Revenue Maximization with a Single Sample☆

Peerapong Dhangwatnotai[1], Tim Roughgarden[2], Qiqi Yan[3]

Stanford University

Abstract

This paper pursues auctions that are *prior-independent*. The goal is to design an auction such that, whatever the underlying valuation distribution, its expected revenue is almost as large as that of an optimal auction tailored for that distribution. We propose the prior-independent Single Sample mechanism, which is essentially the Vickrey-Clarke-Groves (VCG) mechanism, supplemented with reserve prices chosen at random from participants' bids. We prove that under reasonably general assumptions, this mechanism simultaneously approximates all Bayesian-optimal mechanisms for all valuation distributions. Conceptually, our analysis shows that even a single sample from a distribution — some bidder's valuation — is sufficient information to obtain near-optimal expected revenue.

Keywords: Auctions, approximation, revenue-maximization, prior-independence

☆An extended abstract of this paper appeared in the Proceedings of the 11th ACM Conference on Electronic Commerce, June 2010.

Email addresses: `pdh@cs.stanford.edu` (Peerapong Dhangwotnotai), `tim@cs.stanford.edu` (Tim Roughgarden), `qiqiyan@cs.stanford.edu` (Qiqi Yan)

[1]Supported in part by the ONR Young Investigator Award of the second author.

[2]Supported in part by NSF grants CCF-0448664 and CCF-1016885, an ONR Young Investigator Award, an ONR PECASE Award, an AFOSR MURI grant, and an Alfred P. Sloan Fellowship.

[3]Supported by a Stanford Graduate Fellowship. Current address: Google Research, 1600 Amphitheatre Pkway, Mountain View, CA, 94043, `qiqiyan@google.com`.

1. Introduction

The optimal reserve price for a single-item auction is a function of the distribution of the bidders' valuations. In more complex settings, such as with multiple goods, the optimal selling procedure depends on the underlying valuation distributions in still more intricate ways. What if good prior information is expensive or impossible to acquire? What if a single procedure is to be re-used several times, in settings with different or not-yet-known bidder valuations? Can we avoid auction designs that depend on the details of the assumed distribution, in the spirit of "Wilson's Doctrine" (Wilson, 1987)? Are there more robust mechanisms, that are guaranteed to be near-optimal across a range of environments?

This paper pursues auctions that are *prior-independent.*

> *The goal is to design an auction such that, whatever the underlying valuation distribution, its expected revenue is almost as large as that of an optimal auction tailored for that distribution.*

For example, consider a single-item auction with n bidders with valuations drawn i.i.d. from a distribution F. The Vickrey or second-price auction is prior-independent, because its description is independent of F. For well-behaved distributions, the revenue-maximizing auction is the Vickrey auction, supplemented with a reserve price (Myerson, 1981). This reserve price depends on F, and optimal single-item auctions are not prior-independent.

Can there be non-trivial revenue guarantees for prior-independent auctions? After all, this is tantamount to a *single* auction being simultaneously near-optimal for *every* valuation distribution F.

1.1. Our Results

We propose the prior-independent *Single Sample* mechanism. This mechanism is essentially the Vickrey-Clarke-Groves (VCG) mechanism, supplemented with reserve prices chosen at random from participants' bids. We prove that under reasonably general assumptions, this mechanism simultaneously approximates all Bayesian-optimal mechanisms for all valuation distributions. Conceptually, our analysis shows that even a single sample from a distribution — some bidder's valuation — is sufficient information to obtain near-optimal expected revenue.

In more detail, we consider n single-parameter bidders. Each bidder has an independent private valuation for "winning", drawn from a distribution that satisfies a standard technical condition.[4] We allow bidder asymmetry through observable attributes. We assume that the valuations of bidders with a common

[4]Without any restriction on the tails of the valuation distributions, no prior-independent auction has a non-trivial revenue guarantee. To see why, consider a single-item auction with n bidders and valuations drawn i.i.d. from the following distribution F_p, for a parameter p: a bidder has valuation p with probability $1/n^2$, and valuation 0 otherwise. For every p, the optimal auction for F_p has expected revenue proportional to p/n. A prior-independent auction

attribute are drawn i.i.d. from a distribution that is unknown to the seller.[5] Bidders with different attributes can have valuations drawn (independently) from completely different distributions. For example, based on (publicly observable) eBay bidding history, one might classify bidders into "bargain-hunters", "typical", and "aggressive." The assumption is that bidders in the same class are likely to have similar valuations, but nothing is assumed about how their valuations for a given item are distributed. We assume that the environment is *non-singular*, meaning that there is no bidder with a unique attribute.[6]

Feasible allocations are described by a collection of bidder subsets, each representing a set of bidders that can simultaneously win in the auction. For example, in a single-item auction, the subsets are the singletons and the empty set. In combinatorial auctions with single-minded bidders, feasible subsets correspond to bidders seeking mutually disjoint bundles.[7] We consider only "downward-closed" environments, where every subset of a feasible set is again feasible.

Our first main result is that, for every non-singular downward-closed environment in which every valuation distribution has a monotone hazard rate (as defined in Section 2.3), the expected revenue of the prior-independent Single Sample mechanism is at least a constant fraction of the expected optimal welfare (and hence revenue) in that environment. The approximation factor is $\frac{1}{4} \cdot \frac{\kappa-1}{\kappa}$ when there are at least $\kappa \geq 2$ bidders of every present attribute, and our analysis of our mechanism is tight (for a worst-case distribution) for each κ. This factor is $\frac{1}{8}$ when $\kappa = 2$ and quickly approaches $\frac{1}{4}$ as κ grows. This gives, as an example special case, the first revenue guarantee for combinatorial auctions with single-minded bidders outside of the standard Bayesian setup with known distributions (Ledyard, 2007; Hartline and Roughgarden, 2009).

For our second main result, we weaken our assumptions about the valuation distributions but add additional restrictions to the structure of the feasible sets. Precisely, we consider *matroid* environments, where bidders satisfy a substitutes condition (Section 2.1). Examples of such environments include k-unit auctions and certain matching markets. Here, we again prove an approximation factor of $\frac{1}{4} \cdot \frac{\kappa-1}{\kappa}$, assuming only that every valuation distribution is regular — a condition that is weaker than the monotone hazard rate condition and permits distributions with heavier tails. When all bidders have a common attribute and thus have i.i.d. valuations, we modify the mechanism and analysis to achieve an approximation factor of $\frac{1}{2}$ for every $\kappa \geq 2$.

---

essentially has to "guess" the value of p — since bids are almost always zero, they almost never provide any information about F_p — and cannot have expected revenue within a constant factor of p/n for every F_p.

[5]We only consider dominant-strategy incentive-compatible mechanisms. Thus, the valuation distributions can also be unknown to the mechanism participants.

[6]No prior-independent auction has a non-trivial approximation guarantee when there is a bidder with a unique attribute. The argument is similar to the one above for arbitrary valuation distributions; see also Goldberg et al. (2006).

[7]In such an auction, there are n bidders and m goods with unit supply. Each bidder i wants a publicly known subset S_i of goods — for example, a set of geographically clustered wireless spectrum licenses — and has a private valuation v_i for it.

Third, we extend the Single Sample mechanism to make use of multiple samples and provide better approximation guarantees when κ is large. Specifically, provided κ is sufficiently large — at least a lower bound that is polynomial in ϵ^{-1} and *independent* of the underlying valuation distributions — we show how to improve the above approximation factors of $\frac{1}{4}\frac{\kappa-1}{\kappa}$, $\frac{1}{4}\frac{\kappa-1}{\kappa}$, and $\frac{1}{2}$ to $\frac{1}{e}(1-\epsilon)$, $\frac{1}{2}(1-\epsilon)$, and $(1-\epsilon)$, respectively. (Here e denotes $2.718\ldots$.)

1.2. Motivation: The Bulow-Klemperer Theorem

To develop intuition for our techniques, and more generally the possibility of good prior-independent auctions, we review a well-known result of Bulow and Klemperer (1996). This result concerns single-item auctions and states that, for every $n \geq 1$ and valuation distribution F that is regular in the sense of Section 2.3, the expected revenue of the Vickrey auction with $n+1$ bidders with valuations drawn i.i.d. from F is at least that of a revenue-maximizing auction with n such bidders.

First, we observe that the Bulow-Klemperer theorem is an interesting revenue guarantee for a prior-independent auction: with one extra bidder, the prior-independent Vickrey auction is as good as the revenue-maximizing auction tailored to the underlying distribution.

Next, we give a novel interpretation of the Bulow-Klemperer theorem when $n = 1$. Fix a valuation distribution F. The optimal auction for one bidder simply posts a monopoly price — a price p that maximizes $p \cdot (1-F(p))$. In the Vickrey auction, each of the two bidders contributes the same expected revenue. Each bidder effectively faces a reserve price equal to the other bidder's valuation — a random reserve price drawn from F. Thus, the Bulow-Klemperer theorem with $n = 1$ is equivalent to the following statement: *for a bidder with a valuation drawn from a regular distribution F, the expected revenue from a random posted price drawn from F is at least half that from an optimal posted price.*[8] At least in single-item auctions, a random reserve price is an effective surrogate for an optimal one.

1.3. The Main Ideas

Our general results are proved in two parts. The interface between the two is the VCG mechanism with "lazy" monopoly reserves (VCG-L). This mechanism is prior-dependent, in that the valuation distribution F_i of bidder i is known. The VCG-L mechanism first runs the VCG mechanism to obtain a tentative set of winning bidders, and then removes every bidder i with valuation below the monopoly price for F_i.

The first part of our proof approach establishes conditions under which the VCG-L mechanism with monopoly reserves has near-optimal expected revenue. We do this using different arguments for each of the first two main results. We also show that there is no common generalization of these two results, in

[8] See also Lemma 3.6 for a direct, geometric proof of this statement.

that the VCG-L mechanism with monopoly reserves does not have near-optimal expected revenue in every downward-closed environment with regular valuation distributions.

The second part of our proof approach shows that the expected revenue of the Single Sample mechanism is close to that of the VCG-L mechanism with monopoly reserves. Since the Single Sample mechanism uses random reserves and the VCG-L mechanism uses monopoly reserves, this is essentially a generalization of the Bulow-Klemperer argument in Section 1.2.

Our third main result, which modifies the Single Sample mechanism to give better bounds as the number of bidders of every attribute tends to infinity, improves the analysis in the first part of the proof approach above. A weak version of this result, which does not give quantitative bounds on the number of bidders required, can be derived from the Law of Large Numbers. To prove our distribution-independent bound on the number of bidders required, we show that there exists a set of "quantiles" that is simultaneously small enough that concentration bounds can be usefully applied, and rich enough to guarantee a good approximation for every regular valuation distribution. Our arguments rely on a geometric characterization of regular distributions.

1.4. Related Work

Most of the vast literature on revenue-maximizing auctions studies designs tailored to a known distribution over bidders' private information (see, e.g., Krishna (2002)). Here, we mention only the works related to approximation guarantees for prior-independent auctions. Neeman (2003) considers single-item auctions with i.i.d. bidders, and quantifies the fraction of the optimal welfare extracted as revenue by the (prior-independent) Vickrey auction, as a function of the number of bidders. Segal (2003) and Baliga and Vohra (2003) prove asymptotic optimality results for certain prior-independent mechanisms when bidders are symmetric, goods are identical, and the number of bidders is large. As discussed in Section 1.2, the main result in Bulow and Klemperer (1996) is a revenue guarantee for a prior-independent auction. For more general results in the same spirit — that welfare-maximization with additional bidders yields expected revenue (almost) as good as in an optimal mechanism — see Dughmi et al. (2012); Hartline and Roughgarden (2009). Since the conference version of the present work Dhangwatnotai et al. (2010), there have been a number of follow-up papers in prior-independent auctions and related topics; see Azar et al. (2013); Chawla et al. (2013); Devanur et al. (2011); Fu et al. (2013); Roughgarden and Talgam-Cohen (2013); Roughgarden et al. (2012).

We only consider dominant-strategy incentive-compatible mechanisms and make no assumptions about participants' knowledge of the valuation distributions. An alternative approach is to assume that these distributions are known to the participants but unknown to the seller, and to design Bayesian incentive-compatible mechanisms. Caillaud and Robert (2005), for example, give a simple ascending-price contest for a single-item that has a revenue-optimal Bayes-Nash equilibrium. Such mechanisms demand much more from the strategic participants than dominant-strategy implementations; see also the discussion in Bergemann and Morris (2005).

Valuation distributions are used in the analysis, but not in the design, of prior-independent auctions. In *prior-free* auction design, distributions are not even used to evaluate the performance of an auction — the goal is to design an auction with good revenue for *every* valuation profile, rather than in expectation. A key challenge in prior-free auction design, first identified by Goldberg et al. (2006), is to develop a useful competitive analysis framework. Goldberg et al. (2006) proposed a *revenue benchmark* approach, which has been applied successfully to a number of auction settings. The idea is to define a real-valued function on valuation profiles that represents an upper bound on the maximum revenue achievable by any "reasonable" auction. The most well-studied such benchmark is $\mathcal{F}^{(2)}$ for digital goods auctions — auctions with unlimited supply and unit-demand bidders — defined for each valuation profile as the maximum revenue achievable using a common selling price while selling to at least two bidders (Goldberg et al., 2006).

Approximation in this revenue benchmark framework is strictly stronger than the simultaneous approximation goal pursued in the present paper. This connection is made explicit in Hartline and Roughgarden (2008) and is pursued further by Devanur and Hartline (2009); Hartline and Roughgarden (2009); Hartline and Yan (2011); Leonardi and Roughgarden (2012). Advantages of our prior-independent guarantees over the known prior-free results include the ability to handle asymmetric (non-i.i.d.) bidders and more general environments; better approximation factors; and simpler mechanisms.

2. Preliminaries

This section reviews standard terminology and facts about Bayesian-optimal mechanism design. We encourage the reader familiar with these to skip to Section 3.

2.1. Environments

An environment is defined by a set E of bidders, and a collection $\mathcal{I} \subseteq 2^E$ of feasible sets of bidders, which are the subsets of bidders that can simultaneously "win". For example, in a k-unit auction with unit-demand bidders, $\mathcal{I}$ is all subsets of E that have size at most k. We assume that the set system $(E, \mathcal{I})$ is *downward-closed*, meaning that if $T \in \mathcal{I}$ and $S \subseteq T$, then $S \in \mathcal{I}$. Each bidder has a publicly observable *attribute* that belongs to a known set A. We assume that each bidder with attribute a has a private *valuation* for winning that is an independent draw from a distribution F_a. We sometimes denote an environment by a tuple $Env = (E, \mathcal{I}, A, (a_i)_{i \in E}, (F_a)_{a \in A})$. Every subset $T \subseteq E$ of bidders induces a subenvironment in a natural way, with feasible sets $\{S \cap T\}_{S \in \mathcal{I}}$.

Some of our results concern the special case of a *matroid environment*, in which bidders satisfy a substitutes condition. Precisely, the set system $(E, \mathcal{I})$ is a *matroid* if $\mathcal{I}$ is non-empty and downward-closed, and if whenever $S, T \in \mathcal{I}$ with

$|T| < |S|$, there is some $i \in S \setminus T$ such that $T \cup \{i\} \in \mathcal{I}$. This last condition is called the "exchange property" of matroids. See, e.g., Oxley (1992). Examples of matroid environments include digital goods (where $\mathcal{I} = 2^E$), k-unit auctions (where $\mathcal{I}$ is all subsets of size at most k), and certain unit-demand matching markets (corresponding to transversal matroids). Combinatorial auctions with single-minded bidders, where feasible sets correspond to sets of bidders desiring mutually disjoint bundles, induce downward-closed environments that are not generally matroids.

An environment is *non-singular* if there is no bidder with a unique attribute, and is *i.i.d.* if every bidder has the same attribute. An environment is *regular* or *m.h.r.* if every valuation distribution is a regular distribution or an m.h.r. distribution (as defined in Section 2.3), respectively.

2.2. Truthful Mechanisms

Name the bidders $E = \{1, 2, \ldots, n\}$. A (deterministic) mechanism $\mathcal{M}$ comprises an *allocation rule* $\mathbf{x}$ that maps every bid vector $\mathbf{b}$ to a characteristic vector of a feasible set (in $\{0,1\}^n$), and a *payment rule* $\mathbf{p}$ that maps every bid vector $\mathbf{b}$ to a non-negative payment vector in $[0, \infty)^n$. We insist on individual rationality in the sense that $p_i(\mathbf{b}) \leq b_i \cdot x_i(\mathbf{b})$ for every i and $\mathbf{b}$. We assume that every bidder i aims to maximize its quasi-linear utility $u_i(\mathbf{b}) = v_i \cdot x_i(\mathbf{b}) - p_i(\mathbf{b})$, where v_i is its private valuation for winning. We call a mechanism $\mathcal{M}$ *truthful* if for every bidder i and fixed bids $\mathbf{b}_{-i}$ of the other bidders, bidder i maximizes its utility by setting its bid b_i to its private valuation v_i. Since we only consider truthful mechanisms, in the rest of the paper we use valuations and bids interchangeably.

The *efficiency* or *welfare* of the outcome of a mechanism is the sum of the winners' valuations, and the *revenue* is the sum of the winners' payments. By individual rationality, the revenue of a mechanism outcome is bounded above by its welfare.

A well-known characterization of truthful mechanisms in single-parameter settings (Myerson, 1981; Archer and Tardos, 2001) states that a mechanism $(\mathbf{x}, \mathbf{p})$ is truthful if and only if the allocation rule is monotone — that is, $x_i(b_i', \mathbf{b}_{-i}) \geq x_i(\mathbf{b})$ for every i, $\mathbf{b}$, and $b_i' \geq b_i$ — and the payment rule is given by a certain formula involving the allocation rule. We often specify a truthful mechanism by its monotone allocation rule, with the understanding that it is supplemented with the unique payment rule that yields a truthful mechanism. For deterministic mechanisms like those studied in this paper, the payment of a winning bidder is simply the smallest bid for which it would remain a winner.

For example, the *VCG mechanism*, which chooses the feasible set $S \in \mathcal{I}$ that maximizes the welfare $\sum_{i \in S} v_i$, has a monotone allocation rule and can be made truthful using suitable payments.

Two variants of the VCG mechanism are also important in this paper. Let r_i be a *reserve price* for bidder i. The *VCG mechanism with eager reserves* $\mathbf{r}$ *(VCG-E)* works as follows, given bids $\mathbf{v}$: (1) remove all bidders i with $v_i < r_i$; (2) run the VCG mechanism on the remaining bidders to determine the winners; (3)

charge each winning bidder i the larger of r_i and its VCG payment in step (2). In the *VCG mechanism with lazy reserves* $\mathbf{r}$ *(VCG-L)*, steps (1) and (2) are reversed.

Both the VCG-E and VCG-L mechanisms are feasible and truthful in every downward-closed environment. The two mechanisms are equivalent in sufficiently simple environments — see Corollary 3.4 — but are different in general. For instance, suppose $E = \{1, 2, 3\}$, the maximal feasible sets are $\{1, 2\}$ and $\{3\}$, $v_1 = v_2 = 2$, $v_3 = 3$, $r_1 = 3$, and $r_2 = r_3 = 0$. Then, bidder 3 is the winner in the VCG-E mechanism, while bidder 2 is the winner in the VCG-L mechanism. In general, the welfare of the VCG-E mechanism is at least that of the VCG-L mechanism. Each mechanism has larger revenue than the other in some cases.

2.3. Bayesian-Optimal Auctions

Let F be (the cumulative distribution function of) a valuation distribution. For simplicity, we assume that the distribution is supported on a closed interval $[l, h]$, and has a positive and smooth density function on this interval. When convenient, we assume that $l = 0$; a simple "shifting argument" shows that this is the worst type of distribution for approximate revenue guarantees. The *virtual valuation function* of F is defined as $\varphi_F(v) = v - \frac{1}{h(v)}$, where $h(v) = \frac{f(v)}{1-F(v)}$ is the *hazard rate function* of F. This paper works with two different common assumptions on valuation distributions. A *regular* distribution has, by definition, a nondecreasing virtual valuation function. A *monotone hazard rate (m.h.r.)* distribution has a nondecreasing hazard rate function. Several important distributions (exponential, uniform, etc.) are m.h.r.; intuitively, these are distributions with tails no heavier than the exponential distribution. Regular distributions include all m.h.r. distributions along with some additional distributions with heavier tails, such as some power-law distributions.

Myerson (1981) characterized the expected revenue-maximizing mechanisms for single-parameter environments using the following key lemma.

Lemma 2.1 (Myerson's Lemma) *For every truthful mechanism* $(\mathbf{x}, \mathbf{p})$*, the expected payment of a bidder* i *with valuation distribution* F_i *satisfies*

$$E_{\mathbf{v}}[p_i(\mathbf{v})] = E_{\mathbf{v}}[\varphi_{F_i}(v_i) \cdot x_i(\mathbf{v})].$$

Moreover, this identity holds even after conditioning on the bids $\mathbf{v}_{-i}$ *of the bidders other than* i*.*

In words, the (conditional) expected payment of a bidder is precisely its (conditional) expected contribution to the virtual welfare. It follows that if the distributions are regular, then a revenue-maximizing truthful mechanism chooses a feasible set S that maximizes the virtual welfare $\sum_{i \in S} \varphi_{F_i}(v_i)$. The role of regularity is to ensure that this allocation rule is indeed monotone; otherwise, additional ideas are needed (Myerson, 1981).

3. Revenue Guarantees with a Single Sample

In this section, we design a prior-independent auction that simultaneously approximates the optimal expected revenue to within a constant factor in every non-singular m.h.r. single-parameter environment, and in every non-singular regular matroid environment. Section 3.1 defines our mechanism. Section 3.2 introduces some of our main analysis techniques in the simpler setting of i.i.d. matroid environments — here, we also obtain better approximation bounds. Section 3.3 gives an overview of our general proof approach. Sections 3.4 and 3.5 prove our approximation guarantees for m.h.r. downward-closed and regular matroid environments, respectively. Section 3.6 shows that there is no common generalization of these two results, in that the Single Sample mechanism does not have a constant-factor approximation guarantee in regular downward-closed environments. Section 3.7 discusses computationally efficient variants of our mechanism.

3.1. The Single Sample Mechanism

We propose and analyze the *Single Sample* mechanism: we randomly pick one bidder of each attribute to set a reserve price for the other bidders with that attribute, and then run the VCG-L mechanism (Section 2.2) on the remaining bidders.

Definition 3.1 (Single Sample) Given a non-singular downward-closed environment $Env = (E, \mathcal{I}, A, (a_i)_{i \in E}, (F_a)_{a \in A})$, the *Single Sample* mechanism is the following:

(1) For each represented attribute a, pick a *reserve bidder* i_a with attribute a uniformly at random from all such bidders.

(2) Run the VCG mechanism on the sub-environment induced by the non-reserve bidders to obtain a preliminary winning set P.

(3) For each bidder $i \in P$ with attribute a, place i in the final winning set W if and only if $v_i \geq v_{i_a}$. Charge every winner $i \in W$ with attribute a the maximum of its VCG payment computed in step (2) and the reserve price v_{i_a}.

The Single Sample mechanism is clearly prior-independent — that is, it is defined independently of the F_a's — and it is easy to verify that it is truthful. Section 4 shows how to use multiple samples to obtain better approximation factors when there are more than two bidders with each represented attribute.

3.2. Warm-Up: I.I.D. Matroid Environments

To introduce some of our primary analysis techniques in a relatively simple setting, we first consider matroid environments (recall Section 2.1) in which all bidders have the same attribute (i.e., have i.i.d. valuations).

Theorem 3.2 (I.I.D. Matroid Environments) *For every i.i.d. regular matroid environment with at least $n \geq 2$ bidders, the expected revenue of the Single Sample mechanism is at least a $\frac{1}{2} \cdot \frac{n-1}{n}$ fraction of that of an optimal mechanism for the environment.*

The factor of $(n-1)/n$ can be removed with a minor tweak to the mechanism (Remark 3.7).

What's so special about i.i.d. regular matroid environments? One important point is that the VCG mechanism can be implemented in a matroid environment via the greedy algorithm: bidders are considered in nonincreasing order of valuations, and a bidder is added to the winner set if and only if doing so preserves feasibility, given the previous selections. Now, recall that a *monopoly reserve price* of a valuation distribution F is a price in $\text{argmax}_p[p \cdot (1 - F(p))]$. The following proposition follows immediately from Myerson's Lemma, the fact that the greedy algorithm maximizes welfare in matroid environments, and the fact that the virtual valuation function is order-preserving when valuations are drawn i.i.d. from a regular distribution. See, e.g., Dughmi et al. (2012) for additional details.

Proposition 3.3 *In every i.i.d. regular matroid environment, the VCG-E mechanism with monopoly reserves is a revenue-maximizing mechanism.*

The matroid assumption also allows us to pass from eager to lazy reserves.

Corollary 3.4 *In every i.i.d. regular matroid environment, the VCG-L mechanism with monopoly reserves is a revenue-maximizing mechanism.*

Proof: We use that the VCG mechanism admits a greedy implementation. With a common reserve price r, it makes no difference whether bidders with valuations below r are thrown out before or after running the greedy algorithm. Thus in matroid environments, the VCG-L and VCG-E mechanisms with an anonymous reserve price are equivalent. ■

Proving an approximate revenue-maximization guarantee for the Single Sample mechanism thus boils down to understanding the two ways in which it differs from the VCG-L mechanism with monopoly reserves — it throws away a random bidder, and it uses a random reserve rather than a monopoly reserve. The damage from the first difference is easy to control.

Lemma 3.5 *In expectation over the choice of the reserve bidder, the expected revenue of an optimal mechanism for the environment induced by the non-reserve bidders is at least an $\frac{n-1}{n}$ fraction of the expected revenue of an optimal mechanism for the original environment.*

Proof: Condition first on the valuations of all bidders and let S denote the winners under the optimal mechanism for the full environment. Since the reserve bidder is chosen independently of the valuations, each bidder of S is a

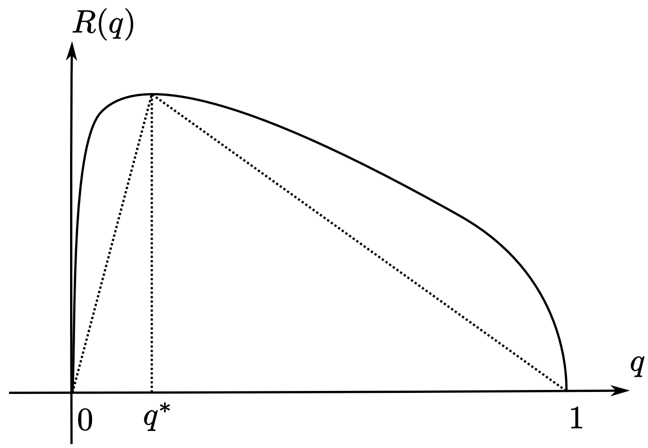

Figure 1: The revenue function in probability space of a regular distribution.

non-reserve bidder with probability $\frac{n-1}{n}$. By the linearity of expectation, the expected virtual welfare — over the choice of the reserve bidder, with all valuations fixed — of the (feasible) set of non-reserve bidders of S is $\frac{n-1}{n}$ times that of S, and the expected maximum-possible virtual welfare in the sub-environment is at least this. Taking expectations over bidder valuations and applying Myerson's Lemma (Lemma 2.1) completes the proof. ■

The crux of the proof of Theorem 3.2 is to show that a random reserve price serves as a sufficiently good approximation of a monopoly reserve price. The next key lemma formalizes this goal for the case of a single bidder. Its proof uses a geometric property of regular distributions. To explain it, for a distribution F, define the *revenue function* by $\widehat{R}(p) = p(1 - F(p))$, the expected revenue earned by posting a price of p on a good with a single bidder with valuation drawn from F. Define the *revenue function in probability space* R as $R(q) = q \cdot F^{-1}(1 - q)$ for all $q \in [0, 1]$, which is the same quantity parameterized by the probability q of a sale. An example of a revenue function in probability space is shown in Figure 1. One can check easily that the derivative $R'(q)$ equals the virtual valuation $\varphi_F(p)$, where $p = F^{-1}(1 - q)$. Regularity of F thus implies that $R'(q)$ is nonincreasing and hence R is concave. Also, assuming that the support of F is $[0, h]$ for some $h > 0$ — recall Section 2.3 — we have $R(0) = R(1) = 0$.

Lemma 3.6 *Let F be a regular distribution with monopoly price r^* and revenue function $\widehat{R}$. Let v denote a random valuation from F. For every nonnegative number $t \geq 0$,*

$$\mathbf{E}_v\left[\widehat{R}(\max\{t, v\})\right] \geq \frac{1}{2} \cdot \widehat{R}(\max\{t, r^*\}). \tag{1}$$

Inequality (1) makes precise and extends the novel interpretation of the Bulow-Klemperer theorem described in Section 1.2. When $t = 0$, the inequality states that a random reserve price, drawn from the valuation distribution, obtains expected revenue at least half that of the optimal (monopoly) reserve price. We require that this statement hold more generally in the presence of an exogenous additional reserve price t, which in our applications arises from competition from other bidders.

Proof: First suppose that $t = 0$, so that the claim is equivalent to the assertion that the expectation of $R(q)$ is at least half of $R(q^*)$, where $q = 1 - F(v)$ and $q^* = 1 - F(r^*)$ are the quantiles corresponding to v and r^*, respectively. Since q is uniformly distributed on $[0,1]$, $\mathbf{E}_q[R(q)]$ equals the area under the curve defined by R. By concavity, this area is at least the area of the triangle in Figure 1, which is $\frac{1}{2} \cdot 1 \cdot R(q^*) = \frac{1}{2} \cdot \widehat{R}(r^*)$.

If $0 < t < r^*$, then the right-hand side of (1) is unchanged. The left-hand side can only be higher: the only difference is to sometimes use a selling price t that, by concavity, is better than the previous selling price v. Finally, if $t > r^*$ and $q_t = 1 - F(t)$, then the right-hand side of (1) is $R(q_t)/2$; and the left-hand side is a convex combination of $R(q_t)$ (when $v \le t$) and the expected value of $R(q)$ when q is drawn uniformly from $[q_t, 1]$, which by concavity is at least $R(q_t)/2$. ■

We prove Theorem 3.2 by extending the approximation bound in Lemma 3.6 from a single bidder to all bidders and blending in Lemma 3.5.

Proof of Theorem 3.2: Condition on the choice of the reserve bidder j. Fix a non-reserve bidder i and condition on all valuations except those of i and j. Recall that j, as a reserve bidder, does not participate in the VCG computation in step (2) of the Single Sample mechanism. Thus, there is a "threshold" $t(\mathbf{v}_{-i})$ for bidder i such that i wins if and only if its valuation is at least $t(\mathbf{v}_{-i})$, in which case its payment is $t(\mathbf{v}_{-i})$.

With this conditioning, we can analyze bidder i as in a single-bidder auction, with an extra external reserve price of $t(\mathbf{v}_{-i})$. Let r^* and $\widehat{R}$ denote the monopoly price and revenue function for the underlying regular distribution F, respectively. The conditional expected revenue that i contributes to the revenue-maximizing mechanism in the sub-environment of non-reserve bidders is $\widehat{R}(\max\{t(\mathbf{v}_{-i}), r^*\})$. The conditional expected revenue that i contributes to the Single Sample mechanism is $\mathbf{E}_{v_j}\left[\widehat{R}(\max\{t(\mathbf{v}_{-i}), v_j\})\right]$. Since v_i, v_j are independent samples from the regular distribution F, Lemma 3.6 implies that the latter conditional expectation is at least 50% of the former. Taking expectations over the previously fixed valuations of bidders other than i and j, summing over the non-reserve bidders i and applying linearity of expectation, and finally taking the expectation over the choice of the reserve bidder j and applying Lemma 3.5 proves the theorem. □

Remark 3.7 (Optimized Version of Theorem 3.2) We can improve the approximation guarantee in Theorem 3.2 from $\frac{1}{2} \cdot \frac{n-1}{n}$ to $\frac{1}{2}$ by making the following minor modification to the Single Sample mechanism. Instead of discarding the reserve bidder j, we include it in the VCG computation in step (2) of the Single Sample mechanism. An arbitrary other bidder h is used to set a reserve price v_h for the reserve bidder j. Like the other bidders, the reserve bidder is included in the final winning set W if and only if it is chosen by the VCG mechanism in step (2) and also has a valuation above its reserve price ($v_j \geq v_h$). Its payment is then the maximum of its VCG payment and v_h.

We claim that, for every choice of a reserve bidder j, a non-reserve bidder i, and valuations $\mathbf{v}$, bidder i wins with bidder j included in the VCG computation in step (2) if and only if it wins with bidder j excluded from the computation. To prove this, recall that the VCG mechanism can be implemented via a greedy algorithm in matroid environments. If $v_i \leq v_j$, then i cannot win in either case (it fails to clear the reserve); and if $v_i > v_j$, then the greedy algorithm considers bidder i before j even if the latter is included in the VCG computation.

This claim implies that the expected revenue from non-reserve bidders in the original and modified versions of the Single Sample mechanism is the same. In the modified version, the obvious analog of Lemma 3.5 for a single bidder and Lemma 3.6 imply that the reserve bidder also contributes, in expectation, a $\frac{1}{2} \cdot \frac{1}{n}$ fraction of the expected revenue of an optimal mechanism. Combining the contributions of the reserve and non-reserve bidders yields an approximation guarantee of $\frac{1}{2}$ for the modified mechanism.

Remark 3.8 (Tightness of Analysis) The bound of $\frac{1}{2}$ in Remark 3.7, and hence also the bound in Lemma 3.6, is tight in the worst case. To see this, consider a digital goods auction with two bidders, and a regular valuation distribution F whose revenue function in probability space is essentially a triangle (cf., Figure 1). For example, the distributions $F(v) = 1 - \frac{1}{v+1}$ on $[0, H)$ and $F(H) = 1$ provide a matching lower bound as $H \to \infty$.[9]

3.3. Proof Framework

Relaxing the matroid or i.i.d. assumptions of Section 3.2 introduce new challenges in the analysis of the Single Sample mechanism. The expected revenue-maximizing mechanism becomes complicated — nothing as simple as the VCG mechanism with reserve prices. In addition, eager and lazy reserve prices are not equivalent.

Our general proof framework hinges on the VCG-L mechanism with monopoly reserves, which we use as a proxy for the optimal mechanism. The analysis proceeds in two steps:

1. Prove that the expected revenue of the VCG-L mechanism with monopoly reserves is close to that of an optimal mechanism.

[9] If a distribution with a continuous density function is desired, the point mass at H can be spread out over a very short interval centered at H.

2. Prove that the expected revenue of the Single Sample mechanism is close to that of the VCG-L mechanism with monopoly reserves in the subenvironment induced by the non-reserve bidders.

Given two such approximation guarantees, we can combine them with a generalized version of Lemma 3.5, as in the proof of Theorem 3.2, to show that the expected revenue of the Single Sample mechanism is a constant fraction of that of the optimal mechanism. Section 3.2 implemented this plan for the special case of i.i.d. regular matroid environments, where the VCG-L mechanism with monopoly reserves is optimal.

The arguments in Section 3.2 essentially accomplish the second step of the proof framework, with an approximation factor of 2, for all regular downward-closed non-singular environments. The harder part is the first step. The next two sections establish such approximation guarantees under two incomparable sets of assumptions, via two different arguments: m.h.r. downward-closed environments, and regular matroid environments.

For m.h.r. downward-closed environments, we prove that the expected revenue of the VCG-L mechanism with monopoly reserves is at least a $1/e$ fraction of that of an optimal mechanism (Theorem 3.11). This implies that the expected revenue of the Single Sample mechanism is at least a $\frac{1}{2e} \cdot \frac{\kappa-1}{\kappa}$ fraction of that of an optimal mechanism when there are at least $\kappa \geq 2$ bidders of every present attribute (Theorem 3.12). Via an optimized analysis, we also prove an approximation factor of $\frac{1}{4} \cdot \frac{\kappa-1}{\kappa}$ (Theorem 3.14). This factor is $\frac{1}{8}$ when $\kappa = 2$ and quickly approaches $\frac{1}{4}$ as κ grows.

For regular matroid environments, we prove that the expected revenue of the VCG-L mechanism with monopoly reserves is at least half that of an optimal mechanism (Theorem 3.17), which in turn implies an approximation guarantee of $\frac{1}{4}\frac{\kappa-1}{\kappa}$ for the Single Sample mechanism (Theorem 3.18).

3.4. M.H.R. Downward-Closed Environments

We now implement the proof framework outlined in Section 3.3 for m.h.r. downward-closed environments. We carry out the arguments for expected welfare, rather than expected revenue, because this gives a stronger result. We first generalize Lemma 3.5 to non-i.i.d. environments.

Lemma 3.9 *For every m.h.r. downward-closed environment with at least $\kappa \geq 2$ bidders of every present attribute, the expected optimal welfare in the subenvironment induced by non-reserve bidders is at least a $(\kappa-1)/\kappa$ fraction of that in the original environment.*

The proof of Lemma 3.9 is essentially the same as that of Lemma 3.5, with valuations assuming the role previously played by virtual valuations. In contrast to Remark 3.7, discarding reserve bidders before the VCG computation in step (2) of Definition 3.1 is important for the analysis of the Single Sample mechanism in non-matroid environments.

Analogous to Lemma 3.6, we require a technical lemma about the single-bidder case to establish step 1 of our proof framework.

Lemma 3.10 *Let F be an m.h.r. distribution with monopoly price r^* and revenue function $\widehat{R}$. Let $V(t)$ denote the expected welfare of a single-item auction with a posted price of t and a single bidder with valuation drawn from F. For every nonnegative number $t \geq 0$,*

$$\widehat{R}(\max\{t, r^*\}) \geq \frac{1}{e} \cdot V(t). \tag{2}$$

Proof: Let s denote $\max\{t, r^*\}$. Recall that, by the definition of the hazard rate function, $1 - F(x) = e^{-H(x)}$ for every $x \geq 0$, where $H(x)$ denotes $\int_0^x h(z)dz$. Note that since $h(z)$ is non-negative and nondecreasing, $H(x)$ is nondecreasing and convex. We can write the left-hand side of (2) as $s \cdot (1 - F(s)) = s \cdot e^{-H(s)}$ and, for a random sample v from F,

$$\begin{aligned} V(t) &= \Pr[v \geq t] \cdot \mathbf{E}[v \mid v \geq t] \\ &= e^{-H(t)} \cdot \left[t + \int_t^\infty e^{-(H(v) - H(t))} dv \right]. \end{aligned} \tag{3}$$

By convexity of the function H, we can lower bound its value using a first-order approximation at s:

$$H(v) \geq H(s) + H'(s)(v - s) = H(s) + h(s)(v - s) \tag{4}$$

for every $v \geq 0$. There are now two cases. If $t \leq r^* = s$, then $h(s) = 1/s$ since r^* is a monopoly price.[10] Starting from (3) and using that H is nondecreasing, and then substituting (4) yields

$$\begin{aligned} V(t) &\leq \int_0^\infty e^{-H(v)} dv \\ &\leq \int_0^\infty e^{-(H(s) + \frac{v-s}{s})} dv \\ &= e \cdot s \cdot e^{-H(s)}. \end{aligned}$$

If $r^* \leq t = s$, then the m.h.r. assumption implies that $h(s) \geq 1/s$ and (4) implies that $H(v) \geq H(t) + (v - t)/t$ for all $v \geq t$. Substituting into (3) gives

$$\begin{aligned} V(t) &\leq e^{-H(t)} \cdot \left[t + \int_t^\infty e^{-(H(t) + \frac{v-t}{t} - H(t))} dv \right] \\ &\leq e^{-H(t)} \cdot \int_0^\infty e^{-\frac{v-t}{t}} dv \\ &= e \cdot s \cdot e^{-H(s)}, \end{aligned}$$

where in the second inequality we use that $e^{-(v-t)/t} \geq 1$ for every $v \leq t$. ■

[10] One proof of this follows from the first-order condition for the revenue function $p(1-F(p))$; alternatively, applying Myerson's Lemma to the single-bidder case shows that $r^* = \varphi_F^{-1}(0)$ and hence $r^* - 1/h(r^*) = \varphi_F(r^*) = 0$.

Lemma 3.10 implies that the expected revenue of the VCG-L mechanism with monopoly reserves is at least a $\frac{1}{e}$ fraction of the expected optimal welfare in every downward-closed environment with m.h.r. valuation distributions.

Theorem 3.11 (VCG-L With Monopoly Reserves) *For every m.h.r. downward-closed environment, the expected revenue of the VCG-L mechanism with monopoly reserves is at least a $\frac{1}{e}$ fraction of the expected efficiency of the VCG mechanism.*

Proof: Fix a bidder i and valuations $\mathbf{v}_{-i}$. This determines a winning threshold t for bidder i under the VCG mechanism (with no reserves). Lemma 3.10 implies that the conditional expected revenue obtained from i in the VCG-L mechanism with monopoly reserves is at least a $1/e$ fraction of the conditional expected welfare obtained from i in the VCG mechanism (with no reserves). Taking expectations over $\mathbf{v}_{-i}$ and then summing over all the bidders proves the theorem. ■

Considering a single bidder with an exponentially distributed valuation ($F(v) = 1 - e^{-v}$) shows that the bounds in Lemma 3.10 and Theorem 3.11 are tight in the worst case.

Theorem 3.11 establishes step 1 of our proof framework. The arguments in Section 3.2 now imply that the expected revenue of the Single Sample mechanism is almost half that of the VCG-L mechanism with monopoly reserves (step 2). Precisely, mimicking the proof of Theorem 3.2, with Lemma 3.9 replacing Lemma 3.5, gives the following result.

Theorem 3.12 (Single Sample Guarantee #1) *For every m.h.r. downward-closed environment with at least $\kappa \geq 2$ bidders of every present attribute, the expected revenue of the Single Sample mechanism is at least a $\frac{1}{2e} \cdot \frac{\kappa-1}{\kappa}$ fraction of the expected optimal welfare in the environment.*

We can improve the guarantee in Theorem 3.12 by optimizing jointly the two single-bidder guarantees in Lemmas 3.10 (step 1) and 3.6 (step 2). This is done in the next lemma.

Lemma 3.13 *Let F be an m.h.r. distribution with monopoly price r^* and revenue function $\widehat{R}$, and define $V(t)$ as in Lemma 3.10. Let v denote a random valuation from F. For every nonnegative number $t \geq 0$,*

$$\mathbf{E}_v\left[\widehat{R}(\max\{t, v\})\right] \geq \frac{1}{4} \cdot V(t). \qquad (5)$$

Proof: Define the function H as in the proof of Lemma 3.10, and recall from that proof that $V(t)$ can be written as in (3). We show that the left-hand side of (5) is at least 25% of that quantity.

Consider two i.i.d. samples v_1, v_2 from F. We interpret v_2 as the random reserve price v in (5) and v_1 as the valuation of the single bidder. The left-hand side of (5) is equivalent to the expectation of a random variable that is equal to t if $v_2 \leq t \leq v_1$, which occurs with probability $F(t) \cdot (1 - F(t))$; equal to v_2 if $t \leq v_2 \leq v_1$, which occurs with probability $\frac{1}{2}(1 - F(t))^2$; and equal to zero, otherwise. Hence,

$$\begin{aligned}
\mathbf{E}_v\left[\widehat{R}(\max\{t,v\})\right] &\geq \frac{1}{2}\left(F(t) \cdot (1 - F(t)) \cdot t\right. \\
&\quad \left. + (1 - F(t))^2 \cdot \mathbf{E}[\min\{v_1, v_2\} \mid \min\{v_1, v_2\} \geq t]\right) \\
&= \frac{1}{2}(1 - F(t)) \cdot \Big(t \cdot F(t) + (1 - F(t)) \cdot \\
&\qquad\qquad \left[t + e^{2H(t)} \int_t^\infty e^{-2H(v)} dv\right]\Big) \\
&\geq \frac{1}{2}(1 - F(t)) \cdot \left[t + e^{H(t)} \int_t^\infty e^{-H(2v)} dv\right] \qquad (6) \\
&= \frac{1}{4}(1 - F(t)) \cdot \left[2t + \int_{2t}^\infty e^{-(H(v) - H(t))} dv\right] \\
&\geq \frac{1}{4}(1 - F(t)) \cdot \left[t + \int_t^\infty e^{-(H(v) - H(t))} dv\right], \qquad (7)
\end{aligned}$$

where in (6) and (7) we use that H is non-negative, nondecreasing, and convex. Comparing (3) and (7) proves the lemma. ■

We then obtain the following optimized version of Theorem 3.12.

Theorem 3.14 (Single Sample Guarantee #2) *For every m.h.r. downward-closed environment with at least $\kappa \geq 2$ bidders of every present attribute, the expected revenue of the Single Sample mechanism is at least a $\frac{1}{4} \cdot \frac{\kappa - 1}{\kappa}$ fraction of the expected optimal welfare in the environment.*

The proof of Theorem 3.14 is the same as that of Theorem 3.2, with the following substitutions: the welfare of the VCG mechanism (with no reserves) plays the previous role of the revenue of the VCG-L mechanism with monopoly reserves; Lemma 3.13 replaces Lemma 3.6; and Lemma 3.9 takes the place of Lemma 3.5.

Remark 3.15 (Theorem 3.14 Is Tight) Our analysis of the Single Sample mechanism is tight for all values of $\kappa \geq 2$, as shown by a digital goods environment with κ bidders with valuations drawn i.i.d. from an exponential distribution ($F(v) = 1 - e^{-v}$): the expected optimal welfare is κ, and a calculation shows that the expected revenue of the Single Sample mechanism is $(\kappa - 1)/4$.

Since the revenue of every mechanism is bounded above by its welfare, we have the following corollary.

Corollary 3.16 *For every m.h.r. environment with at least $\kappa \geq 2$ bidders of every present attribute, the expected revenue of the Single Sample mechanism is at least a $\frac{1}{4} \cdot \frac{\kappa - 1}{\kappa}$ fraction of that of the optimal mechanism for the environment.*

3.5. Regular Matroid Environments

This section proves an approximation guarantee for the Single Sample mechanism under assumptions incomparable to those in Section 3.4, namely for regular matroid environments. We again follow the proof framework outlined in Section 3.3, step 1 of which involves proving an approximation bound for the VCG-L mechanism with monopoly reserves. In Section 3.4 we proved the stronger statement that the expected revenue of this mechanism is at least a constant fraction of the optimal expected *welfare*. No mechanism achieves this stronger guarantee with regular valuation distributions, so we use a different line of argument.

Hartline and Roughgarden (2009) proved that the expected revenue of the VCG-E mechanism with monopoly reserves (Section 2.2) is at least half that of an optimal mechanism in regular matroid environments. The VCG-E and VCG-L mechanisms do not coincide in matroid environments unless all bidders face a common reserve price (cf., Corollary 3.4), and the results of Hartline and Roughgarden (2009) have no obvious implications for the VCG-L mechanism with monopoly reserves in matroid environments with non-i.i.d. bidders. We next supplement the arguments in Hartline and Roughgarden (2009) with some new ideas to prove an approximation guarantee for this mechanism.

Theorem 3.17 (VCG-L With Monopoly Reserves) *For every regular matroid environment, the expected revenue of the VCG-L mechanism with monopoly reserves is at least a $\frac{1}{2}$ fraction of that of an optimal mechanism.*

Proof: Consider a regular matroid environment. For a valuation profile $\mathbf{v}$, let $W(\mathbf{v})$ and $W'(\mathbf{v})$ denote the winning bidders in the VCG-L mechanism with monopoly reserves and in the optimal mechanism, respectively. We claim that

$$E_{\mathbf{v}}\left[\sum_{i \in W(\mathbf{v}) \setminus W'(\mathbf{v})} \varphi_i(v_i)\right] \geq 0 \tag{8}$$

and

$$E_{\mathbf{v}}\left[\sum_{i \in W(\mathbf{v})} p_i(\mathbf{v})\right] \geq E_{\mathbf{v}}\left[\sum_{i \in W'(\mathbf{v}) \setminus W(\mathbf{v})} \varphi_i(v_i)\right], \tag{9}$$

where $\mathbf{p}$ denotes the payment rule of the VCG-L mechanism with monopoly reserves. Given these two claims, Lemma 3.9 in Hartline and Roughgarden (2009) immediately implies the theorem.[11]

[11] The proof goes as follows. First, using (8), the virtual welfare of the optimal mechanism from bidders in $W(\mathbf{v}) \cap W'(\mathbf{v})$ is at most that of the virtual welfare of the VCG-L mechanism with monopoly reserves. Second, using (9), the virtual welfare of the optimal mechanism from bidders in $W'(\mathbf{v}) \setminus W(\mathbf{v})$ is at most the revenue of the VCG-L mechanism with monopoly reserves. Finally, applying Myerson's Lemma (Lemma 2.1) completes the proof.

The inequalities (8) and (9) almost correspond to the definition of "commensurate" in Hartline and Roughgarden (2009, Definition 3.8), but our second inequality is weaker. Nonetheless, the proof of Lemma 3.9 in Hartline and Roughgarden (2009) carries over unchanged.

Inequality (8) holds because the VCG-L mechanism with monopoly reserves allocates to a bidder only if its valuation is at least its monopoly reserve. Since a monopoly reserve r_i^* satisfies $\varphi_i(r_i^*) = 0$ and the virtual valuation function is nondecreasing (by regularity), a bidder wins only when it has a non-negative virtual valuation.

Proving inequality (9) requires a more involved argument. Fix a valuation profile $\mathbf{v}$ and let $W''(\mathbf{v})$ denote the winning bidders under the VCG mechanism with zero reserve prices. Because the greedy algorithm maximizes welfare in matroid environments, the winners $W(\mathbf{v})$ under the VCG-L mechanism with monopoly reserves are precisely the bidders of $W''(\mathbf{v})$ that have nonnegative virtual valuations.

The exchange property of matroids (Section 2.1) implies that all maximal feasible sets have equal size and, since $W''(\mathbf{v})$ must be maximal, that we can choose a subset $S \subseteq W''(\mathbf{v}) \setminus W'(\mathbf{v})$ such that $S \cup W'(\mathbf{v})$ and $W''(\mathbf{v})$ have the same size. Since the winners $W'(\mathbf{v})$ under the optimal mechanism maximize the virtual welfare, all bidders of S have nonpositive virtual values.

We next use a non-obvious but well-known property of matroids (see e.g., Schrijver (2003, Corollary 39.12a)): given two feasible sets of equal size, such as $W''(\mathbf{v})$ and $S \cup W'(\mathbf{v})$, there is a bijection f from $(S \cup W'(\mathbf{v})) \setminus W''(\mathbf{v})$ to $W''(\mathbf{v}) \setminus (S \cup W'(\mathbf{v}))$ such that, for every bidder i in the domain, $W''(\mathbf{v}) \setminus \{f(i)\} \cup \{i\}$ is a feasible set. Since $S \subseteq W''(\mathbf{v})$, the domain of f is simply $W'(\mathbf{v}) \setminus W''(\mathbf{v})$. Since the VCG mechanism chooses a welfare-maximizing set, the threshold bid (and hence the payment) of a winning bidder $f(i)$ in the range of the function f is at least v_i. Summing over all bidders in $W'(\mathbf{v}) \setminus W''(\mathbf{v})$ and using that f is a bijection, the revenue of the VCG mechanism is at least

$$\sum_{i \in W'(\mathbf{v}) \setminus W''(\mathbf{v})} v_i \geq \sum_{i \in W'(\mathbf{v}) \setminus W''(\mathbf{v})} \varphi_i(v_i), \tag{10}$$

where the inequality follows from the definition of a virtual valuation. Because $W'(\mathbf{v})$ contains no bidders with negative virtual valuations and all bidders in $W''(\mathbf{v}) \setminus W(\mathbf{v})$ have nonpositive virtual valuations, the right-hand side of (10) equals

$$\sum_{i \in W'(\mathbf{v}) \setminus W(\mathbf{v})} \varphi_i(v_i).$$

The expected revenue of the VCG mechanism (with no reserves) is at least the expected value of this quantity. Since the VCG-L mechanism with monopoly reserves differs from the VCG mechanism only by excluding bidders with negative virtual valuations, Lemma 2.1 implies that its expected revenue is at least that of the VCG mechanism. This completes the proof of inequality (9) and the theorem. ■

An approximation guarantee for the Single Sample mechanism follows as in the proof of Theorem 3.12, with Theorem 3.17 replacing Theorem 3.11.

Theorem 3.18 (Single Sample Guarantee) *For every regular matroid environment with at least $\kappa \geq 2$ bidders of every present attribute, the expected revenue of the Single Sample mechanism is at least a $\frac{1}{4} \cdot \frac{\kappa-1}{\kappa}$ fraction of that of an optimal mechanism for the environment.*

3.6. Counterexample for Regular Downward-Closed Environments

We now sketch an example showing that a restriction to m.h.r. valuation distributions (as in Section 3.4) or to matroid environments (as in Section 3.5) is necessary for the VCG-L mechanism with monopoly reserves and the Single Sample mechanism to have constant-factor approximation guarantees. The following example is adapted from Hartline and Roughgarden (2009, Example 3.4).

For n sufficiently large, consider two "big" bidders and n "small" bidders $1, 2, \ldots, n$. The feasible subsets are precisely those that do not contain both a big bidder and a small bidder; this is not a matroid environment. Fix an arbitrarily large constant H. Each big bidder's valuation is deterministically $\frac{1}{2}n\sqrt{\ln H}$, so the expected revenue of an optimal mechanism is clearly at least $n\sqrt{\ln H}$. The small bidders' valuations are i.i.d. draws from the distribution $F(v) = 1 - \frac{1}{v+1}$ on $[0, H)$ and $F(H) = 1$. This distribution F is regular — or can be made so with a minor perturbation, as in Remark 3.8 — but is not m.h.r. For n sufficiently large, the sum of the small bidders' valuations is tightly concentrated around $n \ln H$.

We complete the sketch for the VCG-L mechanism with monopoly reserves; the argument for the Single Sample mechanism is almost identical. The VCG mechanism almost surely chooses all small bidders as its preliminary winner set, with a threshold bid of zero for each. The expected revenue extracted from each small winner, via its monopoly reserve H, is at most 1.[12] Thus, the expected revenue of the VCG-L mechanism with monopoly reserves is not much more than n, which is arbitrarily smaller than the optimal expected revenue as $H \to \infty$.

3.7. Computationally Efficient Variants

In the second step of the Single Sample mechanism, a different mechanism can be swapped in for the VCG mechanism. One motivation for using a different mechanism is computational efficiency (although this is not a first-order goal in this paper). For instance, for combinatorial auctions with single-minded bidders — where feasible sets of bidders correspond to those desiring mutually disjoint bundles of goods — implementing the VCG mechanism requires the solution of a packing problem that is NP-hard, even to approximate.

[12] A subtle point is that each small bidder's valuation is now drawn at random from F, conditioned on the event that the VCG mechanism chose all of the small bidders. But since the small bidders are chosen with overwhelming probability (for large n and H), the probability that a given small bidder is pivotal is vanishingly small, so it still contributes at most 1 to the expected revenue of the mechanism.

For example, the proof of Theorem 3.14 evidently implies the following more general statement: if step (2) of the Single Sample mechanism uses a truthful mechanism guaranteed to produce a solution with at least a $1/c$ fraction of the maximum welfare, then the expected revenue of the corresponding Single Sample mechanism is at least a $\frac{1}{4c}\frac{\kappa-1}{\kappa}$ fraction of the expected optimal welfare (whatever the underlying m.h.r. downward-closed environment). For knapsack auctions — where each bidder has a public size and feasible sets of bidders are those with total size at most a publicly known budget — we can substitute the polynomial-time, $(1+\epsilon)$-approximation algorithm by Briest et al. (2005). For combinatorial auctions with single-minded bidders, we can use the algorithm of Lehmann et al. (2002) to obtain an $O(\sqrt{m})$-approximation in polynomial time, where m is the number of goods. This factor is essentially optimal for polynomial-time approximation, under appropriate computational complexity assumptions (Lehmann et al., 2002).

4. Revenue Guarantees with Multiple Samples

This section modifies the Single Sample mechanism to achieve improved guarantees via an increased number of samples from the underlying valuation distributions, and provides quantitative and distribution-independent polynomial bounds on the number of samples required to achieve a given approximation factor.

4.1. Estimating Monopoly Reserve Prices

Improving the revenue guarantees of Section 3 via multiple samples requires thoroughly understanding the following simpler problem: Given an accuracy parameter ϵ and a regular distribution F, how many samples m from F are needed to compute a reserve price r that is $(1-\epsilon)$-*optimal*, meaning that $\widehat{R}(r) \geq (1-\epsilon)\cdot\widehat{R}(r^*)$ for a monopoly reserve price r^* for F? Recall from Section 3.2 that $\widehat{R}(p)$ denotes $p\cdot(1-F(p))$. We pursue bounds on m that depend *only* on ϵ and not on the distribution F — such bounds do not follow from the Law of Large Numbers and must make use of the regularity assumption.

Given m samples from F, renamed so that $v_1 \geq v_2 \geq \cdots \geq v_m$, an obvious idea is to use the reserve price that is optimal for the corresponding empirical distribution, which we call the *empirical reserve*:

$$\operatorname*{argmax}_{i\geq 1}\ i\cdot v_i. \tag{11}$$

Interestingly, this naive approach does *not* in general give distribution-independent polynomial sample complexity bounds. Intuitively, with a heavy-tailed distribution F, there is a constant probability that a few large outliers cause the empirical reserve to be overly large, while a small reserve price has much better expected revenue for F.

Our solution is to forbid the largest samples from acting as reserve prices, leading to a quantity we call the *guarded empirical reserve* (with respect to an accuracy parameter ϵ):

$$\operatorname*{argmax}_{i \geq \epsilon m} i \cdot v_i. \tag{12}$$

We use the guarded empirical reserve to prove distribution-independent polynomial bounds on the sample complexity needed to estimate the monopoly reserve of a regular distribution.

Lemma 4.1 (Estimating the Monopoly Reserve) *For every regular distribution F and sufficiently small $\epsilon, \delta > 0$, the following statement holds: with probability at least $1-\delta$, the guarded empirical reserve (12) of $m \geq c(\epsilon^{-3}(\ln \epsilon^{-1} + \ln \delta^{-1}))$ samples from F is a $(1-\epsilon)$-optimal reserve, where c is a constant that is independent of F.*

Proof: Set $\gamma = \epsilon/11$ and consider m samples $v_1 \geq v_2 \geq \cdots \geq v_m$ from F. Define "q-values" by $q_t = 1 - F(v_t)$ and $q^* = 1 - F(r^*)$, where r^* is a monopoly price for F. Since the q's are i.i.d. samples from the uniform distribution on $[0,1]$, the expected value of the quantile q_t is $t/(m+1)$, which we estimate by t/m for simplicity. An obvious approach is to use Chernoff bounds to argue that each q_t is close to this expectation, followed by a union bound. Two issues are: for small t's, the probability that t/m is a very good estimate of q_t is small; and applying the union bound to such a large number of events leads to poor probability bounds. In the following, we restrict attention to a carefully chosen small subset of quantiles, and take advantage of the properties of the revenue functions of regular distributions to get around these issues.

First we choose an integer index sequence $0 = t_0 < t_1 < \cdots < t_L = m$ in the following way. Let $t_0 = 0$ and $t_1 = \lfloor \gamma m \rfloor$. Inductively, if t_i is defined for $i \geq 1$ and $t_i < m$, define t_{i+1} to be the largest integer in $\{1, \ldots, m\}$ such that $t_i < t_{i+1} \leq (1+\gamma)t_i$. If $m = \Omega(\gamma^{-2})$, then $t_i + 1 \leq (1+\gamma)t_i$ for every $t_i \geq \gamma m$ and hence such a t_{i+1} exists. Observe that $L \approx \log_{1+\gamma} \frac{1}{\gamma} = O(\gamma^{-2})$ and $t_{i+1} - t_i \leq \gamma m$ for every $i \in \{0, \ldots, L-1\}$.

We claim that, with probability 1, a sampled quantile q_t with $t \geq \gamma m$ differs from t/m by more than a $(1 \pm 3\gamma)$ factor only if some quantile q_{t_i} with $i \in \{1, 2, \ldots, L\}$ differs from t_i/m by more than a $(1 \pm \gamma)$ factor. For example, suppose that $q_t > \frac{(1+3\gamma)t}{m}$ with $t \geq \gamma m$; the other case is symmetric. Let $i \in \{1, 2, \ldots, L\}$ be such that $t_i \leq t \leq t_{i+1}$. Then

$$q_{t_{i+1}} \geq q_t > \frac{(1+3\gamma)t}{m} \geq \frac{(1+3\gamma)t_i}{m} \geq \frac{(1+3\gamma)t_{i+1}}{(1+\gamma)m} \geq \frac{(1+\gamma)t_{i+1}}{m},$$

as claimed.

We next claim that the probability that q_{t_i} differs from t_i/m by more than a $(1 \pm \gamma)$ factor for some $i \in \{1, 2, \ldots, L\}$ is at most $2Le^{-\gamma^3 m/4}$. Fix $i \in \{1, 2, \ldots, L\}$. Note that $q_{t_i} > (1+\gamma)\frac{t_i}{m}$ only if less than t_i samples have q-values at most $(1+\gamma)\frac{t_i}{m}$. Since the expected number of such samples is $(1+$

$\gamma)t_i$, Chernoff bounds (e.g., Motwani and Raghavan (1995)) imply that the probability that $q_{t_i} > (1+\gamma)\frac{t_i}{m}$ is at most

$$\exp\{-\gamma^2 t_i/3(1+\gamma)\} \le \exp\{-\gamma^2 t_i/4\} \le \exp\{-\gamma^3 m/4\},$$

where the inequalities use that γ is at most a sufficiently small constant and that $t_i \ge \gamma m$ for $i \ge 1$. A similar argument shows that the probability that $q_{t_i} < (1-\gamma)\frac{t_i}{m}$ is at most $\exp\{-\gamma^3 m/4\}$, and a union bound completes the proof of the claim. If $m = \Omega(\gamma^{-3}(\log L + \log \delta^{-1})) = \Omega(\epsilon^{-3}(\log \epsilon^{-1} + \log \delta^{-1}))$, then this probability is at most δ.

Now condition on the event that every quantile q_{t_i} with $i \in \{1, 2, \ldots, L\}$ differs from t_i/m by at most a $(1 \pm \gamma)$ factor, and hence every quantile q_t with $t \ge \gamma m$ differs from t/m by at most a $(1 \pm 3\gamma)$ factor. We next show that there is a candidate for the guarded empirical reserve (12) which, if chosen, has good expected revenue. Choose $i \in \{0, 1, \ldots, L-1\}$ so that $t_i/m \le q^* \le t_{i+1}/m$. Define t^* as t_i if $q^* \ge 1/2$ and t_{i+1} otherwise. Assume for the moment that $q^* \le 1/2$. By the concavity of revenue function in probability space $R(q)$ — recall Section 3.2 — $R(q_{t_{i+1}})$ lies above the line segment between $R(q^*)$ and $R(1)$. Since $R(1) = 0$, this translates to

$$R(q_{t^*}) \ge R(q^*) \cdot \frac{1 - q_{t_{i+1}}}{1 - q^*} \ge R(q^*) \cdot \frac{1 - (1+3\gamma)\left(\frac{t_i}{m} + \gamma\right)}{1 - \frac{t_i}{m}} \ge (1 - 5\gamma) \cdot R(q^*),$$

where in the final inequality we use that $\frac{t_i}{m} \le \frac{1}{2}$ and γ is sufficiently small. For the case when $q^* \ge \frac{1}{2}$, a symmetric argument (using $R(0)$ instead of $R(1)$ and q_{t_i} instead of $q_{t_{i+1}}$) proves that $R(q_{t^*}) \ge (1 - 5\gamma) \cdot R(q^*)$.

Finally, we show that the guarded empirical reserve also has good expected revenue. Let the maximum in (12) correspond to the index $\hat{t}$. Since $\hat{t}$ was chosen over t^*, we have $\hat{t} \cdot v_{\hat{t}} \ge t^* \cdot v_{t^*}$. Using that each of $q_{\hat{t}}, q_{t^*}$ is approximated up to a $(1 \pm 3\gamma)$ factor by $\hat{t}/m, t^*/m$ yields

$$R(q_t) = q_{\hat{t}} v_{\hat{t}} \ge \frac{(1-3\gamma)\hat{t}}{m} v_{\hat{t}} \ge \frac{(1-3\gamma)t^*}{m} v_{t^*} \ge \frac{1-3\gamma}{1+3\gamma} q_{t^*} v_{t^*} = \frac{1-3\gamma}{1+3\gamma} R(q_{t^*})$$

and hence

$$R(q_t) \ge \frac{(1-5\gamma)(1-3\gamma)}{1+3\gamma} R(q^*) \ge (1 - 11\gamma) R(q^*).$$

Since $\gamma = \epsilon/11$, the proof is complete. ■

Remark 4.2 (Optimization for M.H.R. Distributions) There is a simpler and stronger version of Lemma 4.1 for m.h.r. distributions. We use a simple fact, first noted in Hartline et al. (2008, Lemma 4.1), that the selling probability q^* at the monopoly reserve r^* for an m.h.r. distribution is at least $1/e$. Because of this, we can take the parameter t_1 in the proof of Lemma 4.1 to be $\lfloor m/e \rfloor$ instead of $\lfloor \gamma m \rfloor$ without affecting the rest of the proof. This saves a γ factor in the exponent of the bound on the probability that some q_{t_i} is not well approximated by t_i/m, which translates to a new sample complexity bound of

$m \geq c(\epsilon^{-2}(\ln \epsilon^{-1} + \ln \delta^{-1}))$, where c is some constant that is independent of the underlying distribution. Also, this bound remains valid even for the empirical reserve (11) — the guarded version in (12) is not necessary.

4.2. The Many Samples Mechanism

In the following *Many Samples mechanism*, we assume that an accuracy parameter ϵ is given, and use m to denote the sample complexity bound of Lemma 4.1 (for regular valuation distributions) or of Remark 4.2 (for m.h.r. distributions) corresponding to the accuracy parameter $\frac{\epsilon}{3}$ and failure probability $\frac{\epsilon}{3}$. The mechanism is only defined if every present attribute is shared by more than m bidders.

(1) For each represented attribute a, pick a subset S_a of m *reserve bidders* with attribute a uniformly at random from all such bidders.

(2) Run the VCG mechanism on the sub-environment induced by the non-reserve bidders to obtain a preliminary winning set P.

(3) For each bidder $i \in P$ with attribute a, place i in the final winning set W if and only if v_i is at least the guarded empirical reserve r_a of the samples in S_a. Charge every winner $i \in W$ with attribute a the maximum of its VCG payment computed in step (2) and the reserve price r_a.

We prove the following guarantees for this mechanism.

Theorem 4.3 (Guarantees for Many Samples) *The expected revenue of the Many Samples mechanism is at least:*

(a) a $(1-\epsilon)$ fraction of that of an optimal mechanism in every i.i.d. regular matroid environment with at least $n \geq 3m/\epsilon = \Theta(\epsilon^{-4} \log \epsilon^{-1})$ bidders;

(b) a $\frac{1}{2}(1-\epsilon)$ fraction of that of an optimal mechanism in every regular matroid environment with at least $\kappa \geq 3m/\epsilon = \Theta(\epsilon^{-4} \log \epsilon^{-1})$ bidders of every present attribute;

(c) a $\frac{1}{e}(1-\epsilon)$ fraction of the optimal expected welfare in every downward-closed m.h.r. environment with at least $\kappa \geq 3m/\epsilon = \Theta(\epsilon^{-3} \log \epsilon^{-1})$ bidders of every present attribute.

Bidders with i.i.d. and exponentially distributed valuations show that part (c) of the theorem is asymptotically optimal (as is part (a), obviously).

Proof: The lower bound $\kappa \geq 3m/\epsilon$ on the number of bidders of each attribute implies that at most an $\epsilon/3$ fraction of all bidders are designated as reserve bidders. Lemmas 3.5 and 3.9 imply that the expectation, over the choice of reserve bidders, of the expected revenue of an optimal mechanism for and the expected optimal welfare of the subenvironment induced by the non-reserve bidders are at least a $(1-\frac{\epsilon}{3})$ fraction of those in the full environment.

Now condition on the choice of reserve bidders, but not on their valuations. Fix a non-reserve bidder i, and condition on the valuations of all other non-reserve bidders. Let t denote the induced threshold bid for i and r^* a monopoly price for the valuation distribution F of i. The conditional expected revenue obtained from i using the price $\max\{r^*, t\}$ is precisely that obtained by the VCG-L mechanism with monopoly reserves for the subenvironment induced by the non-reserve bidders.

The Many Samples mechanism, on the other hand, uses the price $\max\{r, t\}$, where r is the guarded empirical reserve of the reserve bidders that share i's attribute. By Lemma 4.1 and our choice of m, r is $(1 - \frac{\epsilon}{3})$-optimal for F with probability at least $1 - \frac{\epsilon}{3}$. Concavity of the revenue function (cf., Figure 1) and an easy case analysis shows that, whenever r is $(1 - \frac{\epsilon}{3})$-optimal, the conditional expected revenue from i with the price $\max\{r, t\}$ is at least a $(1 - \frac{\epsilon}{3})$ fraction of that with the price $\max\{r^*, t\}$, for every value of t. Since valuations are independent of each other and the choice of the reserve bidders, the expected revenue from i in the Many Samples mechanism, conditioned on the choice of reserve bidders and on the valuations of the other non-reserve bidders, is at least a $(1 - \frac{\epsilon}{3})^2 \geq (1 - \frac{2}{3}\epsilon)$ fraction of that of the VCG-L mechanism with monopoly reserves. Removing the conditioning on the valuations of other non-reserve bidders; summing over the non-reserve bidders; and removing the conditioning on the choice of reserve bidders shows that the expected revenue of the Many Samples mechanism is at least a $(1 - \frac{2}{3}\epsilon)$ fraction of that of the VCG-L mechanism with monopoly reserves on the subenvironment induced by the non-reserve bidders. The three parts of the theorem now follow from Corollary 3.4, Theorem 3.17, and Theorem 3.11, respectively. ■

Remark 4.4 (Case Study: Digital Goods Auctions) Our results in this section have interesting implications even in the special case of digital goods auctions. We note that there is no interference between different bidders in such an auction, so the general case of multiple attributes reduces to the single-attribute i.i.d. case (each attribute can be treated separately).

The Deterministic Optimal Price (DOP) digital goods auction offers each bidder i a take-it-or-leave offer equal to the empirical reserve of the other $n - 1$ bidders. The expected revenue of the DOP auction converges to that of an optimal auction as the number n of bidders goes to infinity, provided valuations are i.i.d. samples from a distribution with bounded support (Goldberg et al., 2006) or from a regular distribution (Segal, 2003). However, the number of samples required in these works to achieve a given degree of approximation depends on the underlying distribution F.

As an alternative, consider the variant of DOP that instead uses the guarded empirical reserve (12) of the other $n-1$ bidders to formulate a take-it-or-leave-it offer for each bidder. Our Lemma 4.1 implies a *distribution-independent* bound for this auction: provided the number of bidders is $\Omega(\epsilon^{-3} \log \epsilon^{-1})$, its expected revenue is at least $(1 - \epsilon)$ times that of the optimal auction.

Acknowledgments

We thank Jason Hartline for proposing the term "prior-independent auctions", for the observation in Remark 3.7, and for a number of other helpful comments. We also thank the anonymous referees for their suggestions.

References

Archer, A., Tardos, É., 2001. Truthful mechanisms for one-parameter agents. In: Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS). pp. 482–491.

Azar, P. D., Daskalakis, C., Micali, S., Weinberg, S. M., 2013. Optimal and efficient parametric auctions. In: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). pp. 596–604.

Baliga, S., Vohra, R., 2003. Market research and market design. Advances in Theoretical Economics 3, article 5.

Bergemann, D., Morris, S., 2005. Robust mechanism design. Econometrica 73 (6), 1771–1813.

Briest, P., Krysta, P., Vöcking, B., 2005. Approximation techniques for utilitarian mechanism design. In: Proceedings of the 37th ACM Symposium on Theory of Computing (STOC). pp. 39–48.

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

Caillaud, B., Robert, J., 2005. Implementation of the revenue-maximizing auction by an ignorant seller. Review of Economic Design 9 (2), 127–143.

Chawla, S., Hartline, J. D., Malec, D. L., Sivan, D., 2013. Prior-independent mechanisms for scheduling. In: Proceedings of the 45th ACM Symposium on Theory of Computing (STOC). pp. 51–60.

Devanur, N., Hartline, J. D., 2009. Limited and online supply and the Bayesian foundations of prior-free mechanism design. In: Proceedings of the 10th ACM Conference on Electronic Commerce (EC). pp. 41–50.

Devanur, N., Hartline, J. D., Karlin, A. R., Nguyen, T., 2011. A prior-independent mechanism for profit maximization in unit-demand combinatorial auctions. In: Proceedings of 7th Workshop on Internet & Network Economics. pp. 122–133.

Dhangwatnotai, P., Roughgarden, T., Yan, Q., 2010. Revenue maximization with a single sample. In: Proceedings of the 11th ACM Conference on Electronic Commerce (EC). pp. 129–138.

Dughmi, S., Roughgarden, T., Sundararajan, M., 2012. Revenue submodularity. Theory of Computing 8, 95–119.

Fu, H., Hartline, J. D., Hoy, D., 2013. Prior-independent auctions for risk-averse agents. In: Proceedings of the 14th ACM Conference on Electronic Commerce (EC). pp. 471–488.

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

Hartline, J. D., Mirrokni, V. S., Sundararajan, M., 2008. Optimal marketing strategies over social networks. In: 17th International World Wide Web Conference. pp. 189–198.

Hartline, J. D., Roughgarden, T., 2008. Optimal mechanism design and money burning. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC). pp. 75–84.

Hartline, J. D., Roughgarden, T., 2009. Simple versus optimal mechanisms. In: Proceedings of the 10th ACM Conference on Electronic Commerce (EC). pp. 225–234.

Hartline, J. D., Yan, Q., 2011. Envy, truth, and optimality. In: Proceedings of the 12th ACM Conference on Electronic Commerce (EC). pp. 243–252.

Krishna, V., 2002. Auction Theory. Academic Press.

Ledyard, J. O., 2007. Optimal combinatoric auctions with single-minded bidders. In: Proc. 8th ACM Conf. on Electronic Commerce (EC). pp. 237–242.

Lehmann, D., O'Callaghan, L. I., Shoham, Y., 2002. Truth revelation in approximately efficient combinatorial auctions. Journal of the ACM 49 (5), 577–602.

Leonardi, S., Roughgarden, T., 2012. Prior-free auctions with ordered bidders. In: Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC). pp. 427–434.

Motwani, R., Raghavan, P., 1995. Randomized Algorithms. Cambridge University Press.

Myerson, R., 1981. Optimal auction design. Mathematics of Operations Research 6 (1), 58–73.

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

Oxley, J. G., 1992. Matroid Theory. Oxford.

Roughgarden, T., Talgam-Cohen, I., 2013. Optimal and near-optimal mechanism design with interdependent values. In: Proceedings of the 14th ACM Conference on Electronic Commerce (EC). pp. 767–784.

Roughgarden, T., Talgam-Cohen, I., Yan, Q., 2012. Supply-limiting mechanisms. In: Proceedings of the 13th ACM Conference on Electronic Commerce (EC). pp. 844–861.

Schrijver, A., 2003. Combinatorial Optimziation: Polyhedra and Efficiency. Springer.

Segal, I., 2003. Optimal pricing mechanisms with unknown demand. American Economic Review 93 (3), 509–529.

Wilson, R. B., 1987. Game-theoretic approaches to trading processes. In: Bewley, T. (Ed.), Advances in economic theory: Fifth world congress. Cambridge, pp. 33–77.