

Prior-independent Multi-Parameter Mechanism Design

Nikhil Devanur¹, Jason Hartline² *, Anna Karlin³, and Thach Nguyen³

¹ nikdev@microsoft.com, Microsoft Research

² hartline@eecs.northwestern.edu, Northwestern University

³ {karlin,ncthach}@cs.washington.edu, University of Washington **

Abstract. In a unit-demand multi-unit multi-item auction, an auctioneer is selling a collection of different items to a set of agents each interested in buying at most unit. Each agent has a different private value for each of the items. We consider the problem of designing a truthful auction that maximizes the auctioneer’s profit in this setting. Previously, there has been progress on this problem in the setting in which each value is drawn from a known prior distribution. Specifically, it has been shown how to design auctions tailored to these priors that achieve a constant factor approximation ratio [2, 5]. In this paper, we present a prior-independent auction for this setting. This auction is guaranteed to achieve a constant fraction of the optimal expected profit for a large class of, so called, “regular” distributions, without specific knowledge of the distributions.

1 Introduction

In a unit-demand multi-unit multi-item auction (UMMA), there are n agents and a seller selling a set of m items. The seller has a supply of m_j units of each item j . Each agent, say the i -th, has a private value v_{ij} for item j , and is only interested in purchasing one unit. The seller runs an auction to determine whom to sell to and at what prices. The auction (or mechanism) takes as input a bid b_{ij} from each agent, and based on the collection of bids, determines a feasible⁴ allocation of items to agents and a price to charge each agent. The question we consider here is how to design a truthful auction for this unit-demand setting that maximizes the seller’s profit.

This is an example of a *multi-parameter* mechanism design problem. While single parameter truthful mechanism design is reasonably well-understood, the understanding of truthful multi-parameter mechanism design is still very much in its infancy. In particular, when the objective of the mechanism designer is something other than maximizing social welfare, we know very little.

* Supported in part by NSF CAREER Award CCF-0846113.

** Part of this work was done while the authors were visiting the IAS, Hebrew University, Israel.

⁴ An allocation is feasible if each agent is allocated at most one item and if no more than m_j items of type j are sold.

Among multi-parameter mechanism design problems, the problem of designing profit maximizing mechanisms for UMMAs has received the most attention [2,5] and has yielded the greatest breakthroughs so far. The main results in this area so far concern Bayesian mechanism design, in which each agent’s values v_{ij} are drawn from *known* prior distributions F_{ij} . In this setting, the goal is, given knowledge of the priors, to design a truthful mechanism which maximizes the seller’s expected profit, where the expectation is taken over the random draws from the prior distributions. For example, Chawla, Hartline, Malec and Sivan [5], and independently (in somewhat different settings) Bhattacharya, Goel, Gollapudi and Munagala [2] have shown how to design truthful mechanisms which are guaranteed to obtain a constant fraction of the optimal expected profit. In addition, Cai, Daskalakis and Weinberg [4] have recently shown how to design PTASes for some special cases of the problem. For a large class of Bayesian combinatorial auction settings, where the priors are known to the mechanism designer, Alaei [1] gives a general framework for approximately reducing the mechanism design problem for multiple buyers to single buyer subproblems, which applies to revenue problems such as the one we consider here.

Inspired by [8], we present a “prior-independent” mechanisms for this problem. By prior-independent, we mean two things: first, that there exist prior distributions from which the agents’ values are drawn, and, second, that the mechanism designer has no knowledge of these priors. Thus, the mechanism has to work well, that is, guarantee a constant fraction of the expected profit achieved by the optimal mechanism tailored to the particular prior distributions, without any knowledge of these priors, and no matter what they happen to be, as long as the distributions satisfy a relevant “regularity” condition. In an independent and contemporaneous work, Roughgarden, Talgam-Cohen and Yan [12] shows that, for settings that are very similar to ours, a simple “welfare maximization with supply reduction” mechanism is also a prior-independent constant approximation mechanism.

Our main theorem is the following:

Theorem. *Consider a UMMA setting where for each item j , v_{ij} is drawn independently from an arbitrary regular distribution F_j . There is an efficiently implementable, truthful mechanism \mathcal{M} that, with no knowledge of the F_j ’s, achieves*

$$\mathbb{E}_{\mathbf{v}} [\mathcal{M}(\mathbf{v})] \geq \frac{1}{8} (\mathbb{E}_{\mathbf{v}} [\text{OPT}(\mathbf{v})]).$$

Here OPT is the optimal deterministic mechanism tailored to the priors F_j .⁵

To prove this theorem, we build on a number of ideas from previous works. First, we take advantage of a reduction from Chawla et al [5] that shows that the optimal expected profit achievable in the unit-demand auction setting is upper bounded by the optimal expected profit achievable in a certain single-parameter

⁵ In [6], Chawla, Malec and Sivan show that for UMMA, the profit of the optimal randomized and deterministic mechanisms are within a constant factor of each other. Therefore, our mechanism also obtain a constant factor of the profit of the optimal randomized mechanism.

variant of the problem. We then show how to design a prior-independent mechanism for this related single-parameter variant and, also, how to convert this mechanism back to a multi-parameter mechanism. To design a prior-independent mechanism for the single-parameter variant, we use three ideas. First, we take advantage of our understanding of the optimal mechanism in single-parameter settings, namely the Myerson mechanism [11]. Second, we relax the unit-demand constraint and instead design a mechanism for a relaxed global supply constraint. The effect of this relaxation is to convert the feasibility constraint on the subset of simultaneously served agents from a matroid intersection constraint to the much easier to handle matroid constraint. Finally, we use a Bulow-Klemperer [3] style result due to Hartline and Roughgarden [10] that shows that in single-parameter matroid settings, if each agent is duplicated, and only one of each pair of duplicates is served in any allocation, then VCG is a 2-approximation to Myerson’s optimal mechanism. Putting these ideas together, we are able to design a prior-independent mechanism for the single-parameter variant of the problem. The conversion back to a multi-parameter mechanism consists of an attempt to simulate the single-parameter mechanism by offering each agent a menu of prices and letting the agent choose his favorite item. However, this simulation is not (and cannot be) faithful because of the differences between the single parameter and multi-parameter setting. Thus we need to show that in expectation not too much revenue is lost, which we prove by taking advantage of the interchangeability of the random variables v_{ij} , $1 \leq i \leq n$ (that is a consequence of the fact that they are independent draws from the same distribution).

Other results. We present a simpler mechanism that obtain a constant factor approximation to the optimal mechanism in the case where there is exactly one unit of each item. Depends on the number of agents and items, the approximation factor of this mechanism can be better than the mechanism in our main result.

We also obtain Bulow-Klemperer type result for some special cases of the problem. In particular, we show that VCG with duplicates approximates the optimal mechanism when either there is no constraint on the supplies of the items, or the distributions of the values satisfy the “monotone hazard rate condition”. For these results, we allow each value v_{ij} to be drawn from a different distribution. However, due to the space limit, these results, as well as most proofs, are deferred to the full paper.

2 Preliminaries

2.1 Settings and Definitions

We define UMMA environments. In such an environment, a seller has k items, with m_j units of item j . Each agent i has, as her private information, a valuation v_{ij} for obtaining each item j and would like to buy at most one unit. We will assume that the values v_{ij} are drawn independently from underlying distributions F_{ij} . Formally,

Definition 1. A UMMA environment E is a tuple $(N, M, \mathcal{S}, \mathbf{F})$ where

- $N = \{1, 2, \dots, n\}$ is the set of bidders
- $M = \{1, 2, \dots, m\}$ is the set of items; there are m_j units of item type j .
- \mathcal{S} is the collection of possible allocations given the supply constraints and the unit-demand constraint. Each $S \in \mathcal{S}$ is a set of pairs (i, j) , which represents the assignment of items to agents. Thus, for each bidder i , each set S contains at most one pair containing i , and for each item j , each set S contains at most m_j pairs containing j .
- $\mathbf{F} = \prod_{i \in N, j \in M} F_{ij}$ where F_{ij} is the distribution of v_{ij} . Defining \mathbf{F} to be a product distribution is equivalent to assuming that all values are independent.

In this paper, we will assume that $F_{ij} = F_j$ for all i . For \mathbf{v} that is a valuation profile drawn from \mathbf{F} , we call the tuple $(N, M, \mathcal{S}, \mathbf{v})$, sometimes abbreviated to (E, \mathbf{v}) , an instance of E .

Similarly, a *single-parameter environment* is defined by a tuple $(N, \mathcal{S}, \mathbf{F})$ where N is the set of agents, \mathcal{S} is the feasible set system (i.e., the subsets of agents N that can be simultaneously served) and \mathbf{F} is the distribution of the valuation profiles. This environment corresponds to the scenario where a seller is offering a service (or goods), and the subsets in \mathcal{S} are the sets of agents that can feasibly be served simultaneously. (For example, if we are describing a t -unit auction, then \mathcal{S} consists of all subsets of agents of size at most t .) We will only consider scenarios where the set system is downward-closed, that is, every subset of a feasible set is also feasible. In this setting, each agent $i \in N$, has a value v_i for being served, where v_i is drawn from prior distribution F_i .

A mechanism takes as input a set of bids from the agents, where in the UMMA environment $b_{i,j}$ is agent i 's bid for item j , and in the single-parameter setting b_i is agent i 's bid for service. The mechanism then outputs an allocation and payments. The outcome of a (deterministic) mechanism on a UMMA instance consists of a set of (item, bidder) pairs represented by an 0/1 allocation vector \mathbf{x} and a payment vector \mathbf{p} . Here $x_{ij} = 1$ if and only if item j is assigned to i and p_i represents the amount bidder i has to pay. Similarly for a single-parameter mechanism, the outcome is a set of winning agents and a payment vector. Again we will use the allocation vector \mathbf{x} as an indicator for which agents are served: $x_i = 1$ if agent $i \in N$ is served and, again, p_i is agent i 's payment.

Given the outcome of a mechanism, the utility of a single-parameter agent i is defined by $u_i = x_i v_i - p_i$, while the utility of a unit-demand multi-parameter agent i is $u_i = \sum_j x_{ij} v_{ij} - p_i$. We assume that agents act to maximize their utility, and we will focus on the design of truthful mechanisms.

Definition 2. A mechanism is truthful if each bidder i maximizes her utility by bidding her true values, no matter what other agents do.

As we have already discussed, our goal will be to design truthful mechanisms for the UMMA environment that maximize the expected profit of the auctioneer without knowledge of the priors from which agent's values are drawn. The fact

that a prior-independent mechanism does not use information about the priors means that for any two environments E and E' that differ only on the distributions from which the agents' values are drawn, a prior-independent mechanism does not distinguish between (E, \mathbf{v}) and (E', \mathbf{v}) for any valuation profile \mathbf{v} which is in the support of both distributions.

In the rest of this section, we review a number of important prior results that we will be using.

2.2 VCG

The VCG mechanism [7, 9, 13] is a truthful mechanism for maximizing social welfare in the various environments we consider in this paper. (It also applies much more generally.)

In the single-parameter environments we are discussing, the VCG mechanism takes as input a vector of bids \mathbf{b} and chooses as its output the feasible set S that maximizes social welfare, i.e. $\sum_{i \in S} b_i$. The payment of an agent is its threshold bid, the minimum value it could have bid and still been part of the winning set.

2.3 Myerson's optimal mechanism for single-parameter environments

We will rely heavily on Myerson's optimal mechanism [11] for profit maximization in single-parameter environments with known priors. This result assumes that agent's valuations are drawn from a product distribution $\mathbf{F} = F_1 \times F_2 \times \dots \times F_n$. Thus, the agents' values are independently (but not identically) distributed.

Given a value v_i drawn from the distribution F_i , the *virtual value* corresponding to v_i , denoted by $\phi_{F_i}(v_i)$ is defined by

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}.$$

F_i is *regular* if the function ϕ_i is monotone nondecreasing, and a product distribution $\mathbf{F} = \prod_i F_i$ is *regular* if each F_i is regular. The class of regular distributions is very large and includes many common distributions such as exponential and normal distributions.

Myerson's result is then the following.

Theorem 1 ([11]). *Let $(N, \mathcal{S}, \mathbf{F})$ be a single-parameter environment, where \mathbf{F} is a regular product distribution. For any truthful mechanism for this environment, characterized by an allocation and a payment rules \mathbf{x} and \mathbf{p} , we have*

$$\mathbb{E}_{\mathbf{v}} \left[\sum_i p_i(\mathbf{v}) \right] = \mathbb{E}_{\mathbf{v}} \left[\sum_i \phi_i(v_i) x_i(\mathbf{v}) \right].$$

The Myerson mechanism Mye for regular distributions is a truthful mechanism that optimizes the quantity inside the expectation on the right hand side point-wise. In other words, given a set of bids \mathbf{b} as input, Myerson selects as winners the feasible subset S such that $\sum_{i \in S} \phi_{F_i}(b_i)$ is maximized. This mechanism maximizes the expected profit among truthful mechanisms.

2.4 Reduction from UMMA environments to single-parameter environments

Definition 3. Given a UMMA environment $E = (N, M, \mathcal{S}, \mathbf{F})$, the representative environment $\text{Rep}(E)$ of E is a single-parameter environment represented by the tuple $(N', \mathcal{S}', \mathbf{F})$ where

- $N' = \{ij : i \in N, j \in M\}$,
- Each set $S' \in \mathcal{S}'$ is constructed by taking a set $S \in \mathcal{S}$ and replacing each pair (i, j) by the agent ij .

Each single-parameter agent ij is a representative of the unit-demand agent i .

Chawla et al [5] show that for UMMA, the optimal revenue in the representative environment upper bounds the optimal revenue of the original environment.

Lemma 1 (Corollary of Lemma 5 in [5]). Let $E = (N, M, \mathcal{S}, \mathbf{F})$ be a UMMA environment and let $\text{Rep}(E)$ be its representative environment. Also, let OPT be the optimal deterministic mechanism for E . We have

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [\text{OPT}(E, \mathbf{v})] \leq \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [\text{Mye}(\text{Rep}(E), \mathbf{v})]$$

2.5 Bulow-Klemperer type results

We first review the concept of duplicates.

Definition 4. Given a single parameter environment E , the environment with duplicates $\text{Dup}(E)$ is obtained by adding a new agent i' for each agent i in E such that:

- The value of i and i' are drawn from the same distribution.
- A feasible set in $\text{Dup}(E)$ is constructed by taking a feasible set in E and replace some of the agents by their duplicates.

Hartline and Roughgarden [10] prove the following results:

Lemma 2 (Theorem 4.4 in [10]). Suppose $E = (N, \mathcal{S}, \mathbf{F})$ is a single-parameter environment where \mathcal{S} is a matroid set system (the feasible sets are independent sets in a matroid on N) and \mathbf{F} is a regular product distribution. Then the expected revenue of VCG on $\text{Dup}(E)$ is at least $1/2$ the expected revenue of Mye on E , i.e.

$$\mathbb{E}_{\mathbf{u} \sim \mathbf{F} \times \mathbf{F}} [\text{VCG}(\text{Dup}(E), \mathbf{u})] \geq \frac{1}{2} \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [\text{Mye}(E, \mathbf{v})].$$

3 Prior-independent mechanism for UMMA

We design a prior-independent mechanism that approximates, in expectation, the revenue of the optimal mechanism for UMMA. Since there is no known characterization of the optimal mechanism for UMMA, we will make use of Lemma 1 and design a mechanism that approximates the revenue of Myerson's optimal auction on the representative environment, via a sequence of reductions using a few intermediate environments. To introduce the elements of this process, we start by considering a very simple special case.

3.1 Unit-demand multi-item auction with unit supply

As a warm-up, we consider the case where there is exactly one unit of each item, and m , the number of items, is at most the number of agents, n . Let E be the original unit-demand environment, $\text{Rep}(E)$ be the representative environment of E . An important intermediate environment in our reduction is obtained by relaxing the unit-demand constraint in $\text{Rep}(E)$.

Definition 5. *Let E be a single-parameter environment where the agents can be partitioned into t groups such that at most one agent in each group can be served. Then $\text{Global}(E)$ is the environment where this constraint is replaced by the constraint that overall, at most t agents can be served.*

In particular, the environment $\text{Global}(\text{Rep}(E))$, or in short, $\text{G.R}(E)$, is obtained by replacing the unit-demand constraint in $\text{Rep}(E)$ by the constraint that in total, at most n representatives can be served.

In the special case where $m < n$, the global constraint that at most n representatives can be served is subsumed by the supply constraint. Therefore, $\text{G.R}(E)$ is equivalent to a combination of m independent single-unit auctions. For each single-unit auction, Bulow and Klemperer [3] show that the second-price auction obtains at least $\frac{n-1}{n}$ times the expected profit of Mye. This implies $\sum_j \text{SPA}_j(\text{G.R}(E)) \geq \frac{n-1}{n} \text{Mye}(\text{G.R}(E))$, where SPA_j is second price auction on the representatives interested in item j . Hence, it suffices to design a mechanism that simulates these second price auctions.

The straightforward approach is to offer to sell to each agent every item at a price equal to the highest bid of other agents for that item and ask her to choose her favorite one, as described in Fig. 1

Mechanism \mathcal{M}_1 for unit-demand multi-item auctions with unit supply

Offer each agent i a price menu \mathbf{p}_i where $p_{ij} = \max_{i' \neq i} v_{i'j}$, and ask her to choose her favorite item.

Fig. 1. A mechanism for unit-demand multi-item auction where there is exactly one unit of each item.

It is immediate that \mathcal{M}_1 is truthful and outputs a feasible allocation. To analyze the the revenue of \mathcal{M}_1 , let p_j be the second highest bid for item j and i_j be the highest bidder for item j . Moreover, let ξ_j be the event that $i_\ell \neq i_j$ for all $\ell \neq j$. Then if ξ_j happens, \mathcal{M}_1 gets at least p_j from item j . Therefore,

$$\begin{aligned} \mathbb{E}[\mathcal{M}_1(E)] &\geq \sum_j p_j \Pr[\xi_j] = \sum_j p_j \left(\frac{n-1}{n}\right)^{m-1} \\ &= \left(\frac{n-1}{n}\right)^{m-1} \sum_j \text{SPA}_j(\text{G.R}(E)) = \left(\frac{n-1}{n}\right)^m \text{Mye}(\text{G.R}(E)) \\ &\geq \left(\frac{n-1}{n}\right)^m \text{Mye}(\text{Rep}(E)) \geq \left(\frac{n-1}{n}\right)^m \text{OPT}(E) \end{aligned}$$

While we start with the assumption that $m \leq n$ to motivate the decompose of $G.R(E)$ into m single-item auctions, the above analysis is independent of this assumption. This yields the following theorem.

Theorem 2. *For unit-supply unit-demand multi-item auction with m items and n agents, \mathcal{M}_1 approximates the revenue of the optimal mechanism within a factor of $(\frac{n}{n-1})^m$.*

In particular, when $m = O(n)$, \mathcal{M}_1 is a constant approximation to the optimal mechanism. When $m \leq n$, the approximation ratio is at most 4 for $n \geq 2$, and converges to e when n tends to ∞ .

Remark 1. There is a mechanism, which is a combination of \mathcal{M}_1 and the mechanism \mathcal{M} described in the next section, which approximates the expected revenue of the optimal mechanism for unit-supply unit-demand multi-item auction within a factor of $2(\frac{n}{n-1})^{n+1}$, even when m is much larger than n . This approximation ratio is worse than that of \mathcal{M} when n is small. However, as n tends to ∞ , it approaches $2e$.

3.2 Unit-demand multi-unit multi-item auction

We turn to the general case where there are more than one unit of each items. For the ease of representation, we will assume that n , the number of agents, is even. (If the number of agents is odd, we can simply discard one agent at a small loss in revenue.)

We use the same approach as in the previous section: simulating a prior-independent mechanism for the single-parameter environment by offering a price menu to each agent. To this end, we have to construct a prior-independent single-parameter mechanism for $\text{Rep}(E)$. Lemma 2 gives us a starting point. To use this result, we introduce duplicates. We restrict $\text{Rep}(E)$ to half of the agents and use the remaining agents as duplicates. Moreover, since the resulting environment is not a matroid environment, we will relax the unit-demand constraint to a global constraint, as discussed in the previous section, to transform it into a matroid environment.

Formally, we make use of the following intermediate environments:

- $H.R(E)$ (an abbreviation of $\text{Half}(\text{Rep}(E))$) is the environment obtained by restricting $\text{Rep}(E)$ to the set of representatives $\{ij : 1 \leq i \leq (i + n/2)\}$.
- $G.H.R(E)$ (an abbreviation of $\text{Global}(\text{Half}(\text{Rep}(E)))$) is the environment obtained by relaxing the unit-demand constraint in $H.R(E)$ to the constraint that, overall, at most $n/2$ representatives can be served.
- $D.G.H.R(E)$ (an abbreviation of $\text{Dup}(\text{Global}(\text{Half}(\text{Rep}(E))))$) is the environment obtained by adding a duplicate for each representative in $G.H.R(E)$. We use the representatives discarded by Half as the duplicates, i.e., $D.G.H.R(E)$ contains the representatives of all agents in N , and for each $i \in \{1, 2, \dots, n/2\}$ and each j , ij and $(i + n/2)j$ are duplicates of each other.

We will show that $\text{VCG}(\text{D.G.H.R}(E))$ is a good approximation of $\text{OPT}(E)$ and then design a multi-parameter mechanism \mathcal{M} that approximates $\text{VCG}(\text{D.G.H.R}(E))$. The chain of reductions is summarized as follows

$$\begin{aligned} \mathbb{E}[\text{OPT}(E)] &\leq \mathbb{E}[\text{Mye}(\text{Rep}(E))] \leq 2\mathbb{E}[\text{Mye}(\text{H.R}(E))] \leq 2\mathbb{E}[\text{Mye}(\text{G.H.R}(E))] \\ &\leq 4\mathbb{E}[\text{VCG}(\text{D.G.H.R}(E))] \leq 8\mathbb{E}[\mathcal{M}(E)] \end{aligned} \tag{1}$$

This chain is the proof of our main theorem.

Lemma 1 already gives us the first inequality.

Intuitively, the revenue Mye gets from $\text{Rep}(E)$ is at most the revenue it gets from $\text{H.R}(E)$ and the environment obtained by restricting $\text{Rep}(E)$ to the other half of the representatives. Since these two restricted environment are identical, the optimal revenue in $\text{Rep}(E)$ is at most twice the optimal revenue in $\text{H.R}(E)$. Hence the second inequality holds.

The third inequality follows from the fact that $\text{Mye}(\text{G.H.R}(E))$ optimizes the virtual surplus over a relaxed set of constraints as compared to $\text{Mye}(\text{H.R}(E))$.

The fourth inequality follows from Lemma 2 and the fact that $\text{D.G.H.R}(E)$ is a matroid environment.

It remains to describe \mathcal{M} and prove the last inequality. As discussed, \mathcal{M} would offer each agent i a price menu \mathbf{p}_i and ask her to choose her favorite item. The question is how to determine p_{ij} for each i and j . Since we would like to simulate $\text{VCG}(\text{D.G.H.R}(E))$, the straightforward answer is to set p_{ij} to the VCG price of representative ij .

However, this straightforward approach does not work, as the VCG price of ij may be determined by the value of another representative of i ; hence the menu offered to i is not independent of her bid. This complication stems from the fact that the VCG price of ij is computed by comparing the welfare of other representatives when (i) ij is included and (ii) ij is excluded from the environment. The important observation is that if in (ii), instead of excluding only ij , we excluded all representatives of i , the price menu would be independent of i 's bids. In another word, p_{ij} should be the externality that i would impose on other representatives of $\text{D.G.H.R}(E)$ by taking item j . This leads to our mechanism \mathcal{M} , detailed in Fig. 2.

\mathcal{M} is clearly truthful. The following two lemmas complete (1) and the proof of our main theorem.

Lemma 3. *\mathcal{M} outputs a feasible allocation.*

Proof. The unit-demand constraint is automatically satisfied because each agent is asked to choose one item. On the other hand, \mathcal{M} offers each item j to at most m_j agents at prices smaller than their bids for it. Therefore, at most m_j agents would buy j and the supply constraint is satisfied. \square

Lemma 4. *The expected revenue of \mathcal{M} is at least $1/2$ the expected revenue of $\text{VCG}(\text{D.G.H.R}(E))$.*

To prove this lemma, we first give a condition so that \mathcal{M} and $\text{VCG}(\text{D.G.H.R}(E))$ get the same revenue from an agent.

Mechanism \mathcal{M} for UMMA ^a

For each agent i , do the following

1. Compute a price menu \mathbf{p}_i , where p_{ij} is the externality i would impose on other representatives in $\text{D.G.H.R}(E)$ by taking item j . In another word, let E_{-i} be $\text{D.G.H.R}(E)$ with all representatives of i removed, then p_{ij} is the maximum of three quantities:
 - the value of ij 's duplicate,
 - the value of m_j -th winner of item j in E_{-i} , i.e. the smallest value among the winners if we are to sell m_j unit of item j and nothing else.
 - the value of the $\frac{n}{2}$ -th winner in E_{-i} .
2. Offer \mathbf{p}_i to agent i and ask her to choose her favorite item.

^a This description of the mechanism assumes the absence of ties. When ties are present, extra steps are required to make sure that \mathcal{M} and VCG ($\text{D.G.H.R}(E)$) break ties in the same way. The detailed mechanism is deferred to the full paper.

Fig. 2. A prior-independent, truthful mechanism for UMMA.

Lemma 5. *Consider welfare maximization in environment $\text{D.G.H.R}(E)$ with and without representative ij . If representative ij is served in the former and none of i 's representatives ij' are served in the latter then the payment of representative ij in VCG ($\text{D.G.H.R}(E)$) equals that of agent i in \mathcal{M} .*

Proof. First, in \mathcal{M} , $p_{ij'} > v_{ij'}$ for any $j' \neq j$. This is because surplus maximization without ij failed to assign an item to ij' , so the externality from serving ij' (thus, the payment i must make for item j) must be greater than $v_{ij'}$. Hence, i would not buy any item $j' \neq j$.

Second, the payment of i for item j in \mathcal{M} is the externality when all of i 's representatives are removed, whereas the payment in VCG for ij is the externality when just ij is removed. By the assumption of the lemma, even when we just remove ij , surplus maximization chooses not to serve another representative of i , so these externalities are the same. \square

Based on this lemma, we can now prove Lemma 4.

Proof (of Lemma 4). Let us condition on the set of values drawn from each distribution and the pairing of values given by the duplicates, i.e., from each distribution F_j draw $n/2$ pairs of values, but defer the decision of which representatives belong to which agents until later. Given this conditioning, VCG ($\text{D.G.H.R}(E)$) is deterministic, i.e., both the winning representatives and their payments are fixed⁶.

⁶ For example, the winning set can be calculated as follows: (i) choose the highest-valued representative in each pair; (ii) among the representatives chosen in (i), choose the m_j highest-valued ones for each item j ; and (iii) among the representatives

We argue that \mathcal{M} 's revenue from each item j is at least $1/2$ of VCG(D.G.H.R(E))'s revenue from it. To this end, fix the representatives that win copies of item j and let representative ij be one of them. Now consider (as in the statement of Lemma 5) finding the surplus maximizing allocation in D.G.H.R(E) with ij removed. Since D.G.H.R(E) is a matroid environment, ij will be replaced by some other representative $i'j'$ and all of VCG's other winners will remain winners. This process allocates at most $n/2$ units of items other than item j to representatives of at most $n/2$ distinct agents. While we have conditioned on the representatives that win units of item j , the agents whose representatives win the other items have not yet been fixed. We now consider realizing the assignment of these other representatives to agents. The probability that agent i is assigned one of these (at most) $n/2$ representatives is at most $1/2$. Hence, the assumption of Lemma 5 holds for representative ij with probability at least $1/2$. Therefore, ij 's expected contribution to \mathcal{M} 's revenue is at least half its contribution to VCG's revenue.

The lemma follows. □

References

1. S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *Proc. 52nd IEEE Symp. on Foundations of Computer Science*, 2011.
2. S. Bhattacharya, G. Goel, S. Gollapudi, and K. Munagala. Budget constrained auctions with heterogeneous items. In *Proc. 41st ACM Symp. on Theory of Computing*, 2010.
3. J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86:180–194, 1996.
4. Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. On optimal multidimensional mechanism design. *SIGecom Exchanges*, 10(2):29–33, 2011.
5. S. Chawla, J. Hartline, D. Malec, and B. Sivan. Sequential posted pricing and multi-parameter mechanism design. In *Proc. 41st ACM Symp. on Theory of Computing*, 2010.
6. S. Chawla, D. Malec, and B. Sivan. The power of randomness in bayesian optimal mechanism design. In *ACM Conference on Electronic Commerce*, pages 149–158, 2010.
7. E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
8. P. Dhangwatnotai, T. Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *Proc. 12th ACM Conf. on Electronic Commerce*, 2010.
9. T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
10. J. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce*, 2009.
11. R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

chosen in (ii), chose the $n/2$ highest-valued ones. The payments can be calculated by keeping track of the minimum value a representative needs to have so that she is not discarded in each step.

12. T. Roughgarden, I. Talgam-Cohen, and Q. Yan. Prior-independence without sampling. Manuscript, 2011.
13. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance*, 16:8–37, 1961.