Submitted to *Operations Research*
manuscript OPRE-2015-09-521

Robust Auctions for Revenue via Enhanced Competition

Tim Roughgarden, Inbal Talgam-Cohen
Department of Computer Science, Stanford University, Stanford, CA 94305, tim@cs.stanford.edu, italgam@stanford.edu

Qiqi Yan
Google Inc., Mountain View, CA 94043, contact@qiqiyan.com

Most results in revenue-maximizing mechanism design hinge on "getting the price right" — selling goods to bidders at prices low enough to encourage a sale, but high enough to garner non-trivial revenue. This approach is difficult to implement when the seller has little or no a priori information about bidder valuations, or when the setting is sufficiently complex, such as matching markets with heterogeneous goods. In this paper we apply a robust approach to designing auctions for revenue. Instead of relying on prior knowledge regarding bidder valuations, we "let the market do the work" and let prices emerge from competition for scarce goods.

We analyze the revenue guarantees of one of the simplest imaginable implementations of this idea: first, enhance competition in the market by increasing demand (or alternatively by limiting supply); second, run a standard second-price (Vickrey) auction. In their renowned work, Bulow and Klemperer (1996) apply this method to markets with single goods. As our main result, we give the first application beyond single-parameter settings, proving that simultaneously for many valuation distributions this method achieves expected revenue at least as good as the optimal revenue in the original market.

Our robust and simple approach provides a handle on the elusive optimal revenue in multi-item matching markets, and shows when the use of welfare-maximizing Vickrey auctions is justified even if revenue is a priority. By establishing quantitative trade-offs, our work provides guidelines for a seller in choosing among two different revenue-extracting strategies: sophisticated pricing based on market research, or advertising to draw additional bidders.

Key words: Bidding and auctions, pricing, matchings

1. Introduction

1.1. Auctions for Revenue: Three Challenges

Consider a set of m indivisible goods for sale, and the problem of matching them to n buyers with private values, each of whom wants no more than a single good. This problem has been studied extensively with respect to the goal of maximizing economic efficiency (see, e.g., the classic work of Demange et al. (1986)); in this paper we focus on an alternative important goal — maximizing the seller's revenue.

To demonstrate our setting, consider a for-profit travel website selling overnight accommodation, faced with the task of assigning m available rooms to n buyers. Each buyer needs a single room for the night, and has different private values for different rooms based on their location, size etc. Uncertainty of the seller regarding buyer values is captured by a Bayesian model, in which the values for every good $j \in [m]$ are assumed to be independent draws from a distribution F_j that satisfies a standard regularity condition. The fact that F_j is common to all buyers makes our model *symmetric* with respect to buyers (but not with respect to goods). The seller wishes to maximize her expected revenue by designing a deterministic auction, in which no buyer can do better than to participate and reveal his true values (i.e., *dominant strategy* truthful).[1]

When $m = 1$, that is, when there is a single good on the market, Myerson (1981) characterizes the revenue-optimal deterministic truthful auction under the assumption that the distribution F_1 from which values for the good are drawn is fully known to the seller. The optimal auction in this case turns out to be the well-known second-price auction (Vickrey 1961), with an additional reserve price r tailored to the distribution F_1. The resulting auction is very simple: the bidders report their values to the seller, the bidder with the highest bid above r wins, and the winner pays the second-highest bid above r if there is one, or r otherwise. Myerson's characterization of optimal mechanisms also applies to markets with multiple copies (*units*) of the single good, where each bidder seeks at most one copy. More generally, it applies to all *single-parameter* markets, in which every bidder can either win or lose and has a single private value for winning.[2]

Research efforts to extend Myerson's seminal work have spanned several directions. The direction that has attracted perhaps the most attention aims to generalize the optimal auction characterization *beyond the* $m = 1$ *case, to multi-parameter markets* (see, e.g., Daskalakis et al. 2017, and references within). While much progress has been made, it has also become clear that even in simple markets the optimal auction can be extremely complex; in particular, there is no known characterization of the revenue of the optimal truthful auction for the matching market setting described above. A second important direction, which has become known as "Wilson's doctrine" (Wilson 1987), is to design *robust* auctions for revenue that do not depend on the seller's knowledge of the value distributions. A third direction is inspired by the *simplicity* of Myerson's auction — a second-price auction with reserve — and aims to design similarly simple auctions for revenue in more general settings (see, e.g., Hartline and Roughgarden 2009).

In this paper we contribute to all three goals above by applying a robust approach to revenue maximization. We design mechanisms that are robust, simple, and guaranteed to work well for a variety of market environments, including multi-parameter matching markets. Our mechanisms are based on the natural idea of enhancing bidder competition for the goods by adding competing bidders in the manner of Bulow and Klemperer (1996),[3] and then running a variant of the Vickrey auction. Despite avoiding any reference to the value distributions, the expected revenue achieved by the simple mechanisms with enhanced competition exceeds the expected revenue of the optimal mechanisms tailored to the distributions. Like the original Bulow-Klemperer result, our approach sheds light on the trade-off among alternative seller strategies, by quantifying how many additional buyers are required in order to replace the need to rigorously learn the preferences of existing buyers, and to incorporate these preferences into a complex selling mechanism.

Our treatment of robustness in the remainder of the introduction is intuitive; for a definition and comparison to other robustness notions, as well as a discussion of our focus on deterministic and dominant strategy truthful mechanisms, see Section 2.

1.2. Enhanced Competition via Augmented Demand: Bulow-Klemperer-Type Theorems

A well-known result of Bulow and Klemperer establishes the following:

THEOREM 1 **(Bulow and Klemperer (1996))**. *When selling a single good to bidders whose values are i.i.d. draws from a distribution satisfying regularity,*[4] *the expected revenue of the revenue-optimal mechanism with n bidders is at most that of the Vickrey auction with $n+1$ bidders.*

In other words, when the demand is augmented by adding a single additional bidder competing for the good, the simple Vickrey auction achieves at least the maximum revenue possible with the original demand. This is achieved robustly, despite being oblivious to the value distribution, whereas the optimal Myerson mechanism for n bidders depends on this knowledge to set the reserve price. A further advantage of using Vickrey for revenue is that the seller also guarantees optimal social welfare.

The Bulow-Klemperer theorem attracts attention because it gives a simple answer to the following question: when should a seller invest in learning the value distribution so as to implement the optimal auction, and when should she work to recruit an additional buyer? Both are common business practices — drawing more participants can be achieved e.g. by advertising. Conceptually, the Bulow-Klemperer result rejects the standard model of an auction environment in which the demand is exogenous, treating it instead as an endogenous part of the mechanism design problem. It provides theoretical justification for treating bidder participation as a first-order concern when designing auctions for revenue, as indeed done in practice: "[T]he rules are rarely even a first-tier concern in setting up and running a complex auction. Several other issues are more important. One such issue is marketing: an auction cannot succeed without participants" (Milgrom 2004).

Due to its significance, it is important to understand how generally the Bulow-Klemperer theorem holds and what are its limitations. By a *Bulow-Klemperer-type theorem* we refer to any theorem which states that we can get as much revenue in expectation as from the optimal mechanism by running a variant of the Vickrey auction after augmenting the market with additional bidders. Bulow-Klemperer-type theorems have already been established for many single-parameter settings: When there are k units of the good for sale, the expected revenue of the revenue-optimal mechanism with n bidders is at most that of the Vickrey auction with $n+k$ bidders (Bulow and Klemperer 1996). A similar result holds when there are k units for sale but constraints on which bidders can win them – see Section 6.1. To our knowledge, Bulow-Klemperer-type theorems have not been attempted before beyond single-parameter buyers. The reason for this is that despite extensive research, the optimal revenue from selling multiple goods is ill-understood.

1.3. Main Contribution

Our main contribution is in *formulating and establishing robust revenue guarantees in auctions via competition enhancement (in particular increased demand), for a variety of markets including those for which the optimal mechanism remains unknown.* We show that under minimal regularity assumptions, the simple and robust mechanism of Vickrey with additional bidders achieves good revenue in significantly more complex settings than thought before. As our main result, we prove the following first generalization of Bulow and Klemperer's theorem (Theorem 1) to multi-parameter markets:

THEOREM 2 **(Bulow-Klemperer-Type Theorem for Matching Markets (Informal))**. *For every matching market with n bidders and m goods, assuming symmetry and regularity, the expected revenue of the Vickrey auction with m additional bidders is at least the optimal expected revenue in the original market.*

The formal statement appears below (see Theorem 3). Note that the symmetry assumption in this theorem is among bidders, not goods. That is, values of different bidders for the same good are i.i.d. samples from the same distribution, but different goods have different value distributions. Bidder symmetry is assumed also in the original Bulow-Klemperer theorem, and is a realistic model in practical applications: the seller knows she is selling very different kinds of goods, but sees the bidders — whose identities and characteristics she does not known — as homogeneous (Chung and Ely 2007).

In addition to Theorem 2, we prove Bulow-Klemperer-type theorems that achieve better guarantees for matching markets with more supply than demand ($n \leq m$), and that apply to *asymmetric* markets where bidders' values for a good may belong to different distributions (see Section 6).

Proving Bulow-Klemperer-type theorems for matching markets is the most technically challenging component of this work. The analysis of expected revenue in a multi-parameter setting is non-trivial due to dependency issues: intuitively, the competition for item j determines its price, but depends on the buyers' values for the other items. We overcome this dependency challenge via a technique from the analysis of randomized algorithms known as "the principle of deferred decision", combined with combinatorial stability properties of optimal matchings in bipartite graphs (see Section 5.1 for details).

1.4. Alternatives to Augmented Demand

To increase competition, i.e., demand relative to supply, one can either increase the demand as discussed above, or alternatively reduce the supply. The difference between these two approaches to robust revenue maximization is that one requires augmenting resources whereas the other requires the ability to withhold resources, and this translates to a difference in their revenue guarantees.

The idea of limiting supply to drive up prices is well-established (see, e.g., Vohra and Krishnamurthi 2013). Our contribution is in providing *a quantitative connection between increasing demand, limiting supply and maximizing revenue*. This connection is necessarily relative (involving an approximation factor) rather than absolute, and in this sense weaker than Bulow-Klemperer-type theorems. Interestingly, the approximation guarantees of the supply-limiting mechanisms we design follow from a general reduction (Reduction 5 in Section 5) to our Bulow-Klemperer-type theorems, thus formulating the idea that augmenting demand and limiting supply are two sides of the same coin.

Reducing supply requires that the seller credibly commit to doing so, which may be difficult in some cases (see discussion in Section 5.1). However, when the seller lacks the necessary information to set reserve prices, and/or the revenue-optimal mechanism is unknown, and when recruiting additional bidders is difficult, controlling the supply ("setting quantities") may be a viable alternative to setting prices. In such cases, our results show how to balance between the negative effect on revenue from limiting supply (less goods to sell), and the desired positive effect (enhanced competition), by setting the supply to be a constant fraction of the market size.[5] This approach works well simultaneously for many value distributions, guaranteeing a constant fraction of the optimal expected revenue that is independent of the size of the market (cf. inverse-logarithmic fractions guaranteed by other methods — see, e.g., the analysis in Sandholm and Likhodedov (2015)). Furthermore, since the limitation on supply depends on how "large" and competitive the market is to begin with, in large enough markets no limitation is needed, and the approximation loss asymptotically disappears.

In Section 5 we apply the limited supply approach to two kinds of markets — single- and multi-parameter. An example of the first kind is the market for an artwork of which unlimited copies can be printed. Limiting the supply in such a market amounts to the common practice of producing a limited edition of prints (which can then be auctioned). An example of the second is an auction for multiple one-of-a-kind paintings, in which the auction house offers less paintings for sale than the number of bidders it anticipates.

1.5. Organization

In Section 2 we formulate prior-independent robustness and survey related literature. Section 3 presents our model and preliminaries. Section 4 contains our main results on increasing demand in matching markets. Section 5 discusses limiting supply and large market results as alternatives to increasing demand, for both single-parameter multi-unit markets as well as for matching markets. Extensions appear in Section 6, and Section 7 concludes.

2. Prior-Independent Robustness and Related Work

2.1. Robustness Formalization

Robustness, or *detail-freeness*, is a long-time goal of mechanism and market design. It requires mechanisms that "perform well" for a "large range" of economic environments regardless of their precise details. Any formalization should thus specify a measure of good performance and the relevant range of environments. We now specify these for the notion of *prior-independent* robustness.

Good performance. Let us focus for simplicity on the single good case, and consider first a particular distribution F from which the buyers' i.i.d. values are drawn. Denote by OPT_F the optimal expected revenue that a *truthful deterministic* mechanism can achieve *with full knowledge of* F, and let $\alpha \in (0,1]$ be an approximation factor. We say that a mechanism is α*-optimal* with respect to F if its expected revenue is at least $\alpha\,\mathrm{OPT}_F$.[6] Now let $\mathcal{F}$ be a set of value distributions, or *priors*. A mechanism is *robustly* α*-optimal* with respect to $\mathcal{F}$ if for every $F \in \mathcal{F}$ it is α-optimal with respect to F (equivalently we say it achieves an α*-approximation* to the optimal expected revenue). Thus to "perform well", the mechanism's expected revenue must either exceed or approximate the optimal expected revenue *simultaneously for every distribution in a class of distributions* $\mathcal{F}$.

Rephrasing the results above in this language, Theorem 2 establishes that Vickrey with m additional bidders is robustly 1-optimal, and Theorem 5 states that Vickrey with supply limit $n/2$ is robustly α-optimal for a constant $1/\alpha$.

Range. We set the "large range" of distributions $\mathcal{F}$ to include all distributions that satisfy *regularity*, i.e., whose tail is no heavier than that of a power law distribution (see Section 3 for a precise definition). This standard assumption in auction theory is analogous to the assumption of downward-sloping marginal revenue in monopoly theory (Bulow and Klemperer 1996), and without it no robust revenue guarantees are possible (negative examples appear in Section 3).

Discussion of the Prior-Independence Formalization Prior-independent robustness is an interesting mixture of average- and worst-case performance guarantees: On one hand performance is measured in expectation over random inputs, on the other it is measured in the worst case over all distributions in $\mathcal{F}$. While there is an underlying assumption that values are distributed according to a prior, the mechanism must be entirely independent of this prior. We now briefly discuss the rationales for adding this worst-case component to the classic average-case (Bayesian) approach, i.e., why it makes sense to measure the performance of a robust mechanism against the optimal mechanism with prior knowledge, and why good performance is required for such a large range of distributions (rather than, say, incorporating partial information about the prior to narrow down the range and enable even better performance). We also discuss our choice to focus on deterministic and dominant strategy truthful mechanisms.

First note that since our main results are positive, the choice of OPT_F as a benchmark and $\mathcal{F}$ as the distribution range strengthen our results: for a large range of distributions on which little is assumed, we are able to compete with an ambitious benchmark, even in a challenging environment with multiple items for which the optimal mechanism remains elusive. Choosing OPT_F as a benchmark can also facilitate informed decisions as to how much a seller should invest in obtaining information about F. An alternative approach to robustness, which does not require competing with OPT_F for every distribution $F \in \mathcal{F}$, instead aims to *maximize the minimum* expected revenue over all $F \in \mathcal{F}$. This approach is infeasible in our matching setting as it requires first solving the open problem of finding the optimal mechanism for multi-item markets (for example if $\mathcal{F}$ contains only a single distribution).

The justifications for avoiding dependence on the value distribution also apply to dependence on the *class* of relevant value distributions. Choosing $\mathcal{F}$ to include all regular distributions makes our mechanisms applicable in situations where the seller has no information on which distributions are relevant, such as when the seller or product is new on the market, when the distributions change over time, or when achieving such information is too expensive or inherently noisy. In hindsight, this choice of $\mathcal{F}$ has led to simple and natural mechanisms, thus reinforcing our approach.

Discussion of Restricting to Deterministic Dominant Strategy Truthful Mechanisms We choose as our benchmark the expected revenue OPT_F of the best deterministic and dominant strategy truthful mechanism. For single-parameter settings, our results actually hold for a stronger benchmark, which includes all randomized and Bayesian truthful mechanisms as well. (Bayesian truthfulness is a relaxed truthfulness notion which requires that no buyer can do better *in expectation over the other buyers' values* than to participate and bid truthfully in the auction.) This stronger guarantee follows directly from the properties of single-parameter settings, which were utilized already by Myerson — his deterministic dominant strategy truthful mechanism is optimal among all randomized Bayesian truthful mechanisms. In fact, Gershkov et al. (2013) show that the equivalence of dominant strategy and Bayesian implementations holds for any single-parameter social choice environment. For multi-parameter settings, this equivalence no longer holds, and randomization and/or Bayesian truthfulness give strictly more power to the mechanism (Thanassoulis 2004, Manelli and Vincent 2006, Yao 2016). However in matching markets this extra power is limited — Chawla et al. (2010b) show that the expected revenue from the optimal deterministic mechanism is within a small constant factor of the expected revenue from the optimal randomized mechanism, and Cai et al. (2016) generalize this to deterministic and dominant strategy truthful mechanisms as compared to randomized Bayesian truthful ones. Thus our results guarantee a constant-factor approximation of the stronger benchmark for matching markets.

It is interesting to study how much more competition is needed in multi-parameter settings to surpass (rather than approximate) the stronger benchmark (in this context see our follow-up work described in Section 2.2). However, we argue that the most relevant mechanisms to measure performance against are deterministic and dominant strategy truthful. In other words, requiring many additional extra bidders to reach the stronger benchmark is usually unjustified. Almost all mechanisms applied in real-life auctions — including multi-parameter (combinatorial) ones — do not involve any form of lotteries; thus randomized mechanisms are restricted to almost purely theoretical interest. As for Bayesian truthfulness, for a seller to choose such a mechanism she would have to rely on common knowledge of the prior distributions among the buyers. The seller is often much better positioned to learn these priors than the buyers are (each buyer sees only his own value), and if the seller herself is debating the cost and/or feasibility of learning these priors (as is our underlying assumption in this paper), it is implausible that she would choose to rely on the buyers' knowledge of them. For a mathematical formalization of this line of argument see the work of Chung and Ely (2007), who establish the theoretical foundations of dominant strategy truthfulness.[7] We conclude that in many settings the performance of deterministic and dominant strategy truthful mechanisms is the appropriate benchmark to compete against.

2.2. Robustness in the Literature

Robust mechanism design is widely studied within the economics, operations research and computer science literatures. In economics, Wilson (1987) advocates for mechanisms independent of any details that are not common knowledge among the buyers and seller, noting that repeated weakening of common knowledge assumptions is the only way in which the theory will approximate reality. In operations research, Scarf (1958) writes: "we may have reason to suspect that the future demand will come from a distribution that differs from that governing past history in an unpredictable way", and Bertsimas and Thiele (2014) describe the need for a non-probabilistic theory as "pressing". In computer science, the dominant paradigm of worst-case analysis extends to mechanism design — the expectancy is that mechanisms, like algorithms, should work well across a range of settings, and there is a "general mistrust" in a designer's ability to accurately capture real-world distributions (Nisan 2014). We now survey previous work within these literatures, to which our work contributes by achieving robustness and simplicity for multi-parameter environments.

Prior-Independent Robustness for Single-Parameter Markets

Bulow-Klemperer-type results. Bulow and Klemperer (1996) initiated the study of augmented demand, focusing on single-item and single-parameter *multi-unit* markets (see also Kirkegaard (2006) for a simple proof). Dughmi et al. (2012) generalize their results to markets with *matroid-based constraints*, and Hartline and Roughgarden (2009) achieve similar guarantees for *asymmetric* bidders with heterogeneous value distributions. Sivan and Syrgkanis (2013) develop a Bulow-Klemperer-type result for *irregular* distributions that are a convex combination of regular ones. Fu et al. (2015) compare the Bulow-Klemperer approach to alternative prior-independence methods.

Sampling-based results. A natural approach to prior-independent robustness is to instantiate the optimal mechanism of Myerson (1981) with an empirical distribution sampled from the buyers' bids. This approach is asymptotically optimal as the size of the market goes to infinity (Segal 2003, Baliga and Vohra 2003), and also enables the seller to incorporate any available information regarding the class of relevant prior distributions. Dhangwatnotai et al. (2010) design simple mechanisms that require only a single sample, and are robustly α-optimal for a small constant $1/\alpha$ and markets of any size.

Other Robustness Notions for Single-Parameter Markets In prior-*free* (as opposed to prior-*independent*) mechanism design, prior distributions are used neither in the design nor in the performance evaluation of the auction. Thus OPT_F is no longer an appropriate benchmark. Neeman (2003) compares the revenue of the English auction (an ascending-price version of Vickrey) to the optimal social welfare benchmark (which upper-bounds the optimal revenue). Goldberg et al. (2006) define a benchmark based on the optimal auction within a restricted subclass of auctions.

The relation between different robustness notions has also been studied (see, e.g., Hartline and Roughgarden 2014).

(Non-Robust) Revenue Maximization for Multi-Parameter Markets In recent years there have been significant advances in understanding prior-*dependent* optimal mechanisms for multi-parameter markets. One line of work focuses on characterizing optimal mechanisms and establishing their complexity — they are now known to necessarily involve randomization, infinitely-many outcomes, non-monotonicity of revenue in values, and computational hardness (Cai et al. 2013, Daskalakis et al. 2013, Hart and Reny 2015, Daskalakis et al. 2014), even for extremely simple markets (Hart and Nisan 2012). Another line of work focuses on designing simpler but only approximately-optimal mechanisms. Among the prominent works in this direction (see e.g. Li and Yao 2013, Babaioff et al. 2014, Sandholm and Likhodedov 2015), the most related to ours is by Chawla et al. (2010a), who give an upper bound on the optimal expected revenue for matching markets. They achieve a prior-*dependent* 1/6.75-approximation for matching markets with multiple units and asymmetric buyers, and also a 3/32-approximation for the more general environment of "graphical matroid with unit-demand buyers". Our techniques utilize one of their reductions as described in Section 3.3.

Robustness for Multi-Parameter Markets

Prior-independent approximation results. In contemporaneous work with an early version of this paper, Devanur et al. (2011) independently consider a similar set of problems, but using different mechanisms and analysis. They design approximation mechanisms that are arguably more complex and less natural since they are based on auctions with carefully-constructed price menus (as opposed to enhanced competition). Following our early version, Azar et al. (2014) study matching markets in which partial information about the value distributions is available to the seller in the form of samples. Goldner and Karlin (2016) design prior-independent approximation mechanisms for multi-parameter markets where bidders are *additive* (rather than unit-demand).

Robust optimization. Bandi and Bertsimas (2014) apply a robust optimization approach to optimal mechanism design in multi-item markets. Their model differs from ours in several important aspects, including consideration of additive valuations rather than matching markets, and divisible rather than indivisible goods. A main goal of their paper is orthogonal to ours, namely to study budgets and correlated values in mechanism design (they also address auctions without budget constraints but in their setting these reduce to single-item markets). Still it is interesting to compare the robustness notions in their work and ours: They model the seller's knowledge about values (based say on historical bidding data) by an *uncertainty set*, and then optimally solve the resulting robust optimization problem. Simulations establish significant improvement in single-item settings upon Myerson's mechanism when tailored to an *inaccurate* prior distribution. Our goal in single-item settings is to surpass or approximate the performance of Myerson's mechanism tailored to the accurate distribution, without knowing what this distribution is.

Follow-up work on Bulow-Klemperer-type results. In recent work with the second author, Eden et al. (2016) extend our results in two ways — they prove Bulow-Klemperer-type results that hold for additive valuations subject to downward-closed constraints (these include the matching markets we study here as well as markets with additive bidders), and they focus on the stronger benchmark of the optimal mechanism among all randomized, Bayesian truthful mechanisms (see discussion in Section 2.1). Competing against the stronger benchmark involves enhancing competition by more additional bidders (namely $n+3(m-1)$ additional bidders in total).

Simplicity of the Vickrey Auction The Vickrey or second-price auction is a simple and ubiquitous auction format, common for example in the context of sponsored search and online advertising (see, e.g., Lahaie et al. 2007, Celis et al. 2014). Hartline and Roughgarden (2009) seek conditions on single-parameter markets such that the simple Vickrey auction with (prior-dependent) reserves achieves near-optimal revenue (see also Dughmi et al. 2012). The extension of the Vickrey auction to multi-item markets, called the VCG mechanism (Vickrey 1961, Clarke 1971, Groves 1973), is arguably not as simple (Ausubel and Milgrom 2006, Rothkopf 2007). Yet in the context of matching markets, many of the complications of VCG do not occur: communicating the bids and running the auction are both computationally tractable, and our competition enhancement methods ensure that the revenue does not collapse. Chawla et al. (2013) analyze the VCG mechanism's revenue performance in a job scheduling context, and some of our techniques are inspired by their analysis.

3. Preliminaries

In Section 3.1 we describe our model, in Section 3.2 we review the basics of optimal mechanism design, and in Section 3.3 we present two technical tools — the regularity assumption and representative environments.

3.1. Model

An auction environment (or market) has as set $\{1,\dots,m\}$ of m goods (or items) for sale to a set $\{1,\dots,n\}$ of n bidders (or buyers). As a convention we use the index i for bidders and j for items. Throughout we make the distinction between items and units, where the latter are copies of the same item (so bidders have the same value for them).

Multi-Unit and Other Single-Parameter Environments A single-parameter environment is defined by a non-empty collection $\mathcal{I} \subseteq 2^{[n]}$, containing sets of bidders who can simultaneously *win* and are called *feasible allocations*. Every feasible allocation subset is itself feasible, i.e., $([n],\mathcal{I})$ is a *downward-closed* set system. Every bidder belongs to at least one feasible allocation. Every bidder i has a private value $v_i \in [0,\infty)$ for winning, drawn independently at random from a distribution F_i with a smooth density function f_i that is positive over a nonzero interval support (the environment is "single-parameter" in that the value for winning is fully captured by v_i). We say that single-parameter bidders are *i.i.d.* (or symmetric) if their value distributions are identical, and the environment is i.i.d. if the bidders are i.i.d. We assume a standard risk-neutral quasi-linear utility model, in which a bidder's utility for winning is his value minus his payment, and bidders aim to maximize their expected utilities.

A multi-unit (or *k-unit*) environment is a single-parameter environment in which a subset of bidders is a feasible allocation if and only if its size is at most k. This models k units for sale to bidders interested in at most one unit. We will sometimes impose an additional supply limit of $\ell \leq k$, restricting feasible allocations to size at most ℓ. Multi-unit environments are not to be confused with multi-*item* ones, in which the goods are heterogeneous. For a survey on multi-unit environments see (Nisan 2014).

A *matroid* environment is a single-parameter environment in which the set system $([n],\mathcal{I})$ of bidders and feasible allocations forms a *matroid* (see, e.g., Oxley 1992). k-unit environments are a special case of matroid environments (corresponding to the k-uniform matroid). See Section 6.1 for further details.

Multi-Item Matching Environments In a matching environment there are m different items for sale with one unit available of each item. Feasible allocations are all matchings of items to bidders (each bidder wins at most one item and each item is allocated to at most one bidder) — this models unit-demand bidders. We will sometimes impose an additional supply limit of $\ell \leq m$, restricting the matchings to size at most ℓ. Every bidder i has a private value $v_{i,j} \in [0,\infty)$ for winning item j, which is drawn independently at random from a distribution $F_{i,j}$ with a smooth density function $f_{i,j}$ positive over a nonzero interval support (a matching environment is thus *multi-parameter*). We again assume a risk-neutral quasi-linear utility model. We say the bidders are i.i.d. (or symmetric) if $F_{i,j}$ does not depend on the identity of bidder i, i.e., each item j has an associated distribution F_j and $F_{i,j} = F_j$. In other words, for every item j the values $\{v_{i,j}\}_{i\in[n]}$ are i.i.d. samples from F_j. Note that different items j, j' have different distributions $F_j, F_{j'}$, as necessary for applications with heterogeneous items (e.g., the travel website example in Section 1). For a treatment of *asymmetric* bidders see Sections 6.2 and 6.3.

3.2. Optimal Mechanism Design

Mechanisms By the revelation principle, without loss of generality we may restrict attention to direct mechanisms, which receive a vector of bids $\mathbf{b}$. In the single-parameter case $\mathbf{b} \in \mathbb{R}^n_{\geq 0}$ where b_i is bidder i's bid for winning, and in the matching case $\mathbf{b} \in \mathbb{R}^{nm}_{\geq 0}$ where $b_{i,j}$ is bidder i's bid for winning item j. We focus on deterministic mechanisms, comprised of:

1. An allocation rule $\mathbf{x} = \mathbf{x}(\mathbf{b})$, which maps a bid vector $\mathbf{b}$ to a feasible allocation; in the single-parameter case $\mathbf{x} \in \{0,1\}^n$, where $x_i = x_i(\mathbf{b})$ indicates whether bidder i wins, and in the matching case $\mathbf{x} \in \{0,1\}^{nm}$, where $x_{i,j} = x_{i,j}(\mathbf{b})$ indicates whether bidder i wins item j.

2. A payment rule $\mathbf{p} = \mathbf{p}(\mathbf{b})$, which maps a bid vector $\mathbf{b}$ to a payment vector. The payment vector $\mathbf{p}$ belongs to $\mathbb{R}^n_{\geq 0}$, where $p_i = p_i(\mathbf{b})$ is the payment charged to bidder i.

Fixing a bid vector $\mathbf{b}$, the mechanism's *welfare* in the single-parameter case is $\sum_i x_i v_i$, and in the matching case $\sum_{i,j} x_{i,j} v_{i,j}$. The mechanism's *revenue* is $\sum_i p_i$. Bidder i's utility in the single-parameter case is $x_i v_i - p_i$, and in the matching case $\sum_j x_{i,j} v_{i,j} - p_i$. A mechanism is (dominant strategy) *truthful* if for every bidder i and bid profile $\mathbf{b}_{-i}$ of the other bidders, i maximizes his utility by participating and bidding truthfully, i.e., bidding $b_i = v_i$ in the single-parameter case and $b_{i,j} = v_{i,j}$ for all j in the matching case. All the mechanisms we study are truthful, so from now on we no longer distinguish between bids and values and use v_i or $v_{i,j}$ to denote both. We will mainly be interested in a mechanism's *expected* revenue $\mathbb{E}_{\mathbf{v}}[\sum_i p_i]$, where $\mathbf{p} = \mathbf{p}(\mathbf{v})$ and the expectation is taken over i.i.d. values drawn from the value distributions.

The revenue benchmark against which we measure the performance of our mechanisms is the following: By optimal expected revenue we mean the maximum expected revenue over all (dominant strategy) truthful, deterministic mechanisms. Chawla et al. (2010b) show that for matching environments, the expected revenue from the optimal deterministic mechanism is within a constant factor of the expected revenue from the optimal randomized mechanism. Thus our results for deterministic mechanisms apply to randomized mechanisms up to a constant factor.

Maximizing Welfare and the Vickrey Auction The general form of the Vickrey auction is called the VCG mechanism, and it works for any market whether single-parameter or multi-item. VCG is remarkable in being both truthful and welfare-maximizing for every value profile $\mathbf{v}$. Its allocation rule chooses a feasible allocation that maximizes welfare; its payment rule charges every bidder i a payment equal to i's *externality* — the difference in the maximum welfare *of the other bidders* when i does not participate in the auction and when i does participate in it.

In the context of matching environments, the VCG allocation rule can be implemented as a maximum weighted matching over a bipartite graph, where vertices on one side are the bidders, vertices on the other side are the items, and the weight of every edge (i,j) is $v_{i,j}$ (Bertsekas 1991).

The payment rule also solves bipartite matching problems to compute the payments. For single-parameter k-unit environments, Vickrey's allocation rule finds k bidders with highest values, and for matroid environments it uses a simple greedy algorithm to find a feasible allocation with highest welfare (and similarly for the Vickrey payment rules).

For our supply-limiting mechanisms, we add to the VCG or Vickrey mechanisms a supply limit ℓ and denote them by $\mathrm{VCG}^{\leq \ell}$ and $\mathrm{Vic}^{\leq \ell}$, respectively.

Maximizing Revenue and Myerson's Mechanism For single-parameter environments, Myerson (1981) characterized the optimal mechanism that maximizes expected revenue. (In fact, Myerson showed an even stronger result — his mechanism maximizes the expected revenue over all Bayesian truthful, randomized mechanisms!) Let F be a regular distribution with density f (see Section 3.3 for discussion of regularity). Define its *virtual value* function $\phi_F : \mathbb{R}_{\geq 0} \to \mathbb{R}$ to be $\phi_F(v) = v - \frac{1-F(v)}{f(v)}$. Myerson showed the following.

LEMMA 1 **(Myerson)**. *Given a single-parameter environment and a truthful mechanism* $(\mathbf{x}, \mathbf{p})$, *for every bidder* i *and value profile* $\mathbf{v}_{-i}$ *of the other bidders,* $\mathbb{E}_{v_i \sim F_i}[p_i(\mathbf{v})] = \mathbb{E}_{v_i \sim F_i}[x_i(\mathbf{v})\phi_{F_i}(v_i)]$.

Myerson's lemma says that in expectation over bidder i's value, his payment is equal to his virtual value when he is allocated. By summing over all bidders, this lemma implies that in expectation over the value profile, maximizing the revenue is equivalent to maximizing the total virtual value of allocated bidders, a quantity known as the *virtual surplus*. *Myerson's mechanism* maximizes expected revenue by finding the feasible allocation with maximum virtual surplus. For example, in a k-unit environment this will be the set of $\leq k$ bidders with the highest *positive* virtual values.

3.3. Technical Tools

Regularity We say that bidders are regular if their values are drawn from regular distributions.

DEFINITION 1 (REGULAR DISTRIBUTION). A distribution F is regular if its virtual value function is monotone non-decreasing.

Most commonly-studied distributions are regular, including the uniform, exponential and normal distributions, and distributions with log-concave densities. The assumption that bidders are regular is standard in optimal mechanism design and is necessary for designing good prior-independent mechanisms, as demonstrated by Dhangwatnotai et al. (2010): Fix a value z and a number n of bidders, and define an irregular, long-tailed value distribution F_z such that the probability for z is $1/n^2$ and otherwise the value is zero. Consider a single-item environment with n bidders whose values are drawn from F_z. The optimal auction has expected revenue at least z/n. But any prior-independent truthful auction essentially has to "guess" the value of z, since the probability that the non-winning bids provide information about z is small. Thus its expected revenue cannot be within a constant factor of z/n for every F_z.

Representative Environments A representative environment is the "single-parameter counterpart" of a matching environment. Consider a matching environment with m items, n symmetric bidders and value distributions $\{F_j\}_{j=1}^m$. The corresponding representative environment has the same m items, but nm single-parameter bidders — every bidder in the matching environment has m *representatives* in the representative environment. The jth representative of bidder i is only interested in item j and has a value $v_{i,j} \sim F_j$ for winning it. Every allocation in the representative environment can be translated to an allocation in the matching environment — if the jth representative of i wins, then item j is allocated to bidder i in the matching environment — and vice versa. An allocation in the representative environment is *feasible* if the corresponding allocation in the matching environment forms a matching, meaning that only one representative per bidder wins.

Intuitively, the representative environment is more competitive than the matching one, since representatives of the same bidder compete against each other on who will be the winner. Thus the expected revenue achievable in the representative environment should be at least the optimal expected revenue in the matching environment. Chawla et al. (2010a) formalize this intuition by showing that any truthful mechanism M for the matching environment translates to a truthful mechanism M^{rep} for the representative environment, such that the expected revenue of M^{rep} is only higher. Roughly this is by translating the allocation rule of M to an allocation rule in the representative environment as above, and viewing the payment rule of M as a price menu, whose prices are exceeded in the representative environment by charging every representative the minimum value it needs to bid in order to win.

LEMMA 2 **(Chawla et al. (2010a))**. *The expected revenue of M^{rep} in the single-parameter representative environment is at least the expected revenue of M in the matching environment.*

4. Main Result: Augmenting Demand in Matching Markets

In this section we prove a Bulow-Klemperer-type theorem for matching environments — the first generalization of Bulow and Klemperer (1996) to a multi-item market. Recall what we mean by i.i.d. bidders in a matching environment: different items have different distributions, but independence is both across bidders and across items.

THEOREM 3 **(Bulow-Klemperer-Type Theorem for Matching Markets)**. *For every matching environment with i.i.d. regular bidders and m items, the expected revenue of the VCG mechanism with m additional bidders is at least the optimal expected revenue in the original market. In other words, VCG with m additional bidders is robustly optimal.*

Theorem 3 provides a simple handle on the unknown optimal expected revenue in matching markets. For example, in a market with two goods for sale, the best achievable revenue is at most what VCG can achive with two more bidders. Note that in markets with plentiful supply, i.e. markets in which $m \gg n$, the demand augmentation that is required is substantial. In Section 4.5 we present an alternative Bulow-Klemperer-type theorem with weaker requirements for this case.

4.1. Overview of the Proof

The proof is divided into two parts. In Section 4.2 we identify an upper bound on the optimal expected revenue in the original environment, and a lower bound on the revenue of the VCG mechanism in the augmented environment with m more bidders. These bounds are relatively simple to analyze and are already similar, though not identical, in form. In Section 4.3 we carefully relate the two bounds to establish the theorem.

Our proof is based on the following ideas. We first observe there is a simple upper bound on the optimal expected revenue in the matching environment — *the expected revenue from running m Vickrey auctions to sell each of the m goods to m separate sets of $n+1$ representatives*, who are single-parameter bidders only interested in one particular good (Lemma 3). Our goal is now to show that VCG with a total of m additional multi-parameter bidders does just as well in terms of revenue.

Recall that in the VCG mechanism, the winner of a certain good pays the externality he inflicts upon other bidders, which includes in particular the "damages" he causes the losing bidders who are not allocated any good by the mechanism. Thus, the payment for every good j is at least the highest value for j among the losers. In the augmented matching environment to which VCG is applied, it is guaranteed that there will be n losers, since there are m goods and $n+m$ bidders. The expected revenue from running VCG on the augmented environment is thus at least *the expected welfare from running m Vickrey auctions to allocate each of the m goods separately to the n losers* (Lemma 4). This lower bound is similar to the above upper bound.

The remaining challenge is a dependency issue — by definition, the losers are likely to have lower values for the goods than the $n+1$ representatives. We use the combinatorial structure of maximum weighted matchings to show that a bidder's values conditional on him losing in the VCG mechanism are, while lower, not likely to be *significantly* so compared to the unconditional case. Thus the losers' damages are enough to cover the expected revenue from the representatives.

On a technical level what we show is that, quite remarkably, the only thing that can be deduced about a bidder's value for an item j from his losing the auction completely is that it is lower than the value of the winner of item j. We establish this by introducing an auxiliary selling mechanism for item j, conceptually and revenue-wise half-way between selling the item separately and selling it as part of the VCG mechanism. The auxiliary mechanism runs a maximum weighted matching algorithm as in VCG, but defers the sale of item j until all other goods have been sold and exactly $n+1$ bidders remain unallocated. Thus, by construction, these bidders' values for item j are unaffected by the dependency issue described above.

4.2. Basic Upper and Lower Bounds

Upper bound. To upper bound the expected optimal revenue we use the following notation: let $\text{Vic}_j(\eta)$ be the expected revenue from selling item j to η i.i.d. single-parameter representatives with value distribution F_j using the Vickrey auction. Then:

LEMMA 3 **(Upper Bound on Optimal Expected Revenue)**. *For every matching environment with n i.i.d. regular bidders and m items, $\sum_j \text{Vic}_j(n+1)$ is at least the optimal expected revenue in the original market.*

Proof. Let $\{F_j\}_{j=1}^m$ be the regular value distributions of the matching environment. By Lemma 2, the optimal expected revenue in the matching environment is upper-bounded by the optimal expected revenue in its single-parameter counterpart, the corresponding representative environment. Recall that in the representative environment there are n single-parameter representatives per item j, whose values for j are i.i.d. draws from F_j. The representatives are grouped in n sets of size m corresponding to the original bidders in the matching environment, and feasibility constraints ensure that at most one representative from each set wins.

We now relax these feasibility constraints to get a new single-parameter environment in which the optimal expected revenue has only increased. Relaxing feasibility only increases the optimal expected revenue since by Meyerson's lemma (Lemma 1), it is equal in expectation to the optimal virtual surplus, and clearly the optimum subject to the constraints is bounded from above by the optimum when these are relaxed.

The new environment is equivalent in terms of revenue to a collection of m single-item environments, where in the j-th environment item j is auctioned to its n single-parameter representatives (values are i.i.d. draws from the regular distribution F_j). By Bulow and Klemperer's result (Theorem 1), the optimal expected revenue from the j-th environment is upper-bounded by $\text{Vic}_j(n+1)$. Summing up over all items completes the proof. □

Lower bound. We now turn to the VCG mechanism applied to the augmented environment, whose revenue is the sum of VCG payments for the items. The next lemma lower-bounds the VCG payment for item j.

Lemma 4 **(Lower Bound on VCG Revenue)**. *For every matching environment, the VCG payment for item j is at least the value of any unallocated bidder for j.*

Proof. Say bidder i wins item j. The VCG payment for j is equal to the externality that i imposes on the rest of the bidders by winning j. In particular, i prevents an unallocated bidder i' from being allocated j. Thus the payment is at least the value of i' for j. □

4.3. Relating the Upper and Lower Bounds via Deferred Allocation

The upper and lower bounds above share a similar form. On the one hand, by definition of the Vickrey auction, the upper bound $\text{Vic}_j(n+1)$ on the expected revenue from separately auctioning off item j is equal to the expected *second-highest value for j among $n+1$ bidders with values drawn independently from F_j*. On the other hand, the lower bound on the VCG payment for item j in the augmented environment is equal to the *highest value for j among n unallocated bidders with values drawn independently from F_j*, where we use that in the augmented environment only m out of $n+m$ bidders are allocated. From this it may appear as if we have already shown that the lower bound exceeds the upper bound. However, a dependency issue arises — conditioned on the event that a bidder in the augmented environment is unallocated by VCG, his value for item j is no longer distributed like a random sample from F_j. We address this issue by introducing a *deferred allocation* selling procedure.

Algorithm 1 describes our selling procedure for item j.

Algorithm 1 Selling Item j by Deferred Allocation

Given a matching environment with $n+m$ bidders and m items, and an item j:

1. Find a maximum matching of all $m-1$ items other than j to the bidders. Let U be the set of $n+1$ bidders who remain unallocated.
2. Run the Vickrey auction to sell item j to bidder set U.

The following two claims show how deferred allocation resolves the dependency issue; namely, how the revenue from selling item j via the deferred allocation procedure bridges between the upper and lower bounds in Lemmas 3 and 4. The relation is also depicted in Figures 1a to 1c.

Claim 1 **(Deferred Allocation and Upper Bound)**. *The revenue from selling item j by deferred allocation (Algorithm 1) is equal in expectation to* $\text{Vic}_j(n+1)$.

Proof. The revenue from selling item j to bidder set U by the Vickrey auction is the second-highest value of a bidder in U for j. Since we exclude item j in step (1) of the deferred allocation procedure and allocate it only in step (2), the allocation in step (1) does not depend on the bidders' values for j. Therefore, the values of the unallocated bidders in U for item j are independent random samples from F_j. The expected second-highest among $n+1$ values drawn independently from F_j is equal to $\text{Vic}_j(n+1)$. $\square$

To relate the revenue from deferred allocation to the lower bound in Lemma 4 we need the following stability property.

LEMMA 5 **(Stability of Maximum Matching)**. *Consider an augmented matching environment with $n+m$ bidders and m items. Let bidder set U be as defined in Algorithm 1. If VCG is run on this environment, the set of bidders left unallocated is U with at most one bidder removed.*

Proof. First note that in the matching instances we consider, we may assume there is a unique maximum weighted matching. This holds with probability 1 since the weights are sampled from distributions as described in Section 3.1.[8]

The following is a well-known stability property of maximum weighted matchings (Lovász and Plummer 2009, Lemma 3.2.2): In a complete weighted bipartite graph with $n+m$ nodes on one side and $m-1$ nodes on the other, consider the maximum weighted matching of size $m-1$. Now add a node to the short side of the graph and find the maximum weighted matching of size m. The set of matched nodes on the long side of the graph remains the same up to a single additional node.

The augmented matching environment corresponds to a complete bipartite graph with bidders on one side and items on the other, with the bidders' values for the items as edge weights. Algorithm 1 finds the maximum weighted matching of size $m-1$ in this graph with item j removed. VCG finds the maximum weighted matching of size m in this graph including item j. The lemma follows by applying the above stability property. $\square$

CLAIM 2 **(Deferred Allocation and Lower Bound)**. *Given an augmented matching environment with $n+m$ bidders and m items, the revenue from selling item j by deferred allocation (Algorithm 1) is at most the VCG payment for item j.*

Proof. The revenue from selling item j by deferred allocation is the second-highest value of a bidder in U for j. Let i_1, i_2 be the two bidders in U who value item j the most. By definition, these bidders are left unallocated by the deferred allocation procedure, and by Lemma 5, one of them (say i_1) is also unallocated by the VCG mechanism. Recall that an unallocated bidder's value for item j gives a lower bound on the VCG payment for j (Lemma 4). So the VCG payment for j is at least $v_{i_1,j}$, which in turn is at least the second-highest value of a bidder in U for item j. $\square$

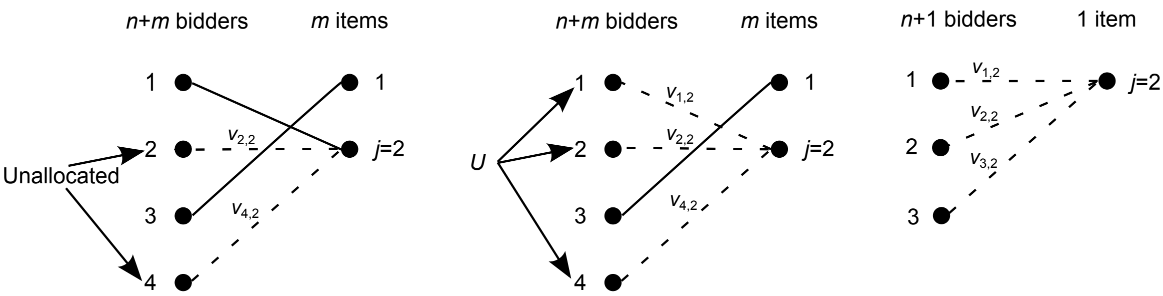

(a) VCG (b) Deferred allocation (c) Vickrey

Figure 1 **Relating bounds by deferred allocation**

Let $n = m = 2$ and item $j = 2$.

(a) **VCG**: Solid edges correspond to the maximum matching. The payment for j is $\geq \max\{v_{2,2}, v_{4,2}\}$ (Lemma 4), where $v_{2,2}, v_{4,2}$ are *not* i.i.d. samples from F_2 given that bidders 2 and 4 are unallocated.

(b) **Deferred allocation**: Solid edges correspond to the maximum matching excluding j. Bidders unallocated in (a) are a subset of the unallocated set U (Lemma 5). Since j is sold to U using Vickrey, the payment for j is the 2nd-highest among $v_{1,2}, v_{2,2}, v_{4,2}$, where these are i.i.d. samples from F_2.

(c) **Vickrey**: The payment for j is the 2nd-highest among $v_{1,2}, v_{2,2}, v_{3,2}$, where these are i.i.d. samples from F_2.

Comparing (a) to (b) and (b) to (c) shows that:

- the payment for j in (a) is at least the payment for j in (b) (Claim 2); and

- in expectation the payment for j in (b) equals the payment for j in (c) (Claim 1).

4.4. Proof of Theorem 3

Putting everything together, we can now complete the proof of the Bulow-Klemperer-type theorem for matching markets.

Proof. We need to show that for every matching environment with n i.i.d. regular bidders and m items, the expected revenue of the VCG mechanism with m additional bidders is at least the optimal expected revneue. By Claim 2, the VCG payment for item j in the augmented environment is at least the revenue from selling item j by deferred allocation, which by Claim 1 is equal in expectation to $\text{Vic}_j(n+1)$. Summing up over all items, the total expected VCG revenue in the augmented environment is at least $\sum_j \text{Vic}_j(n+1)$, and by Lemma 3 this upper-bounds the optimal expected revenue in the original environment. □

4.5. The $m \geq n$ Case

In matching markets where items are more plentiful than bidders, the following Bulow-Klemperer-type theorem provides an alternative to Theorem 3, in which the required demand augmentation is n instead of m bidders.[9] Two additional differences in comparison to previous Bulow-Klemperer-type theorems are that the VCG mechanism is required to be supply-limiting, and the revenue guarantee is an approximation.

Theorem 4 **(Bulow-Klemperer-Type Theorem for Matching with $m \geq n$).** *For every matching environment with n i.i.d. regular bidders and $m \geq n$ items, the expected revenue of the* VCG$^{\leq n}$ *mechanism with n additional bidders is at least an n/m-fraction of the optimal expected revenue in the original market. In other words,* VCG$^{\leq n}$ *with n additional bidders is robustly n/m-optimal.*

Proof. The proof is similar to that of Theorem 3, with the following adjustments.

Consider the bounds in Section 4.2 above. While the upper bound on the optimal expected revenue in Lemma 3 holds and is sufficient, the lower bound on VCG payments in Lemma 4 holds but needs to be strengthened. In the augmented environment, VCG$^{\leq n}$ allocates n out of the m items to n out of the $2n$ bidders. Therefore the VCG payment for item j not only exceeds the value of any unallocated bidder for j, but also exceeds the value of any unallocated bidder for any unallocated item. We shall refer to the highest such value among unallocated bidders and items as the *global* lower bound on VCG payments, and denote it by G.

Now to relate the bounds as in Section 4.3, we use a slightly modified deferred allocation procedure (Algorithm 2).

Algorithm 2 Selling Item j by Deferred Allocation: The Case of $m \geq n$

Given a matching environment with $2n$ bidders and m items, and an item j:

1. Find a maximum matching of $n-1$ items other than j to the bidders. Let U be the set of $n+1$ bidders who remain unallocated.
2. Run the Vickrey auction to sell item j to bidder set U.

Observe that Claims 1 and 2 continue to hold when Algorithm 1 is replaced by Algorithm 2. For Theorem 3 these claims were sufficient to complete the proof, by the following chain of arguments: All items are allocated by VCG in the augmented environment (since it is welfare-maximizing and there are more bidders than items); the VCG payment for item j is at least the revenue from selling j by deferred allocation (by Claim 2); the deferred allocation revenue is equal in expectation to $\text{Vic}_j(n+1)$ (by Claim 1); and $\sum_j \text{Vic}_j(n+1)$ is at least the optimal expected revenue (by Lemma 3). For Theorem 4 we need an additional charging argument — and an approximation factor — since only n out of m items are allocated by VCG$^{\leq n}$.

For every item $j \in [m]$ there are two cases:

1. If j is allocated, then the VCG payment for j is at least the revenue from selling j by deferred allocation (Claim 2).

2. If j is not allocated, then the VCG payment for any allocated item j' is at least the global lower bound G. A straightforward adaptation of the argument in Claim 2 shows that G is an upper bound on the revenue from selling j by deferred allocation (Algorithm 2).

To complete the proof, we charge the VCG payments for the n allocated items against the aggregate revenue from selling each of the m items by deferred allocation, where the latter is equal in expectation to $\sum_j \mathrm{Vic}_j(n+1)$. This leads to the approximation factor of $\frac{n}{m}$. □

5. Alternatives to Augmenting Demand: Limiting Supply and Large Markets

As discussed in Section 1, limiting supply has a similar effect on competition as increasing demand. In this section, the intuitive relation between supply, demand and competition is made precise: we prove approximation guarantees for the Vickrey auction with limited supply by reducing the market to one with augmented demand, and then applying an appropriate Bulow-Klemperer-type theorem. If a market is "large" — i.e., if its supply is limited to begin with in comparison to its demand — our results imply that Vickrey with no supply limit guarantees asymptotically optimal revenue. We begin with an overview of the results and techniques in this section (Section 5.1), then address multi-unit markets (Section 5.2) and matching markets (Section 5.3). For extensions to constraints or asymmetric bidders see Section 6.

5.1. Overview of Our Approach

To demonstrate our approach consider the following example:

Example 1. A *digital goods market* is a multi-unit market with $k = n$ units.

Digital goods markets generalize the example of markets for artwork prints described in Section 1. Consider now the following generic supply-limiting mechanism, which can be applied to many market settings including digital goods:

Mechanism 3 Supply-Limiting Mechanism

1. Set a supply limit $\ell = n/2$ equal to half the number of bidders.
2. Run the Vickrey auction subject to supply limit ℓ.

The "Vickrey auction subject to supply limit ℓ" in Step 2 is a simple variation on the standard second-price auction: in the digital goods example it assigns copies to the ℓ buyers with highest bids (even though there are enough copies for all buyers), and charges them each the $(\ell+1)$th highest bid. The resulting mechanism is simple and natural, and does not rely on knowledge of value distributions. It does rely on the ability of the seller to withhold supply, which may be infeasible in some markets (cf. Coase 1972) but available in others (see the art example above). (A related notion is the ability of some firms to slow down their "service rate" even if the technological cost of speeding up is zero; this can be used to signal quality through queues and thus increase revenue (Debo et al. 2012).)

To analyze the revenue of Mechanism 3 applied to Example 1, we use the following reduction:

Reduction 4 Digital Goods

---

0. **Start with original market with n buyers and n units**

Denote the optimal expected revenue by OPT.

1. **Restrict to market with $n/2$ buyers and supply limit $n/2$**

The optimal expected revenue is OPT /2, by subadditivity of revenue in the buyers (Lemma 6 below) and since the supply limit has no effect here.

2. **Augment to get market with n buyers and supply limit $n/2$**

The expected revenue of the Vickrey auction is OPT /2, by the Bulow-Klemperer-type theorem for multi-unit markets applied to the restricted market.

We now explain how Reduction 4, together with an appropriate Bulow-Klemperer-type theorem, establishes that Mechanism 3 guarantees half the optimal revenue in expectation when applied to Example 1. The argument is as follows: Starting with the *original market*, define a new market by dropping half of the bidders and setting a supply limit of $\ell = n/2$. Consider the resulting *restricted market* with half of the original bidders and corresponding supply limit. One can show that if we were to restrict the optimal mechanism to run on this market instead of the original one, its expected revenue would have been at least half of its original expected revenue. Now conceptually add back the $n/2$ removed bidders but without changing the supply to get the *augmented market*, and run the Vickrey auction. It follows from the Bulow-Klemperer theorem for multi-unit markets that the expected revenue is at least as high as the optimal expected revenue for the restricted market. Therefore the supply-limiting mechanism guarantees at least half of the optimal expected revenue in the original market.

To extend our approach to multi-item matching markets, we first need to define a notion of limiting supply where the supply is heterogeneous rather than homogeneous. Subjecting a multi-item auction to a *supply limit* ℓ means that no more than ℓ goods can be assigned, but places no limitation on which of the total goods to allocate. This lets the market "choose" which part of the supply to limit, and is in line with our robust approach (a seller with no knowledge of how the values for the different goods are distributed cannot make this choice without risking a big loss in revenue).

The simple supply-limiting mechanism we designed for multi-unit markets (Mechanism 3) is now well-defined for multi-item markets as well, and we can prove the following theorem (see Theorem 8 for a formal statement):

Theorem 5 **(Supply-Limiting Mechanism for Matching Markets (Informal))**. *For every matching market with $n \geq 2$ bidders and m goods, the expected revenue of Mechanism 3 is at least a constant fraction of the optimal expected revenue.*

Theorem 5 is interesting in that it shows a simple robust mechanism that achieves a constant fraction of the optimal expected revenue, independently of the size of the market. Whether this is the best possible fraction by such a mechanism with no augmented resources is left as an open problem. The proof of Theorem 5 is via the following generalization of Reduction 5, instantiated with the appropriate Bulow-Klemperer-type theorems. Details appear in Section 5.3.

Reduction 5 Approximation Guarantees via Bulow-Klemperer-Type Theorems

0. **Start with original market with n buyers**

Denote the optimal expected revenue by OPT.

1. **Restrict to market with $< n$ buyers and supply limit ℓ**

The optimal expected revenue is a constant fraction of OPT, by subadditivity.

2. **Augment to get market with n buyers and supply limit ℓ**

The expected revenue of the Vickrey auction is a constant fraction of OPT, by a suitable Bulow-Klemperer-type theorem.

5.2. Warm-up: Alternatives to Augmenting Demand for Multi-Unit Markets

The following is a slightly generalized version of Bulow and Klemperer (1996) (generalized to include a possible supply limit):

Theorem 6 **(Bulow-Klemperer-Type Theorem for Multi-Unit Markets)**. *For every k-unit environment with i.i.d. regular bidders and supply limit ℓ, the expected revenue of the Vickrey auction with $\min\{k,\ell\}$ additional bidders is at least the optimal expected revenue in the original market. In other words, Vickrey with $\min\{k,\ell\}$ additional bidders is robustly 1-optimal.*

Instantiating the general reduction (Reduction 5) with the above Bulow-Klemperer-type theorem (Theorem 6) yields the following (for simplicity of presentation assume the number of bidders n is even).[10]

THEOREM 7 **(Mechanism with No Additional Bidders for Multi-Unit Markets)**. *For every k-unit environment with $n \geq 2$ i.i.d. regular bidders, the expected revenue of the supply-limiting mechanism* $\text{Vic}^{\leq n/2}$ *is at least a* $\max\{\frac{1}{2}, \frac{n-k}{n}\}$*-fraction of the optimal expected revenue. In other words, Vickrey with supply limit $n/2$ is robustly α-optimal for $\alpha = \max\{\frac{1}{2}, \frac{n-k}{n}\}$.*

The supply limit of $n/2$ only "kicks in" when the number of units k exceeds $n/2$, and in this case we get a $1/2$-approximation. If k is no more than $n/2$ to begin with, the competition is inherently high and Vickrey (with no supply limit) provides an $\frac{n-k}{n}$-approximation. In particular, in a large market where $k \ll n$, Vickrey is robustly optimal in the limit.

Proof by Submodularity of Revenue

Proof of Theorem 7. We instantiate Reduction 5 as follows. To go from the original market to the restricted market, remove $\min\{\frac{n}{2}, k\}$ bidders from the original market, and if $k > \frac{n}{2}$ set a supply limit of $\ell = \frac{n}{2}$.

Analysis. We first claim that the restriction of the original market maintains at least a fraction of $\max\{\frac{1}{2}, \frac{n-k}{n}\}$ of the optimal expected revenue in the original market. This is because, as shown by Dughmi et al. (2012), the expected optimal revenue as a function of the bidder set is *submodular*.[11] Revenue submodularity means decreasing marginal returns to the expected revenue as more bidders are added, so the first $\max\{\frac{n}{2}, n-k\}$ bidders already capture at least a $\max\{\frac{1}{2}, \frac{n-k}{n}\}$-fraction of the optimal expected revenue. Limiting the supply to $\frac{n}{2}$ when $k > \frac{n}{2}$ has no effect since in this case the number of bidders remaining in the restricted environment is $\frac{n}{2}$.

We can now apply the Bulow-Klemperer-type theorem for multi-unit markets (Theorem 6) to the restricted environment. In the first case, $k > \frac{n}{2}$ and the restricted environment has $\frac{n}{2}$ bidders, k units and supply limit $\ell = \frac{n}{2}$. In the second case, $k \leq \frac{n}{2}$ and the restricted environment has $n-k$ bidders, k units and no supply limit (i.e., $\ell = k$). In both cases, by Theorem 6 the expected revenue of Vickrey with $\min\{\frac{n}{2}, k\}$ additional bidders is at least the optimal expected revenue in the restricted environment. So running Vickrey with $\min\{\frac{n}{2}, k\}$ additional bidders on the restricted environment is a $\max\{\frac{1}{2}, \frac{n-k}{n}\}$-approximation to the optimal expected revenue in the original environment. But this is equivalent to running the supply-limiting mechanism $\text{Vic}^{\leq n/2}$ on the original environment, completing the proof. □

Tightness The approximation factor in Theorem 7 is asymptotically tight.

PROPOSITION 1. *For every $0 < \gamma < 1$, consider the supply-limiting mechanism* $\text{Vic}^{\leq \gamma n}$*. There exists an n-unit environment with n i.i.d. regular bidders such that the expected revenue of* $\text{Vic}^{\leq \gamma n}$ *is at most a $(\frac{1}{2} + o(1))$-fraction of the optimal expected revenue.*

Proof. Consider first the case that $1/n \leq \gamma \leq 1/2$, i.e., the supply limit is severe. Let the value distribution F be the uniform distribution over the support $[1, 1+\epsilon]$ for a sufficiently small parameter $\epsilon = \epsilon(n)$. The optimal expected revenue is roughly n, while $\text{Vic}^{\leq \gamma n}$ can extract as revenue at most $\gamma n(1+\epsilon) \leq n/2 + o(1)$.

Now suppose $1/2 < \gamma \leq \frac{n-1}{n}$. For sufficiently large H, let the value distribution be $F(z) = \frac{z}{1+z}$ over the support $[0, H]$ with a point mass of $\frac{1}{1+H}$ at H.[12] The optimal expected revenue is at least the expected revenue achieved by offering a posted price H to every one of the n bidders, which extracts $H(1 - \frac{H}{1+H}) = \frac{H}{1+H} \approx 1$ from every bidder in expectation. In comparison, the expected revenue in $\text{Vic}^{\leq \gamma n}$ comes from the $(\gamma n + 1)$st highest bid. This bid is concentrated around $z = \frac{1-\gamma}{\gamma}$, the value of z such that $F(z) = 1 - \gamma$. So VCG achieves an expected revenue of roughly $\frac{1-\gamma}{\gamma}\gamma n = (1-\gamma)n < \frac{n}{2}$.
□

5.3. Alternatives to Augmenting Demand for Matching Markets

Instantiating the general reduction (Reduction 5) with one of the Bulow-Klemperer-type theorems for matching yields the following (for simplicity of presentation assume the number of bidders n is even).

THEOREM 8 **(Mechanism with No Additional Bidders for Matching Markets)**.
For every matching environment with $n \geq 2$ i.i.d. regular bidders and m items, let $\alpha = \max\{\frac{n-m}{n}, \min\{\frac{1}{2}, \frac{n}{4m}\}\}$. Then the expected revenue of the supply-limiting mechanism $\text{VCG}^{\leq n/2}$ is at least an α-fraction of the optimal expected revenue. In other words, VCG with supply limit $n/2$ is robustly α-optimal.

Achieving a good revenue guarantee becomes more difficult as the number of items m grows relatively to the number of bidders n, since the inherent competition among the bidders is split across different items. Accordingly, the fraction α in Theorem 8 depends on the parameters n and m of the environment as follows:

- If $m \leq \frac{n}{2}$ then $\alpha = \frac{n-m}{n}$, i.e., the approximation gets better as m becomes smaller, and for $m = \frac{n}{2}$ we get $\alpha = \frac{1}{2}$.
- If $m \geq \frac{n}{2}$ then $\alpha = \frac{n}{4m}$, and in particular for $m = n$ we get $\alpha = \frac{1}{4}$.

Note that when $m \leq n/2$, the supply limit of $\text{VCG}^{\leq n/2}$ has no effect, that is, the revenue guarantee is achieved by simply applying the VCG mechanism. In particular, in a large market where $m \ll n$, VCG is robustly optimal in the limit.

For the opposite case of $m \gg n$, Theorem 8 does not state a constant approximation. However it still holds in this case that VCG with a supply limit is robustly α'-optimal where $1/\alpha'$ is a constant, albeit with a more involved proof (Yan 2012). Theorem 8 also applies without change to *multi-unit* matching markets, in which there are k_j copies of every item j, and a total of $m = \sum_j k_j$ units overall.

Proof by Revenue Subadditivity The following lemma is used to prove Theorem 8 and may also be of independent interest. It states that in any market environment — including one with asymmetric bidders whose values are drawn independently but not identically — the optimal expected revenue achievable from bidder sets S, T separately is at least the optimal expected revenue achievable from their union. Let $\mathrm{OPT}(\cdot)$ map a bidder set to its optimal expected revenue. Then:

LEMMA 6 **(Subadditivity of Optimal Expected Revenue in Bidder Set)**. *For every auction environment with bidder subsets S and T, $\mathrm{OPT}(S) + \mathrm{OPT}(T) \geq \mathrm{OPT}(S \cup T)$.*

Proof. It is not hard to see that $\mathrm{OPT}(\cdot)$ is monotone, so without loss of generality we can assume that S and T are disjoint. Let M be the optimal mechanism for $S \cup T$. For every value profile $\mathbf{v}_T$ of the bidders in T, we define the mechanism $M_{\mathbf{v}_T}$, which gets bids from the bidders in S and simulates M by using $\mathbf{v}_T$ as the bids of bidders in T. By an averaging argument, there exists a vector $\mathbf{v}_T$ such that mechanism $M_{\mathbf{v}_T}$'s expected revenue is at least the part of the optimal expected revenue of mechanism M that is charged to the bidders in S. On the other hand, the expected revenue of $M_{\mathbf{v}_T}$ is bounded above by $\mathrm{OPT}(S)$. Similarly, the part of the optimal expected revenue that is charged to the bidders in T is bounded above by $\mathrm{OPT}(T)$. This completes the proof. □

A corollary of the subadditivity lemma is that removing bidders from an i.i.d. environment until an α-fraction of the original bidders remains maintains an α-fraction of the optimal expected revenue. By symmetry:

COROLLARY 1. *For every auction environment with n i.i.d. bidders and for every integer c that divides n, $\mathrm{OPT}(n/c) \geq \frac{1}{c}\mathrm{OPT}(n)$.*

We can now prove Theorem 8 by instantiating the general reduction (Reduction 5).

Proof of Theorem 8. Assume first that $m \leq n/2$. We instantiate Reduction 5 as follows: to go from the original market to the restricted market, remove m bidders from the original market. By Corollary 1, this restriction on the original market maintains at least an $\frac{n-m}{n}$-fraction of the original optimal expected revenue. We can now apply the Bulow-Klemperer-type theorem for matching markets (Theorem 3) to the restricted market, which has $n - m$ bidders and m items. The expected revenue of VCG with m additional bidders is at least the optimal expected revenue in the restricted market. But this is equivalent to running VCG on the original market, completing the proof for $m \leq n/2$.

Now assume that $m \geq n/2$. We instantiate Reduction 5 as follows: to go from the original market to the restricted market, remove $n/2$ bidders from the original market. By Corollary 1, this restriction on the original market maintains at least a $\frac{1}{2}$-fraction of the original optimal expected revenue. We can now apply the Bulow-Klemperer-type theorem for matching markets with more items than bidders (Theorem 4) to the restricted market, which has $n/2$ bidders and $m \geq n/2$ items. The expected revenue of VCG$^{\leq n/2}$ with $n/2$ additional bidders is at least an $\frac{n}{2m}$-fraction of the optimal expected revenue in the restricted market, and so an $\frac{n}{4m}$-fraction of the original optimal expected revenue. But this is equivalent to running VCG$^{\leq n/2}$ on the original market, completing the proof. □

6. Extensions and Limitations

In this section we extend our results to markets with a matroid constraint on who can win simultaneously and to asymmetric markets. In Section 6.4 we discuss limitations.

6.1. Matroid Markets

Recall that a matroid environment is a single-parameter environment in which the set system $([n],\mathcal{I})$ of bidders and feasible allocations forms a *matroid*. A matroid satisfies three axioms (Oxley 1992): (A1) $\emptyset \in \mathcal{I}$, (A2) $\mathcal{I}$ is downward-closed, and (A3) if $S,T \in \mathcal{I}$ and $|S| < |T|$ then there is a bidder $t \in T \setminus S$ that can be added to S such that $S \cup \{t\} \in \mathcal{I}$. The *rank* ρ of a matroid is the size of its maximal independent sets or *bases*, and the *packing number* κ of a matroid is its maximum number of disjoint bases. We will use the fact that an *intersection* of a matroid $([n],\mathcal{I})$ with a *u-uniform* matroid is a new matroid $([n],\mathcal{I}')$, in which a set of bidders S belongs to $\mathcal{I}'$ if and only if $S \in \mathcal{I}$ and $|S| \leq u$.

Example 2. A *job scheduling market* is a multi-unit market where the units are slots for running jobs on a machine, and a subset of bidders is feasible if each bidder's job can be matched to a slot between its arrival time and deadline.

Job scheduling markets are an example of matroid environments. In Example 2, if there are four jobs arriving at time 0 of which two must be finished by time 1 and two must be finished by time 2, then the rank is $\rho = 2$ and the packing number is $\kappa = 2$.

Dughmi et al. (2012) show a Bulow-Klemperer-type result for matroid environments:

Theorem 9 **(Bulow-Klemperer-Type Theorem for Matroid Markets)**. *For every matroid environment with i.i.d. regular bidders, the expected revenue of the Vickrey auction with an additional basis of bidders is at least the optimal expected revenue in the original market. In other words, Vickrey with an additional basis of bidders is robustly optimal.*

Applied to the job scheduling example (Example 2), Theorem 9 implies the following: When the best schedule is able to feasibly match $\rho \leq k$ bidders to the k available job slots, the expected revenue of the revenue-optimal mechanism with n bidders is at most that of the Vickrey auction with $n+\rho$ bidders.

We now combine Theorem 9 with our general reduction to prove the following result, in which the approximation depends on the inherent amount of competition in the market, as measured not by the number of bidders n but rather by their packing number κ. For simplicity of presentation assume the rank ρ is even.

THEOREM 10 **(Mechanism with No Additional Bidders for Matroid Markets)**. *For every matroid environment with $n \geq 2$ i.i.d. regular bidders, rank ρ and packing number κ, let $\ell = \rho/2$ if $\kappa = 1$ and $\ell = \rho$ otherwise. Then the expected revenue of the supply-limiting mechanism* $\mathrm{Vic}^{\leq \ell}$ *is at least a* $\max\{\frac{1}{4}, \frac{\kappa-1}{\kappa}\}$*-fraction of the optimal expected revenue. In other words,* $\mathrm{Vic}^{\leq \ell}$ *is robustly* $\max\{\frac{1}{4}, \frac{\kappa-1}{\kappa}\}$*-optimal.*

Note that when $\kappa \geq 2$, the supply limit of $\ell = \rho$ has no effect, and Vickrey is robustly optimal in the limit as $\kappa \to \infty$.

Proof. We instantiate Reduction 5 as follows: If $\kappa = 1$, intersect the original matroid with a $\frac{\rho}{2}$-uniform matroid to get a new matroid $([n], \mathcal{I}')$ with rank $\rho' = \frac{\rho}{2}$ and packing number $\kappa' \geq 2$. Otherwise, if $\kappa \geq 2$, simply set the new matroid $([n], \mathcal{I}')$ to be $([n], \mathcal{I})$. To go from the original market to the restricted market, remove from the original market a basis of bidders of size ρ' according to the matroid $([n], \mathcal{I}')$, and set the matroid of the restricted market to be $([n], \mathcal{I}')$.

Analysis: If $\kappa = 1$, intersecting with the uniform matroid maintains at least a $\frac{1}{2}$-fraction of the original optimal expected revenue. Removing a basis of bidders maintains at least a $\frac{\kappa'-1}{\kappa'}$-fraction. We can now apply the Bulow-Klemperer-type theorem for matroids (Theorem 9) to the restricted market with $n - \rho'$ bidders and matroid $([n], \mathcal{I}')$. The expected revenue of Vickrey with an additional basis of bidders is at least the optimal expected revenue in the restricted market, and so a $\max\{\frac{1}{4}, \frac{\kappa-1}{\kappa}\}$-fraction of the original optimal expected revenue. But this is equivalent to running $\mathrm{Vic}^{\leq \ell}$ on the original market, completing the proof. □

6.2. Asymmetric Bidders: Augmenting Demand

An *attribute-based* environment is a k-unit environment with n bidders, each of whom has a publicly-observable attribute $a = a(i)$ that determines a non-publicly-known value distribution F_a (Dhangwatnotai et al. 2010). Bidders' values in an attribute-based environment are thus independently but not identically distributed. Attributes enable the incorporation of prior information into our model regarding which bidders are alike, while still avoiding assumptions about the value distributions themselves. In fact, our results can be interpreted as an encouragement to invest in this particular kind of prior information, which entails grouping similar bidders together rather than learning distributions. Examples of attributes are bidding styles such as “bargain-hunter” or “aggressive” on `eBay.com`, or in sponsored search and online advertising, advertiser features such as location. Throughout we assume *non-singular* attribute-based environments, where no bidder’s attribute is unique. I.e., for every attribute a, let n_a denote the number of bidders in the environment with attribute a; then $n_a > 0 \implies n_a \geq 2$.

In this section we prove a Bulow-Klemperer-type theorem for attribute-based environments.[13] Let $\text{Vic}^{\leq \ell_a}$ be the Vickrey mechanism with a *local* supply limit ℓ_a for every a, which limits the number of bidders with attribute a who can win simultaneously. The proof of the following theorem uses a *commensuration* argument of Hartline and Roughgarden (2009), and applies the FKG inequality (Alon and Spencer 2008) to solve dependency issues.

Theorem 11 **(Bulow-Klemperer-Type Theorem for Asymmetric Markets)**. *For every attribute-based environment with n_a regular bidders per attribute and k units, the expected revenue of the $\text{Vic}^{\leq n_a}$ auction with $\min\{n_a, k\}$ additional bidders per attribute is at least a $\frac{1}{2}$-fraction of the optimal expected revenue in the original market. In other words, $\text{Vic}^{\leq n_a}$ with $\min\{n_a, k\}$ additional bidders per attribute is robustly $\frac{1}{2}$-optimal.*

Proof. Let $W^{\text{OPT}} = W^{\text{OPT}}(\mathbf{v}), W^{\text{Vic}} = W^{\text{Vic}}(\mathbf{v})$ denote the winning bidders chosen by the optimal mechanism in the original environment and by $\text{Vic}^{\leq n_a}$ in the augmented environment, respectively, given a value profile $\mathbf{v}$ of both original and augmenting bidders. Hartline and Roughgarden (2009) show that to prove a 1/2-approximation it suffices to establish two commensuration conditions among the two mechanisms:

(C1) $\mathbb{E}_{\mathbf{v}}[\sum_{i \in W^{\text{Vic}} \setminus W^{\text{OPT}}} \phi_i] \geq 0$,

(C2) $\mathbb{E}_{\mathbf{v}}[\sum_{i \in W^{\text{Vic}} \setminus W^{\text{OPT}}} p_i(\mathbf{v})] \geq \mathbb{E}_{\mathbf{v}}[\sum_{i \in W^{\text{OPT}} \setminus W^{\text{Vic}}} \phi_i]$,

where ϕ_i is the virtual value of bidder i.

The proof of (C2) in (Hartline and Roughgarden 2009, Lemma 4.5) holds in our setting. In contrast, proving (C1) in our setting turns out to be technically challenging due to dependencies among the random bidder sets W^{OPT} and W^{Vic}. We use an auxiliary allocation procedure (Algorithm 6), and rely on the fact that in our setting, $\text{Vic}^{\leq n_a}$ applies a simple greedy algorithm: it rejects all but the top n_a bidders per attribute, and allocates the units to the $\leq k$ highest remaining bidders.

For the remainder of the proof, fix an attribute a, and let B_a be the set of $n_a + \min\{n_a, k\}$ bidders with attribute a in the augmented environment. Fix the values of the original bidders in B_a as well as the values of all bidders with different attributes, and let $\mathbf{v}_a$ denote the (random) value profile of the augmenting bidders in B_a. Let $W_a^{\text{OPT}} = B_a \cap W^{\text{OPT}}$ denote the bidders in B_a who win in the optimal mechanism, and let $W_a^{\text{Vic}} = W_a^{\text{Vic}}(\mathbf{v}_a) = B_a \cap W^{\text{Vic}}$ denote the bidders in B_a who win in

Algorithm 6 Auxiliary Allocation Procedure

In the augmented environment, given W_a^{OPT}, allocate the k units such that the welfare is maximized subject to the constraint that all bidders in W_a^{OPT} win.

$\text{Vic}^{\leq n_a}$. We can now define our auxiliary procedure in Algorithm 6. Let $W_a^{\text{Aux}} = W_a^{\text{Aux}}(\mathbf{v}_a)$ denote the bidders in B_a who win in this procedure, and observe $W_a^{\text{OPT}} \subseteq W_a^{\text{Aux}} \subseteq B_a$.

We now establish two claims, the first of which relates the auxiliary procedure to $\text{Vic}^{\leq n_a}$ and the second of which relates it to the optimal mechanism.

Claim 3. *For every a and every value profile $\mathbf{v}_a$ of the augmenting bidders with attribute a, the bidders in $W_a^{\text{Vic}} \setminus W_a^{\text{Aux}}$ have non-negative virtual values.*

Proof of Claim 3. Fix $\mathbf{v}_a$ and consider the allocation of $\text{Vic}^{\leq n_a}$ in comparison to that of the auxiliary procedure. $\text{Vic}^{\leq n_a}$ is free to replace bidders in W_a^{OPT}. Since $\text{Vic}^{\leq n_a}$ is greedy, each replacement from B_a will have a higher value than the replaced bidder in W_a^{OPT}, and therefore (using regularity) also a higher virtual value. The proof follows by noticing that all bidders in W_a^{OPT} have non-negative virtual values (Lemma 1). □

Claim 4. *For every a, for every $h \leq \min\{n_a, k\}$, in expectation over the value profile $\mathbf{v}_a$ of the augmenting bidders with attribute a, summing over the h highest virtual values of the bidders in $W_a^{\text{Aux}} \setminus W_a^{\text{OPT}}$ results in a non-negative total virtual value.*

Proof of Claim 4. Given $\mathbf{v}_a$, denote by $\psi_{(1)}, \psi_{(2)}, \ldots$ (where $\psi_{(i)} = \psi_{(i)}(\mathbf{v}_a)$) the virtual values of the bidders in $B_a \setminus W_a^{\text{OPT}}$, sorted in decreasing order of both values and virtual values. Let $\mathbb{1}_{(i)} = \mathbb{1}_{(i)}(\mathbf{v}_a)$ indicate whether the i-th bidder in $B_a \setminus W_a^{\text{OPT}}$ wins in the auxiliary procedure, and let $q_{(i)}$ be the probability that $\mathbb{1}_{(i)} = 1$ over a random choice of $\mathbf{v}_a$. The expected sum of the highest $\min\{n_a, k\}$ virtual values of the bidders in $W_a^{\text{Aux}} \setminus W_a^{\text{OPT}}$ can be written as $\sum_{i \leq \min\{n_a,k\}} \mathbb{E}_{\mathbf{v}_a}[\psi_{(i)} \cdot \mathbb{1}_{(i)}]$. To complete the proof of Claim 4 it is sufficient to show that this expression is non-negative.

Notice that like $\text{Vic}^{\leq n_a}$, the auxiliary procedure is greedy in nature. Thus if the i-th bidder in $B_a \setminus W_a^{\text{OPT}}$ wins in the auxiliary procedure, and the values of bidders in $B_a \setminus W_a^{\text{OPT}}$ increase, then the i-th highest bidder will still win. Formally, for two value profiles $\mathbf{v}_a \leq \mathbf{v}'_a$ we have $\mathbb{1}_{(i)}(\mathbf{v}_a) \leq \mathbb{1}_{(i)}(\mathbf{v}'_a)$ for every i. This positive correlation between how high the values and virtual values are and between the probability of winning in the auxiliary procedure allows us to apply the FKG inequality (Alon and Spencer 2008) as follows:

$$\sum_{i \leq \min\{n_a,k\}} \mathbb{E}_{\mathbf{v}_a}[\psi_{(i)} \cdot \mathbb{1}_{(i)}] \geq \sum_{i \leq \min\{n_a,k\}} \mathbb{E}_{\mathbf{v}_a}[\psi_{(i)}] \cdot \mathbb{E}_{\mathbf{v}_a}[\mathbb{1}_{(i)}]$$

$$= \sum_{i \le \min\{n_a,k\}} \mathbb{E}_{\mathbf{v}_a}[\psi_{(i)}] \cdot q_{(i)}$$
$$= \sum_{i \le \min\{n_a,k\}} \left(\mathbb{E}_{\mathbf{v}_a} \left[\sum_{i'=1}^{i} \psi_{(i')} \right] \cdot \left(q_{(i)} - q_{(i+1)} \right) \right),$$

where the first inequality is by FKG, and where we set $q_{(i+1)}$ for $i = \min\{n_a, k\}$ to 0.

It is not hard to see that $q_{(i)}$ is decreasing in i. Therefore it suffices to prove that $\sum_{i'=1}^{i} \mathbb{E}_{\mathbf{v}_a}[\psi_{(i')}] \ge 0$ for every $i \le \min\{n_a, k\}$. This is the sum of the expected virtual values of the top i bidders in $B_a \setminus W_a^{\mathrm{OPT}}$. Observe that the sum of the expected virtual values of any i augmented bidders in $B_a \setminus W_a^{\mathrm{OPT}}$ equals 0, and there are at least $\min\{n_a, k\}$ such bidders. It follows that this sum for the top i bidders is nonnegative. $\square$

We now use Claims 3 and 4 to complete the proof of (C1). Still holding attribute a fixed, rewrite $\mathbb{E}_{\mathbf{v}_a}[\sum_{i \in W_a^{\mathrm{Vic}} \setminus W_a^{\mathrm{OPT}}} \phi_i]$ as $\mathbb{E}_{\mathbf{v}_a}[\sum_{i \in W_a^{\mathrm{Vic}} \setminus W_a^{\mathrm{Aux}}} \phi_i] + \mathbb{E}_{\mathbf{v}_a}[\sum_{i \in W_a^{\mathrm{Vic}} \cap (W_a^{\mathrm{Aux}} \setminus W_a^{\mathrm{OPT}})} \phi_i]$. The left-hand side is non-negative by Claim 3. We consider two cases for the right-hand side, which by greediness of the auxiliary procedure and $\mathrm{Vic}^{\le n_a}$ are the only possible cases:

1. $(W_a^{\mathrm{Aux}} \setminus W_a^{\mathrm{OPT}}) \subseteq W_a^{\mathrm{Vic}}$: In this case $W_a^{\mathrm{Vic}} \cap (W_a^{\mathrm{Aux}} \setminus W_a^{\mathrm{OPT}}) = W_a^{\mathrm{Aux}} \setminus W_a^{\mathrm{OPT}}$, and so by Claim 4 the sum of virtual values is non-negative in expectation.

2. $W_a^{\mathrm{Vic}} \subseteq W_a^{\mathrm{Aux}}$: In this case, $W_a^{\mathrm{Vic}} \setminus W_a^{\mathrm{OPT}}$ is a subset of the highest $\min\{n_a, k\}$ bidders in $W_a^{\mathrm{Aux}} \setminus W_a^{\mathrm{OPT}}$, and so by Claim 4 the sum of virtual values over this subset is non-negative in expectation.

We have shown that $\mathbb{E}_{\mathbf{v}_a}[\sum_{i \in W_a^{\mathrm{Vic}} \setminus W_a^{\mathrm{OPT}}} \phi_i] \ge 0$. Taking expectation and summing over all attributes we get $\mathbb{E}_{\mathbf{v}}[\sum_{i \in W^{\mathrm{Vic}} \setminus W^{\mathrm{OPT}}} \phi_i] \ge 0$, completing the proof of (C1) and Theorem 11. $\square$

6.3. Asymmetric Bidders: Limiting Supply

Consider an attribute-based environment as defined in Section 6.2. For simplicity of presentation assume that n_a, the number of bidders with attribute a, is even for every a. Recall that $\mathrm{Vic}^{\le n_a/2}$ is the Vickrey mechanism with a local supply limit $n_a/2$ for every a, meaning that no more than half the bidders with the same attribute can win simultaneously. We now show that $\mathrm{Vic}^{\le n_a/2}$ is a good supply-limiting mechanism for attribute-based environments. Note that it is considerably simpler than Myerson's optimal mechanism for asymmetric markets, which requires computing different virtual value functions for different attributes.

THEOREM 12 **(Supply-Limiting Mechanism for Asymmetric Markets)**. *For every attribute-based environment with n_a regular bidders per attribute, the expected revenue of the supply-limiting mechanism $\mathrm{Vic}^{\le n_a/2}$ is at least a $\frac{1}{4}$-fraction of the optimal expected revenue. In other words, $\mathrm{Vic}^{\le n_a/2}$ is robustly $\frac{1}{4}$-optimal.*

Proof. We instantiate Reduction 5 as follows: to go from the original market to the restricted market, remove $\min\{\frac{n_a}{2}, k\}$ bidders with attribute a from the original market. By submodularity (Dughmi et al. 2012), this restriction on the original market maintains at least a $\frac{1}{2}$-fraction of the original optimal expected revenue (since we removed at most half of the bidders). We can now apply the Bulow-Klemperer-type theorem for asymmetric markets (Theorem 11) to the restricted market, which has $\max\{\frac{n_a}{2}, n_a - k\}$ bidders per attribute and k units. The expected revenue of $\mathrm{Vic}^{\leq n_a/2}$ with $\min\{\frac{n_a}{2}, k\}$ additional bidders is at least a $\frac{1}{2}$-fraction of the optimal expected revenue in the restricted market. But this is equivalent to running $\mathrm{Vic}^{\leq n_a/2}$ on the original market, completing the proof. □

6.4. Limitations

Theorem 3 is tight in the sense that fewer than $n+m$ buyers may fail to guarantee the optimal revenue in the original environment: Consider the special case in which $n=1$ and all m items are identical with values drawn from a point-mass distribution. Then unless at least m more bidders are added, the revenue achieved by the VCG mechanism is zero (there is no competition for the m items among the $\leq m$ bidders).

It was recently observed by Eden et al. (2016) that with arbitrarily correlated buyers, the VCG mechanism with additional buyers cannot compete with the optimal revenue in the original environment.

7. Conclusion and Discussion

Robustness has long been recognized as an important design principle in optimization; here we apply it to optimal mechanism design. We study enhanced competition as a means for designing robust and simple auctions, whose revenue guarantees provably exceed or approximate those of the optimal auction, even when the latter is not well understood. Our main contributions are as follows:

- The problem of designing optimal auctions in matching markets is challenging even given full distributional information and regardless of robustness considerations. Yet the robust approach of prior-independence can help us get a better understanding of the optimal auction.
- Prior-independence is similar to a standard approach in computationally-hard optimization: let a polynomial-time algorithm (possibly with augmented resources) compete with what can be achieved by an algorithm with unlimited running time. Here we let a prior-independent mechanism compete with the revenue that can be achieved by a mechanism with access to full distributional information.

- To achieve prior-independence we develop mechanisms based on competition enhancement, primarily by increasing demand (as well as by the complementary approach of limiting supply, related to the former via a general reduction). Our main result is that Vickrey with m more bidders guarantees good revenue in both single- and multi-parameter markets and for many value distributions. As an alternative to adding bidders, limiting the number of allocations to half the size of the market and running Vickrey has good revenue guarantees in a wide range of settings. The mechanisms we design are computationally tractable, and when they involve approximation factors, these often improve as the inherent competition in the market grows. They also take advantage of the combinatorial structure in the market — e.g., the properties of matching play a key role in our analysis.
- A revenue-driven seller may deliberate between acquiring information in order to carefully set prices for existing buyers, and between drawing more potential buyers in order to drive prices up by making what she's selling harder to win. Our results quantify the trade-off between these strategies by answering the question of how many more buyers are needed to replace information acquisition, generalizing the well-known Bulow-Klemperer result beyond multi-unit markets.
- Our approach is flexible and extends to markets with allocation constraints and to markets with asymmetries among the buyers. This demonstrates that robustness and simplicity are achievable even in complicated settings, and in fact are advisable when approaching such settings.

There are two main future research directions arising from our work. First, it is an interesting and challenging direction to study the best possible revenue guarantee subject to robustness. For example, given x additional bidders, what is the optimal prior-independent mechanism? What is the optimal prior-independent mechanism without introducing additional bidders? Second, our techniques may be useful for more general multi-parameter markets. For instance, markets with gross substitutes preferences have combinatorial structure that can possibly be utilized using our methods; and markets with positive correlation (*affiliation*) among the bidders have inherently more competition that can potentially be put to use.

Endnotes

1. See Section 3.2 for a definition. By the revelation principle, this requirement is without loss of generality for a seller seeking a dominant strategy implementation.
2. Myerson's characterization also extends to *asymmetric* environments, where there are multiple distributions $\{F_1^i\}_{i\in[n]}$ and bidder i's value for the good (good 1) is drawn from his distribution F_1^i. In this case the informational burden on the seller is heavier, as she needs to have full knowledge of the distributions of all buyers, and even getting only approximately close to optimal revenue requires many samples from every distribution (Cole and Roughgarden 2014).

3. Note that while we enhance competition among buyers in the market, we do not assume that it becomes a "competitive market" with "price-taker" buyers, and our theorems hold for markets of any size (as opposed to asymptotically).

4. We remark that the *regularity* constraint on the distribution is standard; for a bidder whose value is drawn from the distribution, regularity means that the *revenue curve* describing the trade-off between selling to the bidder often at a low price and selling less often at a higher price is concave. This is satisfied by all common distributions (uniform, normal, power-law, etc.), and without it no result along the lines of Theorem 1 is possible — see Section 3 for details.

5. One can also aim for certain trade-offs between revenue and welfare by setting suitable supply limits.

6. α-optimality is similar to an average-case approximation guarantee in combinatorial optimization, which holds for inputs drawn from a known distribution; the only difference is that the OPT_F benchmark measures the performance of a truthful mechanism rather than an algorithm.

7. It is interesting to note however that there are mechanisms in the literature that do not require the market-maker to know the prior distributions and yet are not dominant strategy truthful — for example the *buyer's bid double auction* (Satterthwaite and Williams 1989).

8. It is not hard to adapt the proof to the case in which there are multiple maximum weighted matchings, to show that one possible allocation of VCG run on the augmented matching environment leaves unallocated a set of bidders equal to U with at most one bidder removed.

9. Similarly, for every $\eta \in [n, m]$ there is a Bulow-Klemperer-type theorem with η additional bidders. This does not improve the guarantees in Section 5.

10. If n is odd, one can first remove a bidder from the environment, losing at most a $1/n$-fraction of the optimal expected revenue.

11. Recall that a function f from sets of bidders to $\mathbb{R}$ is *submodular* if for every two sets $S \subset T$ and every bidder $i \notin T$ it holds that $f(S \cup \{i\}) - f(S) \geq f(T \cup \{i\}) - f(T)$. Equivalently, the marginal contribution of a bidder to the value of f is decreasing.

12. This example uses a distribution which has a constant amount of probability mass concentrated at the highest point H in its support; it can be "smoothed" by spreading out the point mass at H over a small interval.

13. It is not hard to show that the Bulow-Klemperer-type theorem for asymmetric matroid environments of (Hartline and Roughgarden 2009, Theorem 4.4) applies to attribute-based environments. This theorem requires augmenting the demand with an additional "duplicate" bidder for every original bidder, and adding the constraint that at most one of each such pair wins simultaneously. Our version in Theorem 11 utilizes the fact that many of the bidders in an attribute-based environment are symmetric — namely all those with the same attribute — in order to avoid the pair constraints, and when k is relatively small requires less bidders to be added to the environment.

Acknowledgments

An early version of this work appeared as Roughgarden et al. (2012). This work was presented as an invited talk in the Matching and Market Design Session in honor of Al Roth at the INFORMS 2014 Annual Meeting. The authors wish to thank the participants of the microeconomics and IEOR seminars at Stanford and Berkeley for their helpful comments, as well as the anonymous reviewers for their invaluable suggestions and insights. This work was supported in part by NSF grant CCF-1016885, an ONR PECASE Award, and an AFOSR MURI grant, and received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No. 708935. The second author gratefully acknowledges the support of the Hsieh Family Stanford Interdisciplinary Graduate Fellowship.

References

Alon N, Spencer JH (2008) *The Probabilistic Method* (Wiley).

Ausubel LM, Milgrom PR (2006) The lovely but lonely Vickrey auction. Cramton P, Shoham Y, Steinberg R, eds., *Combinatorial Auctions*, chapter 1, 57–95 (Boston, MA, USA: MIT Press).

Azar PD, Kleinberg R, Weinberg SM (2014) Prophet inequalities with limited information. *Proc. 25th ACM Symp. on Discrete Algorithms (SODA)*, 1358–1377.

Babaioff M, Immorlica N, Lucier B, Weinberg SM (2014) A simple and approximately optimal mechanism for an additive buyer. *Proc. 55th IEEE Symp. on Foundations of Computer Science (FOCS)*, 21–30.

Baliga S, Vohra R (2003) Market research and market design. *Advances in Theoretical Economics* 3(1):Article 5.

Bandi C, Bertsimas D (2014) Optimal design for multi-item auctions: A robust optimization approach. *Mathematics of Operations Research* 39(4):1012–1038.

Bertsekas DP (1991) *Linear Network Optimization* (MIT Press).

Bertsimas D, Thiele A (2014) *Robust and Data-Driven Optimization: Modern Decision Making Under Uncertainty*, chapter 5, 95–122 (INFORMS PubsOnline), tutORials in Operations Research.

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

Cai Y, Daskalakis C, Weinberg SM (2013) Understanding incentives: Mechanism design becomes algorithm design. *Proc. 54th IEEE Symp. on Foundations of Computer Science (FOCS)*, 618–627.

Cai Y, Devanur NR, Weinberg SM (2016) A duality based unified approach to Bayesian mechanism design. *Proc. 48th ACM Symp. on Theory of Computing (STOC)*, 926–939.

Celis LE, Lewis G, Mobius MM, Nazerzadeh H (2014) Buy-it-now or take-a-chance: Price discrimination through randomized auctions. *Management Science* 60(12):2927–2948.

Chawla S, Hartline JD, Malec DL, Sivan B (2010a) Multi-parameter mechanism design and sequential posted pricing. *Proc. 42nd ACM Symp. on Theory of Computing (STOC)*, 311–320.

Chawla S, Hartline JD, Malec DL, Sivan B (2010b) The power of randomness in Bayesian optimal mechanism design. *Games and Economic Behavior* 91:297–317.

Chawla S, Hartline JD, Malec DL, Sivan B (2013) Prior-independent mechanisms for scheduling. *Proc. 45th ACM Symp. on Theory of Computing (STOC)*, 51–60.

Chung K, Ely C (2007) Foundations of dominant-strategy mechanisms. *Review of Economic Studies* 74:447–476.

Clarke EH (1971) Multipart pricing of public goods. *Public Choice* 11:17–33.

Coase R (1972) Durability and monopoly. *Journal of Law and Economics* 15(1):143–149.

Cole R, Roughgarden T (2014) The sample complexity of revenue maximization. *Proc. 46th ACM Symp. on Theory of Computing (STOC)*, 243–252.

Daskalakis C, Deckelbaum A, Tzamos C (2013) Mechanism design via optimal transport. *Proc. 14th ACM Conf. on Economics and Computation*, 269–286.

Daskalakis C, Deckelbaum A, Tzamos C (2014) The complexity of optimal mechanism design. *Proc. 25th ACM Symp. on Discrete Algorithms (SODA)*, 1302–1318.

Daskalakis C, Deckelbaum A, Tzamos C (2017) Strong duality for a multiple-good monopolist. *Econometrica* To appear.

Debo LG, Parlour CA, Rajan U (2012) Signaling quality via queues. *Management Science* 58(5):876–891.

Demange G, Gale D, Sotomayor M (1986) Multi-item auctions. *The Journal of Political Economy* 94(4):863–872.

Devanur NR, Hartline JD, Karlin AR, Nguyen CT (2011) Prior-independent multi-parameter mechanism design. *Proc. 7th Conference on Web and Internet Economics (WINE)*, 122–133.

Dhangwatnotai P, Roughgarden T, Yan Q (2010) Revenue maximization with a single sample. *Games and Economic Behavior* 91:318–333.

Dughmi S, Roughgarden T, Sundararajan M (2012) Revenue submodularity. *Theory of Computing* 8(1):95–119.

Eden A, Feldman M, Friedler O, Talgam-Cohen I, Weinberg SM (2016) The competition complexity of auctions: A Bulow-Klemperer result for multi-dimensional bidders, working paper.

Fu H, Immorlica N, Lucier B, Strack P (2015) Randomization beats second price as a prior-independent auction. *Proc. 16th ACM Conf. on Economics and Computation*, 323.

Gershkov A, Goeree JK, Kushnir A, Moldovanu B, Shi X (2013) On the equivalence of Bayesian and dominant strategy implementation. *Econometrica* 81(1):197–220.

Goldberg A, Hartline JD, Karlin A, Saks M, Wright A (2006) Competitive auctions. *Games and Economic Behavior* 55(2):242–269.

Goldner K, Karlin AR (2016) A prior-independent revenue-maximizing auction for multiple additive bidders. *Proc. 12th Conference on Web and Internet Economics (WINE)*, 160–173.

Groves T (1973) Incentives in teams. *Econometrica* 41:617–631.

Hart S, Nisan N (2012) Approximate revenue maximization with multiple items. *Proc. 13th ACM Conf. on Economics and Computation*, 656.

Hart S, Reny PJ (2015) Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics* 10(3):893–922.

Hartline JD, Roughgarden T (2009) Simple versus optimal mechanisms. *Proc. 10th ACM Conf. on Economics and Computation*, 225–234.

Hartline JD, Roughgarden T (2014) Optimal platform design, working paper.

Kirkegaard R (2006) A short proof of the Bulow-Klemperer auctions vs. negotiations result. *Journal of Economic Theory* 28(2):449–452.

Lahaie S, Pennock DM, Saberi A, Vohra RV (2007) Sponsored search auctions. Nisan N, Roughgarden T, Tardos E, Vazirani VV, eds., *Algorithmic Game Theory*, chapter 28 (Cambridge University Press).

Li X, Yao ACC (2013) On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences* 110(28):11232–11237.

Lovász L, Plummer MD (2009) *Matching Theory* (American Mathematical Society).

Manelli AM, Vincent DR (2006) Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory* 127(1):1–35.

McEachern WA (2012) *Economics: A Contemporary Introduction* (Cengage Learning).

Milgrom P (2004) *Putting Auction Theory to Work* (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* 43:214–238.

Nisan N (2014) Algorithmic mechanism design: Through the lens of multi-unit auctions. Young P, Zamir S, eds., *Handbook of Game Theory*, chapter 9, 477–516 (North-Holland).

Oxley JG (1992) *Matroid Theory* (Oxford).

Rothkopf MH (2007) Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. *Operations Research* 55(2):191–197.

Roughgarden T, Talgam-Cohen I, Yan Q (2012) Supply-limiting mechanisms. *Proc. 13th ACM Conf. on Economics and Computation*, 844–861.

Sandholm T, Likhodedov A (2015) Automated design of revenue-maximizing combinatorial auctions. *Operations Research* 63(5):1000–1025.

Satterthwaite MA, Williams SR (1989) The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. *The Review of Economic Studies* 56(4):477–498.

Scarf HE (1958) A min-max solution of an inventory problem. Arrow KJ, Karlin S, Scarf HE, eds., *Studies in the Mathematical Theory of Inventory and Production*, chapter 12, 201–209 (Stanford University Press).

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

Sivan B, Syrgkanis V (2013) Vickrey auctions for irregular distributions. *Proc. 9th Conference on Web and Internet Economics (WINE)*, 422–435.

Thanassoulis JE (2004) Haggling over substitutes. *Journal of Economic Theory* 117(2):217–245.

Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance* 16:8–37.

Vohra RV, Krishnamurthi L (2013) *Principles of Pricing* (Cambridge University Press).

Wilson RB (1987) Game-theoretic analyses of trading processes. *Advances in Economic Theory: Fifth World Congress* 33–70.

Yan Q (2012) Prior-independence: A new lens for mechanism design. PhD Thesis, Stanford University.

Yao AC (2016) On solutions for the maximum revenue multi-item auction under dominant-strategy and Bayesian implementations, working paper.