Monografias em Ciência da Computação, No. 13/05 ISSN: 0103-9741
Editor: Prof. Carlos José Pereira de Lucena April, 2005

On Behalf the Seller and Society: a bicriteria mechanism for unit demand combinatorial auctions

Claudson Bornstein[1], Eduardo S. Laber and Marcelo A. F. Más

[1] Programa de Engenharia de Sistemas e Computação - UFRJ

cfb@cos.ufrj.br, laber@inf.puc-rio.br, marcelomas@inf.puc-rio.br

Abstract. This work focuses on obtaining truthful mechanisms that aim at maximizing both the revenue and the economic efficiency (social welfare) for the unit-demand combinatorial auction problem (UDCAP), in which a set of k items is auctioned to a set of n consumers. Although each consumer bids on all items, no consumer can purchase more than one item in the UDCAP. We present a framework for devising poly-time randomized competitive truthful mechanisms that can be used to either favor economic efficiency or revenue.

Keywords: Combinatorial Auctions, Approximation Algorithms, Randomized algorithms.

Resumo: O objetivo deste trabalho é obter mecanismos truthful que maximizem tanto a receita quanto a eficiência econômica (bem-estar social) do problema de leilão combinaório com demanda unitária (UDCAP), no qual um conjunto de k itens são oferecidos em leilão a um conjunto de n consumidores. Embora cada consumidor faça uma oferta para cada item, nenhum consumidor pode comprar mais de um item no UDCAP. Apresentamos um framework para projetar mecanismos aleatorizados competitivos truthful que podem ser executados em tempo polinomial e que podem ser usados tanto para favorecer eficiência econômica quanto receita.

Palavras-chave: Leilões Combinatórios, Algoritmos de Aproximação, Algoritmos Aleatorizados.

In charge for publications:
Rosane Teles Lins Castilho
Assessoria de Biblioteca, Documentação e Informação
PUC-Rio Departamento de Informática
Rua Marquês de São Vicente, 225 - Gávea
22453-900 Rio de Janeiro RJ Brasil
Tel. +55 21 3114-1516 Fax: +55 21 3114-1530
E-mail: bib-di@inf.puc-rio.br

1 Introduction

Auction mechanism design has long been a field of interest in the Economics and Game Theory communities. With the rise in electronic commerce and high-profile auctions such as the Google IPO and FCC spectrum auctions, the field of mechanism design for auctions has drawn a lot of attention from theoretical Computer Science researchers in recent years. This paper deals with auctions in situations where no previous knowledge of the bid distributions is known, so that in fact the traditional Bayesian approach that relies on prior knowledge is not applicable.

We study single-round sealed-bid combinatorial auctions: one seller wants to auction k distinct items to a set of n consumers, and the auctioneer uses a known mechanism to both allocate the items to consumers and to determine the sale price of each item or set of items (bundles). It is assumed there is no communication or coalition amongst consumers (sealed-bid auction).

The seller and consumers have somewhat conflicting interests, wherein lies the interest in the problem. Each consumer has a *valuation*, i.e., a price expressed in some common currency, for each bundle. The valuation is privately held by the consumer, and is never made public. Each consumer wants to maximize his or her own *utility*, that is, the difference between his private valuation for the bundles purchased as a result of the auction, and the actual price paid for these bundles. The consumer can misrepresent his (her) true valuation for each bundle should he believe it will raise his expected utility. On the other hand, the auctioneer wants to maximize the total revenue of the auction, that is, the sum of all prices paid as a result of the auction.

This game between seller and consumers is avoided with a suitable mechanism design: one that maximizes each consumer's profit when the consumer bids his true valuation on each bundle. These are known as *truthful* or *incentive compatible* mechanisms, and have been extensively studied.

The *economic efficiency* or *social welfare* of an auction is defined as the sum of the valuations that each consumer attributes to the bundles he acquires. It relates to the social value of the auction and often enough, say, on the FCC spectrum auctions, maximizing the efficiency is more important than maximizing the revenue even from the auctioneer's perspective. The famous Vickrey-Clarke-Groves(VCG) mechanisms [17, 5, 12] are early truthful mechanisms and allow the allocation of items so as to maximize total efficiency while giving incentive for truth-telling.

This paper studies the unit-demand combinatorial auction problem *UD-CAP* [13] which consists of designing a mechanism for an auction in which multiple distinct items are being sold. Even though consumers bids on all items they can only purchase a single item.

Our Results. Let $\mathbf{v}$ be the matrix defining the consumers' valuations. We use $\mathcal{T}(\mathbf{v})$ to denote the maximum possible revenue obtained by an 'omniscient' auctioneer on a truthful auction. On the other hand, $\mathcal{F}(\mathbf{v})$ denotes the maximum possible revenue obtained by the same auctioneer, on a truthful auction, under the constraint that a single price must be used to sell all items. We use s to denote $\min\{n, k\}$ where n is the number of consumers and k is the number of items in the auction.

Unlike most recent work on truthful auction mechanisms, we design mechanisms that simultaneously achieve 'high' revenue and 'high' economic efficiency with high probability. Our high-probability results depend on the weak assumption that a small number of consumers cannot generate revenue close to $\mathcal{T}(\mathbf{v})$ (a constant factor of $\mathcal{T}(\mathbf{v})$).

We devise a randomized truthful mechanism that achieves $\Omega(\mathcal{T}(\mathbf{v})/\ln s)$ revenue and $\Omega(\mathcal{T}(\mathbf{v}))$ efficiency with high probability.[1] In addition, we show that for every randomized truthful auction mechanism $\mathcal{A}$ there exists a valuation matrix $\mathbf{v}$ such that the ratio between $\mathcal{T}(\mathbf{v})$ and the expected revenue achieved by $\mathcal{A}$ is $\Omega(\ln s)$, which is matched by our mechanism.

Our mechanism can be viewed as a particular implementation of a more general approach that shows a trade-off between economic efficiency and revenue. Our approach is a generalization of the scheme employed in [11] for unlimited-supply single-item auctions and consists of dividing consumers in two groups, and using one group's bids to estimate a suitable number of items to sell to the consumers of the other group. It then uses a VCG mechanism with this limited number of items to decide both the allocation and the sale prices.

By adjusting our estimate of how many items should be sold we can favor maximizing efficiency or revenue. If we implement our mechanism so as to favor revenue we obtain an auction mechanism that achieves an $\Omega(\mathcal{F}(\mathbf{v}))$ revenue with high probability. We note that proving an $\Omega(\mathcal{F}(\mathbf{v}))$ bound is stronger than proving an $\Omega(\mathcal{T}(\mathbf{v})/\ln s)$ bound, since the inequality $\mathcal{F}(\mathbf{v}) \geq \mathcal{T}(\mathbf{v})/\ln s$ always holds and, in fact, we may even have $\mathcal{F}(\mathbf{v}) = \Omega(\mathcal{T}(\mathbf{v}))$.

By combining this last strategy with the VCG mechanism for UDCAP and with an auction that only sells one item we attain an auction mechanism that achieves $\Omega(\mathcal{F}(\mathbf{v}))$ expected revenue and $\Omega(\mathcal{T}(\mathbf{v}))$ expected efficiency under very weak assumptions. However, in this case we do not have high concentration around the mean.

[1] Throughout this work we do not focus on optimizing constant factors

Finally, we shall mention that our results extend to a more general case where every consumer may purchase a bundle with at most d items, where d is a constant which does not depend on n or k. In this case, however, our mechanism requires exponential time. We defer this discussion to an extended version of this paper.

Related Work. Two natural goals for designing combinatorial auctions mechanisms have been considered in the literature: maximizing the revenue [11, 9, 16, 10, 6, 7, 13] and maximizing the efficiency [14, 1, 2, 4]. Some of the papers that focus on maximizing the efficiency also discuss revenue issues. However, what is usually done is to compare the revenue achieved by the proposed auction mechanisms with that achieved by the VCG mechanism. While VCG mechanisms are widely known and it might be interesting to understand how the revenue of a new auction mechanism relates to it, the revenue achieved by VCG can be low, even 0, which makes it an undesirable benchmark.

With regards to the revenue aspect, our work is related to the series of papers by J. Hartline and others [6, 9, 10, 11, 7, 13]. Most of these papers study the problem of auctioning off an unlimited number of identical copies of a single item. In [9], Goldberg and Hartline handle the problem of auctioning distinct items in unlimited supply, and mention the UDCAP as an open problem.

The paper most closely related to ours is by Guruswami et al. [13], where an interesting mechanism is proposed for the UDCAP. Their mechanism achieves $\Omega(\mathcal{T}(\mathbf{v})/\log h))$ revenue, where h is the ratio between the largest and lowest bid values and consists of randomly selecting reserve prices for a VCG auction. Their approach, however, has some potential disadvantages: it relies on previous knowledge about the range of the bids; the revenue is not highly concentrated around the mean; the revenue can be at a $O(\log h)$ factor from $\mathcal{F}(\mathbf{v})$, even if h is bounded by a polynomial in n. This contrasts with our mechanism that achieves an expected revenue of either $\Omega(\mathcal{F}(\mathbf{v}))$ or $\Omega(\mathcal{T}(\mathbf{v})/\ln s)$ with high probability, depending on whether we want to maximize revenue or economic efficiency. For the sake of fairness, it should be mentioned that, as opposed to ours, their mechanism produces envy-free allocations.

The definition of a suitable benchmark to compare the revenue of a truthful auction with is not trivial as pointed out in [1]. $\mathcal{F}$ (and its variations) has been proved a reasonable metric for evaluating the revenue generated by single-item auction mechanisms [11, 7]. In [13], Guruswami et al. propose to use the envy-free metric (see Section 4), a natural generalization of the $\mathcal{F}$ metric, as a target for the UDCAP mechanisms. We show in Section 4 that for every truthful auction $\mathcal{A}$ there exists a valuation matrix $\mathbf{v}$ such that the ratio between the optimum envy-free revenue for $\mathbf{v}$ and the expected revenue achieved by $\mathcal{A}$ is $\Omega(\ln s)$.

2 Preliminaries

The auctions considered in this paper work as follows: let $C = \{1, 2, \ldots, n\}$ be a set of consumers and let I be a set of k distinct items. We use s to denote $\min\{|C|, |I|\}$. Each consumer submits a bid for each item in I. Based on these bids, the auctioneer allocates at most one item to each consumer and determines the sale price of each allocated item. We denote the bid of consumer i for item j by $b_i(j)$. If j is allocated to i then i buys item j for a sale price p_j, where $p_j \leq b_i(j)$. We disregard the situation in which a number of copies of each item is available, but it can easily be modeled.

We assume that the *mechanism*, that is, both the allocation and the pricing strategy, employed by the auctioneer is publicly known. The following assumptions are made about the consumers:

- For every pair $(i, j) \in C \times I$, consumer i has a private valuation $v_i(j) \geq 0$ for item j. The valuation $v_i(j)$ is the maximum price for which i would be willing to buy item j.
- If consumer i buys item j, then his profit(utility) is $u_i = v_i(j) - p_j$.
- Consumers are rational and will submit bids that try to maximize their utilities.
- Consumers do not collude.
- The consumers in the auction are indistinguishable from the perspective of the auctioneer.

Depending on the mechanism used by the auctioneer, the consumers might be able to increase their utility by presenting bids that misrepresent their valuations. An auction $\mathcal{A}$ is *truthful* if the best strategy for each consumer is to submit his own valuations regardless of his beliefs on the bidding strategies employed by the other consumers. Rational consumers will bid their valuations in truthful auctions.

The input for a deterministic truthful auction $\mathcal{A}$ can be viewed in terms of a weighted complete bipartite graph G among consumers and items. The cost $c(e_{i,j})$ of the edge $e_{i,j}$ (associating the i-th consumer to the j-th item) is $v_i(j)$.

The output of the auction $\mathcal{A}$, for input G, is a pair $(M, \mathbf{p})$, where $M = \cup_{i=1}^{|M|} e_i$ is a matching of G and $\mathbf{p} = (p_1, \ldots, p_{|M|})$ is a vector defining the sale price of every item allocated by M. We must have $p_i \leq c(e_i)$, for $i = 1, \ldots, |M|$. The revenue $\mathcal{R}(\mathcal{A}(G))$ is the sum of the prices assigned to the items of M and the efficiency $\mathcal{W}(\mathcal{A}(G))$ is the sum of the costs of the edges of M. Clearly, $\mathcal{W}(\mathcal{A}(G)) \geq \mathcal{R}(\mathcal{A}(G))$

In this paper, we are interested in *randomized auctions*, which are probability distributions over deterministic auctions. Following [8], we adopt a notion of randomized truthfulness in which a randomized truthful auction is a probability distribution over the set of deterministic truthful auctions.

We use h to refer to the highest bid in the auction or, equivalently, the cost of the most expensive edge of G. Finally, for a consumer i, we use G_{-i} to denote the subgraph induced in G by the removal of i from the set of vertices of G.

2.1 Competitive Framework

We say that the revenue of a randomized auction $\mathcal{A}$ on an input graph G is α revenue-competitive with a metric $\mathcal{M}$ if $\mathrm{E}[\mathcal{R}(\mathcal{A}(G))] \geq \mathcal{M}(G)/\alpha$.

Next we define two metrics that will be used in our competitiveness criteria throughout the paper. Let $H = (V, E)$ be a weighted bipartite graph. Furthermore, let M be a matching in H and let e' be the lowest weight edge of M. We define $\mathcal{F}(M) = |M| \times c(e')$ and $\mathcal{T}(M) = \sum_{e \in M} c(e)$. We use $M_H^{\mathcal{T}}$ ($M_H^{\mathcal{F}}$) to denote the largest matching of H that maximize $\mathcal{T}(\cdot)$ ($\mathcal{F}(\cdot)$). We define $\mathcal{F}_H = \mathcal{F}(M_H^{\mathcal{F}})$ and $\mathcal{T}_H = \mathcal{T}(M_H^{\mathcal{T}})$. For the sake of simplicity, we use $\mathcal{T}$ and $\mathcal{F}$ to refer to $\mathcal{T}_G$ and $\mathcal{F}_G$, respectively.

We also use the concept of competitiveness to evaluate the economic efficiency achieved by our auction mechanisms. However, it only really makes sense to compare the efficiency to $\mathcal{T}$. Thus, we say that a randomized mechanism $\mathcal{A}$ is β efficiency-competitive for an auction G if $\mathrm{E}[\mathcal{W}(\mathcal{A}(G))] \geq \mathcal{T}/\beta$.

We note that for every truthful auction $\mathcal{A}$ and every input graph G, $\mathcal{R}(\mathcal{A}(G)) \leq \mathcal{W}(\mathcal{A}(G)) \leq \mathcal{T}$. As we shall see in the propositions that follow, both $\mathcal{F}$ and $\mathcal{T}$ provide reasonable comparison points for the revenue of truthful auctions.

Proposition 1 *$\mathcal{F} \geq \mathcal{T}/\ln s$, where $s = \min\{|C|, |I|\}$.*

Proof. Let $e_1, \ldots, e_s$ be the edges of $M_G^{\mathcal{T}}$ sorted in non-increasing order of costs and let e_{j^*} be the edge that maximizes $j^* \times c(e_{j^*})$. We have that

$$c(e_i) \leq \frac{j^* \times c(e_{j^*})}{i},$$

for $i = 1, \ldots, s$. Adding all these inequalities we conclude that $j^* \times c(e_{j^*}) \geq \mathcal{T}/\ln s$. Let M be the sub-matching of $M_G^{\mathcal{T}}$ defined by its j^* most expensive edges. Thus, $\mathcal{F} \geq \mathcal{F}(M) \geq \mathcal{T}/\ln s$ ∎

The next proposition shows that there exists a single matching M that has 'high' values for both $\mathcal{T}(\cdot)$ and $\mathcal{F}(\cdot)$.

Proposition 2 *There is a matching M in G such that $\mathcal{F}(M) \geq \mathcal{T}/(2\ln s)$ and $\mathcal{T}(M) \geq \mathcal{T}/2$.*

Proof. Let $e_1, \ldots, e_s$ be the edges of $M_G^{\mathcal{T}}$ sorted in non-increasing order of weights and let i^* be the largest number such that $i^* \times c(e_{i^*}) \geq \mathcal{T}/2\ln s$. The existence of such an i^* is ensured in the proof of the Proposition 1. Define M as $\cup_{j=1}^{i^*} e_j$. Clearly, $\mathcal{F}(M) \geq \mathcal{T}/2\ln s$.

For $j > i^*$, we have that $c(e_j) < \frac{\mathcal{T}}{2j \times \ln s}$. By adding these inequalities we obtain that

$$\sum_{j=i^*+1}^{s} c(e_j) \leq \frac{\mathcal{T} \times (\ln s - \ln i^*)}{2\ln s} \leq \frac{\mathcal{T}}{2}$$

Thus, $\mathcal{T}(M) = \sum_{j=1}^{i^*} c(e_j) \geq \mathcal{T}/2$ ∎

2.2 Approximation matchings

Next, we introduce the concept of an approximation matching for a sequence of matchings. This is used in Section 3.2 as a technical tool to ensure that one of the auctions mechanisms proposed in this paper reaches their expected revenue and efficiency with high probability. Roughly speaking, given a sequence $\mathcal{S}$ of matchings in a graph G, the approximation matching A for $\mathcal{S}$ has the property that for every matching S of $\mathcal{S}$ there is a sub-matching A' of A whose size is at a constant factor of S and, moreover, $\mathcal{F}(A') \geq \min_{S \in \mathcal{S}}\{\mathcal{F}(S)\}/2$ and $\mathcal{T}(A') \geq \min_{S \in \mathcal{S}}\{\mathcal{T}(S)\}$.

For an increasing sequence of integers J, let $min(J) = \min\{j | j \in J\}$, $max(J) = \max\{j | j \in J\}$ and $pred(j)$ be the largest integer of J smaller than j, for $j \in J \setminus min(J)$.

Definition 1 *Let $(M_j)_{j \in J}$ be a sequence of matchings in G, where $|M_j| = j$, for every $j \in J$. We define the sequence $(A_j)_{j \in J}$ as follows:*

$$A_j = \begin{cases} M_j \text{ if } j = min(J) \\ A_{pred(j)} \cup \{e | e \in M_j \text{ and } A_{pred(j)} \cup e \text{ is a matching}\}, \text{ otherwise} \end{cases}$$

We call $A_{max(J)}$ the approximation matching for the sequence $(M_j)_{j \in J}$.

Example 1 *Let* $G = (V_1 \cup V_2, E)$ *be a complete bipartite graph where* $V_1 = \{1,3,5,7,9\}$ *and* $V_2 = \{2,4,6,8,10\}$. *Let us consider the sequence of matchings* M_2, M_3, M_5, *where* $M_2 = \{(1,2),(3,6)\}$, $M_3 = \{(1,2),(3,8),(4,10)\}$ *and* $M_5 = \{(1,2),(3,4),(5,6),(7,8),(9,10)\}$.

Then, we have $A_2 = \{(1,2),(3,6)\}$, $A_3 = \{(1,2),(3,6),(4,10)\}$ *and* A_5, *the approximation matching, is* $\{(1,2),(3,6),(4,10),(7,8)\}$

The following lemma states crucial properties regarding the approximation matching.

Lemma 1 *Let* $(M_j)_{j\in J}$ *be a sequence of matchings in* G, *where* $|M_j| = j$, *for every* $j \in J$. *Furthermore, let* A *be the approximation matching of* $(M_j)_{j\in J}$. *Then, for every* $j \in J$, *there is a sub-matching* A' *of* A *such that: (i)* $\max\{\min(J), j/2\} \leq |A'| \leq 2j$ *; (ii)* $\mathcal{T}(A') \geq \mathcal{T}(M_{min(J)})$ *; (iii)* $\mathcal{F}(A') \geq \min\{\mathcal{F}(M_i)|i \in J\}/2$ *and (iv) If* e' *is the edge of lowest cost in* A', *then* $c(e') \geq c(e)$, *for every edge* e *that belongs to the matching* $A \setminus A'$.

Proof. We refer to Appendix A ■

3 A Truthful Mechanism for the Unit-Demand Combinatorial Auction Problem

In this section we introduce UDCAM , the main auction mechanism proposed in this paper. The auction $\mathcal{A}_l$ defined below is used as a sub-auction of the UDCAM auction. $\mathcal{A}_l$ is a variation of the VCG auction mechanism where the parameter l limits the number of items that can be sold. In fact VCG mechanism is exactly the mechanism $\mathcal{A}_s$.

Mechanism $\mathcal{A}_l$(H: Graph):

1. The allocation is defined by the matching M of H that maximizes $\mathcal{T}(\cdot)$ among all the matchings in H of size l.

2. If M assigns the consumer i to the item j, then the sale price of j is $p_j = v_i(j) - \mathcal{T}(M) + \mathcal{T}(M_{-i})$, where M_{-i} is the matching that maximizes $\mathcal{T}(\cdot)$ among all the matchings in H_{-i} of cardinality l.

The next lemma helps us bound the revenue achieved by auctions that use $\mathcal{A}_l$.

Lemma 2 *Let* H *be a weighted complete bipartite graph in which there is a matching* M' *with* $2l$ *edges and weights at least* y. *Then* $\mathcal{R}(\mathcal{A}_l(H)) \geq l \times y$.

Proof. Let M be the matching determined by $\mathcal{A}_l(H)$. It suffices to argue that $p_j = v_i(j) - \mathcal{T}(M) + \mathcal{T}(M_{-i}) \geq y$, for every consumer i touched by M.

Since $|M'| \geq 2l$, it follows that there is $e \in M'$ such that $M \cup e - e_{ij}$ is a matching in H_{-i}. Thus, $\mathcal{T}(M_{-i}) \geq \mathcal{T}(M \cup e - e_{ij}) \geq \mathcal{T}(M) - v_i(j) + y$ and, as a consequence, $p_j = v_i(j) - \mathcal{T}(M) + \mathcal{T}(M_{-i}) \geq y$ ■

If we could truthfully determine a suitable value l to apply the best mechanism $\mathcal{A}_l$ on any set of input bids, the resulting mechanism would achieve $O(1)$ revenue-competitiveness with $\mathcal{F}$ and $\log s$ efficiency-competitiveness. By simply choosing l randomly from the set $1, 2, \cdots, 2^{\lfloor \log s \rfloor}$, we can build an auction mechanism similar to that obtained in [13].

Instead, we focus on obtaining mechanisms that achieve the expected revenue (and efficiency) with high probability. Our high-probability results depend on not having too few consumers that are important for the auction, that is, not having too few consumers that contribute with a large constant fraction of the revenue and efficiency to the auction. This assumption is a fairly natural extension of the assumption that the highest bid is not much larger than the second highest bid in single-item auctions, which turns out to be a necessary condition for designing competitive mechanisms for single-item auctions [8].

The idea behind UDCAM is the fact that there is a value l, that depends on G, for which $\mathcal{A}_l$ achieves "high" revenue and "high" efficiency. Since we do not know how to compute the optimum l truthfully, the mechanism splits the consumers into two groups and uses one of them to estimate a suitable value of l. At the end, the sub-auction $\mathcal{A}_l$ is run for the consumers of the other group. As shown in the pseudo-code below, the auction is indexed by a function f that maps a weighted bipartite complete graph into an integer. The role of f is determining the value of l for which $\mathcal{A}_l$ is executed. Its definition will determine the welfare, the revenue and the time complexity of the auction mechanism.

Mechanism $UDCAM_f$:

1. Flip a fair coin n times to split the consumers into two groups. Let G_L (G_R) be the bipartite graph induced by the consumers of the left (right) group and the set of all items.

2. Run $\mathcal{A}_{f(G_L)}(G_R)$

For a function f that maps G_L into an integer, the mechanism above is truthful since it belongs to the VCG family.

3.1 The Revenue function

First, we investigate a definition of f that ensures that $UDCAM_f$ is $O(1)$ revenue-competitive with $\mathcal{F}$ and $O(\ln s)$ efficiency-competitive.

Rev(G_L:graph)

1. Let M^1 be the largest (w.r.t. the number of edges) matching in G_L such that $\mathcal{F}(M^1) \geq \mathcal{F}_{G_L}/3$.
2. Return $\lfloor |M^1|/6 \rfloor$.

Recall that we use $M_G^{\mathcal{F}}$ to denote the largest matching of G that maximizes $\mathcal{F}(\cdot)$. The next theorem shows that the probability that $UDCAM_{Rev}$ obtains an expected revenue of $\Omega(\mathcal{F})$ quickly approaches 1 as $|M_G^{\mathcal{F}}|$ grows.

Theorem 1 *There exists a constant K_0 such that $UDCAM_{Rev}$ is both $O(1)$ revenue-competitive with $\mathcal{F}$ and $O(\ln s)$ efficiency-competitive with probability at least $1 - \frac{74}{e^{(|M_G^{\mathcal{F}}|/108)}}$, for every auction graph G where $|M_G^{\mathcal{F}}| \geq K_0$.*

Proof. For every j, let M_j be a matching of size j in G which maximizes $\mathcal{F}(\cdot)$. The matching M_j is said to be *good* if $j \geq |M_G^{\mathcal{F}}|/3$ and $\mathcal{F}(M_j) \geq \mathcal{F}/9$. Let J be the set of integers defined as $J = \{j | M_j$ is a good matching $\}$.

For every $j \in J$, let C_j be the set of consumers of matching M_j. With respect to Step 1 of the UDCAM auction, we define the event $\mathcal{E}_j$ as the event in which the number of consumers of C_j that lie in the left group is at least $j/3$ and at most $2j/3$. Furthermore, let $\mathcal{E} = \bigcup_{j \in J} \mathcal{E}_j$.

In what follows we make some observations under the assumption that $\mathcal{E}$ occurs. Let M' be the sub-matching of $M_G^{\mathcal{F}}$ induced by the consumers of $M_G^{\mathcal{F}}$ that lie in G_L. Then, M' has at least $|M_G^{\mathcal{F}}|/3$ edges and $\mathcal{F}(M') \geq \mathcal{F}/3 \geq \mathcal{F}_{G_L}/3$. This implies that $\mathcal{F}_{G_L} \geq \mathcal{F}/3$ and $|M^1| \geq |M_G^{\mathcal{F}}|/3$.

Since $\mathcal{F}(M^1) \geq \mathcal{F}_{G_L}/3$ it follows that $\mathcal{F}(M^1) \geq \mathcal{F}/9$ and, as a consequence, $|M^1| \in J$. Therefore, there are at least $\lceil |M^1|/3 \rceil$ consumers of $C_{|M^1|}$ in the right group, which implies on the existence of a matching in G_R, say M^2, of size $\lceil |M^1|/3 \rceil$, where every edge costs at least $\mathcal{F}/(9|M^1|)$. Thus, it follows from Lemma 2 that the revenue of $\mathcal{A}_{\lfloor |M^1|/6 \rfloor}(G_R)$ is at least $\lfloor |M^1|/6 \rfloor \times \mathcal{F}/(9|M^1|) = \Omega(\mathcal{F})$. Since Lemma 1 guarantees that $\mathcal{F} \geq \mathcal{T}/\ln s$, it follows that the welfare is $\Omega(\mathcal{T}/\ln s)$.

Now, we obtain a bound on the probability of event $\mathcal{E}$ happening. A direct application of the Chernoff Bound [15] ensures that the probability of event $\mathcal{E}_j$ not happening is at most $2e^{-j/36}$. Applying the union bound we get that the probability of failure of $\mathcal{E}$ is at most

$$\sum_{j\in J} 2e^{-j/36} \leq \sum_{j=\lceil |M_G^{\mathcal{F}}|/3\rceil}^{\infty} 2e^{-j/36} \leq \frac{2e^{-|M_G^{\mathcal{F}}|/108}}{1-e^{-1/36}} \leq \frac{74}{e^{|M_G^{\mathcal{F}}|/108}}$$

■

We note that the condition on the size of $M_G^{\mathcal{F}}$ required in the previous theorem is implied by a perhaps more intuitive condition that states that a small (constant) number of consumers cannot contribute to the overall welfare with more than $\mathcal{T}/\ln s$.

If we do not concern ourselves with achieving high probability we can obtain a simple auction that is $O(1)$ revenue-competitive with $\mathcal{F}$ and $O(1)$ efficiency-competitive for every auction graph G such that $|M_G^{\mathcal{F}}| > 1$.

Theorem 2 *Let Mixed be the auction mechanism that executes one of the following mechanisms with uniform probability: $UDCAM_{Rev}$, $\mathcal{A}_1$ and VCG. Then, Mixed is $O(1)$ revenue-competitive with $\mathcal{F}$ and $O(1)$ efficiency-competitive for every graph G such that $|M_G^{\mathcal{F}}| > 1$.*

Proof. It is well known that VCG mechanism attains the optimal efficiency $\mathcal{T}$. Since VCG is executed with probability 1/3, then the expected welfare is at least $\mathcal{T}/3$.

Let c be the consumer that bids h, the highest bid in the auction, and let h_2 be the highest bid in the auction defined by G_{-c}. Note that the revenue of $\mathcal{A}_1$ is h_2. If, say, $|M_G^{\mathcal{F}}| \geq 500$, the Theorem 1 ensures that $UDCAM_{Rev}$ achieves an expected revenue of $\Omega(\mathcal{F})$. On the other hand, if $1 < |M_G^{\mathcal{F}}| < 500$, then $\mathcal{F} \leq 500h_2$. Since $\mathcal{A}_1(G)$ is executed with probability 1/3, then the expected revenue is $\Omega(\mathcal{F})$. ■

3.2 A Bicriteria Competitive Auction

Now, we investigate a second definition of f that ensures that $UDCAM_f$ is $O(\ln s)$ revenue-competitive with $\mathcal{T}$ and $O(1)$ efficiency-competitive with high probability. The assumptions required are even weaker than those required in Theorem 1. This new function, called *Bic*, estimates the size of the matching in G that satisfies the conditions in Proposition 2.

Bic(G_L: graph)

1. Let $e_1, \ldots, e_{|M_{G_L}^{\mathcal{T}}|}$ be the edges of $M_{G_L}^{\mathcal{T}}$ listed by non-increasing order of costs. Let i^* be the largest integer such that $i^* \times c(e_{i^*}) \geq \mathcal{T}_{G_L}/(2\ln|M_{G_L}^{\mathcal{T}}|)$.

2. Define $M_b^1 = \bigcup_{j=1}^{i^*} e_j$. Return $\lfloor |M_b^1|/12 \rfloor$.

The main result of this section is the following theorem.

Theorem 3 *Let K_1 be the maximum integer such that $\mathcal{T}(M) < \mathcal{T}/8$, for every matching M in G, with $|M| \leq K_1$. $UDCAM_{Bic}(G)$ is $O(\ln s)$ revenue-competitive with $\mathcal{F}$ and $\mathcal{T}$ and $O(1)$ efficiency-competitive with probability at least $1 - \frac{148}{e^{K_1/36}}$.*

The proof consists of showing that with the probability stated above there is a matching M^* in G_R such that: (i) $|M_b^1|/6 \leq |M^*| \leq (4|M_b^1|)/3$; (ii) $\mathcal{F}(M^*) = \Omega(\mathcal{T}/\log s)$ and (iii) $\mathcal{T}(M^*) = \Omega(\mathcal{T})$.

The next lemma shows that the existence of such a matching indeed ensures that the mechanism performs as desired.

Lemma 3 *If there is a matching M^* in G_R that satisfies the properties (i)-(iii) above, then $\mathcal{W}(\mathcal{A}_{Bic(G_L)}(G_R)) = \Omega(\mathcal{T})$ and $\mathcal{R}(\mathcal{A}_{Bic(G_L)}(G_R)) = \Omega(\mathcal{T}/\ln s)$.*

Proof. Let M^2 be the matching of size $\lfloor |M_b^1|/12 \rfloor$ computed by $\mathcal{A}_{Bic(G_L)}(G_R)$. Since M^* has at most $(4|M_b^1|)/3$ edges, the sum of the costs of the $\lfloor |M_b^1|/12 \rfloor$ most expensive edges of M^* is at least $\lfloor |M_b^1|/12 \rfloor \times 3\mathcal{T}(M^*)/(4|M_b^1|)$. It follows that $\mathcal{T}(M^2) = \Omega(\mathcal{T})$.

On the other hand, since $\mathcal{F}(M^*) = \Omega(\mathcal{T}/\ln s)$, then all the edges of M^* cost at least $K\mathcal{T}/(|M^*|\ln s)$, where K is the constant hidden in the $\Omega()$. Since $2|M^2| = 2\lfloor |M_b^1|/12 \rfloor \leq |M_b^1|/6 \leq |M^*|$, it follows from Lemma 2 that the revenue is at least $\lfloor |M_b^1|/12 \rfloor \times K\mathcal{T}/(|M^*|\ln s)$. By using the fact that $|M^*| \leq (4|M_b^1|)/3$, we conclude that the revenue is $\Omega(\mathcal{T}/\ln s)$ ■

Thus, it suffices to prove the existence of such a matching. The following definition is useful in our proofs.

Definition 2 *Given a matching M in G, let C_j be the set of consumers associated to the j most expensive edges of M. Let $\mathcal{E}_j$ be the event where at least $j/3$ consumers of C_j lie in the left group and at least $j/3$ lie in the right one. Finally, let $\mathcal{E}_M = \bigcup_{j=K_1}^{|M|} \mathcal{E}_j$.*

Two properties of $\mathcal{E}_M$ are useful for our analysis: the failure probability of $\mathcal{E}_M$ decreases exponentially with increases in K_1 and if $\mathcal{E}_M$ occurs then the edges of M are "evenly" distributed between G_L and G_R in the sense that the sub-matching of M induced by the consumers that lie in G_L has approximately the same cost (w.r.t $\mathcal{F}$ and $\mathcal{T}$ metrics) of that induced by the consumers that lie in G_R. The following two propositions formalize these observations.

Proposition 3 *The probability of failure of $\mathcal{E}_M$ is at most $74e^{-K_1/36}$*

Proposition 4 *Let M be a matching in G, with $|M| > K_1$. If $\mathcal{E}_M$ occurs then the sub-matching M' of M induced by the consumers of M that lie in G_L (G_R) satisfies the following properties: $\mathcal{F}(M') \geq \mathcal{F}(M)/3$ and $\mathcal{T}(M') \geq \mathcal{T}(M)/3 - \mathcal{T}/24$*

The proof of Proposition 3 follows from a direct application of Chernoff bound [15]. For the proof of Proposition 4 we refer to Appendix B.

Now, we redefine the concept of a good matching. We say that a matching M of G is *good* if $\mathcal{F}(M) \geq \mathcal{T}/(7 \ln s)$ and $\mathcal{T}(M) \geq \mathcal{T}/7$. The existence of at least one good matching is guaranteed by Proposition 2. Let J be an increasing sequence of integers such that $j \in J$ if and only if there is a good matching in G of cardinality j. For every $j \in J$, let M_j be an arbitrary good matching of size j. Note that the definition of good matchings and the assumption over K_1 in Theorem 3 imply that $min(J) > K_1$.

Proposition 5 *Let A be the approximation matching of $(M_j)_{j \in J}$. If the event $\mathcal{E}_{M_G^{\mathcal{T}}} \cup \mathcal{E}_A$ happens then there is a matching in G_R that meets the conditions (i)-(iv).*

Proof. First, we show that if $\mathcal{E}_{M_G^{\mathcal{T}}}$ occurs then there is a good matching of size $|M_b^1|$ in G. Let M' be the sub-matching of $M_G^{\mathcal{T}}$ induced by the consumers of $M_G^{\mathcal{T}}$ that lie in G_L. It follows from Proposition 4 that $\mathcal{T}(M') \geq \mathcal{T}/3 - \mathcal{T}/8 \geq 7 \times \mathcal{T}/24$. Thus, $\mathcal{T}(M_{G_L}^{\mathcal{T}}) \geq \mathcal{T}(M') \geq 7 \times \mathcal{T}/24$. The definition of M_b^1 and Proposition 2 imply that $\mathcal{T}(M_b^1) \geq (7 \times \mathcal{T})/48$ and $\mathcal{F}(M_b^1) \geq 7 \times \mathcal{T}/(48 \ln s)$. Thus, M_b^1 is a good matching in G.

Since there is a good matching in G of size $|M_b^1|$, it follows from Lemma 1 that there is a sub-matching A' of the approximation matching A such that $\max\{min(J), |M_b^1|/2\} \leq |A'| \leq 2|M_b^1|$, $\mathcal{T}(A') \geq \mathcal{T}/7$ and $\mathcal{F}(A') \geq \mathcal{T}/14 \ln s$.

Since $\mathcal{E}_A$ happens and $|A'| \geq min(J) > K_1$, then $\mathcal{E}_{A'}$ also happens. Let A'' be the sub-matching of A' induced by the consumers of A' that lies in G_R. It follows from Proposition 4 and the previous observation about A' that $|M_b^1|/6 \leq |A''| \leq (4|M_b^1|)/3$, $\mathcal{T}(A'') \geq \mathcal{T}/168$ and $\mathcal{F}(A'') \geq \mathcal{T}/42 \ln s$. Thus, A'' meets the conditions (i)-(iv), which establishes our result. ■

Proof of Theorem 3. If the event $\mathcal{E}_{M_G^{\mathcal{T}}} \cup \mathcal{E}_A$ happens, it follows from Proposition 5 and from Lemma 3 that $\mathcal{A}_{Bic(G_L)}(G_R)$ achieves $\Omega(\mathcal{T})$ efficiency and $\Omega(\mathcal{T}/\ln s)$ revenue.

On the other hand, it follows from Proposition 3 that $\mathcal{E}_{M_G^{\mathcal{T}}} \cup \mathcal{E}_A$ fails with probability at most $148e^{-K_1/36}$. Thus, $\mathcal{E}_{M_G^{\mathcal{T}}} \cup \mathcal{E}_A$ happens with probability at least $1 - 148e^{-K_1/36}$. ■

Both $UDCAM_{Rev}$ and $UDCAM_{Bic}$ can be implemented in polynomial time. For details we refer to Appendix C.

We remark that by modifying the choice of i^* in the function Bic we obtain a trade-off between welfare and revenue. For instance, if we select i^* as the largest integer such that $i^* \times c(e_{i^*}) \geq \mathcal{F}_{\mathcal{T}_{G_L}}/12$ we obtain an auction mechanism for which we can prove both efficiency and revenue bounds similar to those given by Theorem 1. We defer a deeper discussion of this issue for an extended version of this paper.

4 Lower Bounds

In this section we show lower bounds on the competitive ratio of truthful mechanisms for the UDCAP.

The notion of bid independence proved very useful when dealing with a single item auction: an auction is bid independent if given a bid vector $\mathbf{b} = (b_1, b_2, \cdots, b_n)$ the sale price for consumer i can be specified as a function $p_i = f(b_{-i})$, that is, consumer i buys the item iff his bid is at or above p_i, but the price p_i does not depend on the actual bid given by i.

Bid independence implies truthfulness for single item auctions, and in fact Goldberg et al [8] prove that (single-item with unlimited copies) auctions are truthful iff they are equivalent to bid independent auctions. This theorem is key to proving their competitive results.

Unfortunately this theorem does not translate directly to the UDCAP problem. Consider the VCG mechanism applied to the UDCAP problem: it is neither bid-independent nor equivalent to a bid-independent auction since when a consumer changes his bids the assignment of consumers to items may change. Thus it cannot be phrased as a bid independent mechanism, and yet this is a truthful mechanism. A weaker version of this theorem holds for UDCAP auctions:

Proposition 6 *Let $G = (V, E)$ and $G' = (V, E)$ be two auction graphs that differ only in the weights of the edges associated with the $i - th$ consumer. A truthful UDCAP mechanism that sells the same item to i on either G or G' must charge the same price from i on both auctions.*

Proof. Follows directly from the definition of truthfulness. ■

Definition 3 *An* envy-free allocation *for an auction graph G is a pair $(M, \mathbf{p})$ that satisfies the following conditions:*

a) M is a matching in G and $\mathbf{p} = (p_1, \ldots, p_{|I|})$ is a vector defining the price of every item in I.

b) Whenever some item j' satisfies $v_i(j') - p_{j'} > 0$ then an item is allocated to consumer i.

c) if M allocates item j to consumer i, then $v_i(j) - p_j \geq v_i(j') - p_{j'}$, for every $j' \in I$.

The intuition behind the definition of envy-free allocations is that all consumers should be content with such an allocation, since given the price vector $\mathbf{p}$, no item would bring profit to a consumer that does not receive any item, whereas all other consumers are allocated to the items that maximize their profit.

For an envy-free allocation $(M, \mathbf{p})$, let $I(M)$ be the set of items allocated by M and let $\mathcal{N}(M, \mathbf{p}) = \sum_{i \in I(M)} p_i$ be the revenue achieved by the allocation. Being $(M^*, \mathbf{p}^*)$ the envy-free allocation which maximizes $\mathcal{N}(\cdot)$ among all the envy-free allocations for G, we define $\mathcal{N}_G = \mathcal{N}(M^*, \mathbf{p}^*)$.

Now, we show that no randomized truthful mechanism is $o(\ln s)$ revenue-competitive with $\mathcal{N}$. As in [7], our proof employs the Yao's minmax principle for online algorithms [3]. We need the following lemma

Lemma 4 *Let H be a spanning subgraph of G such that the lowest weight edge e_h in H and the largest weight edge $e_{G \setminus H}$ in $G \setminus H$ satisfy $c(e_h) \geq c(e_{G \setminus H})$. If the maximum node degree of H is d, then $\mathcal{N}_G \geq \mathcal{T}_H / d$*

Proof. We refer to appendix D. ■

Theorem 4 *No truthful UDCAP mechanism is $o(\ln s)$ revenue-competitive with neither $\mathcal{N}$ nor $\mathcal{T}$.*

Proof. Let us consider an auction with n consumers and n items. Let $\mathbf{D}$ be a probability distribution over the bids of the n consumers. In terms of graphs, $\mathbf{D}$ induces a probability distribution over the costs of the edges of the complete bipartite graph between the n consumers and the n items. Let $\mathcal{A}^*$ be the deterministic truthful auction mechanism that achieves the maximum possible expected revenue for $\mathbf{D}$. The Yao's minmax principle for online algorithms ensures that $E_{\mathbf{D}}[\mathcal{N}_G]/E_{\mathbf{D}}[\mathcal{R}(\mathcal{A}^*(G))]$ is a lower bound on the revenue-competitive ratio of any randomized truthful auction mechanism with $\mathcal{N}$ metric.

Thus, we define a suitable distribution $\mathbf{D}$ and then we give bounds on both $E_{\mathbf{D}}[\mathcal{R}(\mathcal{A}^*(G))]$ and $E_{\mathbf{D}}[\mathcal{N}_G]$. We represent the consumer bids by a matrix $\mathbf{b}$. For $i, j \in \{1, \ldots, n\}$, we use $b_i(j)$ to denote the bid of consumer i for item j. Let $S = \sum_{i=1}^{n} 1/i^2$ and let $i^+ = (i \bmod n) + 1$, for $i = 1, \ldots, n$.

The distribution $\mathbf{D}$ is defined as follows: for $i, x \in \{1, \ldots, n\}$, (i) $Pr[b_i(i) = x] = 1/(S \times x^2)$; (ii) $b_i(i^+) = b_i(i)$ and (iii) $b_i(j) = 1/2$ if $j \neq i$ and $j \neq i^+$. The goal of this distribution is to assure that each consumer offers non-negligible values for a few items and that every item receives non-negligible bids.

Let B be the set of bid matrices that appear with positive probability in $\mathbf{D}$. We give an upper bound on the contribution of a consumer i to $E_{\mathbf{D}}[\mathcal{R}(A^*(G))]$. For $\mathbf{b} \in B$, we use $\mathbf{b}_{-i}$ to denote the submatrix of $\mathbf{b}$ which defines the bids of all consumers but i. Let $B_{-i} = \{\mathbf{b}_{-i} | \mathbf{b} \in B\}$.

Claim 1 *Let $\mathbf{b}' \in B_{-i}$. Then, the contribution of consumer i to $E_{\mathbf{D}}[\mathcal{R}(\mathcal{A}^*(G))]$ due to bid matrices $\mathbf{b}$, for which $\mathbf{b}_{-i} = \mathbf{b}'$, is at most $5 \times Pr[\mathbf{b}']$.*

The proof of this claim employs Proposition 6 to limit the prices that a consumer can pay for either item i or i^+. Due to the space constraints we move the proof for Appendix E. By summing over all $\mathbf{b}' \in B_{-i}$, we get that the total contribution of consumer i is at most 5. Thus, the expected revenue obtained by summing over all consumers is at most $5n$.

Now, we analyze the expected revenue $E_{\mathbf{D}}[\mathcal{N}_G]$ of the best possible envy-free allocation. Let G be a graph that happens with positive probability in the distribution $\mathbf{D}$ and let H be the subgraph of G induced by the edges that connect the consumers to the items for which they offer at least 1 unit. Clearly, H is a spanning subgraph of G and its edges respect the conditions of Lemma 4. Since the maximum node degree in H is 2, then it follows from Lemma 4 that $\mathcal{N}_G \geq \mathcal{T}_H/2$. In addition, it is easy to check that $\mathcal{T}_G = \mathcal{T}_H$. Thus, we have that $E_{\mathbf{D}}[\mathcal{N}_G] \geq E_{\mathbf{D}}[\mathcal{T}_G]/2$. However, one can show that $E_{\mathbf{D}}[\mathcal{T}_G] = \Omega(n \ln n)$ (for details we refer to Appendix F).

Since $E_{\mathbf{D}}[\mathcal{N}_G]/E_{\mathbf{D}}[\mathcal{R}(\mathcal{A}^*(G))] = \Omega(\ln n)$, it follows from Yao's principle that no randomized truthful mechanism can be $o(\ln n)$ revenue-competitive with $\mathcal{N}$ metric. Since $\mathcal{T} \geq \mathcal{N}$ always holds, then no randomized truthful mechanism can be $o(\ln n)$ revenue-competitive with $\mathcal{T}$ metric. ■

References

[1] Aaron Archer, Christos Papadimitriou, Kunal Talwar, and Éva Tardos. An approximate truthful mechanism for combinatorial auctions with

single parameter agents. In *Proceedings of the fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-03)*, pages 205–214, New York, January 12–14 2003. ACM Press.

[2] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi-unit combinatorial auctions. In *Proceedings of Theoretical Aspects of Rationality and Knowledge*, 2003.

[3] Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, Cambridge, 1998.

[4] Patrick Briest, Piotr Krysta, and Berthold Voecking. Approximation techniques for utilitarian mechanism design. In *Proceedings of the ACM Symposium on Theory of Computing*, 2005.

[5] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[6] Amos Fiat, Andrew V. Goldberg, Jason D. Hartline, and Anna R. Karlin. Competitive generalized auctions. In ACM, editor, *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing, Montréal, Québec, Canada, May 19–21, 2002*, pages 72–81, New York, NY, USA, 2002. ACM Press.

[7] Goldberg, Hartline, Karlin, and Saks. A lower bound on the competitive ratio of truthful auctions. In *STACS: Annual Symposium on Theoretical Aspects of Computer Science*, 2004.

[8] A. Goldberg, J. Hartline, A. Karlin, A. Wright, and M. Saks. Competitive auctions. submitted, 2002.

[9] Andrew V. Goldberg and Jason D. Hartline. Competitive auctions for multiple digital goods. In *ESA: European Symposium on Algorithms*, 2001.

[10] Andrew V. Goldberg and Jason D. Hartline. Envy-free auctions for digital goods. In *Proceedings of the 4th ACM Conference on Electronic Commerce (EC-03)*, pages 29–35, New York, June 9–12 2003. ACM Press.

[11] Andrew V. Goldberg, Jason D. Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-01)*, pages 735–744, New York, January 7–9 2001. ACM Press.

[12] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[13] Venkatesan Guruswami, Jason D. Hartline, Anna Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-05)*, 2005.

[14] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, September 2002.

[15] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms.* Cambridge University Press, 1997.

[16] A. Ronen and A. Saberi. Optimal auctions are hard (extended abstract). In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, 2002.

[17] W. Vickrey. Counterspeculation, auctions, and competive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

A Proof of Lemma 1

First, we note that if we obtain a sub-matching A'' of A that satisfies $(i) - (iii)$, then we can define A' as a submatching of A that contains its $|A''|$ most expensive edges. In this case, A' clearly satisfies $(i) - (iv)$. Thus, it suffices to show the existence of a a submatching of A that satisfies $(i) - (iii)$.

The proof consists of showing the existence of a sub-matching A'' of A that satisfies (i) and then arguing that A'' satisfies (ii) and (iii).

Fix $j \in J$. First, we show that $|A_j| \geq j/2$. For the sake of contradiction, we assume that $|A_j| < j/2$. In this case, $A_j \cup e$ is a matching for some edge $e \in M_j$, which contradicts the definition of A_j. If we also have that $|A_j| \leq 2j$ we set $A'' = A_j$. Otherwise, let us consider the minimum i such that $|A_i| > 2j$. Clearly, $min(J) < i \leq j$. Then, $|A_{pred(i)}| \leq 2j$. Furthermore, the definition of sequence $(A_j)_{j \in J}$ implies that $|A_i| \leq |A_{pred(i)}| + i \leq |A_{pred(i)}| + j$. Since $|A_i| > 2j$, it follows that $|A_{pred(i)}| \geq j$. In this case, we set $A'' = A_{pred(i)}$.

Statement (ii). By construction $M_{min(J)} \subseteq A_j$, for every $j \in J$. Thus, $\mathcal{T}(A'') \geq \mathcal{T}(M_{min(J)})$.

Statement (iii). The proof of statement (i) ensures that $A'' = A_j$, for some $j \in J$. Thus, it suffices to show that $\mathcal{F}(A_j) \geq \min\{\mathcal{F}(M_i) | i \in J\}/2$, for every $j \in J$. We use induction on the elements of J. For $j = min(J)$ the result holds since $A_{min(J)} = M_{min(J)}$. Let us assume result the holds for $pred(j)$. Let e' be the lightest edge of A_j. We have two cases:

(a) $e' \in A_{pred(j)}$. Since $|A_j| \geq |A_{pred(j)}|$, it follows that $\mathcal{F}(A_j) \geq \mathcal{F}(A_{pred(j)}) \geq \min\{\mathcal{F}(M_i) | i \in J\}/2$, where the last inequality follows from the induction hypothesis.

(b) $e' \in M_j$. Since $|A_j| \geq |M_j|/2$, we have $\mathcal{F}(A_j) \geq \mathcal{F}(M_j)/2$.

B Proof of Proposition 4

Since $\mathcal{E}_M$ occurs, then at least $\lceil |M|/3 \rceil$ consumers of M lie in G_R. Thus, $\mathcal{F}(M') \geq \mathcal{F}(M)/3$. The proof of the second property makes use of the following technical lemma.

Lemma 5 *Let $B = (b_i)_{i=1}^t$ be a non-increasing sequence of t positive numbers. For $p \leq t$, let $B_p = (b_i)_{i=1}^p$. Furthermore, let K be an integer and let S' be a subset of $\{1, \ldots, t\}$ that satisfies the following condition: for $K \leq j \leq t$,* *at least $j/3$ elements of $\{1,\ldots,j\}$ belong to S'. Then*

$$\sum_{i\in S'} b_i \geq \left(\sum_{i=K}^{t} b_i\right)/3$$

Proof. Through interchanging arguments we can show that $S' = \{i|2K/3 < i \leq K\} \cup \{3i|K/3 < i \leq \lfloor t/3 \rfloor\}$ is the feasible choice for S' that minimizes $\sum_{i\in S'} b_i$.

Since B is non-increasing we have that

$$3\sum_{i\in S'} b_i \geq 3\sum_{i=K/3}^{\lfloor t/3\rfloor} b_{3i} \geq \sum_{i=K}^{t} b_i,$$

which establishes the result. ■

Proof of the proposition. We assume w.l.o.g. that M' is the sub-matching induced by the consumers of M that lie in G_L. Defining B as the sequence of the weights of the edges of M sorted by non-increasing order. The occurrence of $\mathcal{E}_M$ implies that the conditions of Lemma 5 holds. Thus, $\mathcal{T}(M') = (\mathcal{T}(M) - D)/3$, where D is the sum of the costs of the K_1 most expensive edges of M. The hypothesis under matchings of size K_1 implies that $D \leq \mathcal{T}/8$, which completes the result. ■

C Time Complexity Analysis

Here, we argue that the auction mechanisms proposed in the paper run in polynomial time. First, we show that $\mathcal{A}_l(H)$ runs in polynomial time. Given a weighted bipartite graph $H = (V_1 \cup V_2, E)$ we construct a weighted unit-capacity network flow $N = (V \cup \{s,t\}, E \cup E')$, where s (t) is the network source (terminal) and $E' = \{(s,v)|v \in V_1\} \cup \{(v,t)|v \in V_2\}$. The edge of E connecting the consumer i to the item j has cost $-v_i(j)$. On the other hand, all of the edges in E' have weight 0. The matching which maximizes $\mathcal{T}(\cdot)$ among all the matchings of size l in H corresponds to the minimum cost flow in N when the terminal t needs to receive l units of flow from the source s. Thus, $\mathcal{A}_l(H)$ runs in polynomial time.

In order to implement the Rev function we need to evaluate $\mathcal{F}_{G_L}$ and then $|M^1|$. Let $G_L = (V_L, E_L)$. For a positive real value x, let $K(x)$ be the largest matching in G_L where all the edges cost at least x. It is not hard to see that

$$\mathcal{F}_{G_L} = \max_{ij\in E_L}\{v_i(j) \times |K(v_i(j))|\}$$

and that

$$|M^1| = \max_{ij \in E_L} \{|K(v_i(j))|\}$$

constrained to $|K(v_i(j))| \times v_i(j) \geq \mathcal{F}_{G_L}/3$.

D Proof of Lemma 4

We construct an envy-free allocation $(M, \mathbf{p})$ for H. Let $E(H)$ be the edges of H. The allocation M is obtained following the greedy procedure below

$M \leftarrow \emptyset$; $E_{aux} \leftarrow E(H)$.
While E_{aux} is not empty
$e_{i,j} \leftarrow$ edge of E_{aux} with maximum cost (ties are arbitrarily broken).
$M \leftarrow M \cup e_{i,j}$
Remove from E_{aux}:
Every edge that touches i and
Every edge $e_{i',j'}$ such that $e_{i,j'} \in E(H)$.

The price vector $\mathbf{p}$ is defined as $p_j = b_i(j)$ if j is allocated to the consumer i and $p_j = \infty$ if j is not allocated to any consumer. It is easy to check that $(M, \mathbf{p})$ is an envy-free allocation for G.

In order to prove that $\mathcal{N}_G \geq \mathcal{T}_H/d$, let $e_1, e_2, \ldots, e_{|M|}$ be the edges of M sorted in the order they were added to M. Let H_k be the subgraph of H induced by all the edges that are removed from E_{aux} when e_k is added to M. Note that the number of items in H_k is at most d and e_k is the most expensive edge of H_k. Thus we have the following inequality.

$$\mathcal{T}_H \leq \sum_{k=1}^{|M|} \mathcal{T}_{H_k} \leq d \sum_{k=1}^{|M|} c(e_k) \leq d\mathcal{N}_G \blacksquare$$

E Proof of the Claim 1

For a positive integer z, let $[z]$ be the set of the z first positive integers. First, we note that distribution $\mathbf{D}$ ensures that, for $x > 1$, we have

$$Pr[b_i(i) \geq x] = \sum_{i=\lceil x \rceil}^{n} \left(\frac{1}{i^2 \times S} \right) \leq \frac{\int_{\lceil x \rceil - 1}^{n} (1/y^2) dy}{S} \leq \frac{1}{\lceil x \rceil - 1}.$$

This implies that $x \times Pr[b_i(i) \geq x] \leq 2$, for every $x \geq 1$.

For $x \in [n]$, let $\mathbf{b'}_x$ be the bid matrix in which the consumer i bids x on both items i and i^+ and the other consumers bids are given by $\mathbf{b}_{-i}$. Then, let

$$S_i = \{x| \text{ item } i \text{ is sold to consumer } i \text{ when the bid matrix is } \mathbf{b'}_x\}$$

and

$$S_{i^+} = \{x| \text{ item } i^+ \text{ is sold to consumer } i \text{ when the bid matrix is } \mathbf{b'}_x\}.$$

Since $\mathcal{A}^*$ is truthful, Proposition 6 ensures that the item i (i^+) is always sold by the same price, say p_i (p_{i^+}), for every bid matrix $\mathbf{b'}_x$ such that $x \in S_i (x \in S_{i^+})$.

Thus, the contribution of consumer i for the total revenue, given that the other consumers bids are represented by $\mathbf{b'}$, is at most

$$\begin{aligned} Pr[\mathbf{b'}] \times \left(\sum_{\substack{x\in S_i \\ x \geq p_i}} p_i/(S\times x^2) + \sum_{\substack{x\in S_{i^+} \\ x \geq p_{i^+}}} p_{i^+}/(S\times x^2) + \sum_{x \notin S_i \cup S_{i^+}} 1/(2S\times x^2) \right) &\leq \\ Pr[\mathbf{b'}] \times \left(p_i.Pr[b_i(i) \geq p_i] + p_{i^+}.Pr[b_i(i^+) \geq p_{i^+}] + 1/2\right) &\leq 5Pr[\mathbf{b'}], \end{aligned}$$

F Calculations Required in Theorem 4

Here, we prove that $E_{\mathbf{D}}[\mathcal{T}] = \Omega(n \ln n)$. For a graph G that happens with positive probability in $\mathbf{D}$, let M_G be a matching in G, where the consumer i is connected to the item i, that is, $M_G = \bigcup_{i=1}^{n} e_{i,i}$. Thus,

$$E_{\mathbf{D}}[\mathcal{T}_G] \geq E_{\mathbf{D}}[M_G] = \sum_{i=1}^{n} E_{\mathbf{D}}[b_i(i)] = \sum_{i=1}^{n}\sum_{x=1}^{n} x \times (1/(x^2 \times S)) = \Omega(n \ln n),$$

where the last inequality uses ther fact that $S < \pi^2/6$.