

Revenue Submodularity*

Shaddin Dughmi[†]

Tim Roughgarden[‡]

Mukund Sundararajan[§]

August 3, 2010

Abstract

We introduce *revenue submodularity*, the property that market expansion has diminishing returns on an auction’s expected revenue. We prove that revenue submodularity is generally possible only in matroid markets, that Bayesian-optimal auctions are always revenue-submodular in such markets, and that the VCG mechanism is revenue-submodular in matroid markets with i.i.d. bidders and “sufficient competition”. We also give two applications of revenue submodularity: good approximation algorithms for novel market expansion problems, and approximate revenue guarantees for the VCG mechanism with i.i.d. bidders.

1 Introduction

Auctions are often designed for a specific environment. But environments are not always predictable and static: initial expectations might be based on wrong information; bidders might withdraw or bring friends; and the seller can potentially influence the environment directly, for example by attracting new bidders to the market. For these reasons, the way the revenue of an auction *changes* with the underlying environment can be as important as the revenue it achieves in a fixed environment. For example, the VCG mechanism’s lack of revenue monotonicity — the fact that adding new bidders can decrease its revenue — has been widely cited as a “deal breaker” for its possible use as a combinatorial auction (see e.g. Milgrom [17, §2.5.2] and Rastegari et al. [23]).

This paper introduces *revenue submodularity* — essentially, the property that market expansion has diminishing returns on an auction’s expected revenue. For example, in a multi-unit auction with bidders that have unit demand and independent and identically distributed (i.i.d.) valuations, revenue submodularity means that the auction’s expected revenue is a concave function of the number of bidders. In general, an auction is deemed revenue submodular in an environment with

*Some of this work was presented at the 3rd Workshop on Sponsored Search, May 2007, under the title “Is Efficiency Expensive?”, and at the First Conference on Auctions, Market Mechanisms and Their Applications, May 2009, under the title “Revenue Maximization with Efficient Auctions”. A preliminary version of this paper appeared in the Proceedings of the 10th ACM Conference on Electronic Commerce, July 2009.

[†]Department of Computer Science, Stanford University, 460 Gates Building, 353 Serra Mall, Stanford, CA 94305. Supported in part by NSF grant CCF-0448664. Email: shaddin@cs.stanford.edu

[‡]Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, an AFOSR MURI grant, and an Alfred P. Sloan Fellowship. Email: tim@cs.stanford.edu

[§]Google Inc., Mountain View, CA. This work was done while the author was a PhD student in the Stanford Computer Science Department, and supported by NSF Award CCF-0448664 and a Stanford Graduate Fellowship. Email: mukunds@google.com.




Figure 1: Expected revenue of a 10-unit Vickrey auction, as a function of the number of bidders and the reserve price. When the reserve price is .2, the expected revenue is not concave.

potential bidders U if, for every subset $S \subset U$ and bidder $i \notin S$, the increase in the auction’s revenue from supplementing the bidders S by i is at most that of supplementing a set $T \subseteq S$ of bidders by the same additional bidder i .

Natural auctions are not necessarily revenue submodular, even in the simple setting of a multi-unit auction. Figure 1 shows the expected revenue of the Vickrey auction with a reserve price in a 10-unit auction, as a function of the number n of unit-demand bidders with valuations drawn i.i.d. from the uniform distribution on $[0, 1]$.¹ The three curves correspond to the reserve prices $r = .2, .5, .7$. The curve for $r = .5$ — the revenue-maximizing reserve price for this distribution — is noticeably concave. The curve for the high reserve is essentially linear in the range of the plot. The curve for the low reserve $r = .2$, however, is evidently non-concave, with a “kink” between 10 and 15 bidders. The first goal of this work is to understand this phenomenon, by *identifying necessary and sufficient conditions — on environments, valuation distributions, and auctions — such that revenue submodularity holds*.

But why is revenue submodularity an interesting property? We provide two applications in this paper (and anticipate more). The first application is inspired by a famous result of Bulow and Klemperer [4], which states that in multi-unit auctions with i.i.d. bidders, market expansion increases the Vickrey auction’s revenue at least as much as switching to an optimal selling procedure. This idea suggests the *market expansion problem*, which in its simplest form asks: which set of at most k bidders should be recruited to increase a given auction’s revenue as much as possible? Revenue submodularity is the key to achieving a computationally efficient approximation algorithm for this problem. As a second application, we show that revenue submodularity, in conjunction with

¹In this auction, the winners are the highest 10 bidders among those that meet the reserve, and all winners pay the larger of the reserve and the 11th highest bid.

additional conditions, leads to strong quantitative revenue guarantees for the economically efficient VCG mechanism

1.1 Brief Summary of Results

1.1.1 Revenue Submodularity

We first identify the largest class of single-parameter domains for which general revenue submodularity results are possible: *matroid markets*, in which the feasible subsets of simultaneously winning bidders form a matroid (see Section 2 for definitions). Fortunately, matroid markets include several interesting examples, including multi-unit auctions and certain matching markets.

We then prove a number of positive results (Section 3). First is a sweeping result for (Bayesian)-optimal auctions: in every matroid market with independent (not necessarily identical) valuation distributions, the revenue-maximizing auction is revenue-submodular. The fact that the “ $r = .5$ ” curve in Figure 1 is concave is a very special case of this result. The VCG mechanism, on the other hand, enjoys revenue submodularity only under additional conditions, even when valuations are i.i.d. draws from a well-behaved distribution. For example, in a k -unit auction with n unit-demand bidders, the Vickrey auction earns zero revenue when $n \leq k$ and positive revenue when $n \geq k + 1$, a clear violation of revenue submodularity. We identify a natural sufficient condition under which the VCG mechanism is revenue-submodular with i.i.d. bidders, which is a matroid rank condition stating that there is “sufficient competition” in the market. For example, in multi-unit auctions, sufficient competition requires that the number of bidders be at least the number of items. Finally, we prove that revenue submodularity is not a monotone property of the reserve prices used: reserve prices higher than those in an optimal mechanism preserve submodularity (cf., the “ $r = .7$ ” curve in Figure 1), but reserve prices strictly between those in the VCG mechanism (namely, zero) and those in an optimal mechanism always have the potential to destroy revenue submodularity, even when there is sufficient competition in the market (cf., the “ $r = .2$ ” curve in Figure 1).

We obtain reasonably simple and direct proofs of these results by appropriately applying two elegant and powerful techniques: Myerson’s characterization of expected auction revenue in terms of the expected “virtual surplus” of the auction’s allocation; and the submodularity that arises from optimizing a weight function over the independent sets of a matroid.

We also prove in passing that matroid markets are precisely the downward-closed single-parameter domains for which the VCG mechanism is always revenue *monotone*, meaning that additional bidders can only increase the mechanism’s revenue (Section 4).

1.1.2 Application: Market Expansion

Our first application of revenue submodularity is algorithmic and concerns the following problem. Under the constraint that a certain auction mechanism, such as the VCG mechanism, will be deployed, how should the seller recruit new bidders to maximize the auction’s revenue? This problem is clearly faced by sellers on eBay, by companies (like search engines) that run ad actions, and by governments that use spectrum auctions.

We focus on the following version of the *market expansion problem*. The input is a matroid market with a set of potential bidders, a subset of initial bidders, an auction (defined for all induced submarkets), and an expansion budget k . The goal is to recruit a set of at most k new bidders to maximize the expected revenue of the auction on the submarket induced by the original bidders together with the new recruits. The budget models constraints on the seller’s available marketing

resources for recruiting additional bidders. This problem is trivial with i.i.d. bidders in a multi-unit auction, but we prove that it is hard in more general settings. Our main result for this application is that “greedy market expansion” — repeatedly adding the new bidder that (myopically) increases the expected revenue of the auction as much as possible — is a constant-factor approximation algorithm provided the given auction is revenue-submodular over all sets containing the initial bidders. This result also admits several extensions, for example to the variant in which the budget on new recruits is replaced by bidder recruiting costs.

1.1.3 Application: VCG Revenue Is Approximately Optimal

The VCG mechanism maximizes welfare and does not require knowledge about bidders’ valuation distributions. However, its payments are designed to enforce strategyproofness and generate revenue only as a side effect. Are there general conditions under which the VCG mechanism is guaranteed to have good revenue?

Our second application of revenue submodularity is to approximate revenue-maximization guarantees for the VCG mechanism. Specifically, in a matroid market with i.i.d. bidder valuations and “modest competition” — which we formalize using matroid connectivity — the VCG mechanism always obtains a constant fraction of the revenue of an optimal auction. Moreover, the approximation guarantee tends rapidly to 1 as the degree of competition increases. We also extend these guarantees to a standard model of pay-per-click sponsored search auctions. These results suggest an explanation for the persistent use of economically efficient auctions for revenue-maximization problems:² the cost (i.e., revenue loss) of running an efficient auction is typically small and outweighed by the benefits (i.e., economic efficiency and relative simplicity), even for a revenue-maximizing seller.

1.2 Related Work

To the best of our knowledge, we are the first to study revenue submodularity in auctions and to consider market expansion optimization problems. A few other works study revenue guarantees for the VCG mechanism, in simpler settings than ours. Bulow and Klemperer [4] give a sufficient condition on the number of additional bidders required to exceed the (original) optimal revenue in a multi-unit auction, but do not compare the Vickrey and optimal revenue in a fixed environment. Similar results for more general settings appear in [13]. Neeman [19] studies the convergence of revenue to welfare as the number of bidders grows in a single-item auction with i.i.d. bidders from a bounded distribution. Lambert and Shoham [16] study which sponsored search auctions extract the full surplus in the limit, as the number of bidders goes to infinity. Edelman and Schwarz [8] empirically compare the revenue of the VCG and revenue-maximizing auctions in a sponsored search context.

A few other papers study matroids in auction settings but are largely unrelated to our work. Talwar [24] and Karlin et al. [14] study frugality and Cary et al. [5] study profit-maximization in procurement settings. Our proof that matroid markets are precisely those in which the VCG mechanism is always revenue monotone (Theorem 4.1) shares some ideas with arguments in [14, 24] (see also the notes by Hartline [12]), but the techniques in [5, 14, 24] do not seem useful for studying

²For example, current keyword search auctions are modeled after efficient auctions. While these auctions do employ reserve prices, as revenue-maximizing auctions do, these reserves are relatively close to 0 and are apparently meant to discourage irrelevant ads rather than to (directly) maximize revenue — see [21] for more details.

revenue submodularity. Bikhchandani et al. [3] design economically efficient ascending auctions for selling bases of a matroid, but are unconcerned with revenue. For online auctions, where bidders arrive over time, matroid domains are studied in Babaioff et al. [2] and Constantin et al. [6].

2 Preliminaries

This section reviews some standard facts from combinatorial optimization and auction theory that are needed to state and prove our results. Section 2.1 introduces matroids and gives several examples. Section 2.2 states the key facts about matroids that are required in our proofs. Section 2.3 reviews optimal auction design in Bayesian single-parameter environments, as studied in Myerson [18]. We encourage the reader familiar with matroids and auctions to skip ahead to Example 2.11 at the end of the section.

2.1 Matroids

A *set system* consists of a *ground set* U and a collection $\mathcal{I} \subseteq 2^U$ of subsets. We will only be interested in the case where U is finite and both \mathcal{I} and U are non-empty. A *matroid* is a set system (U, \mathcal{I}) that satisfies two properties. First, it is *downward closed*, meaning that if S belongs to \mathcal{I} , then so do all subsets of S . The second condition is the *Exchange Property*, which asserts that whenever $T, S \in \mathcal{I}$ with $|T| < |S|$, there is some $x \in S \setminus T$ such that $T \cup \{x\} \in \mathcal{I}$. Thus T can be extended to a larger set in \mathcal{I} by some element of $S \setminus T$. In a matroid context, the sets of \mathcal{I} are called *independent*, and the maximal such sets are the *bases* of the matroid. The matroid properties easily imply that all bases have equal cardinality. This common size is the *rank* of the matroid.

Matroids model a number of natural auction settings; we mention a few below. In all cases, the ground set of the matroid represents the set of bidders in the auction, and the independent sets of the matroid represent the subsets of bidders that can simultaneously win in the auction.

Example 2.1 (Uniform Matroids) In a *uniform matroid*, the independent sets are the subsets of size at most k , where k is some nonnegative integer. Both matroid properties obviously hold. A uniform matroid models an auction with k identical units of a good and unit-demand bidders. The bases of a uniform matroid are the subsets of size exactly k ; obviously, k is also the rank of the matroid.

Uniform matroids model simple k -unit auctions, where the seller wants to sell at most one item to each bidder.

Example 2.2 (Transversal Matroids) A *transversal matroid* is defined via an undirected bipartite graph (V_1, V_2, E) ; its ground set is V_1 and a subset $S \subseteq V_1$ is independent if and only if the vertices of S can be simultaneously matched to (distinct) vertices of V_2 . The Exchange Property can be proved using an augmenting path argument. The rank of a transversal matroid equals the cardinality of a maximum matching in the corresponding bipartite graph. A transversal matroid represents a matching market, where V_1 is a set of bidders, V_2 is a set of goods, and the edges E specify which goods each bidder is interested in. (Here, each bidder has a common value for the goods in which it is interested.)

Transversal matroids model auction settings where each of the winning bidders is matched to a unique “resource” or “service” from a bidder-specific set of viable resources.

Example 2.3 (Graphic Matroids) A *graphic matroid* is defined by an undirected graph $G = (V, E)$; the ground set is E and the independent sets are the acyclic subsets of E . Such a set system is obviously downward closed, and the Exchange Property can be proved by comparing partitions into connected components. If G is a connected graph, then the bases of the corresponding graphic matroid are the spanning trees of G , and the matroid rank is $|V| - 1$.

We require a standard matroid operation to model the addition (or removal) of new bidders in a market. Given a matroid $M = (U, \mathcal{I})$ and a subset $S \subseteq U$, the *restriction of M to S* is the set system (S, \mathcal{I}_S) , where $\mathcal{I}_S = \{T \subseteq S : T \in \mathcal{I}\}$ is the subsets of \mathcal{I} that lie in S . Every restriction of a matroid to a non-empty set is again a matroid. We sometimes call this a *submatroid* of the original matroid or say that S *induces* the matroid (S, \mathcal{I}_S) .

2.2 Submodularity and Weighted Rank

Suppose we endow every element e of a matroid M with a real-valued *weight* w_e . The weight of a set is then the sum of the weights of its constituent elements. The *weighted rank* of M under w is defined as the maximum weight of one of its independent sets. For a nonnegative weight function the weighted rank is determined by a basis of the matroid; with general weights, non-maximal independent sets can determine the weighted rank.

The weighted rank of a matroid can be computed by the following algorithm which greedily constructs an independent set via a single pass over the elements: (i) sort the elements e_1, \dots, e_n from highest to lowest weight (breaking ties arbitrarily) and initialize $S = \emptyset$; (ii) for each $i = 1, 2, \dots, n$ in turn, if e_i has nonnegative weight and $S \cup \{e_i\} \in \mathcal{I}$, then add e_i to S . For graphic matroids, this algorithm is simply Kruskal’s algorithm (e.g. [15, §2.1]). The correctness of this algorithm for general matroids can be proved using the Exchange Property. Similarly, if the clause “if e_i has nonnegative weight” is omitted, then the corresponding greedy algorithm computes a maximum-weight basis of the matroid.

Submodularity is a set-theoretic analog of concavity, and it is central to this work. We repeat here the formal definition.

Definition 2.4 (Submodular Function) A function $f : 2^U \rightarrow \mathbb{R}$ defined on all subsets of a finite non-empty set U is *submodular* if

$$f(S \cup \{i\}) - f(S) \leq f(T \cup \{i\}) - f(T)$$

for every $T \subseteq S \subset U$ and $i \notin S$.

Our results on revenue submodularity rely on the submodularity of the weighted rank function on the submatroids of a given matroid. For the uniform weight function, this fact is well known.

Proposition 2.5 ([28, Theorem 1.2.3]) For a matroid $M = (U, \mathcal{I})$, let $f(S)$ denote the rank of M restricted to S . Then f is a submodular function.

Proof: Fix sets T, S such that $T \subseteq S \subseteq U$, and an element $i \in U \setminus S$; we show that $f(S \cup \{i\}) - f(S) \leq f(T \cup \{i\}) - f(T)$. The left-hand side is either 0 or 1, and we can assume that it is 1. Let B_T be a basis in the restriction of M to T , and use the Exchange Property to extend B_T to a basis $B_S \supseteq B_T$ in the restriction of M to S . Since $f(S \cup \{i\}) = f(S) + 1$, the Exchange Property implies

that $B_S \cup \{i\}$ is a basis in the restriction of M to $S \cup \{i\}$. Downward closure implies that $B_T \cup \{i\}$ is an independent set of M restricted to $T \cup \{i\}$, so $f(T \cup \{i\}) = f(T) + 1$, as needed. ■

The generalization to weighted rank follows easily.

Corollary 2.6 (Weighted Rank Is Submodular) *For a matroid $M = (U, \mathcal{I})$ and weight function w on U , let $f(S)$ denote the weighted rank of M restricted to S . Then f is a submodular function.*

Proof: First, Proposition 2.5 extends trivially to 0-1 weight functions, since the corresponding weighted rank function of M is the same as the (unweighted) rank function of the restriction of M to the positive-weight elements of U . For a general weight function w , restrict M to the positive-weight elements and label these elements $\{1, 2, \dots, n\}$ so that $w_1 \geq w_2 \geq \dots \geq w_n > 0$. For $i = 1, 2, \dots, n$, let $w^{(i)}$ denote the weight vector in which the elements $\{1, 2, \dots, i\}$ have weight 1 and all other elements have zero weight. Let f_i denote the corresponding (submodular) weighted rank function. Since a nonnegative linear combination of submodular functions is again submodular, $\sum_{i=1}^n (w_i - w_{i+1}) \cdot f_i$ is submodular. (By convention, $w_{n+1} = 0$.) Since the output of the greedy algorithm for maximizing weighted rank depends only on the ordering of the elements' weights (and ties are irrelevant), the independent set that maximizes the weighted rank w.r.t. the weight function w also simultaneously maximizes the weighted rank w.r.t. each of $w^{(1)}, \dots, w^{(n)}$. Thus $f = \sum_{i=1}^n (w_i - w_{i+1}) \cdot f_i$, completing the proof. ■

The following converse to Proposition 2.5 also holds.

Proposition 2.7 ([28, Theorem 1.2.3]) *Let $M = (U, \mathcal{I})$ be a set system and $f(S)$ the largest size of a set of \mathcal{I} that is contained in S . If f is submodular, then M is a matroid.*

There are numerous other characterizations of matroids. We conclude with one useful in the proof of Theorem 4.1.

Proposition 2.8 ([22, Corollary 2.1.5]) *A downward-closed set system (U, \mathcal{I}) with $\mathcal{I} \neq \emptyset$ is a matroid if and only if for every pair A, B of maximal sets in \mathcal{I} and $y \in B$, there is some $x \in A$ such that $A \setminus \{x\} \cup \{y\} \in \mathcal{I}$.*

2.3 Optimal Auction Design

Our auction model is standard (e.g. [18]). There is a population U of bidders, and a set of feasible outcomes, each indicating the “winning” and “losing” bidders in the outcome. For example, in a k -unit auction, there is a feasible outcome for each subset of at most k bidders; in a matching market, feasible outcomes correspond to matchings between bidders and desired goods. The applications we have in mind are *downward closed*; as noted above, this means that for every feasible set of winners, every subset of this set is also feasible.

A bidder i has value 0 for losing, and a *valuation* v_i for winning that is a priori unknown to the auctioneer. Each bidder bids to maximize its payoff $v_i x_i - p_i$, where x_i is 1 if it wins and 0 otherwise, and p_i is its payment to the auctioneer (assumed 0 if $x_i = 0$). A *mechanism* is a specification of an *allocation rule* (the x_i 's) and a *payment rule* (the p_i 's), which together select who wins and who pays what in each bid profile.

The *efficiency* and *revenue* of a mechanism outcome for bidders with valuations v are defined as $\sum_i v_i x_i$ and $\sum_i p_i$, respectively. We usually study the *expected* revenue of a mechanism, under the assumption that the bidders' valuations are independently distributed according to known distributions F_1, \dots, F_n with strictly positive density functions f_1, \dots, f_n . We focus on *strategyproof* mechanisms, in which each bidder is guaranteed to maximize its payoff by revealing its true private valuation to the mechanism, irrespective of the valuations and behavior of the other bidders. Because of this, we use the terms *bids* and *valuations* interchangeably.

In single-parameter problems like those studied in this paper, truthful mechanisms are relatively well understood. An allocation rule can be extended via a (essentially unique) payment rule to a strategyproof mechanism if and only if the rule is *monotone*, meaning that a winner who increases its bid always continues to win (keeping other bids fixed) [18]. For example, in a single-item auction, the “highest bidder wins” rule is monotone; the “second-highest bidder wins” rule is not.

Example 2.9 (The VCG Mechanism) The *VCG mechanism* is defined by the allocation rule that always picks the feasible set with the largest sum of valuations. It is easy to see that this allocation rule is monotone and can therefore be extended to a strategyproof mechanism via suitable payments. These payments are: each winner is charged a price equal to the smallest bid at which it would continue to win (keeping other bids fixed). For a k -unit auction with unit-demand bidders, the VCG mechanism specializes to the Vickrey auction, with all winners paying the $(k+1)$ th highest bid.

Remark 2.10 (The Clarke Pivot Rule) The payments of the general VCG mechanism are defined only up to a bid-independent “pivot term” (see e.g. [17]). In this paper, we study only the *Clarke pivot rule* which normalizes the payment of every losing bidder to zero.

The *virtual valuation* corresponding to a distribution F_i and a valuation v_i is defined as

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

A distribution is *regular* if the corresponding virtual valuation function is increasing over the distribution's support. We note that a virtual valuation can be negative; for example, if $F(v) = v$ on $[0, 1]$, then $\varphi(v) = 2v - 1$.

The importance of virtual valuations is illustrated by the following lemma of Myerson [18, Lemma 3.1]: for *every* mechanism, its expected revenue (over draws from the F_i 's) equals the expected virtual value of its allocation:

$$\mathbf{E}[\text{revenue}] = \int \left(\sum_i \varphi_i(v_i) x_i(v_1, \dots, v_n) \right) f_1(v_1) \cdots f_n(v_n) dv_1 \cdots dv_n. \quad (2)$$

Thus the revenue-maximizing strategyproof mechanism maximizes the expected total virtual value (2) subject to monotonicity of the allocation rule [18].

With regular distributions, the optimal auction simply maximizes the virtual value pointwise (i.e., separately for each valuation profile). Since the virtual valuation functions corresponding to regular distributions are increasing, this defines a monotone allocation rule and yields a strategyproof mechanism once suitable payments are defined.

Finally, we discuss non-regular distributions; this requires some technical concepts, but we require them only for the proof of Theorem 3.1. With such distributions, the allocation rule above

is not monotone and hence cannot be extended to a strategyproof mechanism by any payment rule. To overcome this obstacle, Myerson [18] defined a nondecreasing function called an *ironed virtual valuation*. This function is meant to be a monotone proxy for the virtual valuation function. Maximizing the ironed virtual value of the allocation and breaking ties in a valuation-independent way turns out to be equivalent to maximizing the virtual value of the allocation (2) subject to monotonicity, and therefore results in an optimal auction [18].

Formally, the (nondecreasing) ironed virtual valuation $\bar{\varphi}$ corresponding to a virtual valuation φ is defined by the following procedure.

1. For $q \in [0, 1]$, define $h(q) = \varphi(F^{-1}(q))$.
2. Define $H(q) = \int_0^q h(r)dr$.
3. Define G as the convex hull of H – the largest convex function bounded above by H for all $q \in [0, 1]$.
4. Define $g(q)$ as the derivative of $G(q)$, where defined, and extend to all of $[0, 1]$ by right-continuity.
5. Finally, $\bar{\varphi}(z) = g(F(z))$.

To summarize, in a single-parameter environment with independent valuation distributions, maximizing the expected auction revenue reduces to always selecting the feasible set with maximum total ironed virtual value.

Example 2.11 (Matroid Markets with i.i.d. Bidders) Consider a set U of bidders and suppose that the feasible sets of winners form a matroid. Suppose further that bidders' valuations are i.i.d. draws from a regular distribution F with (increasing) virtual valuation function φ . In this case, the VCG and revenue-maximizing mechanisms are close cousins. The VCG mechanism maximizes the total value of the winners (Example 2.9). In a matroid market, this allocation rule can be implemented by ordering bidders by valuation and running the greedy algorithm from Section 2.2. As discussed above, the optimal mechanism maximizes the total virtual value of the winners. In a matroid market, this corresponds to ordering bidders by virtual valuation, running the greedy algorithm, and halting once the negative virtual valuations are reached. Since bidders' valuations are i.i.d. draws from a regular distribution, the orderings by valuation and by virtual valuation coincide. Thus, the optimal mechanism is nothing more than the VCG mechanism supplemented with the optimal reserve price $r^* = \varphi^{-1}(0)$.

3 Revenue Submodularity

3.1 Optimal Auctions

We first observe that revenue submodularity can be crisply characterized for revenue-maximizing auctions. This characterization also establishes matroid domains as the largest set of domains for which general revenue-submodularity results are possible.

We study the following property of a given domain or market (i.e., a set of bidders and feasible subsets of winners):

(*) for every set of independent valuation distributions for the bidders, the corresponding optimal auction is revenue-submodular.

Theorem 3.1 (Submodularity of Optimal Auctions) *A market has property (*) if and only if it is a matroid market.*

Proof: For the “if” direction, fix a matroid market with bidders U and independent distributions F_1, \dots, F_n for the bidders’ valuations. Condition on the valuations of all bidders in U . For $S \subseteq U$, let $\bar{\varphi}(S)$ denote the maximum sum of ironed virtual valuations (Section 2.3) possessed by an independent set contained in S ; by Myerson’s Theorem [18], the optimal auction for the market induced by S chooses such a set for this valuation profile. Since the market on U is a matroid market, Corollary 2.6 implies that $\bar{\varphi}(S)$ is submodular on U . Taking expectations, the expected sum of the winners’ ironed virtual valuations is submodular on U . (A convex combination of submodular functions is again submodular.) Applying (2) now shows that the optimal auction is revenue-submodular.

For the converse, consider a domain for which the optimal auction is always revenue-submodular, and in particular has this property when every bidder’s valuation is deterministically 1. In this case, the revenue of the optimal auction in a submarket $S \subseteq U$ is the largest size of a feasible set contained in S . In other words, the auction revenue corresponds to the rank function of the set system. This rank function is submodular only when the set system comprises the independent sets of a matroid (Proposition 2.7). ■

3.2 VCG Without a Reserve Price

The plot is thicker for other mechanisms, even in the very special case of multi-unit auctions. (Recall Figure 1.) For example, consider a k -unit auction with n unit-demand bidders. For all $n \leq k$, the Vickrey auction earns zero revenue. There is a sudden jump to positive revenue when $n = k + 1$, a clear violation of revenue submodularity. With non-i.i.d. bidders (e.g., many “small” bidders and few “large” bidders), the same problem can arise even when the total number of bidders is much larger than the number of goods. The best we can hope for with the VCG mechanism is that revenue submodularity kicks in once there is “sufficient competition”. Precisely, we consider i.i.d. bidders and prove submodularity over the *full-rank* sets — sets that contain a basis of the full matroid market M . In a k -unit auction, this corresponds to bidder sets of cardinality at least k .

Theorem 3.2 (Submodularity of VCG) *Fix a matroid market M with bidders U and valuations drawn i.i.d. from a regular distribution F . The expected revenue of the VCG mechanism for induced matroid markets M_S is submodular on the set of full-rank sets $S \subseteq U$.*

Theorem 3.2 is false, even in a single-item auction, if the full-rank assumption is dropped (as we have seen), if the i.i.d. assumption is dropped (Example 3.3), or if the regularity condition is dropped (Example 3.4).

Proof: Condition on the bidders’ valuations. Let C be large enough that the “shifted virtual valuation” $\gamma(v_i) := \varphi(v_i) + C$ is nonnegative for every bidder i . Since the valuations are i.i.d. draws from a regular distribution, the (common) virtual valuation function φ is increasing, and thus γ is a nonnegative weight vector that orders the bidders by valuation.

In a matroid market, all nonnegative weight functions that order the bidders in the same way are maximized by a common (maximal) independent set; this follows from the optimality of the greedy

algorithm, as in the proof of Corollary 2.6. Thus while the VCG mechanism explicitly maximizes the sum of the valuations of the winners, it inadvertently maximizes the sum of their shifted virtual valuations as well. Letting $\gamma(S)$ denote the latter maximum in the submarket induced by S , this observation and Corollary 2.6 establish the submodularity of γ over all subsets of U .

Define $\varphi(S)$ as the total virtual value of the VCG mechanism's allocation in the submarket S . The submodularity of γ translates to submodularity of φ on full-rank sets. To see this, take two full-rank sets A and B with $A \subseteq B$, and a bidder $i \notin B$. By submodularity of γ ,

$$\gamma(A \cup \{i\}) - \gamma(A) \geq \gamma(B \cup \{i\}) - \gamma(B).$$

Now, on full-rank sets S , $\varphi(S) = \gamma(S) - r(M) \cdot C$, where $r(M)$ denotes the rank of the full matroid M . Thus

$$\varphi(A \cup \{i\}) - \varphi(A) \geq \varphi(B \cup \{i\}) - \varphi(B),$$

as claimed.

Finally, taking expectations over the bidders' valuations and applying (2) proves the theorem. ■

Example 3.3 (Necessity of i.i.d. Distributions) Consider a single-item auction where the bidding population consists of two types of bidders, one with value drawn uniformly from $[0, \epsilon]$ (small bidders) for small ϵ and another with value drawn uniformly from $[0, 1]$ (big bidders). Suppose that the market consists initially of small bidders. Consider the revenue of the Vickrey auction. Adding the first big bidder causes an increase in expected revenue of at most ϵ . Adding a second big bidder increases the expected revenue by at least $1/3 - \epsilon$. Thus non-submodularity occurs even on full-rank sets.

Example 3.4 (Necessity of Regular Distributions) Consider a single-item auction where the valuation of each bidder is 1 with probability p , and 0 with probability $1 - p$. (A continuous perturbed version of this distribution also works.) The revenue of the Vickrey auction is 0 with one bidder. Adding the second bidder increases the revenue to p^2 . Adding the third bidder increases the revenue to $p^3 + 3p^2(1 - p)$. For small p , $p^3 + 3p^2(1 - p) - p^2 > p^2 - 0$ and non-submodularity results.

3.3 VCG with an Arbitrary Reserve Price

Consider the VCG mechanism with some reserve price r . Thus far, we have identified conditions for revenue submodularity with zero reserve and with the Myerson reserve r^* , which in matroid markets correspond to the VCG and optimal mechanisms, respectively (Example 2.11). What if a different reserve price is used, either by choice or because of inaccurate statistics? Perhaps surprisingly, there is a big difference between overestimating the optimal reserve price (which never affects submodularity) and underestimating it (which can destroy submodularity, even on full-rank sets).

Theorem 3.5 (VCG with Incorrect Reserve Prices)

- (a) *For every regular distribution F with optimal reserve price $r^* = \varphi^{-1}(0)$, every matroid market with bidders U with valuations drawn i.i.d. from F , and every $r \geq r^*$, the expected revenue of the VCG mechanism with reserve price r is submodular on U .*

- (b) For every $\epsilon \in (0, 1)$, there is a regular distribution F with optimal reserve price r^* and a matroid market for which the expected revenue of the VCG mechanism with reserve price $(1 - \epsilon)r^*$ is not submodular on full-rank sets.

Proof: For part (a), condition on the bidders' valuations. Restrict the independent sets of the matroid to the bidders that meet the reserve price. Since $r \geq r^*$ and F is regular, all such bidders have nonnegative virtual valuations. As in the proof of Theorem 3.2, regularity implies that the VCG mechanism with reserve price r inadvertently maximizes the total virtual value over the independent sets of the restricted matroid. Thus the virtual value of the mechanism's allocation is the weighted rank of a matroid. As in previous proofs, applying Corollary 2.6, taking expectations over valuations, and invoking (2) establishes revenue-submodularity.

Part (b) can be established using the distribution F that is an equal (50/50) mixture of the uniform distribution on $[0, 1]$ and the "equal revenue distribution" with distribution function $1 - 1/x$ on $[1, \infty)$. This distribution is continuous and regular, and its optimal reserve price is 1. Fix an arbitrarily small constant $\epsilon > 0$ and consider a graphic matroid comprising a cycle of length n (for n large) plus one parallel copy e' of one of the edges e . Let $q = (1 + \epsilon)/2$ denote the probability that a sample from F exceeds the reserve price $r = 1 - \epsilon$. When both e and e' are absent, the expected revenue of the VCG mechanism with this reserve is $q(n - 1)(1 - \epsilon)$, with every winner paying the reserve price. For n large, when one of $\{e, e'\}$ is added, the expected revenue increases by $\approx q(1 - \epsilon)$. When both parallel edges are added, there are two relevant cases (for n large). The valuations of e, e' are both above 1 with probability $1/4$, and the additional revenue in this case is the expected minimum of two samples from the equal revenue distribution, which is 2. In the other relevant case, at least one of the valuations of e, e' exceeds the reserve price and at most one of them exceeds 1. This occurs with probability $(1 - (1 - q)^2) - 1/4 = 1/2 + \epsilon - \epsilon^2/4$, and the additional revenue in this case is at least $1 - \epsilon$. A quick calculation shows that the combined extra revenue from the two cases strictly exceeds $2q(1 - \epsilon)$. Since the revenue increase from adding e and e' is more than double that of adding either one individually, we have a violation of revenue-submodularity that involves only full-rank subsets of the matroid. ■

4 Revenue Monotonicity of the VCG Mechanism in Matroids

We note in passing an interesting analog of Theorem 3.1 for the revenue *monotonicity* of the VCG mechanism (with the Clarke pivot rule). Precisely, a mechanism is *revenue monotone* in a single-parameter downward-closed domain (U, \mathcal{I}) if, for every set of bidder valuations v and set $S \subseteq U$, the mechanism's revenue in the full market is at least that in the market induced by S and v .

Theorem 4.1 (Monotonicity of VCG) *The VCG mechanism is revenue monotone in a downward-closed market if and only if the market is a matroid market.*

Proof: For the "if" direction, fix a matroid market $M = (U, \mathcal{I})$ and valuations v for the bidders. By induction, we can consider only sets S that exclude a single bidder e of U . Breaking ties using bidders' names, we can treat the valuations as distinct.

Recall that in a matroid market the VCG mechanism's allocation can be computed by the greedy algorithm (Section 2.2). By the Exchange Property, the greedy algorithm maintains the invariant that, after processing a subset T of the bidders, it has selected the maximum-possible number $r(T)$ of winners. Thus, a bidder i wins if and only if $r(T \cup \{i\}) = r(T) + 1$, where T

is the set of bidders considered before i by the greedy algorithm. Let T_j denote the j bidders of $U \setminus \{i\}$ with the highest valuations. Since r is submodular, there is some $\ell \in 1, 2, \dots, |U|$ so that $r(T_j \cup \{i\}) = r(T_j) + 1$ for all $j < \ell$ and $r(T_j \cup \{i\}) = r(T_j)$ for all $j \geq \ell$. Set p_i equal to the smallest valuation of T_ℓ , or to 0 if $\ell = |U|$. Then, bidder i wins if and only if $v_i > p_i$, in which case it pays p_i (recall Example 2.9).

Now suppose we delete bidder e from M . Let W and W' denote the winners under the VCG mechanism in M and $M \setminus \{e\}$, respectively. For every $i \neq e$, since r is submodular, $r(T_j \setminus \{e\} \cup \{i\}) = r(T_j \setminus \{e\})$ only when $r(T_j \cup \{i\}) = r(T_j)$. Thus, the “threshold index” ℓ for bidder $i \neq e$ is only larger after e ’s deletion, so its “threshold price” p_i is only smaller. Thus, all bidders of $W \setminus \{e\}$ belong to W' and pay lower prices. If W' contains no bidders not in W , then we are done. Otherwise, since $r(M \setminus \{e\}) \leq r(M)$ and the VCG mechanism chooses a basis, we can write $W' = W \setminus \{e\} \cup \{f\}$ for some bidder $f \notin W$. Bidder f contributes revenue at most v_f to the VCG mechanism in $M \setminus \{e\}$. We conclude by noting that the revenue p_e contributed by e to the VCG mechanism in M is at least v_f : if bidder e bids less than v_f , then the VCG mechanism in M will choose the allocation W' instead of W .

For the “only if” direction, consider a non-matroid downward-closed market $M = (U, \mathcal{I})$. By Proposition 2.8, there are maximal sets A, B of \mathcal{I} and an element $y \in B$ such that $A \setminus \{x\} \cup \{y\} \notin \mathcal{I}$ for every $x \in A$. Suppose that the bidders of $A \cup \{y\}$ have valuation 1 and all other bidders have valuation 0. The set A maximizes welfare over the sets of \mathcal{I} (with welfare $|A|$) and, by our choice of A and y , every other set of \mathcal{I} has welfare at most $|A| - 1$. The VCG mechanism thus generates zero revenue in this market. We complete the proof by identifying a submarket in which the VCG mechanism earns strictly positive revenue.

Since \mathcal{I} is downward closed, we have $A \cup \{y\} \notin \mathcal{I}$ and $(A \cap B) \cup \{y\} \in \mathcal{I}$. We can therefore choose a set $A' \supseteq A \cap B$ and an element $x \in A \setminus B$ such that $A' \cup \{x\} \cup \{y\} \notin \mathcal{I}$ and $A' \cup \{y\} \in \mathcal{I}$. In the matroid market induced by $A' \cup \{x, y\}$, there are at least two welfare-maximizing solutions, $A' \cup \{x\}$ and $A' \cup \{y\}$. The VCG mechanism chooses an allocation that includes either x or y (or both), and will collect a payment of 1 from this bidder. ■

Unlike the other results in this paper, Theorem 4.1 is not stated for distributions over bidders’ valuations. Pointwise revenue monotonicity (as in Theorem 4.1) obviously implies expected revenue monotonicity with respect to every distribution over bidders’ valuations.

5 Near-Optimal Market Expansion

Given a market and a mechanism for it, and also an initial submarket, which k additional bidders should be recruited to generate the largest increase in the auction’s revenue? This question is trivial when bidders are indistinguishable, as in a multi-unit auction with i.i.d. bidder valuations: any k additional bidders will do. As the next example shows, this basic optimization problem becomes quite subtle with bidder asymmetries.

Example 5.1 (Expanding a Graphic Matroid) Consider a graphic matroid market $G = (V, U)$, and suppose that the initial submarket S is a spanning tree of G , that $k = 1$, and that the mechanism used is the VCG mechanism. Suppose bidders’ valuations are i.i.d. draws from an exponential distribution with rate 1. Adding a new bidder creates a cycle, say of length ℓ . Once valuations have been sampled, the VCG mechanism will select all bidders but the lowest one i on the cycle, and the other $\ell - 1$ bidders of the cycle will each be charged v_i . (Bidders off the cycle are charged 0.)

The expected revenue of the VCG mechanism (over the random valuations) is $(\ell - 1)/\ell$, since $1/\ell$ is the expected value of the minimum of ℓ independent exponential random variables. Thus, in this instance, the optimal solution to the market expansion problem is to add the edge of $U \setminus S$ that creates the longest cycle.

The market expansion problem is inapproximable without revenue submodularity, for example in matroid markets with the VCG mechanism and without a full-rank initial submarket.

Theorem 5.2 *The market expansion problem for general matroids and the VCG mechanism admits no polynomial-time algorithm with non-zero approximation ratio, unless $P = NP$.*

Proof: Vardy [25] proved that finding a cycle of a given length in a binary matroid — i.e., linearly independent subsets of vectors over the two-element field — is NP -hard. Given such a matroid and the empty initial market, there are k bidders whose recruitment generates positive revenue for the VCG mechanism if and only if the matroid has a cycle of length at most k . ■

Even with revenue submodularity, as in graphic matroids with the revenue-maximizing auction (Theorem 3.1), the market expansion problem is NP -hard.

Theorem 5.3 *Optimal market expansion is NP -hard, even for graphic matroids, i.i.d. valuations, and the revenue-maximizing auction.*

Proof: We provide a reduction from the Hamiltonian cycle problem [10, Problem GT37]. Given a connected graph with n vertices, consider the corresponding graphic matroid, the empty submarket, and a budget of n new bidders to recruit. Bidders' valuations are (say) i.i.d. draws from the exponential distribution with rate 1. By (2), the expected revenue in a given submarket is the expected maximum virtual value of an acyclic subgraph. For a given subgraph with n edges and valuation profile, the maximum virtual value of an acyclic subgraph is simply the sum of the positive virtual valuations, unless these bidders include a cycle, in which case some positive virtual valuations must be thrown out. Hamiltonian cycles (if any) are the submarkets that minimize the expected amount of positive virtual value so wasted, and thus are the submarkets that maximize expected revenue. ■

Revenue submodularity leads directly to a positive result for near-optimal market expansion. By *greedy market expansion*, we mean the heuristic of repeatedly (k times) recruiting the new bidder that increases the expected revenue of an auction the most. Our next result follows from a classic analysis of Nemhauser, Wolsey, and Fisher [20]. They showed that for every nonnegative, monotone, and submodular set function f on a universe U , this greedy heuristic outputs a set S' that is a $(1 - 1/e)$ -approximation to the maximum-value subset of U of size at most k : $f(S') \geq (1 - 1/e) \cdot \max_{|S| \leq k} f(S)$.

Theorem 5.4 *Greedy market expansion is a $(1 - 1/e)$ -approximation algorithm for the market expansion problem whenever the given auction is revenue monotone and revenue submodular on all submarkets containing the initial market.*

For example, since the revenue-maximizing auction with arbitrary (not necessarily regular or i.i.d.) independent distributions is obviously expected revenue monotone, Theorems 3.1 and 5.4 imply that greedy market expansion is a $(1 - 1/e)$ -approximation algorithm in every matroid market.

For the VCG mechanism and bidder valuations that are i.i.d. draws from a regular distribution, Theorems 3.2, 4.1, and 5.4 imply that greedy market expansion is a $(1 - 1/e)$ -approximation algorithm in matroid markets with a full-rank initial market.

Extensions Extensions of Theorem 5.4 are easy to come by. If only an α -approximation algorithm is available for the subroutine that chooses the optimal next bidder to add — for example, due to sampling error in estimating the expected revenue of a mechanism in a given submarket — the approximation bound degrades only to $1 - 1/e^\alpha$ (see e.g. [11]).

Theorem 5.5 *When using a subroutine that is an α -approximation algorithm for the problem of choosing the optimal next bidder, greedy market expansion is a $(1 - 1/e^\alpha)$ -approximation algorithm for the market expansion problem whenever the given auction is revenue monotone and revenue submodular on all submarkets containing the initial market.*

The budget of k can be replaced by an arbitrary matroid constraint on the bidders of $U \setminus S$ without changing the approximation guarantee [27]. For example, the feasible recruitable sets might correspond to assignments of a fixed number of recruiters to different locations subject to geographic constraints (a transversal matroid).

Theorem 5.6 *Greedy market expansion is a $(1 - 1/e)$ -approximation algorithm for the market expansion problem subject to a matroid constraint whenever the given auction is revenue monotone and revenue submodular on all submarkets containing the initial market.*

As a third extension, we can attach a fixed recruiting cost to each bidder e . The objective is then to maximize revenue minus recruiting costs. Adding costs can ruin revenue monotonicity but does not affect submodularity. As long as the revenue submodularity condition in Theorem 5.4 holds and the profit earned when recruiting the entire market is nonnegative, the market expansion problem can be approximated to within a factor of $2/5$ using randomized local search [9].

Theorem 5.7 *There is a randomized local search algorithm that is a $2/5$ -approximation algorithm for the market expansion problem with recruiting costs whenever the given auction is revenue monotone and revenue submodular on all submarkets containing the initial market, and recruiting the entire market guarantees nonnegative profit.*

6 Revenue Guarantees for the VCG Mechanism

6.1 Matroid Markets

Are there interesting conditions under which the VCG mechanism inadvertently yields near-optimal revenue? In matroid markets, even “modest competition” suffices for such a guarantee. We quantify competition via the *packing number* of the matroid, defined as the maximum number of disjoint bases that the matroid contains.

Theorem 6.1 (A Guarantee for the VCG Mechanism’s Revenue) *In every matroid market $M = (U, \mathcal{I})$ with packing number κ and bidders’ valuations drawn i.i.d. from a regular distribution, the expected revenue of the VCG mechanism is at least a $(1 - 1/\kappa)$ fraction of that of the revenue-maximizing mechanism.*

For example, in multi-unit auctions, the packing number is simply the factor by which the number of bidders exceeds the number of goods (rounded down to an integer). The VCG mechanism generates zero revenue when $k \leq n$, so no approximation is possible when the packing number is 1. A packing number of 2 suffices for a constant-factor approximation, and the expected revenue of the VCG mechanism converges rapidly to that of the optimal mechanism as the packing number increases. Similarly to Theorem 3.2, all of the hypotheses in Theorem 6.1 are generally necessary.

A key step in the proof of Theorem 6.1 is a generalization of a result of Bulow and Klemperer [4] that is interesting in its own right. The result resolves the following thought experiment. Suppose a seller initially employs the VCG mechanism in a given market with i.i.d. bidder valuations. Which of the following two options is better for revenue: switching to an optimal auction tailored to the given valuation distribution, or performing a little market expansion? The next lemma shows that expanding a matroid market by a new basis (under the VCG mechanism) is more profitable than switching to an optimal mechanism. For example, for a k -unit auction with n bidders, adding k additional bidders is guaranteed to boost expected revenue beyond that of an optimal auction in the original market. For a matching market with k goods (a transversal matroid), adding k bidders who are collectively willing to accept all k goods achieves the same guarantee.

Lemma 6.2 (Bulow-Klemperer in Matroid Markets) *Let M be a matroid market with bidders U with valuations drawn i.i.d. from a regular distribution. The expected revenue of the VCG mechanism for M is at least that of every optimal mechanism for a matroid market M_S that is induced by a set $S \subseteq U$ that excludes a basis of M .*

To prove Lemma 6.2, we first show that the VCG mechanism is revenue-optimal for a different problem.

Lemma 6.3 *Let M be a matroid market with bidders U with valuations drawn i.i.d. from a regular distribution. The VCG mechanism optimizes the expected revenue over truthful mechanisms that always allocate to a basis of M .*

Proof: Fix a valuation profile. Recall from Example 2.11 that the VCG mechanism can be implemented via the greedy algorithm of Section 2.2 that considers bidders in nonincreasing order of valuation. Recall also that the output of the greedy algorithm depends only on the ordering of the bidders, and that, for i.i.d. valuations drawn from a regular distribution, the bidder orderings by valuation and by virtual valuation coincide. Thus, as in the proof of Theorem 3.2, the VCG mechanism inadvertently maximizes the total virtual value over the bases of the matroid. Taking expectations over the valuation profile and applying (2) completes the proof. ■

Proof of Lemma 6.2: Fix matroids M and M_S that satisfy the conditions of the lemma. The plan is to define an allocation rule that always picks a basis of M and that has expected virtual value equal to that of the optimal auction for M_S . Lemma 6.3 and identity (2) then complete the proof.

To define the allocation rule for M , condition only on the valuations of bidders in S . Let $W \subseteq S$ denote the set of winners in the optimal mechanism for M_S . Since $U \setminus S$ contains a basis, the Exchange Property of matroids implies that W can be extended to a basis of M using bidders of $U \setminus S$. Choose such an extension $X_W \subseteq U \setminus S$ arbitrarily and allocate to the basis $W \cup X_W$ of M .

The expected virtual value earned by this “hybrid mechanism” equals that of the optimal mechanism for M_S plus the conditional expectation $E[\sum_{i \in X_W} \varphi(v_i) | \{v_i | i \in S\}]$. We complete the proof by arguing that this second term is zero. First, for every $i \in X_W$, $E[\varphi(v_i) | \{v_i | i \in$

$S\}} = E[\varphi(v_i)]$ because valuations are independent. Second, for every $i \in X_W$, the unconditional expectation $E[\varphi(v_i)]$ is zero — this can be proved by simple calculation or by applying identity (2) to the single-good, single-bidder auction that always gives the good to the bidder (for free). ■

We now complete the proof of Theorem 6.1 using the revenue submodularity of the optimal mechanism.

Proof of Theorem 6.1: For a matroid market X , let $OPT(X)$ denote the expected revenue of the optimal mechanism. Let B_1, \dots, B_κ denote κ disjoint bases of the given matroid $M = (U, \mathcal{I})$. Obviously, deleting all of these bases from M decreases the expected revenue of the optimal mechanism by at most $OPT(M)$. Since the expected revenue of the optimal mechanism is submodular on subsets of U (Theorem 3.1), we have $\sum_{i=1}^{\kappa} [OPT(M) - OPT(M_{U \setminus B_i})] \leq OPT(M)$. So, there is a basis B_i with

$$OPT(M) - OPT(M_{U \setminus B_i}) \leq OPT(M)/\kappa$$

and hence

$$OPT(M_{U \setminus B_i}) \geq (1 - 1/\kappa) \cdot OPT(M).$$

Lemma 6.2 implies that the expected revenue of the VCG mechanism in M is at least $OPT(M_{U \setminus B_i})$, which completes the proof. ■

For a k -unit n -bidder auction, we can take advantage of the additional symmetry to strengthen Theorem 6.1 using a “fractional packing number”: the expected revenue of the Vickrey auction is at least a $(1 - k/n)$ fraction of that of the optimal auction. We omit the straightforward proof.

7 Revenue Properties of Efficient Keyword Auctions

Search engines use pay-per-click keyword auctions to sell advertising bundled with search results. Our revenue guarantees for the VCG mechanism extend to a standard model of these practically relevant auctions. The main technical observation (Lemma 7.1) that enables this extension is that both the VCG mechanism and the revenue-maximizing auction are revenue-equivalent to randomizations over multi-unit auctions.

7.1 The Model

We study a standard, single-shot model of pay-per-click keyword auctions (see for instance Varian [26] and Edelman et al. [7]). In this model, an auction is run by a search engine on the event of a search query and n bidders (advertisers) compete to have their advertisement displayed in one of k slots. A slot-specific, publicly known parameter called the *click-through rate* Θ_j specifies the probability of a click on an advertisement placed in slot j . Higher slots are assumed to attract more clicks, meaning $\Theta_j \geq \Theta_{j+1}$ for every j .

Each advertiser i has a private valuation v_i for each click on its advertisement — v_i could represent the profit that the advertiser expects to make on a subsequent sale, with the probability of a sale appropriately factored in. The total value realized by allocating advertiser i to slot j is $v_i \cdot \Theta_j$. If a bidder i is allocated x_i clicks, its utility is $x_i \cdot v_i - p_i$, where p_i is the total amount it pays. (The amount paid per click is p_i/x_i .) We assume that the bidders’ valuations are i.i.d. draws from a regular distribution F .

7.2 The Reduction to Multi-Unit Auctions and the Revenue Guarantee

We now show that efficient and revenue-maximizing keyword auctions are both revenue-equivalent to a randomization over k multi-unit auctions, where the j th auction sells j identical units and is chosen with relative probability $\Theta_j - \Theta_{j+1}$. This is shown for efficient auctions in [1, 7, 26]. We prove the lemma more generally for any *truthful sort-by-bid keyword auction with reserve r* , meaning a truthful auction that sets a reserve price r and allocates the slots in order of non-increasing bid.

Lemma 7.1 *The expected revenue from a truthful sort-by-bid keyword auction with reserve price r , k slots, and n bidders with valuations drawn i.i.d. from F , is equal to the weighted sum of expected revenues from k multi-unit auctions, each with reserve price r and n bidders with valuations drawn i.i.d. from F . The i^{th} multi-unit auction sells i objects and has weight $\Theta_i - \Theta_{i+1}$.*

Proof: We prove the result for every possible realization of the valuations. Fix the valuations of the n bidders and index bidders in non-increasing order of bids. Since we have fixed the allocation rule and we assume that the auction is truthful, the payments are uniquely defined (see e.g. [18, Lemma 3]). Specifically, the payment of the i^{th} winning bidder is

$$p_i = \sum_{j=i}^k (\Theta_j - \Theta_{j+1}) \cdot \max(r, v_{j+1}). \quad (3)$$

Let l be the last bidder i with $v_i \geq r$; if there is no such bidder, let $l = 0$. Let $k' = \min(k, l)$. By (3), the total revenue of the auction is

$$\begin{aligned} \sum_{1 \leq i \leq k'} p_i &= \sum_{1 \leq i \leq k'} \sum_{j=i}^k (\Theta_j - \Theta_{j+1}) \cdot \max(r, v_{j+1}) \\ &= \sum_{1 \leq j \leq k} (\Theta_j - \Theta_{j+1}) \cdot \min(l, j) \cdot \max(r, v_{j+1}). \end{aligned} \quad (4)$$

Note that the j^{th} summand in (4) is precisely $(\Theta_j - \Theta_{j+1})$ times the revenue of a truthful j -unit auction with reserve price r . ■

We now use Lemma 7.1 to establish a version of Theorem 6.1, and its extension to fractional packing numbers, for keyword auctions.

Theorem 7.2 *In an n -bidder k -slot pay-per-click keyword auction with bidders' valuations drawn i.i.d. from a regular distribution F , the expected revenue of the VCG mechanism is at least a $(1 - k/n)$ fraction of that of the revenue-maximizing auction.*

Proof: First, recall that the VCG mechanism is a truthful sort-by-bid auction with reserve price 0 (see for instance Varian [26]). Similarly, the optimal keyword auction is a truthful sort-by-bid auction with reserve price $r = \varphi^{-1}(0)$ (see Edelman and Schwarz [8]), where φ is the virtual valuation function for the valuation distribution F .

For $j \leq k$, Let R_{Eff}^j and R_{Opt}^j denote the expected revenue of an efficient and revenue-optimal, respectively, n -bidder j -unit auction with valuations drawn i.i.d. from F . Recall that these two auctions are simply the Vickrey auction and the Vickrey auction with a reserve price $r = \varphi^{-1}(0)$, respectively [18].

By Lemma 7.1, the expected revenue of the VCG mechanism and the revenue-optimal auction for the keyword auction setting are

$$\sum_{1 \leq j \leq k} (\Theta_j - \Theta_{j+1}) R_{Eff}^j.$$

and

$$\sum_{1 \leq j \leq k} (\Theta_j - \Theta_{j+1}) R_{Opt}^j,$$

respectively.

By Theorem 6.1, we have

$$R_{Eff}^j \geq \left(1 - \frac{k}{n}\right) \cdot R_{Opt}^j$$

for every $j \leq k$. Thus the expected revenue of the efficient auction is

$$\sum_{1 \leq j \leq k} (\Theta_j - \Theta_{j+1}) \cdot R_{Eff}^j \geq \left(1 - \frac{k}{n}\right) \sum_{1 \leq j \leq k} (\Theta_j - \Theta_{j+1}) \cdot R_{Opt}^j,$$

which proves the desired approximation guarantee. ■

Theorem 7.2 states that even under moderate competition, when the number of bidders is a small multiple of the number of slots, efficient keyword auctions yield near-optimal revenue. Practically, this is good news for three reasons. First, an optimal auction is not trivial to implement, as it requires knowledge of the valuation distribution F to choose an optimal reserve price. Second, optimal auctions are tailored for monopoly settings. Third, economic efficiency, which is a socially desirable objective, is not necessarily at odds with revenue, which is the natural objective for search engines to optimize.

Remark 7.3 Most real-world keyword auctions are not truthful (see [1, 7, 26]) and can have multiple Nash equilibria. On the other hand, in these non-truthful auctions, there is always a natural equilibrium that is revenue-equivalent to the truthful outcome of the VCG mechanism [1, 7, 26].

Remark 7.4 In some keyword auctions, the number k of slots is in principle infinite. But intuitively, only the slots that receive non-negligible clicks should matter. To make this idea precise, assume that click-through rates fall geometrically with ratio γ . Then, we can show that for every $j \in \{1, \dots, k\}$, the efficient auction $\alpha(j, n, \gamma)$ -approximates the revenue of the optimal auction, where $\alpha(j, n, \gamma) = (1 - j/n)(1 - (j+1)(\gamma^j - \gamma^{j+1}) - \gamma^{j+1})$. For instance, when $\gamma = 1/2$ and $j = 5$ (and k is arbitrarily large), the expected revenue of the VCG mechanism is at least a $(1 - 5/n)(0.89)$ fraction of that of the revenue-maximizing auction.

8 Open Questions

Our work suggests a number of open directions. First, are there additional assumptions (e.g., on bidders' valuations) under which revenue submodularity or an approximate variant holds in non-matroid domains? The same question can be asked about approximation guarantees for the

expected revenue of the VCG mechanism. Finally, we believe that market expansion optimization problems should be studied more broadly. For example, is there a better approximation algorithm than greedy market expansion in matroids? Are there non-trivial approximation algorithms for non-matroid domains, where revenue submodularity — and, for the VCG mechanism, even revenue monotonicity — can fail?

References

- [1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *EC '06: Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 1–7, 2006.
- [2] M. Babaioff, N. Immorlica, and R. D. Kleinberg. Matroids, secretary problems, and online mechanisms. In *SODA '07: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 434–443, 2007.
- [3] S. Bikhchandani, S. de Vries, J. Schummer, and R. V. Vohra. Ascending auctions for integral (poly)matroids with concave nondecreasing separable values. In *SODA '08: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 864–873, 2008.
- [4] J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–94, March 1996.
- [5] M. C. Cary, A. D. Flaxman, J. D. Hartline, and A. R. Karlin. Auctions for structured procurement. In *SODA '08: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 304–313, 2008.
- [6] F. Constantin, J. Feldman, S. Muthukrishnan, and M. Pál. An online mechanism for ad slot reservations with cancellations. In *SODA '09: Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1265–1274, 2009.
- [7] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
- [8] B. Edelman and M. Schwarz. Optimal auction design and equilibrium selection in sponsored search auctions. *American Economic Review*, 100(2):597–602, 2010.
- [9] U. Feige, V. S. Mirrokni, and J. Vondrak. Maximizing non-monotone submodular functions. In *FOCS '07: Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 461–471, 2007.
- [10] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, 1979.
- [11] P. R. Goundan and A. S. Schulz. Revisiting the greedy approach to submodular set function maximization. Preprint, 2009.
- [12] J. D. Hartline. Lectures on frugal mechanism design. EECS 510 Lecture Notes, 2008.

- [13] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *EC '09: Proceedings of the Tenth ACM conference on Electronic Commerce*, pages 225–234, 2009.
- [14] A. R. Karlin, D. Kempe, and T. Tamir. Beyond VCG: Frugality of truthful mechanisms. In *FOCS '05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 615–626, 2005.
- [15] D. C. Kozen. *The Design and Analysis of Algorithms*. Springer-Verlag, New York, 1991.
- [16] N. S. Lambert and Y. Shoham. Asymptotically optimal repeated auctions for sponsored search. In *ICEC '07: Proceedings of the Ninth International Conference on Electronic Commerce*, pages 55–64, 2007.
- [17] P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.
- [18] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [19] Z. Neeman. The effectiveness of English auctions. *Games and Economic Behavior*, 43(2):214–238, 2003.
- [20] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions – I. *Mathematical Programming*, 14(3):265–294, 1978.
- [21] M. Ostrovsky and M. Schwarz. Reserve prices in Internet advertising auctions: A field experiment. Technical report, Stanford University, Graduate School of Business, 2009.
- [22] J. G. Oxley. *Matroid Theory*. Oxford, 1992.
- [23] B. Rastegari, A. Condon, and K. Leyton-Brown. Revenue monotonicity in combinatorial auctions. *SIGecom Exchanges*, 7(1):45–47, 2007.
- [24] K. Talwar. The price of truth: Frugality in truthful mechanisms. In *STACS '03: Proceedings of the 20th Annual Symposium on Theoretical Aspects of Computer Science*, pages 608–619, 2003.
- [25] A. Vardy. The intractability of computing the minimum distance of a code. *IEEE Transactions on Information Theory*, 43(6):1757–1766, 1997.
- [26] H. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.
- [27] J. Vondrak. Optimal approximation for the submodular welfare problem in the value oracle model. In *STOC '08: Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 67–74, 2008.
- [28] D. J. Welsh. *Matroid Theory*. Academic Press, 1976.