Supply-Limiting Mechanisms

TIM ROUGHGARDEN, Department of Computer Science, Stanford University
INBAL TALGAM-COHEN, Department of Computer Science, Stanford University
QIQI YAN, Department of Computer Science, Stanford University

Most results in revenue-maximizing auction design hinge on "getting the price right" — offering goods to bidders at a price low enough to encourage a sale, but high enough to garner non-trivial revenue. Getting the price right can be hard work, especially when the seller has little or no a priori information about bidders' valuations.

A simple alternative approach is to "let the market do the work", and have prices emerge from competition for scarce goods. The simplest-imaginable implementation of this idea is the following: first, if necessary, impose an artificial limit on the number of goods that can be sold; second, run the welfare-maximizing VCG mechanism subject to this limit.

We prove that such "supply-limiting mechanisms" achieve near-optimal expected revenue in a range of single- and multi-parameter Bayesian settings. Indeed, despite their simplicity, we prove that they essentially match the state-of-the-art in prior-independent mechanism design.

Categories and Subject Descriptors: F.0 [**Theory of Computation**]: General

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Mechanism design, revenue optimization, prior-independence

1. INTRODUCTION

1.1. A Matching Problem

Consider the problem of matching agents with unknown preferences to a set of goods for sale, with the goal of maximizing the seller's revenue. For example, a travel website selling hotel accommodation would like to match agents to a set of hotel rooms. Each unit-demand agent only needs one room, and has a different private value for each type of room. Uncertainty about agents' values is modeled by drawing the values from *prior distributions*, with one distribution per good (one distribution for a suite at the Ritz, another for a room at Best Western, and so on). The seller wishes to maximize her expected revenue, but at the same time wants to minimize the resources spent estimating the underlying distributions, as well as the risks associated with getting these distributions wrong and the complexity of the procedures involved in carrying out the sale. In addition, a far-seeing seller might also want to (approximately) maximize social welfare.

Maximizing expected revenue in the matching problem above is difficult even when the value distributions are known. The difficulty stems from the problem's *multi-parameter* nature. The theory of optimal auction design stops short of solving settings

---

This work is supported in part by NSF grant CCF-1016885, an ONR PECASE Award, and an AFOSR MURI grant.
Authors' address: T. Roughgarden, I. Talgam-Cohen and Q. Yan, Department of Computer Science, Stanford University, Stanford, CA 94305.

EC'12, June 4–8, 2012, Valencia, Spain. Copyright 2012 ACM 978-1-4503-1415-2/12/06...$10.00.

in which the description of an agent's preferences requires multiple parameters. Recent breakthrough work of Chawla et al. [2010a] circumvents this limitation by introducing *approximately* optimal mechanisms. These mechanisms make use of a priori knowledge of the value distributions.

Our work focuses on *prior-independent mechanisms*, whose description does not reference any prior distributions, yet for every set of prior distributions satisfying standard assumptions, has expected revenue close to that of the optimal mechanism tailored to the distributions. We present prior-independent mechanisms that have near-optimal expected revenue for a variety of market environments, including the multi-parameter matching problem described above. The importance of prior-independence is discussed by Dhangwatnotai et al. [2010].

Our mechanisms are extremely simple, and are based on the natural idea of artificially *limiting the supply* to increase bidder competition for the goods. Previous prior-independent mechanisms are based largely on some form of random sampling to estimate the prior distributions [Balcan et al. 2005; Dhangwatnotai et al. 2010]. Ours are the first known prior-independent mechanisms for nontrivial multi-parameter settings.[1]

1.2. A Supply-Limiting Mechanism

Before describing our results in more detail, we motivate the approach of limiting supply and develop some intuition for how it works by considering a very simple, single-parameter context — multi-unit auctions. In this context the seller only has a single item, but may have several identical units of it, and the agents are unit-demand. A concrete example is the sale of digital goods such as software licenses. In this particular case, the supply is essentially unlimited.

Now consider the supply-limiting mechanism in Algorithm 1. This mechanism is simple and natural, and does not rely on knowing or sampling distributions. Intuitively, limiting the supply increases competition, and indeed it is not hard to show that this mechanism guarantees approximately-optimal revenue, assuming that bidder valuations are i.i.d. draws from a distribution satisfying a standard regularity assumption. In particular, although the mechanism is oblivious to the underlying distribution, in expectation it achieves at least half of the expected revenue of the optimal mechanism tailored to the distribution. We remark that by using the VCG mechanism,[2] our mechanism also guarantees near-optimal social welfare, even though our weak distributional assumptions allow the revenue and welfare of other mechanism to be very far from each other.

ALGORITHM 1: A Generic Prior-Independent Mechanism

(1) Set a supply limit equal to half of the number of bidders.
(2) Run the VCG mechanism subject to this supply limit.

This paper shows that under minimal regularity assumptions, the simple, prior-independent mechanism above and its revenue guarantee generalize to significantly more complex settings. In other words, we identify settings in which the prior-independent VCG mechanism with limited supply is guaranteed to have near-optimal

[1] Simultaneously and independently, another group obtained similar results using different mechanisms; see Section 1.5 for a detailed discussion.

[2] We assume for now that the reader is familiar with the VCG mechanism, see the Preliminaries section for details.

expected revenue. For the matching problem discussed above, we prove the following theorem.

THEOREM 1.1 (PRIOR-INDEPENDENT MECHANISM FOR MATCHING (INFORMAL)). *For every matching environment with $n \geq 2$ bidders, the expected revenue of the mechanism in Algorithm 1 is at least a constant fraction of the optimal expected revenue.*

The constant fractions we achieve are quite good in many cases, e.g., we achieve a fraction of $\frac{1}{4}$ when the number of bidders exceeds the number of goods.

1.3. Technical Approach: Reduction to Bulow-Klemperer-Type Theorems

Our technical approach to establishing approximation properties of supply-limiting mechanisms is based on a general reduction to "Bulow-Klemperer-type" theorems. To convey the basic idea, we once again illustrate our approach using the simple example of a multi-unit auction for digital goods.

A slightly generalized version of Bulow and Klemperer's well-known result [1996] states that when selling k units of an item to n unit-demand bidders, whose single-parameter values are drawn i.i.d. from a regular distribution, the expected revenue of the VCG mechanism with k additional bidders is at least that of the optimal mechanism without additional bidders (see [Kirkegaard 2006] for a simple proof). We briefly sketch how to use this Bulow-Klemperer theorem to analyze the mechanism in Algorithm 1 in the multi-unit auction context.

Consider first the "halved environment", with half of the n original bidders and a corresponding supply limit of $n/2$. One can show that if we were to restrict the optimal mechanism to run on this sub-environment instead of the original one, its expected revenue is at least half of that of the optimal auction for the original environment. Now we conceptually add back the removed bidders but without changing the supply, and run VCG. It follows from the Bulow-Klemperer theorem that the expected revenue is at least as high as the optimal expected revenue for the halved environment. Therefore the supply-halving mechanism guarantees at least half of the optimal expected revenue on the original environment. In fact, one can achieve other trade-offs between revenue and welfare by setting suitable supply limits.

In general, a Bulow-Klemperer-type theorem states that instead of running the optimal mechanism on the original environment, we can get approximately as much revenue in expectation by running the VCG mechanism on a suitably augmented environment with additional bidders or supply. We make explicit the connection between Bulow-Klemperer-type theorems and prior-independent mechanisms. A sketch of our reduction procedure, which applies to both single- and multi-parameter settings, appears in Reduction 2.

REDUCTION 2: From Prior-Independent Auctions to Bulow-Klemperer-Type Theorems

---

(1) **Restriction**: Restrict the auction environment by dropping bidders and/or limiting supply.
Guarantee: Optimal expected revenue approximately maintained.
Proof: By a general subadditivity property of optimal expected revenue in the bidder set.

(2) **Augmentation**: Augment the restricted environment by adding bidders while maintaining the supply limit.
Guarantee: Expected revenue of VCG approximates optimal expected revenue.
Proof: By a suitable Bulow-Klemperer-type theorem.

To instantiate the reduction in various single-parameter environments, we can use generalizations of the original Bulow-Klemperer result to matroid environments

[Dughmi et al. 2009] and to non-i.i.d. bidders [Hartline and Roughgarden 2009]. For the matching problem, we need to prove the first generalization of the original Bulow-Klemperer theorem to a nontrivial multi-parameter setting.

THEOREM 1.2 (B-K-TYPE THEOREM FOR MATCHING (INFORMAL)). *For every matching environment with n bidders and m goods, the expected revenue of the VCG mechanism with either (1) m additional bidders or (2) $O(n)$ additional bidders and a supply limit n, is at least a constant fraction of the optimal expected revenue in the original environment.*

Our proof is based on the following ideas (see Section 5 for details). We first observe there is a very simple upper bound on the optimal expected revenue for the matching environment — the expected revenue from selling each good separately to $n+1$ bidders. We need to show that VCG with additional bidders does just as well. Recall that in VCG, the winner of a certain good pays for the "damages" incurred by the losers, and in both cases stated in the above theorem it is guaranteed there are $O(n)$ many losers. The technical challenge is a dependency issue — by definition the losers seem likely to have lower values for the good than those of the $n+1$ bidders to which it is sold separately, so that their damages may not be enough to cover the expected revenue. We show this is not the case due to the structure of matchings, specifically the fact that the values of bidders for one particular good play a limited role in VCG's choice of winners, so limited in fact that the only implication on the losers' values is being lower than that of the winner. We formalize these ideas by introducing an auxiliary selling procedure, conceptually and revenue-wise half-way between selling the good separately and selling it as part of VCG, namely we run VCG but defer the sale of the good until exactly $n+1$ bidders remain, so that by construction their values are unaffected by dependency issues.

1.4. Our Results and Organization

Our main result is a collection of approximately-optimal supply-limiting mechanisms for different auction environments with i.i.d. regular bidders, as detailed in Table I. For single-parameter settings, we show supply-limiting mechanisms for k-unit and matroid environments. For multi-parameter settings, we show three versions of a supply-limiting mechanism for matching environments, each with a different approximation guarantee. The choice among these versions should be according to the parameters of the environment at hand. Our results generalize to multi-unit matching environments as well.

1.5. Related Work

Most related to our results are the following. Dughmi et al. investigate conditions under which VCG inadvertently yields near-optimal revenue. They use a generalized Bulow-Klemperer result to show this is guaranteed in matroid environments with sufficient competition in the form of disjoint bases [Dughmi et al. 2009]. Our reduction encompasses and generalizes this result. Hartline and Roughgarden also study conditions for when simple Vickrey-based mechanisms achieve near-optimal revenue and in particular they derive an anonymous-reserve mechanism from one of their Bulow-Klemperer-type results [Hartline and Roughgarden 2009, Theorem 5.1]. This mechanism however is not prior-independent and is inherently limited to single items [Hartline and Roughgarden 2009, Example 5.4]. Chawla et al. use posted-price mechanisms that rely on prior distributions — i.e., mechanisms that are *not* prior-independent — to achieve a 6.75-approximation for the matching setting with multiple units and non-i.i.d. bidders, and also a $\frac{32}{3}$-approximation for a more general environment than what we consider, namely a graphical matroid with unit-demands [Chawla et al. 2010a].

Table I: Summary of Our Results

Environment[a]	*Value Parameters*	*Definition Section*	*Main Result*	*Approximation Factor*	*Mechanism*[b]	*Proof Section*
k-unit	Single	Sect. 2.1	Theorem 3.1	$\min\{2, \frac{n}{n-k}\}$	$\text{VCG}^{\leq n/2}$	Sect. 3
Matroid[c] $\kappa \geq 2$	Single	Sect. 6.1.1	Theorem 6.1	2	VCG	Full version
Matroid[c] $\kappa = 1$	Single	Sect. 6.1.1	Theorem 6.1	4	$\text{VCG}^{\leq \lfloor r/2 \rfloor}$	Full version
Matching[d] $m \leq n/2$	Multiple	Sect. 2.1	Theorem 4.1	2	VCG	Sect. 4.3 Sect. 5 (B-K)
Matching[d] $m \geq n/2$	Multiple	Sect. 2.1	Theorem 4.2	$\frac{4m}{n}$	$\text{VCG}^{\leq n/2}$	Sect. 4.3 Sect. 5 (B-K)
Matching general m	Multiple	Sect. 2.1	Theorem 4.3	27	$\text{VCG}^{\leq n/3}$	Sect. 4.3 Full version (B-K)
Matching k_j-unit general m	Multiple	Sect. 6.2	Theorem 6.5	27	$\text{VCG}^{\leq n/3, \leq \lceil k_j/2 \rceil}$	Full version

[a]Bidders in all environments are assumed to be i.i.d. and regular.
[b]$\text{VCG}^{\leq \ell}$ is the VCG mechanism with supply limit ℓ (and $\text{VCG}^{\leq \ell, \leq \ell_j}$ is the VCG mechanism with global supply limit ℓ and local supply limits $\{\ell_j\}$, defined in Section 6.2).
[c]Matroid has rank r and packing number κ.
[d]Result holds for multi-unit matchings as well (defined in Section 6.2).

In terms of techniques, our limited-supply mechanisms are special cases of maximal-in-range-mechanisms (see, e.g., [Nisan and Ronen 2000]), which can be implemented in dominant strategies. We apply a reduction of Chawla et al. that relates single- and multi-parameter settings ([Chawla et al. 2010a], see Preliminaries). Some of our techniques are inspired by Chawla et al.'s analysis of the VCG mechanism's performance in a job scheduling context [Chawla et al. 2011].

Finally, the paper of Devanur et al. [Devanur et al. 2011] is very closely related to this work. These two independent papers considered essentially the same set of problems and gave similar results, though using different mechanisms. The mechanisms in [Devanur et al. 2011] are arguably quite complicated, certainly more complex and less natural than the supply-limiting mechanisms studied here. On the other hand, the mechanisms in [Devanur et al. 2011] seem to be a bit easier to analyze than supply-limiting mechanisms, and they also achieve better constant factors in the matching problem for the case when m is large. The analyses in the two papers share some common preliminary steps, but at their core are quite different, reflecting the different mechanisms studied. Finally, in the matching problems studied in the original version of the present work, each item was assumed to have unit supply; we were inspired by [Devanur et al. 2011] to pursue the more general multi-unit results presented here.

2. PRELIMINARIES

2.1. Basic Auction Environments

Auction environments are settings in which goods are sold to bidders. We distinguish between *items* — different kinds of goods, and *units* — different copies of the same good. While bidders have the same value for different units of the same item, their values for different items are independent from one another (although possibly drawn from the same distribution). We use the following notation: item j means the j-th kind of good sold in the auction, k_j denotes the number of units available of item j, and $m = \sum_j k_j$ is the total number of units available of all items.

Our results apply to two basic auction environments — in the first there are k multiple units of a single item, and in the second there are m multiple items but only one unit of each. In both cases we will assume there are at least $n \geq 2$ bidders, since prior-independence is impossible with a single bidder (in the sense that for every prior-independent mechanism there is an environment for which the mechanism's approximation guarantee is arbitrarily poor). We now describe the two environments of interest in more detail; for extensions of these environments and generalizations of our results see Section 6.

2.1.1. Single-Parameter k-Unit Environments. In general, a single-parameter auction environment is composed of a set $\{1, \ldots, n\}$ of bidders, and a non-empty collection $\mathcal{I} \subseteq 2^{\{1,\ldots,n\}}$ containing subsets of bidders who can win the auction simultaneously.[3] Subsets in $\mathcal{I}$ are called *feasible allocations*. We assume that every subset of a feasible allocation is also feasible (formally, the set system $(\{1, \ldots, n\}, \mathcal{I})$ is downward-closed). We also assume that every bidder belongs to at least one feasible allocation. Every bidder i has a private value v_i for winning, which is drawn independently at random from a distribution F_i (the environment is called single-parameter since the value for winning is described by one parameter). We say the bidders are *i.i.d.* if their distributions are independent and identical.

A *k-unit* (or *multi-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 (we assume $k \leq n$). This captures the situation where there are k identical units of the same item for sale, and every bidder is *unit-demand*, i.e., interested in buying at most one unit. We can also impose an additional supply limit of $\ell \leq k$, restricting feasible allocations to size at most ℓ. An i.i.d. k-unit environment is one in which the bidders are i.i.d; our supply-limiting results apply to such environments.

2.1.2. Multi-Parameter Matching Environments. A (single-unit) *matching* environment is a multi-parameter environment with n bidders and m different items for sale. We only have one unit of each item available, and a multi-unit version of matching environment will be defined later. Bidders are unit-demand, in the sense that each bidder can only win at most one item. Feasible allocations are all matchings of items to bidders, such that each bidder wins at most one item and each item is assigned to at most one bidder. We can also impose an additional supply limit of $\ell \leq m$, restricting the matchings to size at most ℓ. Bidder i has a private value $v_{i,j}$ for winning item j, which is drawn independently at random from a distribution $F_{i,j}$. We say the bidders are i.i.d. if $F_{i,j}$ does not depend on i, which we can simply denote by F_j. Our supply-limiting results apply to i.i.d. matching environments in which the bidders are i.i.d.

2.2. Mechanisms

Our work focuses on deterministic mechanisms (but applies to randomized mechanisms with only a constant-factor loss).[4] A deterministic mechanism is comprised of

— an allocation rule $\mathbf{x}$, which maps a bid vector $\mathbf{b} \in [0, \infty)^n$ in the single-parameter case or $\mathbf{b} \in [0, \infty)^{nm}$ in the multi-parameter case to a feasible allocation; and
— a payment rule $\mathbf{p}$, which maps a bid vector $\mathbf{b}$ to a payment vector in $[0, \infty)^n$.

[3]This description of a single-parameter environment is sufficient for our purpose; more general definitions can be found in the literature.

[4]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.

We assume a quasi-linear utility model, in which each bidder aims to maximize her value for the chosen allocation minus her payment for it. A mechanism is *truthful* if given any bid profile $\mathbf{b}_{-i}$, bidder i maximizes her utility by being truthful, i.e., bidding $b_i = v_i$ in the single-parameter case and $b_{i,j} = v_{i,j}$ in the multi-parameter case. All the mechanisms we study are both truthful and *individually rational* — the utility for bidding truthfully is always non-negative. From now on we no longer distinguish between bids and values and use v_i or $v_{i,j}$ to denote both.

The famous *VCG mechanism* is of special interest to us [Vickrey 1961; Clarke 1971; Groves 1973]. It maximizes social welfare by choosing a feasible allocation $\mathbf{x}^*$ that maximizes the total value to the bidders, and charges every bidder i a payment equal to i's "externality" — the difference between the maximum total value if i does not participate in the auction and the value of all other bidders when i does participate. In the context of single-item environments the VCG mechanism is called the *Vickrey auction*. We augment the VCG mechanism by adding to it a supply limit of ℓ, such that the total number of allocated units in its chosen allocation is at most ℓ; we denote this supply-limiting VCG mechanism by $\mathrm{VCG}^{\leq \ell}$.

2.3. Optimal Mechanism Design

2.3.1. Myerson's Mechanism for Single-Parameter Environments. For single-parameter environments, Myerson [1981] determined the optimal mechanism that maximizes expected revenue. Given a distribution F with density f, define its *virtual valuation* function to be $\phi_F(v) = v - \frac{1-F(v)}{f(v)}$. Myerson showed the following.

LEMMA 2.1 (MYERSON'S LEMMA). *Given a single-parameter environment and a truthful mechanism* $(\mathbf{x}, \mathbf{p})$*, for every bidder* i *and fixed value profile* $\mathbf{v}_{-i}$ *of the other bidders,* $\mathbb{E}_{v_i \sim F_i}[p_i(\mathbf{v})] = \mathbb{E}_{v_i \sim F_i}[\phi_{F_i}(v_i) \cdot x_i(\mathbf{v})]$.

Myerson's lemma implies that maximizing expected revenue can be reduced to maximizing the expected total virtual value of the bidders, also called *virtual surplus*. *Myerson's mechanism* maximizes virtual surplus and is optimal by the above lemma.

2.3.2. The Regularity Assumption

Definition 2.2 (*Regular Distribution*). A distribution F is regular if its virtual valuation function is monotone non-decreasing.

Most of the commonly-studied distributions are regular, including the uniform, exponential and normal distributions. We say that bidders are regular if their values are drawn from regular distributions. The assumption that bidders are regular is standard in optimal mechanism design, in particular when aiming for simplicity and/or prior-independence. In what follows we assume that all distributions are regular and possess positive smooth density functions.

2.3.3. Representative Environments for Upper-Bounding Optimal Multi-Parameter Revenue. The optimal mechanism for multi-parameter matching environments is currently unknown. Chawla et al. [2010a] introduced the concept of *representative environments* in order to upper-bound the optimal expected revenue in i.i.d. matching environments despite not knowing the optimal mechanism.

Given a matching environment Env with m items, n i.i.d. bidders and value-distributions $\{F_j\}_j$, the representative environment $\mathrm{Env}^{\mathrm{rep}}$ has the same m items but nm single-parameter bidders, one for every pair of original bidder and item (i, j). The m bidders in $\mathrm{Env}^{\mathrm{rep}}$ corresponding to original bidder i are called i's representatives. Like bidder i's value for item j in Env, representative (i, j)'s value $v_{i,j}$ for winning in $\mathrm{Env}^{\mathrm{rep}}$ is drawn independently at random from F_j. Note that every subset S of representatives in $\mathrm{Env}^{\mathrm{rep}}$ corresponds to an allocation in Env — if representative (i, j) is in S then item j is allocated to bidder i in Env. Feasible allocations in $\mathrm{Env}^{\mathrm{rep}}$ are subsets of representatives such that the corresponding allocation in Env is feasible. In particular, since every bidder i in Env is unit-demand, only one of its representatives in $\mathrm{Env}^{\mathrm{rep}}$ can win simultaneously.

Given a truthful mechanism M for Env, its allocation rule can be used to construct a truthful mechanism M^{rep} for $\mathrm{Env}^{\mathrm{rep}}$. The following lemma relates the expected revenue of the two mechanisms. Intuitively, $\mathrm{Env}^{\mathrm{rep}}$ involves more competition than Env since representatives of the same bidder compete with one another, and so the expected revenue of M^{rep} is higher.

LEMMA 2.3 ([CHAWLA ET AL. 2010A]). *The expected revenue of M^{rep} for $\mathrm{Env}^{\mathrm{rep}}$ upper bounds the expected revenue of M for Env.*

3. WARM-UP: A SUPPLY-LIMITING MECHANISM FOR I.I.D. k-UNIT ENVIRONMENTS

In this section we formally prove the following theorem, discussed informally in the introduction, in order to illustrate our general reduction in a simple single-parameter environment. To simplify notation we assume that the number of bidders n is even. This assumption is essentially without loss of generality since 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.

THEOREM 3.1 (SUPPLY-LIMITING MECHANISM FOR I.I.D. k-UNIT ENVIRONMENTS). *For every k-unit environment with $n \geq 2$ i.i.d. regular bidders, the expected revenue of the supply-limiting mechanism $\mathrm{VCG}^{\leq n/2}$ is a $\min\{2, \frac{n}{n-k}\}$-approximation to the optimal expected revenue.*

The proof of Theorem 3.1 using the reduction requires the following slightly generalized version of the original Bulow-Klemperer theorem [1996].

THEOREM 3.2 ((GENERALIZED) BULOW-KLEMPERER THEOREM). *For every k-unit environment with i.i.d. regular bidders and supply limit ℓ, the expected revenue of VCG with $\min\{k, \ell\}$ additional bidders is at least as high as the optimal expected revenue.*

PROOF OF THEOREM 3.1. We instantiate our general reduction (Reduction 2) as follows.

(1) **Restriction**: Remove $\min\{\frac{n}{2}, k\}$ bidders from the environment, and if $k > \frac{n}{2}$ limit the supply to $\frac{n}{2}$ units.
(2) **Augmentation**: Add back $\min\{\frac{n}{2}, k\}$ bidders.

We first claim that the restriction step reduces the expected optimal revenue by only a small constant factor. By submodularity of the expected optimal revenue in the bidder set [Dughmi et al. 2009, Theorem 3.1], removing bidders from the environment reduces the optimal expected revenue by a factor of at most $\min\{2, \frac{n}{n-k}\}$. Limiting the supply to $\frac{n}{2}$ when $k > \frac{n}{2}$ does not affect the optimal expected revenue since the bidders are unit-demand and since in this case the number of bidders in the restricted environment is $\frac{n}{2}$.

We now apply the Bulow-Klemperer theorem as stated in Theorem 3.2 to the restricted environment. If $k > \frac{n}{2}$, the restricted environment has $\frac{n}{2}$ bidders, k units and supply limit $\frac{n}{2}$, while if $k \leq \frac{n}{2}$, the restricted environment has $n - k$ bidders and k units. In both cases, by the Bulow-Klemperer theorem, the expected revenue of VCG with $\min\{\frac{n}{2}, k\}$ additional bidders is at least as high as the optimal expected revenue.

We have shown that running VCG with $\min\{\frac{n}{2}, k\}$ additional bidders on the restricted environment is a $\min\{2, \frac{n}{n-k}\}$-approximation to the optimal expected revenue on the original environment. But this is equivalent to running the supply-limiting mechanism $\text{VCG}^{\leq n/2}$ on the original environment (where the supply limit of $\frac{n}{2}$ is vacuous if $k \leq \frac{n}{2}$) . This completes the proof. $\square$

The approximation factor in the above theorem is asymptotically tight, by a simple proof deferred to the full version of the paper.

PROPOSITION 3.3 (ASYMPTOTIC TIGHTNESS). *For every $0 \leq \rho \leq 1$, there exists an n-unit environment with n i.i.d. bidders whose values are drawn from a regular distribution F such that $\text{VCG}^{\leq \rho n}$ gives in expectation at most $(\frac{1}{2} + o(1))$-fraction of the optimal expected revenue.*

4. A SUPPLY-LIMITING MECHANISM FOR I.I.D. MATCHING ENVIRONMENTS

4.1. Statement and Discussion of Main Theorems

In this section we present our main result — a supply-limiting mechanism for i.i.d. matching environments. More precisely, we present three alternative supply-limiting mechanisms, all VCG-based, with different approximation factors depending on the parameters n and m of the i.i.d. matching environment. The relation between the number of bidders n and total number of items m in the environment at hand determines which supply-limiting mechanism is most suitable for it.

We denote the revenue from the optimal mechanism for n bidders by $\text{OPT}(n)$, and the revenue from the supply-limiting VCG mechanism for n bidders by $\text{VCG}^{\leq \ell}(n)$, sometimes omitting ℓ from the notation when $\ell \geq \min\{n, m\}$. Note that $\text{OPT}(n)$ and $\text{VCG}^{\leq \ell}(n)$ are random variables over the sample space of bidder valuation profiles $\mathbf{v}$. All expectations below are over $\mathbf{v}$.

For simplicity of notation we will assume that $n/2$ (or $n/3$ where appropriate) is integer. If this is not the case, the approximation guarantees below hold up to a small multiplicative factor (the maximum loss in optimal expected revenue from dropping one or two bidders from the environment).

THEOREM 4.1 (2-APPROXIMATION FOR $m \leq n/2$). *For every matching environment with $n \geq 2$ i.i.d. regular bidders and $m \leq n/2$ items, $\mathbb{E}[\text{VCG}(n)] \geq \frac{1}{2}\mathbb{E}[\text{OPT}(n)]$.*

THEOREM 4.2 ($\frac{4m}{n}$-APPROXIMATION FOR $m \geq n/2$). *For every matching environment with $n \geq 2$ i.i.d. regular bidders and $m \geq n/2$ items, $\mathbb{E}[\text{VCG}^{\leq n/2}(n)] \geq \frac{n}{4m}\mathbb{E}[\text{OPT}(n)]$.*

THEOREM 4.3 (27-APPROXIMATION FOR GENERAL n, m). *For every matching environment with $n \geq 3$ i.i.d. regular bidders and m items, $\mathbb{E}[\text{VCG}^{\leq n/3}(n)] \geq \frac{1}{27}\mathbb{E}[\text{OPT}(n)]$.*

Intuitively, achieving good approximation guarantees becomes more difficult as the number of items grows relative to the number of bidders, since the natural competition among the bidders in the environment is diversified across different items. Accordingly, when number of items is less than half of the number of bidders, we show that simply applying VCG achieves a 2-approximation to the optimal expected revenue (Theorem 4.1). When the number of items is more than half of the number of bidders but still proportional to it, applying VCG while artificially limiting the supply to half of the number of bidders achieves a $\frac{4m}{n}$-approximation, in particular a 4-approximation when $m = n$ (Theorem 4.2). Finally, when the number of items is possibly much larger than the number of bidders, limiting the supply still achieves a constant-factor approximation but with a larger constant. We find that setting the supply limit to a third of the number of bidders guarantees a 27-approximation (Theorem 4.3). We believe this approximation factor can be further improved, and leave this as an open problem.

The remainder of the paper is largely dedicated to proving the first two of the above theorems. The proof of Theorem 4.3 is more involved, and its details appear only in the full version of the paper.[5] More specifically, by applying our general reduction, all proofs boil down to proving appropriate Bulow-Klemperer-type theorems. In Section 4.2 we state these theorems and in Section 4.3 we show how the main theorems reduce to them. The proofs of the Bulow-Klemperer-type theorems themselves are the main technical contribution of our work, and appear in Section 5 (for those corresponding to the first two main theorems) and in the full version (for that corresponding to Theorem 4.3).

4.2. Statement and Discussion of Multi-Parameter Bulow-Klemperer-Type Theorems

In order to prove Theorems 4.1 to 4.3 via our general reduction, we need the following corresponding Bulow-Klemperer-type theorems.

THEOREM 4.4 (B-K WITH m MORE BIDDERS). *For every matching environment with n i.i.d. regular bidders and m items,* $\mathbb{E}[\mathrm{VCG}(n+m)] \geq \mathbb{E}[\mathrm{OPT}(n)]$.

THEOREM 4.5 ($\frac{m}{n}$-APPROXIMATE B-K FOR $m \geq n$ WITH n MORE BIDDERS). *For every matching environment with n i.i.d. regular bidders and $m \geq n$ items,* $\mathbb{E}[\mathrm{VCG}^{\leq n}(2n)] \geq \frac{n}{m}\mathbb{E}[\mathrm{OPT}(n)]$.

THEOREM 4.6 (9-APPROXIMATE B-K WITH $2n$ MORE BIDDERS). *For every matching environment with n i.i.d. regular bidders and m items,* $\mathbb{E}[\mathrm{VCG}^{\leq n}(3n)] \geq \frac{1}{9}\mathbb{E}[\mathrm{OPT}(n)]$.

For proofs see Section 5 (Theorems 4.4 and 4.5) and full version (Theorem 4.6).

The first of the above Bulow-Klemperer-type theorems states that for matching environments with m items, the expected revenue of VCG with m additional bidders is at least as high as the optimal expected revenue. This generalizes the original Bulow-Klemperer theorem to the more complex multi-parameter matching setting. If $m \gg n$ however, the required resource augmentation — adding m bidders when originally there are only n — is substantial, which will cause our reduction to give weak bounds.

Our second and third Bulow-Klemperer-type theorems address this issue by requiring the addition of $O(n)$ bidders. This is made possible by using supply-limiting VCG, which restricts the total number of allocated items to at most n out of the m items available, and by relaxing the optimality requirement to approximate-optimality. Theorem 4.5 provides a good approximation factor when m is larger than n but proportional to it. Theorem 4.6 guarantees a constant approximation factor of 9 for any values of n, m.

4.3. Proof of Main Theorems by Applying the General Reduction

We now reduce the three main theorems to the three Bulow-Klemperer-type theorems by instantiating our general reduction. The following lemma is the key and may be of independent interest. It applies to general auction environments (including multi-parameter ones) and states that the optimal expected revenue achievable from two sets of bidders separately exceeds that is achievable from the union of the two sets. A corollary of this subadditivity lemma is that if we remove bidders from an i.i.d. environment until only a constant fraction of the bidders are left, we still maintain a constant fraction of the optimal expected revenue.

---

[5]Available from the authors upon request.

Let $\mathrm{OPT}(S)$ denote the optimal expected revenue achievable from bidder set S.

LEMMA 4.7 (SUBADDITIVITY OF OPTIMAL EXPECTED REVENUE IN BIDDER SET). *For every auction environment with bidder subsets S and T, $\mathbb{E}[\mathrm{OPT}(S)] + \mathbb{E}[\mathrm{OPT}(T)] \geq \mathbb{E}[\mathrm{OPT}(S \cup T)]$.*

PROOF. It is easy to prove that $\mathrm{OPT}(S)$ is monotone in S, and therefore we can assume that S and T are disjoint. Let M be the optimal mechanism for $S \cup T$. For every valuation profile $\mathbf{v}_T$ of bidders in T, we define the following mechanism $M_{\mathbf{v}_T}$. The mechanism $M_{\mathbf{v}_T}$ gets bids from 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 $\mathbb{E}_{\mathbf{v}_S}[M_{\mathbf{v}_T}(S)]$ is at least the part of the optimal expected revenue $\mathbb{E}_{\mathbf{v}_S,\mathbf{v}_T}[M(S \cup T)]$ that comes from S, and the expected revenue of $M_{\mathbf{v}_T}$ is upper-bounded in turn by $\mathbb{E}[\mathrm{OPT}(S)]$. Similarly, the part of the optimal expected revenue that comes from T is upper-bound by $\mathbb{E}[\mathrm{OPT}(T)]$. Summing up we have the desired subadditivity claim. $\square$

COROLLARY 4.8. *For every auction environment with n i.i.d. bidders and for every integer c that divides n, $\mathbb{E}[\mathrm{OPT}(n/c)] \geq \frac{1}{c}\mathbb{E}[\mathrm{OPT}(n)]$.*

4.3.1. Reduction Instantiations

PROOF OF THEOREM 4.1. We need to show $\mathbb{E}[\mathrm{VCG}(n)] \geq \frac{1}{2}\mathbb{E}[\mathrm{OPT}(n)]$ under assumptions of i.i.d. bidders, regularity and $m \leq n/2$. We instantiate our general reduction (Reduction 2) as follows.

(1) **Restriction**: Remove m bidders from the environment.
(2) **Augmentation**: Add back m bidders.

By Corollary 4.8 of the subadditivity property and by monotonicity of the optimal expected revenue, restricting the environment does not hurt the optimal expected revenue too much, i.e., $\mathbb{E}[\mathrm{OPT}(n)] \leq 2\mathbb{E}[\mathrm{OPT}(n/2)] \leq 2\mathbb{E}[\mathrm{OPT}(n-m)]$. Applying the appropriate Bulow-Klemperer-type theorem, Theorem 4.4, to the restricted environment with $n-m$ bidders and m items, gives $\mathbb{E}[\mathrm{OPT}(n-m)] \leq \mathbb{E}[\mathrm{VCG}(n)]$, completing the proof. $\square$

PROOF OF THEOREM 4.2. We need to show $\mathbb{E}[\mathrm{VCG}^{\leq n/2}(n)] \geq \frac{n}{4m}\mathbb{E}[\mathrm{OPT}(n)]$ under assumptions of i.i.d. bidders, regularity and $m \geq n/2$. We instantiate the reduction as follows.

(1) **Restriction**: Remove $n/2$ bidders from the environment.
(2) **Augmentation**: Add back $n/2$ bidders.

As above, the proof is by the inequality chain $\mathbb{E}[\mathrm{OPT}(n)] \leq 2\mathbb{E}[\mathrm{OPT}(n/2)] \leq \frac{4m}{n}\mathbb{E}[\mathrm{VCG}^{\leq n/2}(n)]$, where the first inequality is by Corollary 4.8, and the second inequality is by applying the appropriate Bulow-Klemperer-type theorem (Theorem 4.5) to the restricted environment with $n/2$ bidders and $m \geq n/2$ total units. $\square$

PROOF OF THEOREM 4.3. We need to show $\mathbb{E}[\mathrm{VCG}^{\leq n/3}(n)] \geq \frac{1}{27}\mathbb{E}[\mathrm{OPT}(n)]$ under assumptions of i.i.d. bidders and regularity. We instantiate the reduction as follows.

(1) **Restriction**: Remove $\frac{2}{3}n$ bidders from the environment.
(2) **Augmentation**: Add back $\frac{2}{3}n$ bidders.

As above, the proof is by the inequality chain $\mathbb{E}[\mathrm{OPT}(n)] \leq 3\mathbb{E}[\mathrm{OPT}(n/3)] \leq 27\mathbb{E}[\mathrm{VCG}^{\leq n/3}(n)]$, where the first inequality is by Corollary 4.8, and the second inequality is by applying the appropriate Bulow-Klemperer-type theorem (Theorem 4.6) to the restricted environment with $n/3$ bidders and m items. □

5. PROOF OF BASIC BULOW-KLEMPERER-TYPE THEOREMS FOR I.I.D. MATCHING ENVIRONMENTS

In this section we prove two Bulow-Klemperer-type theorems for i.i.d. matching environments — Theorems 4.4 and 4.5. Theorem 4.6 is much more challenging to prove and is not dealt with in this section.

We begin with the proof of Theorem 4.4, divided into two parts. In Section 5.1 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. The advantage of this step is that these bounds are relatively simple to analyze and are already similar in form, though not identical. In Section 5.2 we carefully relate the two bounds, thus establishing the theorem. In Section 5.3 we show how the proof extends to establish Theorem 4.5 as well.

5.1. Basic Upper and Lower Bounds

Let $\mathrm{Vic}_j(n+1)$ be the revenue from selling item j to $n+1$ i.i.d. bidders with value-distribution F_j using the Vickrey (second-price) auction. We use the concept of representative environment to show that the optimal expected revenue from selling all items to n bidders in an i.i.d. matching environment is upper-bounded by the expected revenue from selling each item separately to $n+1$ single-parameter bidders.

LEMMA 5.1 (UPPER BOUND ON OPTIMAL EXPECTED REVENUE). *For every matching environment with n i.i.d. regular bidders,* $\mathbb{E}[\mathrm{OPT}(n)] \leq \sum_j \mathbb{E}[\mathrm{Vic}_j(n+1)]$.

PROOF. Given the matching environment, consider the corresponding complete bipartite graph with bidders on one side and items on the other, and the bidders' values for items drawn from distributions $\{F_j\}_j$ as edge weights; recall that feasible allocations correspond to matchings. By Lemma 2.3, 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.

We now relax the feasibility constraints, by which we may only increase the optimal expected revenue. We define a new environment in which feasible allocations are all subsets of edges such that at most one edge is incident to an item-node (but unlike a matching, multiple edges can be incident to every bidder-node). Observe that 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 n single-parameter bidders whose values are drawn i.i.d. from the regular distribution F_j. By the original Bulow-Klemperer theorem (Theorem 3.2), the optimal expected revenue from the j-th environment is upper-bounded by $\mathbb{E}[\mathrm{Vic}_j(n+1)]$. Summing up over all items completes the proof. □

The revenue from the VCG mechanism is the sum of VCG payments for allocated items. We lower-bound the VCG payment for an allocated item j.

OBSERVATION 5.2 (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. If bidder i wins item j, then the VCG payment for j is equal to the externality that i imposes on the rest of the bidders by winning j. Since i prevents any unallocated bidder from getting j, the payment is at least the unallocated bidder's value for j. □

In our matching context, the upper and lower bounds above turn out to share a similar form. On one hand, by definition of the Vickrey auction, the upper bound $\mathbb{E}[\mathrm{Vic}_j(n+1)]$ on the expected revenue from separately auctioning item j is equal to the 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. We are using here the fact that since the augmented environment includes m more bidders, all items are allocated and exactly 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, her value for item j is no longer a random sample from F_j. We address this issue in the next section by introducing a *deferred allocation* selling procedure.

5.2. Relating the Upper and Lower Bounds via Deferred Allocation

Algorithm 3 describes a deferred allocation procedure for selling item j.

ALGORITHM 3: Selling Item j by Deferred Allocation

Given a matching environment with $n+m$ bidders and m items:

(1) Find a welfare-maximizing feasible allocation (a maximum matching) of all items other than j to a subset of 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.

We now show how deferred allocation resolves the dependency issue. Consider the revenue from selling item j to bidder set U by the deferred allocation procedure described in Algorithm 3. We use this revenue to relate the upper and lower bounds found in the previous section, as depicted in Figures 1a to 1c.

Claim 5.3 (*Relating to Upper Bound*). The revenue from selling item j by deferred allocation is equal in expectation to $\mathbb{E}[\mathrm{Vic}_j(n+1)]$.

PROOF. Observe that 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 still independent random samples from F_j. The expected second-highest among $n+1$ values drawn independently from F_j is equal to $\mathbb{E}[\mathrm{Vic}_j(n+1)]$. □

To relate to the lower bound in Claim 5.2, we need the following stability property.

Claim 5.4 (*Stability*). For every value profile of the augmented matching environment, the set of bidders left unallocated by VCG is U with at most one bidder removed.

PROOF. Given the augmented matching environment, consider again the corresponding complete bipartite graph with bidders on one side and items on the other, and the bidders' values for items as edge weights. The VCG mechanism finds the maximum-weighted matching in this graph.[6] Our claim is a direct corollary of the following well-known property of matchings (see, e.g., [Wastlund 2008, Lemma 2.2]). Starting with a maximum-weighted matching of size $m-1$, if we add a node (the excluded item j) to one side of the bipartite graph and find the maximum-weighted matching of size m, the set of matched nodes on the other side remains the same up to a single additional node. □

Using this claim we can lower-bound the VCG payment for item j in the augmented environment.

Claim 5.5 (*Relating to Lower Bound*). For every value profile of the augmented matching environment, the VCG payment for item j is at least the revenue from selling item j by deferred allocation.

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 the previous claim, 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 (Observation 5.2). 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. □

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

PROOF OF THEOREM 4.4 (B-K FOR MATCHING WITH m MORE BIDDERS). We need to show that for every matching environment with n i.i.d. regular bidders and m total units, $\mathbb{E}[\mathrm{VCG}(n+m)] \geq \mathbb{E}[\mathrm{OPT}(n)]$. By Claim 5.5, 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 5.3 is equal in expectation to $\mathbb{E}[\mathrm{Vic}_j(n+1)]$. Summing up over all items, the total expected VCG revenue in the augmented environment is at least $\sum_j \mathbb{E}[\mathrm{Vic}_j(n+1)]$, and by Lemma 5.1 this upper-bounds the optimal expected revenue in the original environment. □

5.3. The Case of $m \geq n$

PROOF OF THEOREM 4.5 (B-K FOR MATCHING WITH n MORE BIDDERS). We need to show that for every matching environment with n i.i.d. regular bidders and $m \geq n$ items, $\mathbb{E}[\mathrm{VCG}^{\leq n}(2n)] \geq \frac{n}{m}\mathbb{E}[\mathrm{OPT}(n)]$. The proof is similar to that of Theorem 4.4; here we highlight the necessary adjustments.

The upper bound on the optimal expected revenue remains $\sum_j \mathbb{E}[\mathrm{Vic}_j(n+1)]$ (Lemma 5.1). As for the lower bound, it is no longer the case that in the augmented environment all items are allocated, and so we make use of a generalization of Observation 5.2 — the VCG payment for item j is lower-bounded not only by the value of any unallocated bidder for j itself, but also by the value of any unallocated bidder for any unallocated item. We call the highest of the latter among all unallocated bidders and items the *global* lower bound on VCG payments, and denote it by G. Note that since VCG is now applied with a supply limit of n, exactly n out of the $2n$ bidders in the augmented environment remain unallocated.

We use the modified deferred allocation selling procedure in Algorithm 4; observe that Claims 5.3, 5.4 and 5.5 hold. For Theorem 4.4 these claims were sufficient to com-

[6]We assume there is a unique maximum-weighted matching. This holds with probability 1 as all distributions have smooth density functions.

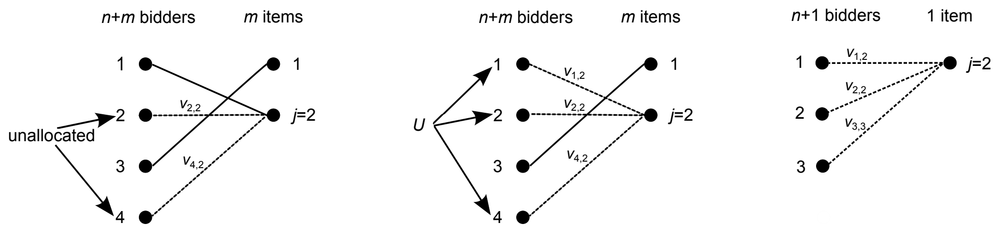

(a) VCG

(b) Deferred allocation for item $j = 2$

(c) Vickrey for item $j = 2$

Fig. 1: Example of relating bounds by deferred allocation ($n = 2, m = 2$).
(a) Applying VCG to augmented matching environment (not all edges are shown). Solid edges correspond to maximum-weighted matching. Payment for item $j = 2$ is at least $\max\{v_{2,2}, v_{4,2}\}$ (Observation 5.2). Values of unallocated bidders are *not* i.i.d. samples from distribution F_2.
(b) Applying deferred allocation for item $j = 2$ to augmented matching environment. Solid edge corresponds to maximum-weighted matching excluding $j = 2$. The set U is a superset of the unallocated bidders in (a) (Claim 5.4). Payment for $j = 2$ is $\max_2\{v_{1,2}, v_{2,2}, v_{4,2}\} \leq \max\{v_{2,2}, v_{4,2}\}$, where $\max_2$ is the second-highest value (Claim 5.5). Values of bidders in U are i.i.d. samples from F_2.
(c) Applying Vickrey for item $j = 2$ to $n+1$ bidders, as part of the upper bound (Lemma 5.1). Values of bidders are i.i.d. samples from F_2 so in expectation we get the same payment for $j = 2$ as from deferred allocation (Claim 5.3).

ALGORITHM 4: Selling Item j by Deferred Allocation — The Case of $m \geq n$

Given a matching environment with $2n$ bidders and m items:

(1) Find a welfare-maximizing feasible allocation (a maximum matching) of $n - 1$ items other than j to a subset of 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.

---

plete the proof, by applying the following chain of arguments: all items are allocated by VCG in the augmented environment; the VCG payment for item j is at least the revenue from selling j by deferred allocation; the deferred allocation revenue is equal in expectation to $\mathbb{E}[\text{Vic}_j(n+1)]$; and $\sum_j \mathbb{E}[\text{Vic}_j(n+1)]$ is at least the optimal expected revenue. However, for Theorem 4.5, we need an additional charging argument, since only n out of m items are allocated by VCG.

(1) If item j is allocated by VCG, then as above the VCG payment for it is at least the revenue from selling j by deferred allocation.
(2) If item j is not allocated by VCG, then the VCG payment for any allocated item $j' \neq j$ is at least the global lower bound G, and so is at least the revenue from selling j by deferred allocation (cf. Claim 5.5).

By the above, we can charge the VCG payments for the n allocated items against the total revenue from selling all m items by deferred allocation, which is equal in expectation to $\mathbb{E}[\text{Vic}_j(n+1)]$, thus obtaining an approximation factor of $\frac{m}{n}$. $\square$

6. ADDITIONAL RESULTS

In this section we list additional results, whose proofs appear in the full version of the paper.

6.1. Single-Parameter Environments

6.1.1. A Supply-Limiting Mechanism for I.i.d. Matroid Environments. A *matroid environment* is a single-parameter environment in which the set system $(\{1,\ldots,n\},\mathcal{I})$ of bidders and feasible allocations forms a matroid (for an exposition on matroids see, e.g., [Oxley 1992]). Recall that the *rank* of a matroid is the size of its bases, and the *packing number* of a matroid is its maximum number of disjoint bases.

THEOREM 6.1 (SUPPLY-LIMITING MECHANISM FOR I.I.D. MATROIDS). *For every matroid environment with $n \geq 2$ i.i.d. regular bidders, rank r and packing number κ*

(*1*) *If $\kappa \geq 2$ then the expected revenue of the VCG mechanism is a 2-approximation to the optimal expected revenue.*
(*2*) *If $\kappa = 1$ then the expected revenue of the supply-limiting mechanism $\mathrm{VCG}^{\leq \lfloor r/2 \rfloor}$ is a 4-approximation to the optimal expected revenue.*

The proof is by instantiating the general reduction, where the restriction and augmentation consist roughly of removing and adding back a basis of bidders, and applying Dughmi et al.'s [2009, Lemma 6.1] Bulow-Klemperer-type result for i.i.d. matroid environments. For completeness we state their result.

THEOREM 6.2 (B-K FOR I.I.D. MATROID ENVIRONMENTS). *For every matroid environment with i.i.d. regular bidders, the expected revenue of VCG with an additional matroid-basis of bidders is at least as high as the optimal expected revenue.*

6.1.2. A Tighter Bulow-Klemperer-Type Theorem for Parallel Multi-Unit Environments. A single-parameter *parallel k_j-unit* (or *parallel multi-unit*) environment consists of a k_j-unit auction for each item j, and these auctions are related by a global supply limit ℓ. To be specific, for each item j, there are $k_j \leq n$ units available, and n unit-demand bidders interested in it, whose values are drawn i.i.d. from a regular distribution F_j. There is a total of $m = \sum_j k_j$ units, of which at most ℓ can be allocated at the same time. In other words, a feasible allocation is a set of bidders containing at most $\min\{n,\ell\}$ bidders overall and at most k_j bidders per item j. We can also impose a local supply limit ℓ_j on the number of units allocated of item j.

A parallel multi-unit environment is a particular case of a matroid environment (the underlying matroid being the intersection of a partition matroid with an ℓ-uniform matroid). As such, the Bulow-Klemperer-type theorem for non-i.i.d. matroid environments by Hartline and Roughgarden [2009, Theorem 4.4] applies to it. However, this theorem requires augmenting the environment with an additional "duplicate" bidder for every original bidder, and adding the constraint that at most one of each pair of duplicates wins simultaneously. This is wasteful, as many of the original bidders in a parallel multi-unit environment are i.i.d., namely, all bidders interested in the same item j. The following theorem shows that it is sufficient to augment the environment with only k_j additional bidders per item j.

THEOREM 6.3 (B-K FOR PARALLEL MULTI-UNIT ENVIRONMENTS). *For every parallel k_j-unit environment with local supply limits $\{\ell_j\}$, the expected revenue of VCG with $\min\{k_j,\ell_j\}$ additional bidders per item j is a 2-approximation to the optimal expected revenue.*

The proof of this theorem uses a commensuration argument developed by Hartline and Roughgarden, and applies the FKG inequality to solve dependency issues. One application of this theorem is in our proofs of supply-limiting results for i.i.d. multi-unit matching environments (see Section 6.2 and full version). Also, we can readily use this Bulow-Klemperer-type theorem to derive a corresponding supply-limiting mechanism for parallel multi-unit environments. However to keep the description simple, in the following we derive a supply-limiting mechanism for a natural subclass of such environments, namely *attribute-based* non-i.i.d. multi-unit environments.

Attribute-Based Environments. Consider a single-parameter k-unit environment, where every bidder has a publicly-observable attribute a, say age bracket, which determines her (regular) distribution F_a. In other words, the values are i.i.d. for bidders with the same attribute, and are independent but not necessarily identically distributed for bidders with different attributes. Furthermore, assume that the environment is *non-singular* in the sense that there is no bidder with a unique attribute. (This setting was first introduced by Dhangwatnotai et al. [2010].)

Now consider the following supply-limiting mechanism. For every set of n_a bidders with the same attribute a (and hence the same distribution), impose a local supply limit of $\min\{k, \frac{n_a}{2}\}$ on the number of units that can be allocated to this bidder set, and run VCG (we assume here that n_a is even; the case of odd n_a can be handled with a small loss).

THEOREM 6.4 (SUPPLY-LIMITING FOR ATTRIBUTE-BASED ENVIRONMENTS). *For every non-singular k-unit environment with attribute-based regular bidders, the expected revenue of the above supply-limiting VCG mechanism is a 4-approximation to the optimal expected revenue.*

PROOF. We instantiate our reduction as follows.

(1) **Restriction**: For every a, remove $\min\{k, \frac{n_a}{2}\}$ bidders with attribute a from the environment, and if $k \geq \frac{n_a}{2}$, limit the supply for bidders with attribute a to $\frac{n_a}{2}$ units.
(2) **Augmentation**: For every a, add back $\min\{k, \frac{n_a}{2}\}$ bidders with attribute a.

In the restriction step, we remove at most half of the bidders and so lose a factor of at most 2 [Dughmi et al. 2009, Theorem 3.1]. Note that limiting the supply has no effect on the revenue, and that after the restriction we essentially have a global supply limit of $\min\{k, \sum_a \frac{n_a}{2}\}$ on the restricted environment. Therefore, even though there is just a single item, we can view this as a parallel multi-unit environment with bidder sets of size $\max\{n_a - k, \frac{n_a}{2}\}$, a global supply limit of $\min\{k, \sum_a \frac{n_a}{2}\}$, and local supply limits of $\min\{k, \frac{n_a}{2}\}$. Now we can apply Theorem 6.3 to conclude that in the augmentation step we lose at most another factor of 2. We have thus shown that running VCG on the restricted environment with $\min\{k, \frac{n_a}{2}\}$ additional bidders per attribute a is a 4-approximation to the optimal expected revenue on the original environment, and this is equivalent to running the supply-limiting mechanism on the original environment, completing the proof. □

We remark that this supply-limiting mechanism is considerably simpler than Myerson's optimal mechanism for this setting, which requires computing different virtual value functions for different attributes.

6.2. Supply-Limiting Mechanisms for I.i.d. Multi-Unit Matching Environments

A k_j*-unit* (or *multi-unit*) *matching* environment is a multi-parameter matching environment with multiple units per item. There are k_j units of every item j and a total of $m = \sum_j k_j$ units. We can also impose an additional global supply limit $\ell \leq m$ on the total number of allocated units, and/or local supply limits $\{\ell_j\}$ where $\ell_j \leq k_j$ is on the number of allocated units of every item j.

Two out of the three supply-limiting mechanisms in Section 4 for i.i.d. matching environments apply directly to i.i.d. multi-unit matching as well. In fact, Theorems 4.1 and 4.2 hold without change for multiple units. Recall that Theorem 4.3 gives a constant approximation guarantee in the challenging case where the number of items m is much larger than the number of bidders n. In order to generalize this theorem to multi-unit matching, we introduce a slightly more general supply-limiting mechanism. Let $\mathrm{VCG}^{\leq \ell, \leq \ell_j}$ be the VCG mechanism with a global supply limit ℓ on the total number of allocated units, and local supply limits $\{\ell_j\}$ on the number of allocated units of every item j. We then have the following multi-unit version of Theorem 4.3.

THEOREM 6.5 (27-APPROXIMATION FOR MULTIPLE UNITS). *For every multi-unit matching environment with $n \geq 3$ i.i.d. regular bidders, m total units and k_j units per item j,* $\mathbb{E}[\mathrm{VCG}^{\leq n/3, \leq \lceil k_j/2 \rceil}(n)] \geq \frac{1}{27}\mathbb{E}[\mathrm{OPT}(n)]$.

We prove this theorem via our general reduction, using the following multi-unit version of the Bulow-Klemperer-type result in Theorem 4.6.

THEOREM 6.6 (9-APPROXIMATE B-K FOR MULTIPLE UNITS). *For every multi-unit matching environment with n i.i.d. regular bidders, m total units and k_j units per item j,* $\mathbb{E}[\mathrm{OPT}(n)] \leq 9\mathbb{E}[\mathrm{VCG}^{\leq n, \leq \lceil k_j/2 \rceil}(3n)]$.

REFERENCES

BALCAN, M., BLUM, A., HARTLINE, J., AND MANSOUR, Y. 2005. Mechanism design via machine learning. In *Proc. 46th IEEE Symp. on Foundations of Computer Science (FOCS)*.

BULOW, J. AND KLEMPERER, P. 1996. Auctions versus negotiations. *American Economic Review 86,* 1, 180–194.

CHAWLA, S., HARTLINE, J., MALEC, D., AND SIVAN, B. 2010a. Sequential posted pricing and multi-parameter mechanism design. In *Proc. 41st ACM Symp. on Theory of Computing (STOC)*.

CHAWLA, S., HARTLINE, J., MALEC, D., AND SIVAN, B. 2011. Private communication.

CHAWLA, S., MALEC, D., AND SIVAN, B. 2010b. The power of randomness in bayesian optimal mechanism design. In *Proc. 12th ACM Conf. on Electronic Commerce (EC)*.

CLARKE, E. H. 1971. Multipart pricing of public goods. *Public Choice 11*, 17–33.

DHANGWATNOTAI, P., ROUGHGARDEN, T., AND YAN, Q. 2010. Revenue maximization with a single sample. In *Proc. 12th ACM Conf. on Electronic Commerce (EC)*.

DUGHMI, S., ROUGHGARDEN, T., AND SUNDARARAJAN, M. 2009. Revenue submodularity. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*. ACM, New York, NY, USA, 243–252.

GROVES, T. 1973. Incentives in teams. *Econometrica 41*, 617–631.

HARTLINE, J. D. AND ROUGHGARDEN, T. 2009. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*.

KIRKEGAARD, R. 2006. A short proof of the bulow-klemperer auctions vs. negotiations result. *Journal of Economic Theory 28,* 2, 449–452.

MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research 6,* 1, 58–73.

NISAN, N. AND RONEN, A. 2000. Computationally feasible vcg mechanisms. In *Proc. 2nd ACM Conf. on Electronic Commerce (EC)*. ACM Press, 242–252.

OXLEY, J. G. 1992. *Matroid Theory*. Oxford.

VICKREY, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance 16*, 8–37.

WASTLUND, J. 2008. Random matching problems on the complete graph. *Elect. Comm. in Probab. 13*, 258–265.