanisms in most cases and thus can't be used to evaluate mechanisms.

One way to achieve this goal is by defining a benchmark and comparing the payments of the mechanism to the benchmark as the ratio of the overpayment. Such a benchmark receives the costs as input and maps these values to a non-negative real number. The goal then is to minimize the ratio of the mechanism's total payment to the value of the benchmark. This formulation was first defined by the pioneering work of Archer and Tardos [2] who call this ratio the *frugality* ratio of a truthful mechanism. In their work, they also propose a benchmark which is the cost of the second cheapest option. However, this measure was specifically designed to study the shortest path auction, and hence it was not extensible to general problems. Later, several attempts were made to improve the benchmark of Archer and Tardos [2], but none of the proposed benchmarks seem to be a right fit for the problem setting (see the full-version for a more detailed review). Perhaps the reason that studying frugal mechanisms is unreasonably hard or sometimes faulty is the unexpected and counterintuitive behavior of the current benchmarks.

Our main contribution is proposing a new benchmark. Our benchmark benefits from two attractive properties which have been lacking in most of the previous proposals.

- *Smoothness:* If a provider changes his bid by a small amount ϵ, the benchmark changes by at most $nO(\epsilon)$ where n is the number of providers.
- *Monotonicity:* If a provider decreases his bid, the benchmark does not increase.

Almost none[2] of the currently known benchmarks satisfy both of the above properties. This makes our proposal unique since both properties are naturally expected of a reasonable benchmark. We use this newer definition to show positive results for many well-studied auctions. We emphasize that the results presented in this work are based on our benchmark and thus are not direct improvements/corrections of the previous works. Nonetheless, given the justifications of our benchmark, we believe our results are strong evidence in support of the efficiency of our mechanisms.

[2]Except the proposal Karlin *et al.* [12] that has been ruled out by Elkind *et al.* [8] due to several undesired computational and non-computational properties.

1.1 Our Benchmark

We model an auction by a set system $(\mathcal{E}, \mathcal{F})$ where $\mathcal{E}$ denotes the set of providers and $\mathcal{F}$ is a collection of all feasible subsets of providers who can perform the task. For instance, in the case of shortest path auction, the auctioneer wishes to buy one feasible set of roads to connect the source to the destination. Therefore, every edge is a provider and every feasible set corresponds to a path from the source to the destination.

As aforementioned, our proposal is based on the price of cooperation. Suppose the auctioneer commits to a feasible set $S \in \mathcal{F}$ and asks the providers to come up with prices. Of course, the prices should be reasonable, hence, no other feasible set $T \in \mathcal{F}$ should be cheaper for the auctioneer. The maximum total payment of the auctioneer in this case is when all the winning providers collude to maximize the total payment. Therefore, we assume the providers cooperate in order to maximize the total payment of the auctioneer. Such prices can be formulated via an LP.

To formulate the above definition, let c denote the true costs of the providers. We define a function $v^*_S(c)$ for an arbitrary feasible set S as the solution to the following LP:

$$
\begin{aligned}
\max \quad & v^*_S(c) = \textstyle\sum_{e \in S} x_e & & (1)\\
s.t. \quad & x_e = c_e & \forall e \notin S & \quad (2)\\
& x_e \geq 0 & \forall e \in S & \quad (3)\\
& \textstyle\sum_{e \in S} x_e \leq \sum_{e' \in T} x_{e'} & \forall T \in \mathcal{F} & \quad (4)
\end{aligned}
$$

Our proposal for the benchmark is the maximum of such values over all choices of feasible sets for the auctioneer.

Justification of the benchmark: Our benchmark identifies as the worst cooperative equilibrium. This explains why we take the maximum over all possible sets. Note that the previous benchmarks were defined based on the cheapest feasible set only because the corresponding LP was infeasible for the rest of the feasible sets. This is not the case for cooperative equilibria and therefore we take the maximum value over all possible feasible sets. One may think of x_e as a candidate payment to provider e. This way, $v^*_S(c)$ formulates the smallest payment of the auctioneer subject to the winning set being a set S. Roughly speaking, the auctioneer has to pay at least $v^*_S(c)$ as there is no better option for him. This intuition justifies the fact that $v^*_S(c)$ is a good candidate lower bound for the payment of the auctioneer for any S. Therefore, we define our benchmark as

$$\mathcal{B}_{(\mathcal{E},\mathcal{F})}(c) = \max_{S \in \mathcal{F}} v^*_S(c). \tag{5}$$

Note that since the winning set is not necessarily the cheapest set, we take the maximum of $v^*_S(c)$ over all feasible sets. This is where our definition differs from that of Elkind *et al.* [8]. To capture the concept of cooperation in this benchmark, we do not enforce the candidate prices to be higher than the costs which might, in some cases, leave the LP infeasible. We provide a complete review of this benchmark, previously existing benchmarks, and their relation and properties in the full version of the paper.

1.2 An Overview of the Results and Techniques

Our main contribution is a novel and well-motivated benchmark for measuring the efficiency of truthful mechanisms. Our benchmark is inspired by the work of Elkind *et al.* [8] in which the authors suggest "the price of cooperation" to define the frugality ratio of mechanisms. We show that our benchmark benefits from two important properties monotonicity and smoothness both of which are naturally expected from a benchmark.

Note that, most of the previous benchmarks lack either of these properties. For instance, the benchmark of [5] ($v(c)$) is not monotone, i.e., an increase in the agents' costs might result in a decrease in the benchmark. Similarly, $\mathsf{TUMax}(c)$ and $\mathsf{TUMin}(c)$ proposed by Elkind *et al.* [8] (worst and best Nash equilibrium of a utility transferable game) satisfy neither of the properties and thus show a counterintuitive behavior in certain situations (see the full-version of the paper for more detail). We show our benchmark is easier to work with, by presenting several positive results for different set systems. Both properties of our benchmark are essential in proving all results of this paper.

1.2.1 Vertex Cover Auctions (Section 3). In a vertex cover auction, we are given a graph G and every provider corresponds to a vertex of the graph. The goal is to hire a set of providers that cover all edges of the graph. In other words, $\mathcal{E} = V(G)$ and a set $S \in \mathcal{F}$ if and only if for every edge $(u, v) \in E(G)$ either $u \in S$ or $v \in S$ (or both).

In Section 3, we analyze the mechanism of Kempe *et al.* [13] for vertex cover auction and prove a constant competitive factor for this mechanism. Kempe *et al.* 's [13] mechanism for vertex cover is based on a spectral analysis of the graph. They define a matrix $\mathsf{WAM}(G)$ based on the adjacency matrix of the graph. $\mathsf{WAM}(G)$ is essentially the adjacency matrix of the graph except that the values are scaled based on the benchmark. Notice that the benchmark of Kempe *et al.* [13] is different from our benchmark, however replacing their benchmark with ours doesn't change the definition of $\mathsf{WAM}(G)$. $\mathsf{WAM}(G)$ only depends on the structure of the graph and is regardless of the bids or costs. Next, they define a weight q_v for every vertex v of the graph and later use the weights to define the winning set of their mechanism. They show the entries of the dominant eigenvector of this matrix are positive real numbers and thus formulate the weight of each vertex based on the dominant eigenvector of $\mathsf{WAM}(G)$. We denote the weights of the vertices with vector q where every q_v is the weight of vertex v.

Finally, based on the weights of the elements, they leverage a variant of the VCG mechanism to purchase the elements. We call such a mechanism a weighted VCG mechanism (WVCG). In a WVCG mechanism, every element is associated with a weight. The winning set of such auctions is the one with the minimum total weighted cost. Based on this selection rule, the prices are set according to Myerson's characterization of truthful mechanisms [14]. Therefore, every WVCG mechanism is truthful. As such, Kempe *et al.* [13]'s mechanism chooses the vertex cover with the minimum total weighted cost and computes the payments based on the Myerson's rule.

In Section 3.3, we give a correct analysis for the above mechanism and show its frugality ratio is bounded by the dominant eigenvalue of $\mathsf{WAM}(G)$. Of course, this proof is based on our benchmark $\mathcal{B}_{(\mathcal{E},\mathcal{F})}(c)$.

Theorem 3.5 (restated). *The frugality ratio of the vertex cover mechanism is bounded by α where α is the dominant eigenvalue of $\mathsf{WAM}(G)$.*

The complete proof of this theorem appears in Section 3.3.

The intuition behind the proof is that we charge the payments to the value of our benchmark. Recall that $\mathcal{B}_{(\mathcal{E},\mathcal{F})}(c) = \max_{S\in\mathcal{F}} v^*_S(c)$ and the value of function $v^*_S(c)$ for a specific set S is determined by variables x_e's (see LP 1). Let S^{w} be the winning set of the mechanism and x be a solution to the LP corresponding to $v^*_{S^{\mathrm{w}}}(c)$. We redistribute the value of x_e for each element $e \notin S^{\mathrm{w}}$ in an optimal solution of $v^*_{S^{\mathrm{w}}}(c)$ in a way that guarantees the following condition: the payment to every provider is at most α times the value distributed to him. To this end, we leverage spectral techniques and exploit properties of the dominant eigenvector of $\mathsf{WAM}(G)$. The redistribution is essentially a double counting argument based on the payments in specific instances of the problem. Theorem 3.5 follows from the conditions of our redistribution.

Kempe *et al.* [13] also show a lower bound on the frugality ratio of any truthful mechanism for the vertex cover auction. More precisely, they show the frugality ratio of any truthful mechanism is at least $\alpha/2$ where α is the largest eigenvalue of $\mathrm{WAM}(G)$. We show a similar analysis works for our benchmark and hence such a lower bound holds in our setting. We extend this lower bound to randomized mechanisms as well.

Theorem 3.6 (restated). *No truthful randomized mechanism can get a frugality ratio better than $\alpha/4$.*

For the complete proof see Section 3.4. The above theorem along with Theorem 3.5 proves a constant competitive factor (4) for this mechanism.

One may ask how is our result different from the work of Chen *et al.* [5]? We emphasize that the result of Chen *et al.* is not constant competitive. None of the previous work provides any constant competitive analysis for any vertex cover mechanism. We present a new analysis for the mechanism of Kempe *et al.* [13] with respect to our benchmark and show it is constant competitive.

1.2.2 Knapsack Auctions (Section 4). Knapsack auctions have been studied in the literature of both economics and computer science [1, 3, 9, 16]. Knapsack auction captures natural economic interactions. For instance, when a company or individual aims at purchasing at least a specific amount of advertisement opportunities, the problem can be modeled via a knapsack auction. In knapsack auctions, every item has an integer volume and an auctioneer wishes to purchase a subset of items having a total volume of at least a certain number. In the knapsack auction, every element e corresponds to a volume v_e and every feasible set is a subset of items that has a total volume of at least k. All the volumes are at least 1, i.e., $v_i \geq 1$ for all items. In other words, $\mathcal{E} = \{1, 2, \ldots, |\mathcal{E}|\}$ and $\mathcal{F}$ includes every subset S of $\mathcal{E}$ such that

$$\sum_{e \in S} v_e \geq k.$$

In Section 4, we present an approximately optimal mechanism for such auctions.

Theorem 4.3 (restated). *There exists a weighted* VCG *mechanism for the knapsack auction with a frugality ratio of at most $\sqrt{2\mathrm{v}_{\max}}$ where $\mathrm{v}_{\max}$ is an upper bound on the volume of the elements.*

Our proposal is again a WVCG mechanism described below.

- Every element e has a weight $w_e = 1/\sqrt{v_e}$ and the winning set is a feasible subset of elements with the smallest total sum of the weighted costs.
- The payment to every element e is set to its critical value, and as such, the mechanism is truthful.

In Section 4 we show the frugality ratio of the above mechanism is bounded by $\sqrt{2\mathrm{v}_{\max}}$. To this end, we consider an instance of the knapsack auction and run the above mechanism to specify the winning set S^{w} and the associated payments to the elements of S^{w}. More precisely, for every element e in the winning set, we define p_e as the payment made by the auctioneer to element e.

Next, we provide a feasible solution for the benchmark LP. To do so, for every element e in the winning set, we define $x_e = p_e/\sqrt{2\mathrm{v}_{\max}}$ and for every element e outside the winning set we set $x_e = c_e$. According to Lemma 4.2, we argue that x is a feasible solution for $v^*_{S^{\mathrm{w}}}(c)$ and thus the total sum of payments made by the mechanism is bounded by $\sqrt{2\mathrm{v}_{\max}}$ times the value of the benchmark. The reader can find the complete proof in Section 4. This result is followed by an almost matching lower bound.

Theorem 4.1 (restated). *No truthful mechanism can guarantee a frugality ratio better than $\sqrt{\mathrm{v}_{\max}}$.*

The counter-example of Theorem 4.1 is essentially equivalent to a star instance in the vertex cover auctions. We translate such an instance of the vertex cover auction into a knapsack auction by putting an item for every vertex and setting the sizes in a way that the edge constraints are implied. Theorem 4.1 along with Theorem 4.3 proves that our mechanism is optimal up to a constant factor $\sqrt{2}$.

2 PRELIMINARIES

A set system is a pair $(\mathcal{E}, \mathcal{F})$ where $\mathcal{E}$ denotes the set of all ground elements and $\mathcal{F} \subset 2^{\mathcal{E}}$ denotes the list of all feasible subsets of the elements. In this setting, an auctioneer wants to purchase a feasible subset of elements $S \in \mathcal{F}$, and each element $e \in \mathcal{E}$ is owned by a selfish strategic agent. For instance, in the case of shortest path auction, every element of the ground set corresponds to an edge of the graph and every feasible set in $\mathcal{F}$ is a path from the source to the target. Similarly, in a vertex cover system, every element corresponds to a vertex of the graph and every feasible set is a vertex cover of the graph.

A set system is called *monopoly-free*, if no element appears in all feasible sets ($\bigcap_{S_i \in \mathcal{F}} S_i = \emptyset$). We assume both the auctioneer and providers have the full information of the set system in advance and the auction is monopoly-free. Every provider $e \in \mathcal{E}$ has a private cost c_e for the service or product he is offering. The auction proceeds as follows:

(1) Every provider submits a bid b_e for his service. At this point, he is aware of the set system, the costs, and the mechanism which the auctioneer uses to run the auction.
(2) The auctioneer receives all of the bids, he will then run a predetermined mechanism $\mathcal{M}$ to select a winning set $S \in \mathcal{F}$ and assign a payment $p_e \geq b_e$ to every element $e \in S$.

A mechanism $\mathcal{M}$ for such an auction is an algorithm that receives a set system $(\mathcal{E}, \mathcal{F})$ and a vector $b = \langle b_1, \ldots, b_{|\mathcal{E}|} \rangle$ of bids as input, and reports a feasible set $S \in \mathcal{F}$ and a payment $p : S \to \mathbb{R}$ to the providers such that $p_e \geq b_e$ for every $e \in S$.

A mechanism is called *truthful*, if it is in every provider's best interest to report his bid b_e equal to the cost of his service c_e. It is known that a mechanism is truthful if and only if no losing agent can become a winner by increasing his cost and the payment of each winner is set to the highest value for which he remains in the winning set. We refer to this as the agent's critical value.

We measure the performance of a mechanism $\mathcal{M}$ by comparing the total payments of an auctioneer who uses $\mathcal{M}$ to a benchmark. Note that, since our focus is on truthful mechanisms, we assume that all of the bids are equal to the true costs and use the terms bid and cost interchangeably. For a set system $(\mathcal{E}, \mathcal{F})$ and a cost vector $c = \langle c_1, c_2, \ldots, c_{|\mathcal{E}|} \rangle$, we denote our benchmark by $\mathcal{B}_{(\mathcal{E}, \mathcal{F})}(c)$ or just $\mathcal{B}(c)$ when the parameters are clear from the context. We review different benchmarks and show the properties of our benchmark in the full-version of the paper.

The frugality ratio of a mechanism $\mathcal{M}$ is formulated as

$$\mathsf{FR}(\mathcal{M}, (\mathcal{E}, \mathcal{F})) = \sup_{c} \frac{P_{\mathcal{M},(\mathcal{E},\mathcal{F})}(c)}{\mathcal{B}_{(\mathcal{E},\mathcal{F})}(c)}$$

where $P_{\mathcal{M},(\mathcal{E},\mathcal{F})}(c)$ denotes the total sum of payments when the auctioneer runs mechanism $\mathcal{M}$ on cost vector c for set system $(\mathcal{E}, \mathcal{F})$. Similarly, the competitive ratio of a set system is defined as

$$\mathsf{FR}(\mathcal{E}, \mathcal{F}) = \inf_{\mathcal{M}} \mathsf{FR}(\mathcal{M}, (\mathcal{E}, \mathcal{F}))$$

where $\mathcal{M}$ is a truthful mechanism.

Finally, we say a mechanism $\mathcal{M}$ is κ-competitive for a set system $(\mathcal{E}, \mathcal{F})$ if

$$\mathsf{FR}(\mathcal{M}, (\mathcal{E}, \mathcal{F})) \leq \kappa \cdot \mathsf{FR}(\mathcal{E}, \mathcal{F}).$$

Throughout this paper, we design competitive mechanisms or mechanisms with bounded frugality ratios for classic set systems such as vertex cover and knapsack. We formally define these set systems as follows.

Definition 2.1. Given a connected graph $G = \langle V(G), E(G) \rangle$, a vertex cover set system is a pair $(\mathcal{E}, \mathcal{F})$ such that $\mathcal{E} = V(G)$ and $\mathcal{F}$ is the set of all vertex covers of G. More precisely, a subset of vertices S is in $\mathcal{F}$ if and only if every edge $e \in E(G)$ is incident to at least a vertex of S.

We use the notation $N(u)$ to refer to the set of the neighbors of a vertex u in a graph.

Definition 2.2. Given a collection U of items, a vector v specifying the volumes of the items in U, and an integer number k, a set system $(\mathcal{E}, \mathcal{F})$ is a knapsack set system, if $\mathcal{E} = U$ and $\mathcal{F}$ contains every subset of S whose total sum of volumes is at least k. All the volumes are integer numbers greater than or equal to 1.

Finally, we define a class of truthful mechanisms for set systems. This class is a generalization of the well-known VCG mechanisms and has been used for many different auctions.

Definition 2.3. A WVCG mechanism $\mathcal{M}$ for a set system $(\mathcal{E}, \mathcal{F})$ is associated to a weight vector $w = \langle w_1, w_2, \ldots, w_{|\mathcal{E}|} \rangle$. For every cost vector c, it selects a feasible set $S \in \mathcal{F}$ with the minimum weighted sum of costs (*i.e.* $\sum_{e \in S} c_e w_e$ is minimized) . The payments are set to the critical values, and hence every WVCG mechanism is truthful.

3 VERTEX COVER

As mentioned before, Kempe *et al.* [13] propose a mechanism for the vertex cover auction and claim their mechanism is constant competitive with regard to their benchmark. However, the proof they provide is not entirely correct. In this section, we first explain the mechanism and show why their proof fails. Next, we provide an alternative proof for the upper bound of the mechanism with regard to our benchmark. Moreover, in Section 3.4, we show the analysis of Kempe *et al.* [13] gives the same lower bound on the frugality ratio of truthful mechanisms for our benchmark. We also extend this lower bound to all randomized truthful mechanisms. This result in addition to the upper bound given in Section 3.3 proves that the proposed mechanism is 4 competitive.

3.1 Mechanism

In this section, we explain the mechanism of Kempe *et al.* [13] for the vertex cover auction, namely $\mathcal{M}^{\mathsf{VC}}$. Let $G = (V, E)$ be the underlying graph and the goal be to purchase a vertex cover of G. Without loss of generality, we assume $V = \{1, 2, \ldots, n\}$ and the cost of each vertex u is c_u. For each node $u \in V$, let $\mathbf{1^u}$ be an n-dimensional vector with all indices equal to 0 except for index u which is equal to 1. We define the externality of node u as follows:

Definition 3.1. The externality of node u denoted by β_u is defined as the value of the benchmark for cost vector $\mathbf{1^u}$. In other words $\beta_u = \mathcal{B}(\mathbf{1^u})$ (recall Equation (5)).

Roughly speaking, β_u reflects the amount of cost incurred to the auctioneer to cover all of the edges of the graph *without* buying u. Let $\mathsf{WAM}(G)$ be a weighted adjacency matrix of graph G where the weight of an edge (u, v) depends on the externality of vertex u. More precisely

$$\mathsf{WAM}(G)_{(u,v)} = \begin{cases} 0, & \text{if } (u, v) \notin E \\ \frac{1}{\beta_u}, & \text{if} (u, v) \in E. \end{cases}$$

We denote the dominant eignevalue of $\mathsf{WAM}(G)$ by α and refer to its corresponding eignevector with q. By the Perron-Frobenius theorem, eigenvalues of $\mathsf{WAM}(G)$ are real and indices of q are non-negative real numbers. Mechanism $\mathcal{M}^{\mathsf{VC}}$ is a weighted VCG (WVCG) mechanism where the weight of every vertex $i \in V$ is $w_i = 1/q_i$. More formally $\mathcal{M}^{\mathsf{VC}}$ is defined as follows.

- The winning set is a vertex cover S with the minimum total sum of weighted costs, i.e, $\sum_{i \in S} w_i c_i$ is minimized.
- The payment of every vertex i is its critical value.

Since we set the payment of each node u to its critical value, the mechanism is truthful.

3.2 Counter-Example to the Argument of Kempe *et al.* [13]

In the following, we give a counter-example to the proof of Kempe *et al.* [13] for showing that $\mathcal{M}^{\mathsf{VC}}$ has a frugality ratio of at most α. The faulty part of their argument is for bounding the sum of the payments of $\mathcal{M}^{\mathsf{VC}}$ over their benchmark (the optimal value of LP 6 which is equal to $v(c)$). They claim that the following equation holds based on the linearity of their benchmark ($v(c)$).

$$v(c) \geq \sum_u c_u v(\mathbf{1^u}).$$

We show that this is not the case by a counter-example. Let us first formulate $v(c)$ by a linear program:

$$\begin{aligned} \max \quad & v(c) = \textstyle\sum_{e \in \mathrm{S^c}} x_e \\ s.t. \quad & x_e = c_e \qquad \forall e \notin \mathrm{S^c} \\ & x_e \geq c_e \qquad \forall e \in \mathrm{S^c} \\ & \textstyle\sum_{e \in \mathrm{S^c}} x_e \leq \sum_{e \in T} x_e \quad \forall T \in \mathcal{F} \end{aligned} \tag{6}$$

In the above formulation, $\mathrm{S^c}$ denotes the cheapest feasible set. Here we give a counter-example in which the optimal solution of the above LP is smaller than $\sum_u c_u v(\mathbf{1^u})$. Kempe *et al.* [13] support their argument with the following incorrect statement: if we define $x^{(u)}$ as the optimal solution of $v(\mathbf{1^u})$ then $x = \sum_u c_u x^{(u)}$ is a feasible solution for $v(c)$. We show in the following that such a solution for $v(c)$ does not necessarily have a value of at least $\sum_u c_u v(\mathbf{1^u})$. The example is shown in Figure 1.

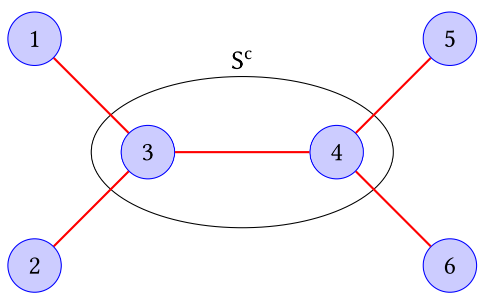

Fig. 1. The counter-example graph. Here all the costs are equal to 1 and the optimal vertex cover is shown by $\mathrm{S^c}$.

The following table illustrates the values of vector x in an optimal solution of $v(\mathbf{1^u})$ for each $u \in \{1, \dots, 6\}$. We refer to these variables as $x^{(u)}$ for each vertex u.

$v(c)$	$x^{(u)}$
$v(\mathbf{1^1}) = 1$	(1, 0, 1, 0, 0, 0)
$v(\mathbf{1^2}) = 1$	(0, 1, 1, 0, 0, 0)
$v(\mathbf{1^3}) = 1$	(1, 0, 1, 0, 0, 0)
$v(\mathbf{1^4}) = 1$	(0, 0, 0, 1, 0, 1)
$v(\mathbf{1^5}) = 1$	(0, 0, 0, 1, 1, 0)
$v(\mathbf{1^6}) = 1$	(0, 0, 0, 1, 0, 1)

All costs in this example are equal to 1, i.e., $c = \langle 1, 1, 1, 1, 1, 1 \rangle$. Moreover, the cheapest feasible sets ($\mathrm{S^c}$) are different for different $x^{(u)}$'s. For instance, when $c = \mathbf{1^1}$, one possible cheapest set is $\{3, 4\}$ whereas when $c = \mathbf{1^3}$, one cheapest set is $\{4, 1, 2\}$. Since $x = \sum_u c_u x^{(u)}$ we have $x = \langle 2, 1, 3, 3, 1, 2 \rangle$. Moreover, the objective function of LP 6 for x is 6 although the optimal solution ($v(c)$) is actually 4 which can be obtained by vector $\langle 1, 1, 2, 2, 1, 1 \rangle$. This shows that they bound the total payment of $\mathcal{M}^{\mathrm{VC}}$ by some value higher than their benchmark. The same counter-example holds even when applying their argument to our benchmark.

3.3 Upper Bound on the Frugality Ratio of $\mathcal{M}^{\mathrm{VC}}$

Recall that α is the dominant eigenvalue of matrix $\mathrm{WAM}(G)$ corresponding to eigenvector q. We show that the frugality ratio of $\mathcal{M}^{\mathrm{VC}}$ is bounded by α with respect to our benchmark.

The main steps for showing such an upper bound on the frugality ratio of $\mathcal{M}^{\mathrm{VC}}$ are the following: recall that the externality of a node u (β_u) is defined as the value of the benchmark for cost vector $\mathbf{1^u}$. We fix the feasible set purchased by the mechanism to be a vertex cover S^{w} and in Lemma 3.2 show an upper bound on the total payment of the mechanism. The upper bound follows from the fact that our mechanism buys a set S^{w} with the minimum total weighted costs and the payments are set according to the Myerson's characterization of the truthful mechanisms. Thus, the payment to every vertex $u \in S^{\mathrm{w}}$ is bounded by a linear expression in terms of the costs of the vertices in $N(u)$. This enables us to prove an upper bound of $(\sum_{u \notin S^{\mathrm{w}}} c_u \beta_u) \cdot \alpha$ on the total payment of the mechanism. Next, in Lemma 3.4 we show a lower bound of $\sum_{u \notin S^{\mathrm{w}}} c_u \beta_u$ for the value of the benchmark $\mathcal{B}(c)$. Notice that $\mathcal{B}(c) \geq v_S^*(c)$ for every $S \in \mathcal{F}$ and therefore $\mathcal{B}(c) \geq v_{S^{\mathrm{w}}}^*(c)$ also holds. Therefore, to show such a lower bound, it only suffices to construct a solution for $v_{S^{\mathrm{w}}}^*(c)$ that has a value of at least $\sum_{u \notin S^{\mathrm{w}}} c_u \beta_u$. The two bounds together imply that the frugality ratio of $\mathcal{M}^{\mathrm{VC}}$ is bounded by α.

We begin by showing an upper bound on the total payment of the mechanism to the vertices of the winning set. In Lemma 3.2, we use the fact that S^{w} is a vertex cover with the minimum total weighted costs and that the mechanism is truthful. We denote the payment to a vertex $u \in S^{\mathrm{w}}$ by p_u.

Observation 3.1. *For every $u \in S^{\mathrm{w}}$, $p_u \leq \sum_{v \in (N(u) \setminus S^{\mathrm{w}})} c_v w_v / w_u$ holds.*

Proof. This observation follows from the fact that the payment of each provider is determined by its critical value. Recall that the critical value or payment of vertex u is the minimum bid for which he remains in the winning set fixing the bids of all other bidders. Now if the critical value of u (p_u) were higher than $\sum_{v \in N(u), v \notin S^{\mathrm{w}}} c_v w_v / w_u$ then WVCG would select all neighbors of u which are not in S^{w} ($N(u) \setminus S^{\mathrm{w}}$) instead of u since $\sum_{v \in N(u), v \notin S^{\mathrm{w}}} c_v w_v$ would be less than $p_u w_u$. □

Lemma 3.2. $\sum_{u \in S^{\mathrm{w}}} p_u \leq \left(\sum_{u \notin S^{\mathrm{w}}} c_u \beta_u \right) \cdot \alpha$.

Proof. This lemma follows from the following inequalities

$$
\begin{aligned}
\sum_{u \in S^{\mathrm{w}}} p_u &\leq \sum_{u \in S^{\mathrm{w}}} \sum_{v \in (N(u) \setminus S^{\mathrm{w}})} c_v w_v / w_u && \text{Observation 3.1} \\
&= \sum_{v \notin S^{\mathrm{w}}} \sum_{u \in N(v)} c_v w_v / w_u && S^{\mathrm{w}} \text{ is a vertex cover} \\
&= \sum_{v \notin S^{\mathrm{w}}} \frac{c_v}{q_v} \sum_{u \in N(v)} q_u && w_u = 1/q_u \\
&= \sum_{v \notin S^{\mathrm{w}}} \frac{c_v \beta_v}{q_v} \sum_{u \in N(v)} q_u / \beta_v \\
&= \sum_{v \notin S^{\mathrm{w}}} c_v \beta_v \Big(\frac{1}{q_v} \sum_{u \in N(v)} q_u / \beta_v \Big) \\
&= \sum_{v \notin S^{\mathrm{w}}} \alpha c_v \beta_v && q \text{ is an eigenvector of WAM(G)} \\
& && \text{and } \alpha \text{ is its corresponding eigenvalue.}
\end{aligned}
$$

□

To complement the result of Lemma 3.2, we aim to prove a lower bound on the value of the benchmark ($\mathcal{B}$). Let S be a vertex cover of G and x be a feasible solution to $v_S^*(c)$. For each subset of nodes $U \subseteq V$, we define $x(U)$ as $\sum_{u \in U} x_u$. We state the following lemma which reinterprets the last constraint of LP 1 (the benchmark LP) for vertex cover auctions.

Lemma 3.3. *Let $\langle \mathcal{E}, \mathcal{F} \rangle$ be a vertex cover auction on graph G and $S \in \mathcal{F}$ be a vertex cover of G. A vector x is a feasible solution of LP 1 for S ($v^*_S(c)$) if and only if*

- *constraints of type (2) and (3) hold for x*
- *and for every independent set $R \subseteq S$ we have $x(R) \leq x(N(R) \setminus S)$.*

Proof. only if part. If x is a feasible solution for $v^*_S(c)$ then clearly constraints of type (2) and (3) hold. Now for every independent set $R \subseteq S$, let T be equal to $(S \setminus R) \cup (N(R) \setminus S)$. T is a vertex cover of G since every edge with one endpoint in R will be covered by either $N(R) \setminus S$ or $S \setminus R$. Moreover, the rest of the edges will be covered by $S \setminus R$. Therefore, by Constraint (4) and the fact that $R \subseteq S$ we have $x(S) = x(S \setminus R) + x(R) \leq x(T)$ which can be written as $x(S \setminus R) + x(R) \leq x(S \setminus R) + x(N(R) \setminus S)$ and hence $x(R) \leq x(N(R) \setminus S)$.

if part. Let $T \in \mathcal{F}$ be an arbitrary vertex cover of G. $S \setminus T$ is an independent set and $N(S \setminus T)$ is a subset of T since otherwise T would not be a vertex cover. Moreover, let $R = S \setminus T$ be an independent set of S. Therefore we have

$$
\begin{aligned}
x(R) &\leq x(N(R) \setminus S) && \text{assumption of the lemma} \\
x(R) + x(S \setminus R) &\leq x(N(R) \setminus S) + x(S \setminus R) && \text{add } x(S \setminus R) \text{ to both sides} \\
x(S) &\leq x(N(R) \setminus S) + x(S \setminus R) && \\
x(S) &\leq x(T) && N(R) \setminus S \text{ and } S \setminus R \\
& && \text{are disjoint subsets of } T
\end{aligned}
$$

and thus constraints of type (4) hold for every vertex cover $T \in \mathcal{F}$ and therefore x meets all constraints of LP 1. □

Now we are ready to prove the lower bound on the value of $\mathcal{B}$. Let c be a vector of costs for the providers and S^{w} be the set of items that $\mathcal{M}^{\mathsf{VC}}$ buys. For every $u \notin S^{\mathrm{w}}$, let set $\mathcal{S}^u$ be a vertex cover such that $\beta_u = \mathcal{B}(\mathbf{1^u}) = v^*_{\mathcal{S}^u}(\mathbf{1^u})$. In other words, $v^*_{\mathcal{S}^u}(\mathbf{1^u})$ is at least as large as $v^*_S(\mathbf{1^u})$ for every vertex cover S of G. If there are multiple vertex covers that maximize $v^*_{\mathcal{S}^u}(\mathbf{1^u})$, we select the one that does not contain u and if there are still more than one such vertex covers, we select one arbitrarily. In the following, we claim that $u \notin \mathcal{S}^u$ holds for all $\mathcal{S}^u$. Another way to interpret Observation 3.2 is that for every vertex u, there always exists a vertex cover S that does not contain u but maximizes $v^*_S(\mathbf{1^u})$.

Observation 3.2. *No vertex cover $\mathcal{S}^u$ contains vertex u.*

Proof. We consider two cases. The first case is where $\mathcal{B}(\mathbf{1^u}) = 0$. In this case for every vertex cover $S \in \mathcal{F}$ we have $v^*_S(\mathbf{1^u}) = 0$ and hence any vertex cover that does not contain u is a candidate for $\mathcal{S}^u$ and thus $\mathcal{S}^u$ does not contain u. Next we investigate the case where $\mathcal{B}(\mathbf{1^u}) > 0$. Note that in this case if $u \in \mathcal{S}^u$ then it implies $\mathcal{B}(\mathbf{1^u}) = 0$ which is self-contradictory. □

Next, we show that $\mathcal{B}(c)$ is at least $\sum_{u \notin S^{\mathrm{w}}} c_u \beta_u$.

Lemma 3.4. *$\mathcal{B}(c) \geq \sum_{u \notin S^{\mathrm{w}}} c_u \beta_u$.*

Proof. To show the lemma, we construct a solution for $v^*_{S^{\mathrm{w}}}(c)$ as follows. Let vector $x^{(u)}$ be a feasible solution that maximizes $v^*_{\mathcal{S}^u}(\mathbf{1^u})$. For every vertex $u \in S^{\mathrm{w}}$, we set $x_u = \sum_{v \notin S^{\mathrm{w}}} c_u x^{(u)}_v$ and for every vertex $u \notin S^{\mathrm{w}}$ we set $x_u = c_u$.

We argue that 1) x is a feasible solution for $v^*_{S^w}(c)$ and 2) the objective value of solution x for $v^*_{S^w}(c)$ is equal to $\sum_{u \notin S^w} c_u \beta_u$. These two together imply that $\mathcal{B}(c) \geq v^*_{S^w}(c) \geq \sum_{u \notin S^w} c_u \beta_u$ and hence the lemma holds. We present the proof of each claim separately.

Proof of the feasibility of x for $v^*_{S^w}(c)$: Notice that constrints of type (2) and (3) trivially hold for solution x. Therefore, by Lemma 3.3, it only suffices to show that for every independent set $R \subseteq S^w$ we have $x(R) \leq x(N(R) \setminus S^w)$ to complete the proof of this part. In what follows, we prove this inequality. Since $x^{(u)}$ is a feasible solution for $v^*_{\mathcal{S}^u}(\mathbf{1^u})$, it follows from Lemma 3.3 that

$$\sum_{v \in R} x^{(u)}_v \leq x^{(u)}(N(R) \setminus \mathcal{S}^u) = \sum_{v \in N(R) \setminus \mathcal{S}^u} \mathbf{1^u}_v \leq 1$$

for solution $x^{(u)}$ of $v^*_{\mathcal{S}^u}(\mathbf{1^u})$ and any independent set $R \subseteq \mathcal{S}^u$. Morever, by scaling both sides of the inequality we get

$$\sum_{v \in R} c_u x^{(u)}_v \leq c_u \tag{7}$$

for every independent set $R \subseteq \mathcal{S}^u$. For any $v \notin (\mathcal{S}^u \cup \{u\})$ we have $x^{(u)}_v = 0$. This implies that if $u \notin S^w$, then $x^{(u)}_v = 0$ holds for every $v \in (S^w \setminus \mathcal{S}^u)$. This, in addition to Inequality (7) implies that for every node $u \notin S^w$ and independent set $R \subseteq S^w$ we have

$$\sum_{v \in R} c_u x^{(u)}_v \leq c_u$$

and therefore

$$\sum_{u \in (N(R) \setminus S^w)} \sum_{v \in R} c_u x^{(u)}_v \leq \sum_{u \in (N(R) \setminus S^w)} c_u. \tag{8}$$

For an independent set $R \subseteq S^w$ and a vertex $u \notin (S^w \cup N(R))$ we have $x^{(u)}_v = 0$ for every $v \in R$ because of the following argument: If $v \notin \mathcal{S}^u$, $x^{(u)}_v = 0$ holds by definition. Otherwise the costs of vertex v and all its neighbors in $\mathbf{1^u}$ are equal to zero and if $x^{(u)}_v > 0$ for a $v \in R$ then we can replace v with all of its neighbors which are not in $\mathcal{S}^u$ and obtain a vertex cover with a smaller value. Therefore we can extend Equation (8) to

$$\left(\sum_{u \in (N(R) \setminus S^w)} \sum_{v \in R} c_u x^{(u)}_v + \sum_{u \notin (S^w \cup N(R))} \sum_{v \in R} c_u x^{(u)}_v \right) \leq \sum_{u \in (N(R) \setminus S^w)} c_u$$

Therefore we have

$$\begin{aligned}
\left(\sum_{u \in (N(R) \setminus S^w)} \sum_{v \in R} c_u x^{(u)}_v + \sum_{u \notin (S^w \cup N(R))} \sum_{v \in R} c_u x^{(u)}_v \right) &\leq \sum_{u \in (N(R) \setminus S^w)} c_u \\
\sum_{v \in R} \left(\sum_{u \in (N(R) \setminus S^w)} c_u x^{(u)}_v + \sum_{u \notin (S^w \cup N(R))} c_u x^{(u)}_v \right) &\leq \sum_{u \in (N(R) \setminus S^w)} c_u \qquad \text{reordering} \\
\sum_{v \in R} \sum_{u \notin S^w} c_u x^{(u)}_v &\leq \sum_{u \in (N(R) \setminus S^w)} c_u \\
\sum_{v \in R} x_v &\leq \sum_{u \in (N(R) \setminus S^w)} x_u \qquad \text{definition of } x \\
x(R) &\leq x(N(R) \setminus S^w)
\end{aligned}$$

Proof of $\sum_{u \in S^w} x_u \geq \sum_{u \notin S^w} c_u \beta_u$: We have

$$\begin{aligned}
\sum_{u \notin S^w} c_u \beta_u &= \sum_{u \notin S^w} \sum_{v \in \mathcal{S}^u} c_u x_v^{(u)} \\
&= \sum_{u \notin S^w} \sum_{v \in S^w} c_u x_v^{(u)} && \text{explained below} \\
&= \sum_{v \in S^w} \sum_{u \notin S^w} c_u x_v^{(u)} && \text{reordering} \\
&= \sum_{v \in S^w} x_v && \text{definition of } x
\end{aligned}$$

The second equation follows from the fact that $x_v^{(u)}$ is equal to zero if v is not in $N(u)$ and that $\mathcal{S}^u$ contains all vertices of $N(u)$ since $u \notin \mathcal{S}^u$. □

Finally, Lemma 3.2 along with Lemma 3.4 shows an upperbound of α on the frugality ratio of $\mathcal{M}^{\mathrm{VC}}$.

THEOREM 3.5 (A COROLLARY OF LEMMAS 3.2 AND 3.4). *The frugality ratio of $\mathcal{M}^{\mathrm{VC}}$ is bounded by α.*

3.4 Lower Bound

In this section we show the frugality ratio of every truthful mechanism is at least $\alpha/4$. This along with Theorem 3.5 proves the competitive factor of $\mathcal{M}^{\mathrm{VC}}$ is bounded by 4. Our proof is similar to [13] except that we extend this lower bound to all randomized mechanisms as well.

Let mechanism $\mathcal{M}$ be an arbitrary truthful randomized mechanism for a vertex cover auction $\langle \mathcal{E}, \mathcal{F} \rangle$ on graph G. We construct a directed graph $\overrightarrow{G}$ as follows. For any edge $(u, v) \in G$, let cost vector c be $c_u = q_u$, $c_v = q_v$ and $c_i = 0$ for all $i \neq u, v$. Now if we run $\mathcal{M}$ on graph G with cost vector c, either u or v appears in the winning set with probability at least $1/2$. If u appears in the winning set with probability at least $1/2$ we add a directed edge from u to v in $\overrightarrow{G}$. Similarly if that's the case with v we add a directed edge from v to u in $\overrightarrow{G}$. Since we iterate over all edges of G, for any edge $(u, v) \in E(G)$ either u has a directed edge to v in $\overrightarrow{G}$ or the vice versa or both.

We associate a weight to every vertex of $\overrightarrow{G}$. More precisely we set the weight of every vertex u equal to q_u. One could show that there exists a node $v \in V(\overrightarrow{G})$ such that

$$\sum_{(u,v) \in E(\overrightarrow{G})} q_u \leq \sum_{(v,u) \in E(\overrightarrow{G})} q_u$$

(for a proof see Lemma 11 of [12]). We define a cost vector c in the following way: $c_u = 0$ for all $u \neq v$ and $c_v = q_v$. Now we run $\mathcal{M}$ on graph G with cost vector c. Note that since $\mathcal{M}$ is truthful, every node u such that $(v, u) \in E(G)$ should be in the winning set. Moreover, again by truthfulness of $\mathcal{M}$, each vertex u such that $(u, v) \in E(\overrightarrow{G})$ receives a payment of at least $\frac{q_u}{2}$.

The rest of the proof follows from the following inequalities that bound the total payment of $\mathcal{M}$.

$$\begin{aligned}
p(\mathcal{M}) &\geq \sum_{(v,u)\in E(\overrightarrow{G})} \frac{1}{2} q_u \\
&\geq \frac{1}{4} \sum_{(v,u)\in E(G)} q_u && \text{by selection of } v \\
&= \frac{1}{4}\beta_v \sum_{(v,u)\in E(G)} \frac{1}{\beta_v} q_u \\
&= \frac{1}{4}\beta_v \alpha \cdot q_v && q \text{ is an eigenvector of } \mathsf{WAM}(G)
\end{aligned}$$

Note that the value of our benchmark for cost vector c equals $q_v \cdot \beta_v$ which concludes that $\mathcal{M}$ has a frugality ratio of at least $\frac{\alpha}{4}$.

Theorem 3.6. *Every truthful mechanism (deterministic or randomized), has a frugality ratio of at least $\alpha/4$.*

4 KNAPSACK AUCTION

In this section, we formally define the knapsack auction and discuss the mechanism we propose for this problem. Let $\mathcal{E}$ be a set of elements, each having a volume v_e. In a knapsack auction, the auctioneer is willing to purchase a subset of items having a total volume of at least k. Therefore, every subset $S \subseteq \mathcal{E}$ is in $\mathcal{F}$ if and only if the total sum of the volumes of the elements in S is at least k. As an example of such auction, consider a company is interested in purchasing some ad slots that guarantee a certain number of user clicks over a period of time. Moreover, we have a number of advertisement companies that offer an ad slot which is guaranteed to be clicked by a certain number of people. Now, the company wishes to run a procurement auction to purchase the ad slots.

We first show it is impossible to design any truthful mechanism for this system with a frugality ratio better than $\sqrt{\mathsf{v}_{\max}}$, where $\mathsf{v}_{\max}$ is the volume of the largest item. We then, propose a mechanism and show the frugality ratio of this mechanism is bounded by $\sqrt{2\mathsf{v}_{\max}}$.

Theorem 4.1. *Let $\mathsf{v}_{\max} = \max_{e\in\mathcal{E}} v_e$ be the volume of the largest item in $\mathcal{E}$. No truthful mechanism can guarantee a frugality ratio better than $\sqrt{\mathsf{v}_{\max}}$.*

Proof. Let $\mathcal{E} = \{1, 2, \ldots, \mathsf{v}_{\max} + 1\}$ contain $\mathsf{v}_{\max} + 1$ elements such that $v_i = 1$ for all $1 \leq i \leq \mathsf{v}_{\max}$ and $v_{\mathsf{v}_{\max}+1} = \mathsf{v}_{\max}$. We set $k = \mathsf{v}_{\max}$ and hence, $\mathcal{F}$ would be the union of all subsets of items whose total sum of volumes is at least k. We show it is impossible to design a truthful mechanism for purchasing a feasible set of items in this instance of knapsack auction that has a frugality ratio better than $\sqrt{\mathsf{v}_{\max}}$.

To this end, assume $\mathcal{M}$ is a truthful mechanism and consider a cost vector $c = \langle 1, 0, 0, \ldots, 0, \sqrt{\mathsf{v}_{\max}}\rangle$ for the elements such that $c_1 = 1$, $c_{\mathsf{v}_{\max}+1} = \sqrt{\mathsf{v}_{\max}}$, and the cost of all other elements is equal to 0. Suppose we run mechanism $\mathcal{M}$ on cost vector c and end up buying a set S of the elements. There are two cases for set S.

- $\mathsf{v}_{\max} + 1 \in S$.
- $\mathsf{v}_{\max} + 1 \notin S$ and hence $i \in S$ for all $1 \leq i \leq \mathsf{v}_{\max}$.

We show in both cases the frugality ratio of $\mathcal{M}$ cannot be less than $\sqrt{\mathsf{v}_{\max}}$. First, assume $\mathsf{v}_{\max} + 1 \in S$. Now, consider a different cost vector c' in which the cost of $\mathsf{v}_{\max} + 1$ is equal to 0 and the cost of all other elements is the same as c. Since $\mathcal{M}$ is truthful, $\mathsf{v}_{\max} + 1$ should also be in the winning set and it should receive a payment of at least $\sqrt{\mathsf{v}_{\max}}$. Therefore, the total payment of the auctioneer in this case is at least $\sqrt{\mathsf{v}_{\max}}$. However, the benchmark $\mathcal{B}(c') = 1$ and thus, the frugality ratio of $\mathcal{M}$ is at least $\sqrt{\mathsf{v}_{\max}}$.

Next, we investigate the case where $\mathsf{v}_{\max} + 1 \notin S$. Let $c^1, c^2, c^3, \ldots, c^{\mathsf{v}_{\max}}$ be $\mathsf{v}_{\max}$ cost vectors such that $c^i_{\mathsf{v}_{\max}+1} = \sqrt{\mathsf{v}_{\max}}$ for all $1 \le i \le \mathsf{v}_{\max}$ and $c^i_i = 1$. Apart from this, all other costs in these vectors are equal to 0. With the same argument, we can conclude that running $\mathcal{M}$ on each of the cost vectors $c^1, c^2, \ldots, c^{\mathsf{v}_{\max}}$, does not lead to buying $\mathsf{v}_{\max} + 1$ unless $\mathsf{FR}(\mathcal{M}) \ge \sqrt{\mathsf{v}_{\max}}$. Now, consider a cost vector c'' such that $c''_{\mathsf{v}_{\max}+1} = \sqrt{\mathsf{v}_{\max}}$ and $c''_i = 0$ for all $1 \le i \le \mathsf{v}_{\max}$. Since $\mathcal{M}$ is truthful and it buys all of the elements $1, 2, \ldots, \mathsf{v}_{\max}$ for all cost vectors $c^1, c^2, \ldots, c^{\mathsf{v}_{\max}}$, it also buys all those elements for cost vector c''. Moreover, each item receives a payment of at least 1. Therefore, the auctioneer makes a payment of at least $\mathsf{v}_{\max}$, where $\mathcal{B}(c'') = \sqrt{\mathsf{v}_{\max}}$. Hence, $\mathsf{FR}(\mathcal{M}) \ge \sqrt{\mathsf{v}_{\max}}$. □

In the rest of this section, we propose a WVCG mechanism $\mathcal{M}^{\mathsf{K}}$ for knapsack auction and bound its frugality ratio by $\sqrt{2\mathsf{v}_{\max}}$. We first define a weight w_e for every item e which is equal to $\frac{\sqrt{v_e}}{v_e}$ and then define $\mathcal{M}^{\mathsf{K}}$ as follows:

- The winning set is a $S \in \mathcal{F}$ such that that minimizes the expression $\sum w_e b_e$. Ties are broken lexicographically.
- The payment of every item $e \in S$ is determined based on its critical value.

In the following, we show the frugality ratio of $\mathcal{M}^{\mathsf{K}}$ is bounded by $\sqrt{2\mathsf{v}_{\max}}$ where $\mathsf{v}_{\max} = \max_{e\in\mathcal{E}} v_e$. Before we proceed to the proof, we state an axillary lemma which we use in the proof of the theorem.

Lemma 4.2. *Let $Y = \{1, 2, \ldots, |Y|\}$ be a set of elements with volumes $v_1, v_2, \ldots, v_{|Y|}$ and costs $c_1, c_2, \ldots, c_{|Y|}$ such that $\sum v_i = b$. If $V \ge \max_{e\in Y} v_e$ and every $v_i \ge 1$ then, for any $1 \le d \le b$, there exists a subset S of elements in Y such that $\sum_{e\in S} v_e \ge d$ and*

$$\sum_{e\in S} c_e/\sqrt{v_e} \le \frac{\sqrt{2dV}}{b}\sum_{e\in Y} c_e.$$

Proof. Without loss of generality, we assume the elements are sorted in increasing order of $\frac{c_e}{v_e}$. In other words,

$$\frac{c_1}{v_1} \le \frac{c_2}{v_2} \le \ldots \le \frac{c_{|Y|}}{v_{|Y|}}.$$

Therefore for every $1 \le i \le |Y|$ we have

$$\frac{\sum_{j=1}^{i} c_j}{\sum_{j=1}^{i} v_j} \le \frac{\hat{c}}{b} \tag{9}$$

where $\hat{c}$ denotes the total cost of items in Y. Now let set $Q = \{1, 2, 3, \ldots, |Q|\}$ be the smallest prefix of items in Y whose total sum of volumes is at least d. Note that, $\sum_{i=1}^{|Q|} v_i \ge d$ but $\sum_{i=1}^{|Q|-1} v_i < d$. Now, let q be the total cost of all items in $Q \setminus \{|Q|\}$. In the rest of the proof we show

$$\sum_{e\in Q} c_e/\sqrt{v_e} \le \frac{\sqrt{2dV}}{b}\hat{c}.$$

Since $\min_{e\in Q\setminus\{|Q|\}} v_e \ge 1$ we have

$$\sum_{e\in W} c_e/\sqrt{v_e} \le q + c_{|Q|}/\sqrt{v_{|Q|}}. \tag{10}$$

Moreover, by Inequality (9) we have

$$q \leq \frac{(d-1)\hat{c}}{b}$$

and

$$q + c_{|Q|} \leq \frac{(d + v_{|Q|} - 1)\hat{c}}{b}.$$

Therefore (10) is maximized when $q = \frac{(d-1)\hat{c}}{b}$ and $c_{|Q|} = \frac{(v_{|Q|})\hat{c}}{b}$ and thus

$$\begin{aligned}
\sum_{e \in W} c_e/\sqrt{v_e} &\leq q + c_{|Q|}/\sqrt{v_{|Q|}} \\
&\leq \frac{(d-1)\hat{c}}{b} + \frac{(v_{|Q|})/\sqrt{v_{|Q|}}\hat{c}}{b} \\
&= \frac{(d-1)\hat{c}}{b} + \frac{\sqrt{v_{|Q|}}\hat{c}}{b} \\
&\leq \frac{(d-1)\hat{c}}{b} + \frac{\sqrt{V}\hat{c}}{b} \\
&= \frac{\hat{c}}{b}(d - 1 + \sqrt{V}) \\
&= \frac{\hat{c}}{b}\sqrt{V}(1 + \frac{d-1}{\sqrt{V}}) \\
&\leq \frac{\hat{c}}{b}\sqrt{2Vd}.
\end{aligned}$$

□

Next we prove $\mathsf{FR}(\mathcal{M}^{\mathsf{K}}) \leq \sqrt{2\mathsf{v}_{\max}}$.

THEOREM 4.3. *Let $(\mathcal{E}, \mathcal{F})$ be a knapsack set system and $\mathsf{v}_{\max} = \max_{e \in \mathcal{E}} v_e$. Then we have*

$$\mathsf{FR}(\mathcal{M}^{\mathsf{K}}) \leq \sqrt{2\mathsf{v}_{\max}}.$$

Proof. Let $(\mathcal{E}, \mathcal{F})$ be a set system and c be a cost vector for the elements of $\mathcal{E}$. Suppose we run $\mathcal{M}^{\mathsf{K}}$ on the bids and determine a set S as the winning set and assign a payment p_e to every element e in S. In the following we show

$$\sum_{e \in S} p_e \leq \sqrt{2\mathsf{v}_{\max}}\mathcal{B}(c).$$

By the definition, we have $\mathcal{B}(c) \geq v_S^*(c)$. Recall that $v_S^*(c)$ is equal to the objective value of the following linear program:

$$\begin{aligned}
\max \quad & v_S^*(c) = \textstyle\sum_{e \in S} x_e & & (11)\\
s.t. \quad & x_e = c_e & \forall e \notin S \quad & (12)\\
& x_e \geq 0 & \forall e \in S \quad & (13)\\
& \textstyle\sum_{e \in S} x_e \leq \sum_{e \in T} x_e & \forall T \in \mathcal{F} \quad & (14)
\end{aligned}$$

Let x be a solution of $v_S^*(c)$ in which for every $e \in S$, $x_e = p_e/\sqrt{2\mathsf{v}_{\max}}$ and for every $e \notin S$, we have $x_e = c_e$. We first show x is a feasible solution of LP (11). Note that all constraints of type (12) and (13) are trivially satisfied. Now let T be a feasible set for constraints of type (14). We define

$X = S \setminus T$ to be the elements of S that are not in T and $Y = T \setminus S$ to be the elements of T that replace X in S. Since x_e is the same for both sides of constraint (14), we can rewrite the inequality as

$$\sum_{e \in X} x_e \leq \sum_{e \in Y} x_e = \sum_{e \in Y} c_e. \tag{15}$$

Now let $\sum_{e \in Y} v_e = b$ be the total volume of the elements in Y and $b' = \sum_{e \in X} v_e$. Note that, since all elements of Y can replace elements of X, then every element $e \in X$ can be replaced by some elements of Y that have a total sum of volumes of at least $v_e - \max(0, b' - b)$. Since we set the price of every item in S by its critical value, the payment of any item $e \in X$ that can be replaced by a subset R_e, (i.e., $R_e \cap S = \emptyset$) is bounded by

$$w_e p_e \leq \sum_{e' \in R_e} w_{e'} c_{e'}$$

and hence

$$p_e / \sqrt{v_e} \leq \sum_{e' \in R_e} c_{e'} / \sqrt{v_{e'}}.$$

Therefore

$$p_e \leq \sum_{e' \in R_e} \sqrt{v_e} c_{e'} / \sqrt{v_{e'}} \tag{16}$$

By applying Equation (16) to Lemma 4.2 we get

$$p_e \leq \sqrt{v_e} \frac{\sqrt{2(v_e - \max(0, b' - b))\mathsf{v}_{\max}}}{b} \sum_{e' \in Y} c'_e$$

Since $\sqrt{v_e}\sqrt{(v_e - \max(0, b' - b)} \leq v_e - \frac{\max(0, b'-b)}{2}$, we have

$$p_e \leq \sqrt{\mathsf{v}_{\max}} \frac{\sqrt{2}(v_e - \frac{\max(0, b'-b)}{2})}{b} \sum_{e' \in Y} c'_e. \tag{17}$$

We sum Equation (17) over all elements $e \in X$ and obtain

$$\sum_{e \in X} p_e \leq \sqrt{2\mathsf{v}_{\max}} \frac{\sum_{e \in X}(v_e - \frac{\max(0, b'-b)}{2})}{b} \sum_{e \in Y} c_e = \sqrt{2\mathsf{v}_{\max}} \frac{b' - |X|\frac{b'-b}{2}}{b} \sum_{e \in Y} c_e. \tag{18}$$

Note that if $|X| = 1$, (16) directly proves

$$\sum_{e \in X} p_e \leq \sqrt{2\mathsf{v}_{\max}} \sum_{e \in Y} c_e$$

otherwise $b' - |X|\frac{b'-b}{2} \leq b$ and by (18) we have

$$\sum_{e \in X} p_e \leq \sqrt{2\mathsf{v}_{\max}} \sum_{e \in Y} c_e.$$

Since we set $x_e = p_e / \sqrt{2\mathsf{v}_{\max}}$, we have

$$\sum_{e \in X} x_e \leq \sum_{e \in Y} c_e.$$

and thus all constraints of type (14) are satisfied. Since x is a feasible solution of LP 11, we have

$$\mathcal{B}(c) \geq v_S^*(c) \geq \sum_{e \in S} x_e = \sum_{e \in S} p_e / \sqrt{2\mathsf{v}_{\max}}$$

and therefore $\mathsf{FR}(\mathcal{M}) \leq \sqrt{2\mathsf{v}_{\max}}$. □

5 ACKNOWLEDGMENT

We would like to thank Nick Gravin for having helpful discussions and his comments on an earlier version of the paper.

REFERENCES

[1] Gagan Aggarwal and Jason D Hartline. 2006. Knapsack auctions. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. Society for Industrial and Applied Mathematics, 1083–1092.

[2] Aaron Archer and Éva Tardos. 2007. Frugal path mechanisms. *ACM Transactions on Algorithms (TALG)* 3, 1 (2007), 3.

[3] Xiaohui Bei, Ning Chen, Nick Gravin, and Pinyan Lu. 2012. Budget feasible mechanism design: from prior-free to bayesian. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM, 449–458.

[4] Martin Bichler and Jayant R Kalagnanam. 2006. Software frameworks for advanced procurement auction markets. *Commun. ACM* 49, 12 (2006), 104–108.

[5] Ning Chen, Edith Elkind, Nick Gravin, and Fedor Petrov. 2010. Frugal mechanism design via spectral techniques. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*. IEEE, 755–764.

[6] Peter Cramton, Yoav Shoham, and Richard Steinberg. 2006. Combinatorial auctions. (2006).

[7] Sudipto Dasgupta and Daniel F Spulber. 1990. Managing procurement auctions. *Information Economics and Policy* 4, 1 (1990), 5–29.

[8] Edith Elkind, Leslie Ann Goldberg, and Paul W Goldberg. 2007. Frugality ratios and improved truthful mechanisms for vertex cover. In *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, 336–345.

[9] Paul W Goldberg and Antony McCabe. 2012. Commodity auctions and frugality ratios. In *Algorithmic Game Theory*. Springer, 180–191.

[10] Srabana Gupta. 2002. Competition and collusion in a government procurement auction market. *Atlantic Economic Journal* 30, 1 (2002), 13–25.

[11] Mireia Jofre-Bonet and Martin Pesendorfer. 2000. Bidding behavior in a repeated procurement auction: A summary. *European Economic Review* 44, 4 (2000), 1006–1020.

[12] 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*. IEEE Computer Society, 615–626.

[13] David Kempe, Mahyar Salek, and Cristopher Moore. 2010. Frugal and truthful auctions for vertex covers, flows and cuts. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*. IEEE, 745–754.

[14] Paul Klemperer. 2004. Auctions: theory and practice. (2004).

[15] Robert H Porter and J Douglas Zona. 1992. *Detection of bid rigging in procurement auctions*. Technical Report. National bureau of economic research.

[16] Yaron Singer. 2010. Budget feasible mechanisms. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*. IEEE, 765–774.