Simple versus Optimal Mechanisms

Jason D. Hartline*
Electrical Engineering and Computer Science
Northwestern University
Evanston, IL, USA
hartline@eecs.northwestern.edu

Tim Roughgarden†
Computer Science
Stanford University
Stanford, CA, UA
tim@cs.stanford.edu

ABSTRACT

The monopolist's theory of optimal single-item auctions for agents with independent private values can be summarized by two statements. The first is from Myerson [8]: the optimal auction is *Vickrey with a reserve price*. The second is from Bulow and Klemperer [1]: it is better to recruit one more bidder and run the Vickrey auction than to run the optimal auction. These results hold for single-item auctions under the assumption that the agents' valuations are independently and identically drawn from a distribution that satisfies a natural (and prevalent) regularity condition.

These fundamental guarantees for the Vickrey auction fail to hold in general single-parameter agent mechanism design problems. We give precise (and weak) conditions under which *approximate* analogs of these two results hold, thereby demonstrating that simple mechanisms remain almost optimal in quite general single-parameter agent settings.

Categories and Subject Descriptors

F.0 [**Theory of Computation**]: General

General Terms

Economics, Theory, Algorithms

Keywords

Auctions; revenue-maximization; Vickrey auction; VCG mechanism; optimal auctions

1. INTRODUCTION

A striking theme in the theory of single-item auctions is that *simple auctions are optimal*. Foremost, Myerson [8] showed that practically prevalent reserve-price-based auctions are indeed expected revenue-maximizing in natural, though stylized, models. This is fortunate as these auctions are simple and easy to optimize (just set the reserve price). That practitioners employ widely reserve-price-based auctions — in settings much more complex that those where they are provably optimal — motivates our first question: *When are reserve-price-based mechanisms approximately optimal?*

Bulow and Klemperer [1] up the ante further by showing that *the reserve price is unnecessary*, in the sense that a seller of a single item earns more revenue from the Vickrey auction with one extra bidder than from the optimal auction with the original bidders. This guarantee for an auction even simpler than Vickrey with an optimal reserve motivates our second question: *How generally does a Bulow-Klemperer-style result approximately hold?*

Thus the theme of this paper is tight approximation guarantees for the expected revenue of simple auctions, relative to that of an optimal auction, in reasonably general environments — much more general than single- or multi-unit auctions with i.i.d. bidders. To explain our results more precisely, we review some standard auction theory terminology.

The Setting

We consider mechanisms that are ex post incentive compatible and individually rational, meaning that it is a dominant strategy for agents to participate in the mechanism and bid their true private values. A basic example for the single-item setting is the *Vickrey auction with a reserve price*: the highest bidder above the reserve price (if any) wins, and it pays the maximum of the second highest bid and the reserve price.

We are interested in settings much more general than single-item auctions, and consider *general single-parameter environments* where each agent has a valuation for receiving service and there is a set system specifying *feasible sets*, i.e., sets of agents that can be served simultaneously. For example, the feasible sets of a k-unit auction are precisely those with cardinality at most k. We focus on the typical case of *downward-closed* environments where every subset of a feasible set is again feasible. Another example of such an environment is a combinatorial auction with single-minded bidders, where feasible sets correspond to subsets of bidders seeking disjoint bundles of goods. The generalization of the Vickrey auction to arbitrary single-parameter environments is the well-known Vickrey-Clarke-Groves (VCG) mechanism. The VCG mechanism selects a feasible set of

*Supported in part by NSF CAREER Award CCF-0846113.
†Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, an AFOSR MURI grant, and an Alfred P. Sloan Fellowship.

EC'09, July 6–10, 2009, Stanford, California, USA.

winners with maximum total value and charges each winner the externality it imposes on the other agents.

As is standard in the Bayesian optimal auction literature, we assume that each agent's valuation is drawn independently from a known distribution. A standard and common assumption on such a distribution is that it is *regular*, a condition equivalent to concavity of the revenue as a function of the probability of sale (for a single agent). An important subclass of regular distributions are those meeting the *monotone hazard rate (MHR)* condition: intuitively, these have tails no heavier than that of an exponential distribution (which has constant hazard rate). Uniform, normal, and exponential distributions meet the MHR condition.

Myerson [8] characterized the optimal auction in all single-parameter environments, and this result is the starting point of our investigations. In the special case of a multi-unit auction with bidders' valuations drawn i.i.d. from a regular distribution, *Vickrey with an anonymous (i.e., the same for each bidder) reserve price* maximizes the seller's expected revenue over all incentive-compatible and individually rational auctions. The reserve price that provides this guarantee is simply the *monopoly price*, i.e., the revenue-optimal price for a single agent. Thus, the only thing a designer must know about such a distribution to implement the optimal auction is this monopoly price.

Vickrey with an anonymous reserve price is not optimal in more general settings — not with irregular distributions, not with non-identical distributions, and not in more general single-parameter environments. Furthermore, the optimal auction in such settings is typically heavily dependent on the exact form of the bidders' valuation distribution(s). The main message of this paper is that in quite general settings *simple auctions* — like the VCG mechanism with an anonymous reserve, or with bidder-specific reserves, or with no reserves but in an expanded market — *provably approximate the optimal expected revenue, to within a small constant factor.*

Results: Approximation via anonymous reserve.

In an eBay auction, the seller is forced to choose an anonymous reserve price. If the seller has information to distinguish the probable valuations of different bidders — based on past history, say — then a revenue-maximizing auction would employ bidder-specific reserves. Can near-optimal revenue still be achieved with an anonymous reserve price?

In Section 5 we derive the following from one of our Bulow-Klemperer-style bounds.

single-item, regular: For single-item settings and independent (non-identical) regular valuation distributions, Vickrey with a suitable anonymous reserve is a 4-approximation (or better) to the optimal auction.

We also give examples that show a lower bound of 2 (which we believe is tight); that no analogous constant-factor approximation is possible using the Vickrey auction and an anonymous reserve price with irregular distributions; and that no constant-factor approximation is possible using the VCG mechanism and an anonymous reserve price in environments where a super-constant number of non-identically distributed bidders can simultaneously win.

Results: Approximation via monopoly reserves.

Revenue guarantees in more general environments require bidder-specific reserve prices. A natural candidate for such reserve prices are the monopoly prices, which are generally distinct for non-identically distributed bidders. We characterize conditions under which the VCG mechanism with (bidder-specific) monopoly reserves is approximately optimal. (The single-item setting was previously studied in [2, 9, 10].) Our main results for this question, which we prove in Section 3, are as follows.

downward-closed, MHR: For every downward-closed set system and independent (not necessarily identical) valuation distributions that satisfy the MHR condition, VCG with monopoly reserve prices is a 2-approximation (or better) to the optimal mechanism.

We also prove that the above bound is tight, and that no constant-factor approximation guarantee holds for regular (not necessarily MHR) valuation distributions.

Our second result on the topic exchanges additional structure on the feasible sets for more general valuation distributions.

matroid, regular: For every matroid set system — see Section 3.2 for a definition — and independent (not necessarily identical) valuation distributions that satisfy the regularity condition, VCG with monopoly reserve prices is a 2-approximation (or better) to the optimal mechanism.

We also prove a matching lower bound, even for single-item auctions with only two bidders; and that no constant-factor guarantee holds with irregular distributions.

Results: Approximation via duplicating agents.

As already mentioned, Bulow and Klemperer [1] proved that, for a single-item auction with bidder valuations drawn i.i.d. from a regular distribution, the Vickrey auction with an additional bidder outperforms the optimal auction without an additional bidder. Under the same distributional assumptions, this result extends to matroid set systems when the additional agents form a basis (i.e., a maximal independent set) in the matroid [4].

What is a reasonable analog of the Bulow-Klemperer theorem when bidders are asymmetric? To obtain a general result that makes minimal assumptions about the bidders' valuations and the structure of the feasible sets, we consider duplicating every bidder and running the VCG mechanism. Each bidder and its duplicate have i.i.d. valuations, are interchangeable within the set system, and cannot be served simultaneously. *When does the revenue of VCG with duplicate bidders approximate the revenue of the optimal auction without duplicates?*

Our main results for this question, which we prove in Section 4, are as follows.

downward-closed, MHR: For every downward-closed set system and independent (not necessarily identical) valuation distributions that satisfy the MHR condition, the VCG mechanism with duplicate bidders achieves a 3-approximation (or better) of the expected revenue of the optimal mechanism in the original environment.

We obtain this result even though the VCG mechanism is not generally revenue monotone in downward-closed set systems, meaning that adding new agents can actually reduce the revenue of the VCG mechanism (see e.g. [3, 4]). We also prove that the above bound is tight, and that no constant-factor approximation guarantee holds for regular (not necessarily MHR) valuation distributions.

matroid, regular: For every matroid set system and independent (not necessarily identical) valuation distributions that satisfy the regularity condition, the VCG mechanism with duplicate bidders achieves a 2-approximation (or better) of the expected revenue of the optimal mechanism in the original environment.

We do not know if the above bound is tight, as our best lower bound is 4/3.

2. PRELIMINARIES

By a *general single-parameter environment*, we mean a set of n bidders and a collection of *feasible sets* of bidders, which represent the subsets of bidders that can be served simultaneously. An environment is *downward-closed* if every subset of a feasible set is again feasible. Each bidder i has a private *valuation* v_i for service. The profile of agent valuations is denoted by $\mathbf{v} = (v_1, \ldots, v_n)$.

A *mechanism* comprises an allocation rule and a payment rule. An *allocation rule* $\mathbf{x}$ is a function from bid profiles to $\{0, 1\}^n$, indicating the winners (1) and losers (0). A *payment rule* $\mathbf{p}$ is a function from bid profiles to n-vectors of non-negative payments. A mechanism is *truthful* if for every bidder i and every fixed set $\mathbf{v}_{-i}$ of bids by the other bidders, the bidder maximizes its utility $v_i \cdot x_i(\mathbf{v}) - p_i(\mathbf{v})$ by bidding its true valuation (as opposed to some false bid $b_i \neq v_i$). This paper studies only truthful mechanisms, and thus we do not distinguish between bidders' true valuations and the bids they submit to the mechanism. A mechanism is *individually rational* if $p_i(\mathbf{v}) \leq v_i \cdot x_i(\mathbf{v})$ for every bidder i and input $\mathbf{v}$, implying that truthful bidders are guaranteed non-negative utility by the mechanism.

The *Vickrey-Clarke-Groves (VCG) mechanism* works as follows. Given the valuations of the agents, the mechanism selects winners to maximize the social *surplus*, i.e., $\sum_{i=1}^{n} v_i \cdot x_i(\mathbf{v})$, subject to feasibility. The payment of a winning bidder i is the lowest bid at which it would continue to win, which is the difference between the surplus of other agents in the optimal allocation and the surplus of an optimal allocation that excludes bidder i. One easily checks that the VCG mechanism is truthful and individually rational.

The VCG mechanism can also be supplemented with *reserve prices* $\mathbf{r}$, with one reserve price per bidder. The corresponding allocation rule first deletes every bidder i with valuation below its reserve r_i, and then invokes the VCG allocation rule on the remaining bidders. The corresponding payment rule invokes the VCG payment rule on the remaining bidders, and charges a winning bidder i the maximum of its VCG payment and its reserve price r_i. This is again a truthful and individually rational mechanism, for any set $\mathbf{r}$ of reserve prices. We denote this mechanism by $\text{VCG}_{\mathbf{r}}$.

The valuations of the agents are drawn from a product distribution $\mathbf{F} = F_1 \times \cdots \times F_n$. As such, the bidders' valuations are independently, but not necessarily identically, distributed. For simplicity, we assume that every distribution F_i has a continuous density function f_i that is strictly positive on the support of the distribution, which we assume is an interval of the non-negative real line.[1] The *hazard rate* of the distribution F_i at a point z in its support is defined as $h_i(z) = f_i(z)/(1 - F_i(z))$. A distribution F satisfies the *monotone hazard rate (MHR)* condition if its hazard rate $h(z)$ is nondecreasing over its support. Many of the most common distributions (exponential, uniform, etc.) satisfy this condition. A weaker condition is *regularity*, which requires only that $z - 1/h(z)$ is nondecreasing over the support of the distribution and thereby allows for heavier tails. A canonical distribution that is regular but not MHR is the *equal-revenue* distribution, defined by $F(z) = 1 - 1/z$ on $[1, \infty)$.

Our objective is to maximize the revenue of the mechanism. The revenue of a mechanism $\mathcal{M} = (\mathbf{x}, \mathbf{p})$ on an input $\mathbf{v}$, denoted by $\mathcal{M}(\mathbf{v})$, is the sum of the payments collected: $\sum_{i=1}^{n} p_i(\mathbf{v})$. Generally, two different mechanisms earn incomparable revenue: one will collect more on some inputs, the other on other inputs. However, for a fixed distribution over valuations, the expected revenues of different mechanisms are absolutely comparable. As is traditional in optimal mechanism design, we assume that the mechanism designer knows this distribution but not the actual valuations of the agents.

Myerson [8] characterized the optimal (i.e., expected revenue-maximizing) mechanism for every single-parameter environment. To state his results, define the *virtual value* $\varphi(v)$ corresponding to a distribution F and valuation v by $\varphi(v) = v - 1/h(v)$. The following lemma states that the expected payment of an agent is equal to its expected contribution to the allocation's virtual value. This result is central to Myerson's analysis, and also to the present paper.

Proposition 2.1 (Myerson's Lemma) *For every mechanism $(\mathbf{x}, \mathbf{p})$, the expected payment of agent i satisfies*

$$\mathbf{E}_{\mathbf{v}}[p_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}[\varphi_i(v_i) x_i(\mathbf{v})] \,.$$

Given Proposition 2.1, the optimal mechanism is intuitively obvious: for every input $\mathbf{v}$, choose the feasible set that maximizes the *virtual surplus*, $\sum_{i=1}^{n} \varphi_i(v_i) \cdot x_i(\mathbf{v})$. This approach gives the optimal (truthful) mechanism if and only if the distributions are regular. (When distributions are not regular, this approach yields a non-truthful mechanism, and a more sophisticated construction of the optimal mechanism is required [8].) We refer to the optimal mechanism in a single-parameter environment as the Myerson mechanism.

Two special cases of Myerson's result are especially illuminating. First, consider an auction to a single bidder with valuation drawn from a regular distribution F. The virtual surplus-maximizing allocation sells to this bidder whenever the bidder's virtual valuation is non-negative. Notice that this is tantamount to making a take-it-or-leave-it offer of $\varphi^{-1}(0)$. This offer price is known as the *monopoly price*. Second, consider an auction of a single item to one of several bidders with valuations drawn i.i.d. from a regular distribution F. The virtual surplus-maximizing allocation sells to the bidder with the maximum positive virtual valuation. Since the bidders are identically distributed, the winner is the bidder with the highest valuation that is at least $\varphi^{-1}(0)$, the monopoly price. Notice that this is tantamount to running the Vickrey auction with the monopoly reserve price.

[1]For convenience, some examples use discrete distributions.

3. VCG WITH MONOPOLY RESERVES

3.1 General Environments and MHR Valuation Distributions

This section proves that using (bidder-specific) monopoly reserve prices and no additional information about bidders suffices for a 2-approximation of the optimal expected revenue in general downward-closed single-parameter environments, assuming that the valuation distributions are independent (not necessarily identical) and satisfy the MHR condition. We also show that no better bound is possible, and that no constant bound is possible if the MHR condition is relaxed to regularity. We begin with a simple lemma.

Lemma 3.1 *Let F be an MHR distribution with monopoly price r and virtual valuation function φ. For every $v \geq r$,*

$$r + \varphi(v) \geq v.$$

PROOF. We can derive

$$r + \varphi(v) = r + v - \tfrac{1}{h(v)} \geq r + v - \tfrac{1}{h(r)} = v,$$

where the first equality follows from the definition of a virtual valuation, the inequality from the MHR assumption and the fact that $v \geq r$, and the final equality from the definition of the monopoly price (Section 2). □

We now give a tight bound on the expected revenue of the VCG mechanism with (bidder-specific) monopoly reserve prices (denoted $\text{VCG}_{\mathbf{r}}$).

Theorem 3.2 *For every downward-closed environment with valuations drawn independently from distributions that satisfy the MHR condition, the expected revenue of VCG with the monopoly reserves is a 2-approximation to the expected revenue of the optimal mechanism.*

PROOF. Let $\mathbf{x}$ be the allocation rule of $\text{VCG}_{\mathbf{r}}$ — of VCG with the monopoly reserve prices $\mathbf{r} = (r_1, \ldots, r_n)$. Let $\mathbf{x}^*$ denote the allocation rule of Myerson. We first note that

$$\sum\nolimits_{i=1}^{n} v_i \cdot x_i(\mathbf{v}) \geq \sum\nolimits_{i=1}^{n} v_i \cdot x_i^*(\mathbf{v}) \qquad (1)$$

for every profile $\mathbf{v}$, as $\text{VCG}_{\mathbf{r}}$ chooses a surplus-maximizing allocation among those that exclude all bidders that fail to meet their reserve prices, and by downward-closure, Myerson outputs one such allocation.

By Myerson's Lemma (Proposition 2.1),

$$\mathbf{E}[\text{VCG}_{\mathbf{r}}(\mathbf{v})] = \mathbf{E}\Big[\sum\nolimits_{i=1}^{n} \varphi_i(v_i) x_i(\mathbf{v})\Big].$$

Since all winners pay at least the reserve price,

$$\mathbf{E}[\text{VCG}_{\mathbf{r}}(\mathbf{v})] \geq \mathbf{E}\Big[\sum\nolimits_{i=1}^{n} r_i x_i(\mathbf{v})\Big].$$

By the linearity of expectation,

$$2\mathbf{E}[\text{VCG}_{\mathbf{r}}(\mathbf{v})] \geq \mathbf{E}\Big[\sum\nolimits_{i=1}^{n} (r_i + \varphi_i(v_i)) x_i(\mathbf{v})\Big]. \qquad (2)$$

By the definition of $\text{VCG}_{\mathbf{r}}$, $x_i(\mathbf{v}) = 0$ whenever $v_i < r_i$. We can therefore complete the proof by writing

$$\begin{aligned} 2\mathbf{E}_{\mathbf{v}}[\text{VCG}_{\mathbf{r}}(\mathbf{v})] &\geq \mathbf{E}\Big[\sum\nolimits_{i=1}^{n} v_i \cdot x_i(\mathbf{v})\Big] \\ &\geq \mathbf{E}\Big[\sum\nolimits_{i=1}^{n} v_i \cdot x_i^*(\mathbf{v})\Big] \\ &\geq \mathbf{E}[\text{Myerson}(\mathbf{v})], \end{aligned}$$

where the first inequality follows from (2) and Lemma 3.1, the second from (1), and the third from the individual rationality of Myerson. □

The bound of 2 in Theorem 3.2 cannot be improved, even in the special case of i.i.d. bidders.

Example 3.3 (Sketch.) We first give a non-i.i.d. example. For n sufficiently large, the bidders $1, 2, \ldots, n$ are "small", with valuations drawn i.i.d. from an exponential distribution with rate 1. Bidder 0 is "big" and its valuation is deterministically equal to $(2 - \epsilon)n/e$, where $e = 2.718\ldots$ and $\epsilon > 0$ is an arbitrarily small constant. The feasible subsets are precisely those that do not contain both the big bidder and a small bidder.

An obvious mechanism is to always allocate to the big bidder and earn revenue $(2-\epsilon)n/e$; the optimal mechanism earns at least this. To estimate $\mathbf{E}_{\mathbf{v}}[\text{VCG}_{\mathbf{r}}(\mathbf{v})]$ in this environment, observe that the monopoly reserves are $(2-\epsilon)n/e$ for the big bidder and 1 for the small bidders. As n grows large, the number of small bidders that meet the reserve is tightly concentrated around n/e, and the average valuation of these bidders is tightly concentrated around 2. The mechanism $\text{VCG}_{\mathbf{r}}$ will, almost surely as $n \to \infty$, allocate to the small bidders that meet their reserve, collect a payment of 1 from each, and earn only half of the revenue obtained by an optimal mechanism.

This example can be extended to an i.i.d. environment. First, the example remains valid even if the exponential distribution used for small bidders is truncated at a sufficiently high value, say H. (We can safely take n much larger than H.) Let F denote this distribution. The main idea for simulating the big bidder with small ones with valuations drawn from F is this: if a sufficiently large group of small bidders participate in a single-item (Vickrey) auction, then both the highest- and second-highest valuations will be nearly the maximum-possible value (with high probability), so both the value and revenue obtained from this group are tightly concentrated around H. The big bidder with valuation $(2-\epsilon)n/e$ can thus be simulated with $\approx (2-\epsilon)n/eH$ independent such groups. The feasible allocations are now those that do not allocate to both an original small bidder and also one of the small bidders used to simulate the big bidder. The logic behind the previous example continues to hold.

Also, the MHR condition is necessary for Theorem 3.2. Even for i.i.d. and regular valuation distributions, no constant approximation factor is possible.

Example 3.4 (Sketch.) Consider one big bidder, n small bidders for large n, and feasible allocations as in Example 3.3. Fix an arbitrarily large constant H. The small bidders' valuations are i.i.d. draws from the equal-revenue distribution on $[1, H]$ (see Section 2). The monopoly price for such bidders is 1. (The distribution can also be perturbed so that 1 is the unique monopoly price.) The big bidder's valuation is deterministically $n\sqrt{\ln H}$; the expected revenue of the optimal mechanism is clearly at least this. For n sufficiently large, the sum of the small bidders' valuations is tightly concentrated around $n \ln H$. $\text{VCG}_{\mathbf{r}}$ almost surely allocates to all small bidders and obtains revenue n.

Modifying this example as in Example 3.3 shows that, even with i.i.d. regular distributions, the expected revenue of $\text{VCG}_{\mathbf{r}}$ is not always a constant fraction of that of an optimal mechanism.

3.2 Matroid Environments and Regular Valuation Distributions

Theorem 3.2 can be extended to the more general class of regular valuation distributions if we further restrict the structure of the environment. Recall that a *matroid* comprises a ground set of elements E and a non-empty collection $\mathcal{I} \subseteq 2^E$ of *independent sets* that satisfy two properties: (1) subsets of independent sets are again independent; (2) given two independent sets I_1 and I_2 with $|I_1| < |I_2|$, an element of $I_2 \backslash I_1$ can be added to I_1 without destroying its independence. A *base* of a matroid is a maximal independent set; all maximal independent sets have the same cardinality.

Talwar [12] showed that the special structure of matroids leads to unusually well-behaved VCG payments; his work was later refined by Karlin et al. [7]. One important property of matroids is the following (e.g. [11, Corollary 39.12a]).

Proposition 3.5 *Let B_1, B_2 be independent sets of size k in a matroid M. Then there is a bijective function $f : B_2 \backslash B_1 \to B_1 \backslash B_2$ such that, for every $i \in B_2 \backslash B_1$, the set $B_1 \backslash \{f(i)\} \cup \{i\}$ is independent in M.*

Also, the following proposition follows immediately from the definition of VCG payments.

Proposition 3.6 *Fix an arbitrary single-parameter environment and bidder valuations $\mathbf{v}$. Suppose that the winners in the VCG mechanism are W, with $i \in W$ and $j \notin W$; and that $W \backslash \{i\} \cup \{j\}$ is a feasible set. Then the VCG payment of i is at least v_j.*

In a *matroid environment*, the feasible sets of winners form a matroid on the bidders. Examples include multi-item auctions (corresponding to a symmetric or uniform matroid), spanning tree auctions (graphic matroids), matchable nodes in a bipartite graph (transversal matroids), and so on. We now show that, in such environments, the revenue guarantee in Theorem 3.2 extends to regular valuation distributions. The following theorem is a direct consequence of Lemmas 3.9 and 3.10 that are stated and proved below.

Theorem 3.7 *For every matroid environment with valuations drawn independently from distributions that satisfy the regularity condition, the expected revenue of VCG with the monopoly reserves is a 2-approximation to the expected revenue of the optimal mechanism.*

We cast the proof of this theorem in a more general context to assist later arguments in the paper.

Definition 3.8 Let $\mathcal{M}$ and $\mathcal{M}'$ be two mechanisms for a given environment. Let $W(\mathbf{v})$ and $W'(\mathbf{v})$ denote the winners in $\mathcal{M}$ and $\mathcal{M}'$, respectively, with the valuation profile $\mathbf{v}$. The mechanism $\mathcal{M}$ *is commensurate with* $\mathcal{M}'$ if

(C1) $\mathbf{E}_{\mathbf{v}}\left[\sum_{i \in W(\mathbf{v}) \backslash W'(\mathbf{v})} \varphi_i(v_i)\right] \geq 0$

and

(C2) $\mathbf{E}_{\mathbf{v}}\left[\sum_{i \in W(\mathbf{v}) \backslash W'(\mathbf{v})} p_i(\mathbf{v})\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in W'(\mathbf{v}) \backslash W(\mathbf{v})} \varphi_i(v_i)\right]$,

where $\mathbf{p}$ denotes the payment rule of $\mathcal{M}$.

The first condition (C1) asserts that the expected virtual value of bidders winning in $\mathcal{M}$ but not $\mathcal{M}'$ is non-negative. This assertion is generally non-trivial even though the unconditional expectation of a bidder's virtual valuation is zero, because of the implicit conditioning on the bidders losing in $\mathcal{M}'$. The second condition (C2) requires that the expected total payment from bidders of $W(\mathbf{v}) \backslash W'(\mathbf{v})$ in $\mathcal{M}$ is at least the expected virtual surplus of $W'(\mathbf{v}) \backslash W(\mathbf{v})$ in $\mathcal{M}'$. In our applications of Definition 3.8, $\mathcal{M}$ will be a variant on the VCG mechanism and $\mathcal{M}'$ will be the optimal mechanism.

Satisfying Definition 3.8 is a sufficient condition for the expected revenue of mechanism $\mathcal{M}$ to be a 2-approximation to that of the mechanism $\mathcal{M}'$.

Lemma 3.9 *If a mechanism $\mathcal{M}$ is commensurate with a mechanism $\mathcal{M}'$, then*

$$\mathbf{E}_{\mathbf{v}}[\mathcal{M}(\mathbf{v})] \geq \frac{1}{2} \cdot \mathbf{E}_{\mathbf{v}}\left[\mathcal{M}'(\mathbf{v})\right].$$

PROOF. We argue separately that

$$\mathbf{E}_{\mathbf{v}}[\mathcal{M}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v}}\left[\sum\nolimits_{i \in W(\mathbf{v}) \cap W'(\mathbf{v})} \varphi_i(v_i)\right] \tag{3}$$

and

$$\mathbf{E}_{\mathbf{v}}[\mathcal{M}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v}}\left[\sum\nolimits_{i \in W'(\mathbf{v}) \backslash W(\mathbf{v})} \varphi_i(v_i)\right]. \tag{4}$$

Adding these and applying linearity of expectation and Myerson's Lemma (Proposition 2.1) yields the theorem: the left-hand side equals $2\mathbf{E}_{\mathbf{v}}[\mathcal{M}(\mathbf{v})]$ and the right-hand side equals $\mathbf{E}_{\mathbf{v}}[\mathcal{M}'(\mathbf{v})]$.

To derive inequality (3), write

$$\mathbf{E}_{\mathbf{v}}[\mathcal{M}(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in W(\mathbf{v})} \varphi_i(v_i)\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in W(\mathbf{v}) \cap W'(\mathbf{v})} \varphi_i(v_i)\right],$$

where the equality follows from Myerson's Lemma and the inequality follows from condition (C1).

To derive inequality (4), let $\mathbf{p}$ denote the payment rule of $\mathcal{M}$ and write

$$\mathbf{E}_{\mathbf{v}}[\mathcal{M}(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in W(\mathbf{v})} p_i(\mathbf{v})\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in W(\mathbf{v}) \backslash W'(\mathbf{v})} p_i(\mathbf{v})\right]$$
$$\geq \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in W'(\mathbf{v}) \backslash W(\mathbf{v})} \varphi_i(v_i)\right],$$

where the equality is the definition of revenue, the first inequality follows from the non-negativity of payments, and the final inequality follows from condition (C2). $\square$

The next lemma completes the proof of Theorem 3.7.

Lemma 3.10 *For matroid environments and regular valuation distributions, the VCG mechanism with monopoly reserve prices,* $\text{VCG}_{\mathbf{r}}$*, is commensurate with the optimal mechanism* Myerson.

PROOF. For the first condition (C1) of Definition 3.8, recall that regularity of the distribution F_i implies that $\varphi_i(v_i) \geq 0$ if and only if $v_i \geq r_i$. Thus, in the mechanism $\text{VCG}_{\mathbf{r}}$, all winners have non-negative virtual valuations with probability 1. The inequality (C1) follows.

We prove the second condition (C2) pointwise (for each $\mathbf{v}$). Let M denote the given matroid and $M_{\geq \mathbf{r}}$ its restriction to the bidders that meet their reserves — equivalently, to the bidders with a non-negative virtual valuation. This restriction is again a matroid. Both $\text{VCG}_{\mathbf{r}}$ and Myerson are defined as maximizers of a non-negative weight function over $M_{\geq \mathbf{r}}$ — valuations and virtual valuations, respectively. Thus, the winners of each form bases (and have the same cardinality) of the matroid $M_{\geq \mathbf{r}}$. Let $f : W'(\mathbf{v}) \backslash W(\mathbf{v}) \to W(\mathbf{v}) \backslash W'(\mathbf{v})$ denote a bijection of the form guaranteed by Proposition 3.5. Applying Proposition 3.6 proves that $p_{f(i)}(\mathbf{v}) \geq v_i$ for every $i \in W'(\mathbf{v}) \backslash W(\mathbf{v})$. Of course, $v_i \geq \varphi_i(v_i)$ by the definition of a virtual valuation. Summing over all $i \in W'(\mathbf{v}) \backslash W(\mathbf{v})$ completes the proof. □

Example 3.11 We sketch a matching lower bound for Theorem 3.7. Let H be a sufficiently large constant and consider a single-item auction with two bidders. The first bidder has valuation deterministically equal to 1, while the second bidder's valuation is drawn from the equal-revenue distribution on $[1, H]$. The monopoly price for the second bidder is 1 — again, the distribution can be perturbed so that this is uniquely optimal — and the mechanism $\text{VCG}_{\mathbf{r}}$ will earn unit revenue on every input.

Now consider the mechanism that: (1) sells the good to the second bidder at price H, if possible; and (2) otherwise sells the good to the first bidder at price 1. The expected revenue of this mechanism is $H \cdot \frac{1}{H} + 1 \cdot (1 - \frac{1}{H})$, which approaches 2 as H tends to infinity.

Remark 3.12 Unlike Examples 3.3 and 3.4, there is no hope of producing an i.i.d. version of Example 3.11. The reason is that the two mechanisms $\text{VCG}_{\mathbf{r}}$ and Myerson coincide in matroid environments with i.i.d. bidder valuations, provided the distribution is regular (see [4]).

4. BULOW-KLEMPERER-TYPE RESULTS

This section uses approximation to extend the Bulow-Klemperer theorem [1] to general (asymmetric) single-parameter environments. The analog of "adding one or more bidders" is not clearly defined in general environments; to obtain a result with minimal assumptions on the environment and bidders' valuations, we consider the expected revenue of the VCG mechanism after duplicating every bidder. Our revenue guarantees with duplicate bidders also have interesting consequences for anonymous reserve prices in single-item auctions *without* duplicate bidders (Theorem 5.1).

Formally, the *duplication* of a single-parameter environment is defined as follows. Each bidder i with valuation distribution F_i is replaced by a pair i, i' of bidders, whose valuations are i.i.d. draws from F_i. The feasible sets of the duplicated environment are defined as those satisfying: (1) at most one bidder from each pair is selected; (2) the set of winners, when naturally interpreted as a set of bidders from the original environment, is a feasible set in that environment.

Notation. Let $\mathbf{v}, \mathbf{v}'$ denote the valuation profile of the $2n$ bidders where $\mathbf{v}$ is that of the originals and $\mathbf{v}'$ is that of the duplicates. For ease of notation we will let $i' = n + i$ denote the index of the duplicate of bidder i. Furthermore, we will refer to the value of i's duplicate as both v_i' and $v_{i'}$. The valuation profile without both i and i's duplicate is $\mathbf{v}_{-i}, \mathbf{v}'_{-i}$.

4.1 General Environments and MHR Valuation Distributions

Our main result in this section relies on the following technical lemma, which we prove in Appendix A.

Lemma 4.1 *Let v_1, v_2 denote two i.i.d. samples from a monotone hazard rate distribution F with virtual valuation function φ, and t a non-negative real number. Then*

$$\mathbf{E}[\max\{\varphi(v_1), \varphi(v_2)\} \mid \max\{v_1, v_2\} \geq t] \\ \geq \tfrac{1}{3} \cdot \mathbf{E}[\max\{v_1, v_2\} \mid \max\{v_1, v_2\} \geq t] .$$

We now show that the VCG mechanism's expected revenue in a duplicated environment is a constant-factor approximation of the maximum expected revenue achievable in the original environment, provided the valuation distributions satisfy the MHR condition.

Theorem 4.2 *For every downward-closed environment with valuations drawn independently from distributions that satisfy the MHR condition, the expected revenue of VCG with duplicates is a 3-approximation to the expected revenue of the optimal mechanism without duplicates.*

Proof. Fix a bidder i and its duplicate i' and bids $\mathbf{v}_{-i}, \mathbf{v}'_{-i}$ for the other bidders of the duplicated environment. By the definition of the VCG mechanism, there is a threshold $t \geq 0$ such that: if $v_i, v_{i'} < t$, then neither i nor i' is allocated to; and if at least one of $v_i, v_{i'}$ exceeds t, then the bidder among i, i' with a higher valuation (and hence a higher virtual valuation) is allocated to. Lemma 4.1 then implies that

$$\mathbf{E}\left[\varphi_i(v_i) \cdot x_i(\mathbf{v}, \mathbf{v}') + \varphi_i(v_{i'}) \cdot x_{i'}(\mathbf{v}, \mathbf{v}') \,\middle|\, \mathbf{v}_{-i}, \mathbf{v}'_{-i}\right] \\ \geq \tfrac{1}{3} \cdot \mathbf{E}\left[v_i \cdot x_i(\mathbf{v}, \mathbf{v}') + v_{i'} \cdot x_{i'}(\mathbf{v}, \mathbf{v}') \,\middle|\, \mathbf{v}_{-i}, \mathbf{v}'_{-i}\right] .$$

Taking expectations over $\mathbf{v}_{-i}, \mathbf{v}'_{-i}$, summing over all pairs of duplicates, and applying linearity of expectation and Myerson's Lemma (Proposition 2.1) yields

$$\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\text{VCG}(\mathbf{v}, \mathbf{v}')\right] \geq \tfrac{1}{3} \cdot \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i=1}^{2n} v_i \cdot x_i(\mathbf{v}, \mathbf{v}')\right] . \quad (5)$$

VCG always picks a surplus-maximizing solution, and the expected maximum-possible surplus in the duplicated environment — the expectation on the right-hand side of (5)) — is clearly at least that in the original environment. This in turn upper bounds the expected revenue of every individually rational mechanism (like Myerson) in the original environment. The theorem follows. □

The factor of 3 in Theorem 4.2 cannot be improved, even in the special case of i.i.d. bidders.

Example 4.3 Consider the following variant of Example 3.3. In the original environment, there are n small bidders with valuations drawn i.i.d. from an exponential distribution with rate 1. The big bidder's valuation is deterministically $(\frac{3}{2} - \epsilon)n$ for a small constant $\epsilon > 0$. The optimal mechanism in this environment obtains expected revenue at least $(\frac{3}{2} - \epsilon)n$.

In the corresponding duplicated environment, the expected maximum and minimum valuation of a pair of small bidders are $\frac{3}{2}$ and $\frac{1}{2}$, respectively. Almost surely as $n \to \infty$, VCG with duplicates will allocate only to small bidders (to achieve welfare $\approx \frac{3}{2}n$) and obtain revenue tightly concentrated around $n/2$.

As in Examples 3.3 and 3.4, this lower bound can be modified to apply to the special case of i.i.d. bidders.

4.2 Matroid Environments and Regular Valuation Distributions

We now show that for matroid environments and regular distributions, the VCG mechanism's expected revenue in a duplicated environment is at least half that of the optimal mechanism in the original environment.

Theorem 4.4 *For every matroid environment with valuations drawn independently from distributions that satisfy the regularity condition, the expected revenue of VCG with duplicates is a 2-approximation to the expected revenue of the optimal mechanism without duplicates.*

Theorem 4.4 follows by combining Lemma 3.9 with the next lemma, which shows that the VCG mechanism in a duplicated environment is commensurate with the optimal mechanism in the original environment, in the sense of Definition 3.8.[2]

Lemma 4.5 *For matroid environments and regular valuation distributions, the VCG mechanism with duplicates is commensurate with the optimal mechanism* Myerson *without duplicates.*

PROOF. We begin with the first requirement (C1) of Definition 3.8. Let $W(\mathbf{v},\mathbf{v}')$ and $W'(\mathbf{v},\mathbf{v}')$ denote the winners in the VCG mechanism (with duplicates) and the optimal mechanism (without duplicates), respectively. By definition, $W'(\mathbf{v},\mathbf{v}')$ is independent of $\mathbf{v}'$ and cannot contain any duplicate bidders.

Condition on $\mathbf{v}$ but not $\mathbf{v}'$; this fixes the value of $W'(\mathbf{v},\mathbf{v}')$. We argue that

$$\mathbf{E}_{\mathbf{v}'}\left[\sum\nolimits_{i \in W(\mathbf{v},\mathbf{v}')\setminus W'(\mathbf{v},\mathbf{v}')} \varphi_i(v_i)\right] \geq 0; \qquad (6)$$

the unconditional inequality in (C1) follows. We prove (6) by showing that the expected combined contribution of each original bidder i and its duplicate i' to the left-hand side is non-negative. Recall that if one of i, i' belongs to $W(\mathbf{v},\mathbf{v}')$, it is the bidder with higher valuation and hence, by regularity, with higher virtual valuation.

First consider an original bidder i that belongs to the winner set $W'(\mathbf{v},\mathbf{v}')$. Since the valuation distributions are regular, the optimal mechanism selects only bidders with a non-negative virtual valuation, so $\varphi_i(v_i) \geq 0$. It follows that the contribution from i, i' to the left-hand side of (6) in this case is non-negative with probability 1: if $i, i' \notin W(\mathbf{v},\mathbf{v}')$ the contribution is zero; otherwise it is

$$\max\{\varphi_i(v_i), \varphi_i(v_{i'})\} \geq \varphi_i(v_i) \geq 0.$$

Now suppose that the original bidder i is not in $W'(\mathbf{v},\mathbf{v}')$. Condition further on the valuations $\mathbf{v}'_{-i}$ of all duplicates other than i', and let $\mathcal{E}$ denote the event that one of i, i' is included in $W(\mathbf{v},\mathbf{v}')$. If $\neg\mathcal{E}$ occurs, then the contribution from i, i' to the left-hand side of (6) is zero. Since $\mathbf{v}, \mathbf{v}'_{-i}$ are fixed, event $\mathcal{E}$ occurs if and only if $v_{i'}$ is at least some non-negative threshold t. In this case, the expected contribution of i, i' is

$$\mathbf{E}_{v_{i'}}[\max\{\varphi_i(v_i), \varphi_i(v_{i'})\}],$$

conditioned on $\mathbf{v}$, $\mathbf{v}'_{-i}$, and $\mathcal{E}$. This is lower bounded by the analogous conditional expectation of $\varphi_i(v_{i'})$, which is equivalent to

$$\mathbf{E}_{v_{i'}}[\varphi_i(v_{i'}) \,|\, \varphi_i(v_{i'}) \geq t]. \qquad (7)$$

Since the unconditional expectation of a virtual valuation is zero and φ_i is nondecreasing (by regularity), the quantity in (7) in non-negative. Taking expectations over whether or not $\mathcal{E}$ occurs, and then over $\mathbf{v}'_{-i}$, completes the argument.

For condition (C2) of Definition 3.8, we proceed as in the proof of Lemma 3.10. We prove the condition pointwise, for each profile $\mathbf{v}, \mathbf{v}'$. Let M' denote the duplication of the given matroid environment, which is itself a matroid environment. The winners $W(\mathbf{v},\mathbf{v}')$ of the VCG mechanism are a basis of M'. The winners $W'(\mathbf{v},\mathbf{v}')$ of the Myerson mechanism can be naturally viewed as an independent set of M'. Choose a subset $S \subseteq W(\mathbf{v},\mathbf{v}')$ of winners of size $|W'(\mathbf{v},\mathbf{v}')|$. Applying Propositions 3.5 and 3.6 as in the proof of Lemma 3.10 shows that

$$\sum_{i \in S\setminus W'(\mathbf{v},\mathbf{v}')} p_i(\mathbf{v},\mathbf{v}') \geq \sum_{i \in W'(\mathbf{v},\mathbf{v}')\setminus S} \varphi_i(v_i),$$

where $\mathbf{p}$ denotes the payment rule of the VCG mechanism in the duplicated environment. Condition (C2) follows from the non-negativity of payments by bidders of $W(\mathbf{v},\mathbf{v}')\setminus S$. □

We suspect that the upper bound in Theorem 4.4 can be improved. Our best lower bound is 4/3.

Example 4.6 The example is similar to Example 3.11. Consider a single-item setting with two bidders. Bidder one's valuation is deterministically 1 while bidder two's valuation comes from a nearly equal-revenue distribution with $F_2(z) = 1 - \frac{1}{z+1}$. (Notice that the virtual valuation of bidder two is $\varphi_2(v_2) = -1$ which is monotone.)

One possible way to auction the item to these two bidders is to offer bidder two the item at price H (which is accepted with probability $\frac{1}{H+1}$) and if declined, offer the item to bidder one at price 1. The expected revenue of this auction is $\frac{H}{H+1} + 1 - \frac{1}{H+1} \approx 2$. The revenue of the optimal auction is at least this.

The Vickrey auction with duplicate bidders has revenue z if the minimum value of bidder two and its duplicate is $z \geq 1$, and otherwise it has revenue 1. Thus, the cumulative distribution function for the revenue is given by

$$\begin{cases} 1 - \frac{1}{(1+z)^2} & \text{if } z \geq 1 \\ 0 & \text{otherwise.} \end{cases}$$

The expected revenue is the integral of one minus the cumulative distribution function on $[0, \infty)$:

$$1 + \int_1^\infty \frac{1}{(1+z)^2} dz = \frac{3}{2}.$$

We conclude that 4/3 is a lower bound on the worst-case revenue ratio of Vickrey with duplicates and the optimal auction, even in single-item settings.

[2]Strictly speaking, we defined this notion only for two mechanisms for a common environment. Definition 3.8 and Lemma 3.9 can be extended easily to the case where one environment contains the other.

5. VICKREY WITH ANONYMOUS RESERVE

We now return to single-item auctions and derive from Theorem 4.4 a constant-factor guarantee for the Vickrey auction with an anonymous reserve price when bidders' valuations are drawn from non-identical regular distributions. (Recall that Myerson [8] implies that Vickrey with an anonymous reserve is revenue-maximizing only when valuations are *i.i.d.* draws from a regular distribution.)

Theorem 5.1 *For every single-item setting with valuations drawn independently from distributions that satisfy the regularity condition, there is an anonymous reserve price such that the expected revenue of VCG with this reserve is a 4-approximation to the expected revenue of the optimal auction.*

PROOF. By Theorem 4.4, the expected revenue of the Vickrey auction (with no reserve) in the duplication of a single-item environment is at least half that of the optimal auction in the original environment. In this duplicated environment, symmetry dictates that the winner of the Vickrey auction is equally likely to be an original bidder or a duplicate. The expected revenue from the original bidders is thus at least a quarter of that of the optimal auction in the original environment.

We can simulate the allocation for and revenue obtained from original bidders in the duplicated environment by running the Vickrey auction, with the original set of bidders, with a random reserve price that is distributed according the maximum valuation. This random reserve takes the place of the maximum valuation of a duplicate bidder in the duplicated environment. The expected revenue of the best anonymous reserve price is at least that obtained by this random (anonymous) reserve price. □

The following example, which is similar to Examples 3.11 and 4.6, shows a lower bound of two that we suspect is tight.

Example 5.2 Consider a single-item auction with two bidders. Bidder one's valuation is deterministically 1. Bidder two's valuation is drawn from the equal-revenue distribution (i.e., $F_2(z) = 1 - 1/z$). Every anonymous reserve price yields expected revenue at most 1. As in Example 3.11, Of course, setting a reserve of H for bidder two and selling to bidder one if bidder two refuses yields expected revenue 1 from bidder two and $1 - 1/H$ from bidder one. As H tends to infinite, this expected revenue approaches 2. Myerson's revenue is at least the revenue of this auction.

The two restrictions in Theorem 5.1 are to single-item settings and to regular distributions, and both are necessary. The necessity of regularity was established by Chawla et al. [2].

Proposition 5.3 ([2]) *There are non-identical irregular distributions such that the expected revenue of Vickrey with the best anonymous reserve is no better than a logarithmic approximation to that of the optimal auction.*

The proof idea is to choose the valuation of bidder i according to the distribution

$$v_i = \begin{cases} 1/i & \text{with probability } 1/n \\ 0 & \text{otherwise.} \end{cases}$$

Chawla et al. [2] also showed that for irregular *identical* distributions and a single-item setting, the expected revenue of Vickrey with a suitable anonymous reserve price is a 4-approximation to the optimal auction.

We conclude the section with an example showing that the restriction to single-item settings is also necessary in Theorem 5.1: in general, the VCG mechanism with the best anonymous reserve price is no better than a logarithmic approximation in the cardinality of the largest set of winners.

Example 5.4 Let $S = \{1, \ldots, k\}$ be the largest set of bidders that can be served simultaneously in some downward-closed environment and consider the following (deterministic) valuations. All bidders not in S have zero valuation. Bidder $i \in S$ has value $1/i$. Bidder-specific reserve prices extract revenue $\sum_{i=1}^{k} 1/i \approx \log k$, whereas every anonymous reserve yields revenue at most 1.

6. CONCLUSIONS

We have used approximation to extend two recommendations from the theory of single-item auctions with i.i.d. bidder valuations to the much more general setting of downward-closed single-parameter agent environments with non-identical distributions. Our first result shows that the VCG mechanism, in conjunction with monopoly reserve prices, gives a 2-approximation of the optimal expected revenue. Our second result is that recruiting a copy of each agent in the environment suffices to lift the expected revenue of the VCG mechanism (with no reserves) to one third of that of the optimal mechanism in the original environment. A version of this second result also implies a constant-factor approximation guarantee for Vickrey with anonymous (and even random) reserve prices for single-item auctions and regular valuation distributions. We conclude that, in many contexts, the VCG mechanism with simple reserve prices is near-optimal in a very practical sense. It would not be difficult for a designer with modest prior knowledge of the valuation distributions to implement one of our proposed approximations.

Looking toward future work, there remain small gaps between our upper and lower bounds in some cases. For Bulow-Klemperer-type results in matroid environments with regular valuation distributions, is the tight approximation guarantee closer to our upper bound of 2 (Theorem 4.4) or our lower bound of 4/3 (Example 4.6)? Relatedly, for the Vickrey auction with an anonymous reserve price in a single-item setting with (non-identical) regular valuation distributions, is the correct answer closer to 4 (Theorem 5.1) or 2 (Example 5.2)? We strongly suspect that both of these upper bounds can be improved. We also suspect that our revenue guarantees for matroid environments (Theorems 3.7 and 4.4) can be improved if the valuation distributions are further restricted to satisfy the monotone hazard rate condition.

More broadly and importantly, our results illuminate a road map for the rigorous study of prior-free revenue-maximization in asymmetric single-parameter environments. (Most prior-free revenue-maximizing mechanism design results to date hold only in unlimited supply problems.) Specifically, Theorems 3.2 and 3.7 motivate the following prior-free benchmark $\mathcal{G}$, defined for every bid vector $\mathbf{v}$ in a given environment: the maximum revenue earned by the VCG mechanism when supplemented with a common reserve price for all agents. More precisely, we define

$$\mathcal{G}(\mathbf{v}) = \max_{r \le v_{(2)}} \mathrm{VCG}_{\mathbf{r}}(\mathbf{v}), \tag{8}$$

where the maximum is over all anonymous reserve prices no larger than the second-highest valuation $v_{(2)}$. The upper bound on the reserve price is needed for standard technical reasons, explained in [5].

Our results in Section 3 demonstrate that the economic motivation for comparing to the benchmark in (8) is strong: a prior-free approximation of $\mathcal{G}$ is guaranteed to *simultaneously approximate* the performance of a Bayesian monopolist for *every* i.i.d. valuation distribution that satisfies the regularity condition (in matroid environments) or the MHR condition (in general downward-closed environments). See Hartline and Roughgarden [6] for detailed discussion of this point.

Designing prior-free approximations of this benchmark $\mathcal{G}$ appears technically challenging. The reason is that $\mathcal{G}$ might be large due to VCG payments (rather than the reserve price), and these arise from competition between bidders. Random sampling approaches (as in [5]) tend to destroy competition, leading to the inaccurate estimation of the optimal reserve price.

For matroid environments, we can prove that randomizing between the VCG mechanism (with no reserve) and a random sampling-type auction achieves a constant-factor approximation of the prior-free benchmark $\mathcal{G}$.

Theorem 6.1 *In every matroid domain, there is a mechanism that 8-approximates the benchmark $\mathcal{G}(\mathbf{v})$ for every input $\mathbf{v}$.*

The mechanism used in Theorem 6.1, as well as simple variants of it, provably fail to achieve a constant-factor approximation of $\mathcal{G}$ in more general environments. We conclude with an open question that should serve as the next major challenge in prior-free revenue-maximizing mechanism design.

Open Question. Design prior-free mechanisms that approximate the benchmark $\mathcal{G}$ for all downward-closed single-parameter environments.

7. REFERENCES

[1] J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–194, 1996.

[2] S. Chawla, J. D. Hartline, and R. D. Kleinberg. Algorithmic pricing via virtual valuations. In *Proc. 9th ACM Conf. on Electronic Commerce (EC)*, pages 243–251, 2007.

[3] R. Day and P. Milgrom. Core-selecting auctions. *International Journal of Game Theory*, 36(3-4):393–407, 2008.

[4] S. Dughmi, T. Roughgarden, and M. Sundararajan. Revenue submodularity. In *these proceedings*.

[5] A. V. Goldberg, J. D. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.

[6] J. D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *Proc. 39th ACM Symp. on Theory of Computing (STOC)*, pages 75–84, 2008.

[7] A. Karlin, D. Kempe, and T. Tamir. Beyond VCG: Frugality in truthful mechanisms. In *Proc. 46th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 615–626, 2005.

[8] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[9] Z. Neeman. The effectiveness of English auctions. *Games and Economic Behavior*, 43(2):214–238, 2003.

[10] A. Ronen. On approximating optimal auctions. In *Proc. 3rd ACM Conf. on Electronic Commerce (EC)*, pages 11–17, 2001.

[11] A. Schrijver. *Combinatorial Optimziation: Polyhedra and Efficiency*. Springer, 2003.

[12] K. Talwar. The price of truth: Frugality in truthful mechanisms. In *Proc. 20th Annual Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 608–619, 2003.

APPENDIX

A. PROOF OF LEMMA 4.1

PROOF OF LEMMA 4.1. First suppose that $t = 0$. Write $v_{\max}$ and $v_{\min}$ for $\max\{v_1, v_2\}$ and $\min\{v_1, v_2\}$, respectively, and μ for the expected value of the given MHR distribution F. As a thought experiment, consider running a standard Vickrey (second-price) auction with bids v_1 and v_2. The expected revenue of this auction is obviously $\mathbf{E}[v_{\min}]$. By Myerson's Lemma (Proposition 2.1), its expected revenue is also the expected virtual value of the allocation, namely $\mathbf{E}[\max\{\varphi(v_1), \varphi(v_2)\}]$. Thus

$$\mathbf{E}_{v_1,v_2}[\max\{\varphi(v_1), \varphi(v_2)\}] = \mathbf{E}_{v_1,v_2}[v_{\min}]\,. \tag{9}$$

Now recall the characterization of μ in terms of the hazard rate $h(x)$ of F:

$$\mu = \int_0^\infty (1 - F(x))dx = \int_0^\infty e^{-H(x)}dx,$$

where $H(x)$ denotes $\int_0^x h(t)dt$. Since h is non-negative and nondecreasing, H is nondecreasing and convex. Similarly,

$$\begin{aligned}\mathbf{E}_{v_1,v_2}[v_{\min}] &= \int_0^\infty e^{-2H(x)}dx \\ &\ge \int_0^\infty e^{-H(2x)}dx = \frac{1}{2}\int_0^\infty e^{-H(x)}dx = \frac{\mu}{2},\end{aligned} \tag{10}$$

where the inequality follows from the fact that H is nondecreasing and convex. Also,

$$\begin{aligned}\mathbf{E}_{v_1,v_2}[v_{\max} \mid v_{\min}] &= v_{\min} + \left(\int_{v_{\min}}^\infty e^{-H(x)}dx\right) \cdot e^{H(v_{\min})} \\ &\le v_{\min} + \int_{v_{\min}}^\infty e^{-H(x - v_{\min})}dx \\ &= v_{\min} + \mu,\end{aligned} \tag{11}$$

where the inequality again holds because H is nondecreasing and convex. Taking expectations in (11) and then using (10) yields

$$\mathbf{E}[v_{\max}] \le \mathbf{E}[v_{\min}] + \mu \le 3 \cdot \mathbf{E}[v_{\min}]\,;$$

combining this with (9) gives the lemma, in the special case where $t = 0$.

Finally, suppose that $t > 0$. By the definition of a virtual valuation, we can rephrase our progress so far as

$$\mathbf{E}_{v_1,v_2}[v_{\max}] - \mathbf{E}_{v_1,v_2}\left[\frac{1}{h(v_{\max})}\right] = \mathbf{E}_{v_1,v_2}\left[v_{\max} - \frac{1}{h(v_{\max})}\right] \geq \frac{1}{3} \cdot \mathbf{E}_{v_1,v_2}[v_{\max}]\,. \quad (12)$$

The distribution of $v_{\max}$, conditioned on the event that $v_{\max} \geq t$, stochastically dominates the unconditional distribution of $v_{\max}$; in other words, $\Pr[v_{\max} \geq s] \leq \Pr[v_{\max} \geq s \,|\, v_{\max} \geq t]$ for every $s \geq 0$. Similarly, using the MHR assumption, the corresponding conditional distribution of $-1/h(v)$ stochastically dominates its unconditional distribution. It follows that the inequality (12), and hence the lemma, continues to hold after conditioning on the event $v_{\max} \geq t$. $\square$