

On Revenue-Maximizing Walrasian Equilibria for Size-Interchangeable Bidders

Enrique Areyan Viqueira^(✉), Amy Greenwald, Victor Naroditskiy,
and Daniels Collins

Brown University, Providence, RI 02912, USA
{eareyan, amy, daniel_collins}@brown.edu, vnarodit@gmail.com

Abstract. We study a market setting in which bidders are single-valued but size-interchangeable, and there exist multiple copies of heterogeneous goods. Our contributions are as follows: (1) providing polynomial-time algorithms for finding a *restricted* envy-free equilibrium with reserve prices (EFER); (2) posing the problem of finding a revenue-maximizing EFER, and running experiments to show that our algorithms perform well on the metrics of revenue, efficiency, and time, without incurring too many violations of the stronger Walrasian equilibrium with reserve (envy-free plus market clearance) conditions.

1 Introduction

In a **centralized combinatorial matching market (CCMM)**, a market maker offers a set of n heterogeneous **goods** to m consumers (or **bidders**), the latter of which are interested in acquiring certain combinations (or **bundles**) of goods. In general, there are multiple copies of each good i , but the total supply N_i of each good is finite. Bidder j 's preferences are captured by a **valuation** function $v_j(\cdot)$ that describes how bidder j values each bundle.

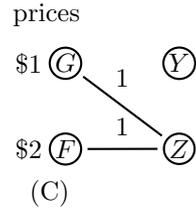
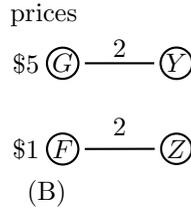
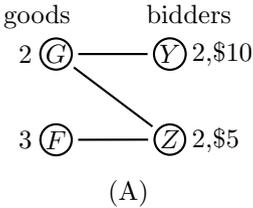
In general, a bidder's valuation function can be an arbitrary function of the set of all bundles. We study a case where bidders are only interested in specific varieties of goods, and we model these interests as edges in a graph connecting bidders only to their goods of interest. Furthermore, in our model, bidder's valuations are **single-valued**, and depend only on the bundle's size, assuming the bundle is a match for the bidder. The value then is either a positive value R_j , if the size of the bundle is at least some threshold I_j , and 0 otherwise.

Our model is motivated by the Trading Agent Competition Ad Exchange game (TAC AdX) [15], which in turn models online ad exchanges in which agents face the challenge of bidding for display-ad impressions needed to fulfill advertisement contracts, after which they earn the amount the advertiser budgeted. Other settings captured by this model include the problem of how to allocate specialized workers to firms, and how to compensate the workers, where each firm requires a certain number of workers to produce an output (a new technology, for instance) that yields a certain revenue.

One well-studied special case of our model is that of (single-valued) **single-minded** consumers [11]. There, bidders are only interested in one particular bundle. Hence their valuation function can be understood as assigning value R_j to that bundle or any superset thereof, and 0 to all other bundles. We call our valuations (single-valued) **size-interchangeable**, because bidders can be satisfied (i.e., achieve value R_j) by any bundle of size I_j that consists of their desired goods. Like single-minded valuations, our interchangeable valuations model complements, since a bidder is not satisfied unless it receives a bundle of sufficient size. Furthermore, our interchangeable valuations model (perfect) substitutes, since any bundle of sufficient size that consists of suitable goods will do.

In this paper, we assume valuations are known to the market maker. Thus, our problem is one of equilibrium computation rather than traditional mechanism design (where values are private). A market **outcome** is an allocation-pricing pair (\mathbf{X}, \mathbf{p}) , where \mathbf{X} describes the assignment of goods to bidders, and \mathbf{p} ascribes prices to goods. While \mathbf{X} is a matrix, we assume \mathbf{p} is a vector, which precludes any form of price discrimination (all copies of the same good must have the same price). Furthermore, we assume **item pricing**, not bundle pricing, so that the price of a bundle is the sum of the prices of all the goods (items) in the bundle. Both of these assumptions—no price discrimination and item pricing—are most natural.

Example 1 (CCMM and possible outcomes). Consider the CCMM in Figure (A). There are two goods, G and F , with 2 copies of good G and 3 copies of good F , and two bidders, Y and Z . Bidder Y wants two copies of good G (as indicated by the edge from G to Y) and values this bundle at 10, and bidder Z ascribes the value 5 to any bundle of size 2 comprised of any combination of G s and F s (also indicated by edges). Possible outcomes of this markets are depicted in Figures (B) and (C).



Outcome (B) allocates 2 copies of good G to bidder Y at a price of \$5 per copy, and 2 copies of good F to bidder Z at a price of \$1 per copy. This outcome results in the optimal social welfare (\$15) and a revenue of \$12.

Outcome (C) allocates to bidder Z only, 1 copy of good G at a price of \$1, and 1 copy of good F at a price of \$2. This outcome results in a social welfare of \$5 and a revenue of \$3.

In an important related setting, bidders have unit-demand valuations, meaning they are interested in at most one good, but may have different valuations for different goods. This is a well-studied setting [1, 2, 6, 7], with important theoretical guarantees. In particular, there always exists a Walrasian equilibrium (WE)

outcome [5] in which bidders are envy-free, meaning they all receive one of their favorite bundles at the set prices, and the market clears, meaning unallocated goods are priced at zero. It follows from this second condition, by the first welfare theorem of economics, that any allocation that is part of any WE outcome maximizes social welfare. In addition, it is possible to find a revenue-maximizing WE among all WE outcomes in polynomial-time (assuming unit demand) [6].

Building on Myerson’s [10] intuition, Guruswami et al. [7] generalized the problem of searching for a revenue-maximizing WE to that of searching for a revenue-maximizing Walrasian equilibrium with reserve prices (WER), where bidders are envy-free and the market clears up to the reserve price: i.e., unallocated goods are priced at the reserve. For bidders with unit-demand valuations a WE always exists [5]; likewise, a WER always exists. In addition, for a fixed reserve price, we can find the revenue-maximizing WER in polynomial-time (using the same approach as in [6]). With this more general solution concept in mind, Guruswami et al. pose the problem of finding a revenue-maximizing WER, for which they propose a polynomial-time approximation algorithm that picks a particular set of candidate reserve prices, generates a revenue-maximizing WER for each, and then returns the revenue-maximizing WER among those considered.

For our model—single-valued, size-interchangeable bidders—we remind the reader that WE do not exist in general; in fact, they are not guaranteed to exist even for single-minded bidders. By relaxing the market clearance condition, we arrive at more general solution concept—an **Envy-Free Equilibrium (EFE)**—which insists only that bidders are envy-free, and which always exists. However, the first welfare theorem does not hold for EFE, so maximal social welfare is not guaranteed by this solution concept.

Departing from the social welfare concern, we instead tackle the competing problem of maximizing seller revenue. Since finding revenue-maximizing WER is APX-hard in a CCMM assuming single-minded bidders [7], we propose a polynomial time heuristic for approximately solving for revenue-maximizing EFER (an EFE in which unallocated goods are priced at a reserve), where we relax the envy-free condition to a **restricted envy-free** condition, which we are able to express as linear constraints. Building on the ideas of Guruswami et al., we then search over a space of carefully chosen reserve prices to find an approximately revenue-maximizing EFER. In particular, whereas Guruswami et al. used this approach to find an approximately revenue-maximizing WER for unit-demand bidders, we apply this same idea to the case of size-interchangeable (and hence, single-minded) bidders.

In sum, our contributions are: (1) providing polynomial-time algorithms for finding a *restricted* envy-free equilibrium with reserve prices (EFER); (2) posing the problem of finding a revenue-maximizing EFER, and running experiments to show that our algorithms perform well on the metrics of revenue, efficiency, and time, without incurring too many violations of the WER (envy-free plus market clearance) conditions.

2 Model and Solution Concepts

We define a **centralized combinatorial matching market** (CCMM) to be an augmented bipartite graph $(U, C, E, \mathbf{N}, \mathbf{I})$ with a set of n types of goods U , a set of m bidders C , a set of edges E from bidders to goods indicating which goods are of interest to which bidders, a supply vector $\mathbf{N} = (N_1, \dots, N_n)$, and a demand vector $\mathbf{I} = (I_1, \dots, I_m)$. That is, there are $N_i > 0$ copies of good $i \in U$, and $I_j > 0$ total goods are demanded by bidder $j \in C$. Note, however, that bidder j only demands copies of goods it is connected to via an edge.

In other words, in a CCMM, bidders are interested in acquiring bundles of goods of at least some fixed size. Note that this is a more general case of the well-studied problem of single-minded consumers [11] where bidders are interested only in one particular bundle of goods. For this reason, we call our valuation function size-interchangeable. Formally:

Definition 1 (*Single-valued, Size-interchangeable valuations*). Given a CCMM $M = (U, C, E, \mathbf{N}, \mathbf{I})$, a bidder j is single-valued, size-interchangeable, if it demands $I_j > 0$ goods among those to which it is connected, and values all such bundles by the function: $v_j(X_j) = R_j > 0$, if $\sum_{i|(i,j) \in E} x_{ij} \geq I_j$, and 0 otherwise. We call R_j the reward attained by j in case its demand I_j is fulfilled.

Definition 2 (*Market*). We call a **market** M a pair consisting of a CCMM $(U, C, E, \mathbf{N}, \mathbf{I})$ and a reward vector $\mathbf{R} = (R_1, \dots, R_m)$.

Given a market M , an **allocation** A is a labeling $x(i, j) \in \mathbb{Z}_{\geq 0}$ of E that represents the number of copies of good i allocated to bidder j . Such an allocation can be represented by a matrix $\mathbf{X} \in \mathbb{Z}_{\geq 0}^n \times \mathbb{Z}_{\geq 0}^m$ where entry $x_{ij} = x(i, j)$. The j th column of an allocation matrix is the **bundle** of goods assigned to bidder j , which we denote by $X_j \in B(\mathbf{N})$, where $B(\mathbf{N}) = \prod_i \{0, 1, \dots, N_i\}$.

A market **outcome** is an allocation-pricing pair (\mathbf{X}, \mathbf{p}) , assigning goods to bidders and per-good prices $p_i \in \mathbb{R}_+$. Given such an outcome, the cost of bundle X_j to bidder j is given by $P_j(X_j) = \sum_i x_{ij} p_i$.

An allocation is **feasible** if, for all i , the total number of goods assigned across bidders is no more than i 's supply: i.e., for all i : $\sum_{j=1}^m x_{ij} \leq N_i$. We use $F \equiv F(M)$ to denote the set of all feasible allocations. In a **feasible outcome** the allocation is feasible.

The **utility** of bidder j is defined as follows: $u_j(\mathbf{X}, \mathbf{p}) = v_j(X_j) - P_j(X_j)$. A standard assumption is that all bidders are utility maximizers, and thus a bidder prefers outcomes with higher utilities.

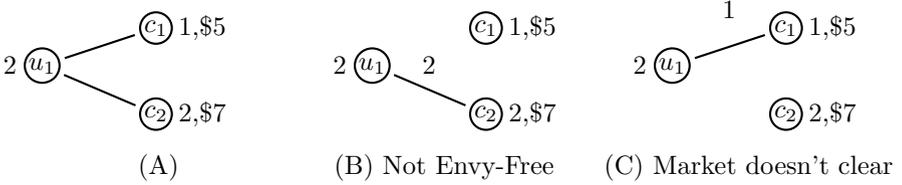
A fundamental market outcome studied in the literature is that of Walrasian Equilibrium (WE) [14], which we define using our notation as follows.

Definition 3 (*Walrasian Equilibrium*). A feasible outcome (\mathbf{X}, \mathbf{p}) is a **Walrasian Equilibrium** (WE) if the following two conditions hold:

1. **Envy-freeness** (EF): There is no bundle X'_j that any bidder j prefers to its assigned bundle X_j , i.e., for all j , $X_j \in \arg \max_{X'_j \in B(\mathbf{N})} \{v_j(X'_j) - P_j(X'_j)\}$.
2. **Market clearance** (MC): Every unallocated good is priced at zero.

The EF condition is a fairness condition; it ensures that the outcome maximizes the utility of every bidder. Note that each bidder is individually rational i.e., $u_j(\mathbf{X}, \mathbf{p}) \geq 0$, since the null allocation is always a feasible allocation. The MC condition, together with EF, implies, by the first welfare theorem of economics, that any allocation that is part of a WE is also welfare-maximizing. However, a WE need not exist in the markets studied in this paper.

Example 2 (Non-existence of WE). Consider the market in Figure (A) with one good and two single-minded bidders. Good u_1 is supplied in $N_1 = 2$ copies, bidder c_1 demands $I_1 = 1$ good, and bidder c_2 demands $I_2 = 2$ goods. Rewards are $R_1 = 5$ and $R_2 = 7$.



There are a total of 6 feasible allocations in this market and none of them are part of a Walrasian Equilibrium. Two such allocations are depicted in (B) and (C). In (B), there is no price p_1 for u_1 at which both bidders would be envy-free. In (C), we must have that $p_1 \geq 3.5$, or otherwise c_2 would have preferred 2 copies from u_1 . But then the market does not clear since there is an unsold copy of u_1 with price greater than 0.

To address this existence problem, we drop the market clearance condition.

Definition 4 (*Envy-Free Equilibrium*). A feasible outcome (\mathbf{X}, \mathbf{p}) is an **Envy-Free Equilibrium (EFE)** if envy-freeness holds.

Note that, in outcome (C) of Example 2, at any price p_1 for u_1 such that $3.5 \leq p_1 \leq 5$, both bidders c_1 and c_2 are envy-free. It follows that this outcome is an EFE.

Unlike in the unit-demand case, where, by the first welfare theorem, a WE implies a welfare-maximizing allocation [12], an EFE (even for unit-demand bidders) does not guarantee a welfare-maximizing allocation. An outcome with a null allocation and prices high enough such that no bidder can afford even a single good is an EFE outcome with 0 welfare. But even for non-null allocations, an EFE outcome can yield low welfare.

Example 3 (An EFE does not imply an efficient allocation). Consider a market M with a single good that is supplied in $N_1 = m - 1$ copies and m bidders, where all bidders demand 1 good (i.e., $I_j = 1$, for all j). Rewards are defined as follow: for $1 \leq j < m : R_j = 1$ and $R_m = 2$.

Consider outcome (\mathbf{X}, \mathbf{p}) , where bidder m is allocated one good at any price p_1 such that $1 < p_1 \leq 2$, and no other bidder is allocated any good. This outcome is an EFE since all bidders are envy-free. However, the welfare of this outcome is 2 for any $m > 1$, since we only satisfy bidder m .

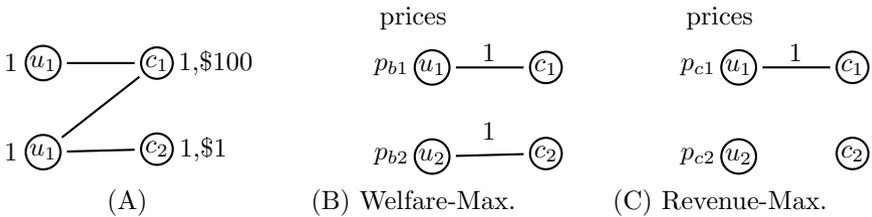
A welfare maximizing allocation for this market is one that satisfies bidder m as well as many other bidders as possible, obtaining welfare $2 + \sum_{j=1}^{m-2} R_j = 2 + m - 2 = m$. Therefore, the welfare of the EFE (\mathbf{X}, \mathbf{p}) , relative to the optimal welfare, approaches zero, as m approaches ∞ .

Although at first glance it may seem disappointing the first welfare theorem does not hold for EFE, it is not a show stopper. Even in the unit-demand case, where the first welfare theorem does hold, there exists the competing, and incompatible goal, of maximizing seller revenue.

In a unit-demand CCMM, if we let $m(< r) = \{j \mid R_j < r\}$, then $|m(< r)|$ is the number of bidders with reward less than the reserve r . Assuming there exists an allocation that satisfies all bidders, by setting a reserve price r we lose at least $R_{\min} |m(< r)|$ welfare, where $R_{\min} = \min_j \{R_j\}$, while we are guaranteed revenue of at least $r |m(< r)^c|$. (Here A^c denotes the complement of set A .)

The following example further illustrate the tradeoff between welfare and revenue, in the case of single-minded bidders.

Example 4 (Welfare-Revenue Tradeoff). Consider the market in Figure (A) and the two different outcomes in Figures (B) and (C)



Outcome (B)'s allocation is welfare maximizing. To support an EFE we must have $0 \leq p_{b2} \leq 1$; otherwise c_2 would not be envy-free. Moreover, $p_{b1} \leq p_{b2}$; otherwise c_1 would have preferred a copy of u_2 . So prices can only be as high as $p_{b1} = p_{b2} = 1$, yielding revenue of 2. Outcome (C)'s allocation is not welfare maximizing. However, in this case, an EFE can be supported by higher prices than those in (B). In particular, $p_{c2} \geq 1$; otherwise c_2 would have preferred a copy of u_2 . Again $p_{c1} \leq p_{c2}$ for the same reasons as in (B). Prices could be as high as $p_{c1} = p_{c2} = 100$, yielding revenue of 100.

Example 4 motivates the introduction of reserve prices as a way to increase revenue while maintaining envy-freeness among market participants. In the previous example, we could set a reserve price of \$2 for u_2 . Doing so would increase revenue from a maximum possible of \$2 (with no reserve price) to \$100. However, by setting reserve prices some bidders are effectively thrown out of the market, so welfare might not be maximized, because any value these bidders bring to the market is lost.

Motivated by this discussion, we generalize the definition of WE so that unallocated goods are priced at some, possibly non-zero, reserve price $r \in \mathbb{R}_+$.

Definition 5 (*Walrasian Equilibrium with reserve r*). A feasible outcome (\mathbf{X}, \mathbf{p}) is a **Walrasian Equilibrium with reserve r (WE r)** if it is a WE with prices at least r , including unallocated goods, which must be priced at exactly r .

Analogously, we augment the definition EFE to an EFE with reserve price r .

Definition 6 (*Envy-Free Equilibrium with reserve r*). A feasible outcome (\mathbf{X}, \mathbf{p}) is an **Envy-Free Equilibrium with reserve r (EFE r)** if it is an EFE with prices at least r .

3 Computation of Envy-Free Equilibria

Chen et al. [8] showed that deciding the existence of WE in a CCMM assuming single-minded bidders is NP-hard. Consequently, we propose a natural restriction on the envy-freeness condition, which lends itself to a polynomial-time computation. With it, we can find a *restricted* WE (outcomes that satisfy restricted EF and MC) in polynomial time in single-minded CCMMs. Furthermore, in size-interchangeable CCMMs, we can find a *restricted* EFE in polynomial time.

Before presenting our algorithm, we formally define **restricted envy-free prices**. Let $|X_j| = \sum_{i=1}^n x_{ij}$ be the size of the bundle assigned to bidder j .

Definition 7 (*Restricted Envy-Free*). A price vector \mathbf{p} is called **restricted envyfree** with respect to a feasible allocation \mathbf{X} if, for all J such that $|X_j| > 0$:

$$X_j \in \arg \max_{X'_j \in B(\mathbf{N}_{|X_j|})} \{v_j(X'_j) - P_j(X'_j)\}$$

where $B(\mathbf{N}_{|X_j|}) = \{\mathbf{0}\} \cup \{X'_j \in B(\mathbf{N}) \mid |X'_j| = |X_j|\}$, i.e., the set of all feasible bundles of size equal to $|X_j|$.

This definition is “restricted” because it assumes an allocation, and then is only concerned with bidders that are allocated (non-zero) bundles in that allocation. Any envy felt by any other bidders is simply ignored.

Another seeming restriction is that even for a bidder j with $|X_j| > 0$, it does not require envy-freeness with respect to all bundles $X'_j \in B(\mathbf{N})$, but only with respect to bundles of the same size as X_j (i.e., $X'_j \in B(\mathbf{N}_{|X_j|})$), and the empty bundle $\mathbf{0}$. This definition might seem overly restrictive, but as we are focused on bidders with single-valued, size-interchangeable valuations, we are likewise concerned with all-or-none allocations, which either allocate to a bidder in full—meaning a bundle of size I_j —or does not allocate at all. Hence, for our purposes the size restriction is not restrictive at all.

Finally, note that restricted envy-free prices always exist. Given an allocation, we can simply set prices equal to zero, and the condition will be satisfied. No one who is allocated would have any envy at zero prices; and the restricted envy-free condition ignores bidders that are unallocated.

Theorem 1. *Given a market M and a feasible allocation \mathbf{X} , the following conditions are necessary and sufficient for \mathbf{p} to be restricted envy-free.*

Individual Rationality: $\forall j \in C : P_j(X_j) \leq v_j(X_j)$.

Compact Condition: $\forall i \in U, j \in C : \text{If } x_{ij} > 0 \text{ then}$

$$\forall k \in U : \text{If } (k, j) \in E \text{ and } x_{kj} < N_k \text{ then } p_i \leq p_k.$$

Proof (Sketch): The Individual Rationality condition ensures that bidders do not pay more than their reward while the Compact Condition ensures that, among all goods assigned to a bidder, she first consumes cheaper goods before consuming more expensive ones. Equivalently, this condition states that prices of goods that are completely consumed are cheaper than those that are only partially consumed. \square

The linear program shown in Algorithm 1, which uses seller revenue as the objective function and the linear conditions that characterize restricted envy-freeness as constraints, can be used to find a set of restricted envy-free prices that maximizes seller revenue.

Algorithm 1. LP *restricted* EFE

Input: Market (U, C, E, N, I, R) and allocation X

Output: A pricing p

maximize $\sum_j \sum_i x_{ij} p_i$

subject to (1) $\forall j \in C : \text{If } |X_j| > 0, \text{ then } p_j(X_j) \leq v_j(X_j)$

(2) $\forall i \in U, \forall j \in C : \text{If } x_{ij} > 0 \text{ then}$

$$\forall k \in U : \text{If } (k, j) \in E \text{ and } x_{kj} < N_k \text{ then } p_i \leq p_k$$

4 Revenue Maximizing Prices

In the remainder of this paper, we will be concerned with finding prices that maximize seller revenue for different market outcomes. We start by defining what a revenue-maximizing problem means for different solution concepts and review algorithms found in the literature to compute these prices in the special case of unit-demand bidders. We then present our algorithm for finding revenue-maximizing EFER in size-interchangeable CCMMs.

Definition 8. *The revenue-maximizing WE problem: Given a CCMM, find a revenue-maximizing WE.*

Gul and Stachetti [6] presented a VCG-inspired [13] polynomial-time algorithm that solves the revenue-maximizing WE problem in unit-demand CCMMs: let $V \in \mathbb{R}_+^n \times \mathbb{R}_+^m$ be the valuation matrix of a market with n items and m unit-demand bidders where entry v_{ij} denotes bidder j 's valuation for good i . Let π denote a maximum weight matching of V , and let $w(V)$ denote the weight of π . Let V_{-i} denote the same valuation matrix, but with good i removed. For each good i , set $p_i = w(V) - w(V_{-i})$. We call this algorithm, which returns (π, p) , **MaxWE**.

Then, building on Myerson's [10] intuition that reserve prices can boost revenue, Guruswami et al. [7] went one step further, essentially generalizing the problem of searching for a revenue-maximizing WE to that of searching for a revenue-maximizing WER.

Definition 9. *The revenue-maximizing WEr problem: Given a CCMM, find a revenue-maximizing WEr.*

Recall from Example 4 that different allocations can support different levels of seller revenue, while still maintaining the envy-freeness property. Algorithm 2 is a high-level strategy for searching among WEr for one that is revenue maximizing. The algorithm searches over different allocations \mathbf{X} , computing WEr prices \mathbf{p} for each, and then outputs a pair (\mathbf{X}, \mathbf{p}) among those seen with maximal seller revenue. The interesting choice, which governs the algorithm's success, is which allocations to search over. Generally speaking, based on some initial allocation, the algorithm determines a set of reserve prices, each of which corresponds to an alternative allocation, whose supporting envy-free prices may or may not yield higher revenue than the others seen.

Algorithm 2. Strategy for finding a revenue-maximizing WEr

Input: Market $M = (U, C, E, N, I, R)$

Output: A pricing \mathbf{p} and an allocation \mathbf{X}

1. Find an initial allocation \mathbf{X} .
2. For all $x_{ij} > 0$:
 - 2.0 Set a reserve price r as a function of x_{ij} .
 - 2.1 Find a WEr (\mathbf{X}, \mathbf{p}) .

Output a pair (\mathbf{X}, \mathbf{p}) among those seen with maximal seller revenue.

In the unit-demand case, Guruswami et al. [7], showed that the following instance of Algorithm 2 finds a revenue-maximizing WEr with revenue at least $\text{OPT}/(2 \ln m)$, where OPT is the revenue of a revenue-maximizing WEr. (Step 1.) Find a maximum weight matching \mathbf{X} of V . (Step 2.) For each valuation r on the edges of \mathbf{X} , compute a WEr as follows: for each good i augment the valuation matrix to include two dummy bidders, each with reward r . Run **MaxWE** on the new valuation matrix to obtain a WE (π, \mathbf{p}) , based on which a new matching π' can be inferred by reallocating goods from dummy bidders to real bidders.

As shown in Example 2, a WE might not exist for a CCMM; thus, a WEr might also not exist. But recall that an EFER always exists. Hence, we define the following problem:

Definition 10. *The revenue-maximizing EFER problem: Given a CCMM, find a revenue-maximizing EFER.*

Like the algorithm of Guruswami et al. [7], our approach (Algorithm 3) to searching for a revenue-maximizing EFER in a size-interchangeable CCMM follows the structure of Algorithm 2. That is, for various choices of r , corresponding to various allocations \mathbf{X}_r , we find an EFER, and then we output an EFER which is revenue-maximizing among all those considered.

More specifically, we first find an allocation \mathbf{X} (Step 1), and then for all $x_{ij} > 0$, we find a restricted EFER (Step 2). Step 2.0 defines a reserve price, and Step 2.1 finds an allocation that respects this reserve price (see Definition 12).

Step 2.2 then invokes a subroutine we call *restricted* EFER, which is a straightforward generalization of Algorithm 1 that finds a revenue-maximizing EFER with restricted envy-free prices in polynomial time. The generalization is simply that this algorithm takes as input a reserve price r , and then includes the additional set of constraints: $\forall i \in U : p_i \geq r$.

Algorithm 3. Revenue-Maximizing EFER (RM-EFER)

Input: Market $M = (U, C, E, \mathbf{N}, \mathbf{I}, \mathbf{R})$

Output: A pricing \mathbf{p} and an allocation \mathbf{X}

1. Find an initial allocation \mathbf{X} .

2. For all $x_{ij} > 0$:

2.0 Set reserve price $r = R_j/x_{ij}$.

2.1 Find an allocation \mathbf{X}_r that respects reserve price r .

2.2 Run *restricted* EFER on M , reserve price r , and allocation \mathbf{X}_r .

Output the pair (\mathbf{X}, \mathbf{p}) with maximum seller revenue.

Allocations in Size-interchangeable CCMMs. Two steps in Algorithm 3 depend on an allocation. The natural place to look are among those of optimal value. A feasible allocation is *optimal* if it maximizes the total of all bidders' rewards.

Definition 11 (*Optimal Allocation*). *An optimal allocation is a solution to the following optimization problem:*

$$\max_{\mathbf{X}} \sum_{j=1}^m R_j y_j, \quad \text{subject to: } \forall i : \sum_{j=1}^m x_{ij} \leq N_i, \quad \forall j : y_j \in \{0, 1\}$$

In the Appendix we present an ILP to compute optimal allocations.

Definition 12 (*Optimal Allocation that respects a reserve price*). *An optimal allocation that respects a reserve price is a solution to the following optimization problem:*

$$\max_{\mathbf{X}} \sum_{j=1}^m (R_j - r I_j) y_j, \quad \text{subject to: } \forall i : \sum_{j=1}^m x_{ij} \leq N_i, \quad \forall j : y_j \in \{0, 1\}$$

The following theorem implies that is unlikely that one can devise an Algorithm that optimizes welfare in size-interchangeable CCMMs in polynomial time.

Theorem 2. *Finding an optimal allocation is NP-hard. (Proof in Appendix)*

Since finding an optimal allocation is NP-Hard, we present a greedy heuristic (Algorithm 4) to find allocations. This algorithm can easily be adapted to produce an allocation that respects reserve price r as follows: given input market M , construct new market M' by removing any bidder j for which $R_j - r I_j < 0$, and setting the reward of the remaining bidders to be $R_j - r I_j$. Now run Algorithm 4 on input M' to obtain an allocation \mathbf{X}' , which we lift up to create an allocation, \mathbf{X}_r , in the original market M that respects reserve prices.

Algorithm 4. Greedy Approximation Algorithm

Input: Market $M = (U, C, E, N, I, R)$
Output: Allocation \mathbf{X}
For all i, j set $x_{ij} = 0$.
foreach $j \in C$ **do**
 Let $U_j = \{i \mid (i, j) \in E \text{ and } \sum_{j=1}^m x_{ij} < N_i\}$.
 if $\sum_{i \in U_j} N_i \geq I_j$ **then**
 foreach $i \in U_j$ **do**
 $x_{ij} = \min\{I_j - \sum_{i=1}^n x_{ij}, N_i - \sum_{j=1}^m x_{ij}\}$.

There are two sources of non-determinism in Algorithm 4: (1) the order in which to loop through bidders and (2) the order in which to loop through goods. One approach is to orders bidders in descending order by rewards per square root of goods demanded, i.e., $R_j/\sqrt{I_j}$, and goods in ascending order of remaining supply. We also experiment with other combinations, e.g., ordering goods in descending order of remaining supply.

5 Experiments

Experimental Setup. Given outcome (\mathbf{X}, \mathbf{p}) , seller revenue $\rho = \sum_j \sum_i x_{ij} p_i$, and total welfare $W = \sum_j R_j y_j$, where $y_j = 1$ in case bidder j is a winner under \mathbf{X} and 0 otherwise. Let OPT_W be the value of a welfare-maximizing allocation. Since we assume bidders are individually rational, seller revenue cannot exceed OPT_W . We thus report metrics of efficiency W/OPT_W , and seller revenue ρ/OPT_W . We also report metrics based on violations of the (unrestricted) envy-freeness and market clearing conditions. Given a market M and outcome (\mathbf{X}, \mathbf{p}) , we define an envy-free violation (EF) as the ratio between the number of bidders that are not envy-free, and the total number of bidders in the market; and define a market clearance violation (MC) as the ratio between the number of goods completely unallocated whose price is greater than zero, and the total number of goods in the market.

All metrics are reported over random markets M drawn from a distribution we call *Random- k -Market*(n, m, p, k). Let $S = \sum_i N_i$ be the total supply of M , and let $D = \sum_j I_j$ be the total demand of M . The supply-to-demand ratio S/D , is a measure of how much over (or under) demanded a market is. A market is over demanded if $S/D < 1$ and under demanded if $S/D > 1$. A random market drawn from *Random- k -Market*(n, m, p, k) over CCMM has n goods and m bidders. The parameter p is the probability that an edge (i, j) is present in E , and thus, the expected number of edges is pnm . Both N_i and I_j are randomly and independently drawn integers between 1 and 10 such that the supply-to-demand ratio is k . Finally, each bidder's reward R_j is uniformly and independently drawn uniformly on $[1, 10]$. From *Random- k -Market*(n, m, p, k) we generate markets with $n, m = 1, \dots, 20$, $p = 0.25, 0.5, 0.75, 1.0$ and $k = 0.25, 0.33, 0.5, 1, 2, 3, 4$. For each metric, we report the average across markets over 100 independent trials. Results are shown in Table 1.

Table 1. Results.

	Efficiency	Revenue	Time (ms)	EF	MC
<i>Under-demanded Singleton CCMM</i>					
CK	1.00	0.00	0.25	0.00	0.00
MaxWE	1.00	0.04	0.41	0.00	0.00
MaxWErApprox	0.85	0.67	12.94	0.00	0.64
LP Optimal	0.86	0.73	15.73	0.00	0.63
<i>Over-demanded Singleton CCMM</i>					
CK	1.00	0.49	4.57	0.00	0.00
MaxWE	1.00	0.80	0.24	0.00	0.00
MaxWErApprox	0.96	0.84	5.69	0.00	0.12
LP Optimal	0.95	0.86	11.47	0.00	0.13
<i>Under-demanded CCMM, $k \in \{2, 3, 4\}$</i>					
CK	0.97	0.08	4.22	0.00	0.07
LP Optimal	0.84	0.65	232.22	0.00	0.43
LP Greedy	0.84	0.65	12.87	0.00	0.49
<i>Over-demanded CCMM, $k \in \{0.25, 0.33, 0.5\}$</i>					
CK	0.79	0.41	2.50	0.00	0.21
LP Optimal	0.93	0.75	199.59	0.02	0.11
LP Greedy	0.89	0.72	9.08	0.02	0.15

Algorithms. Algorithms' names are abbreviated as follows: **CK** refers to the Crawford-Knoer ascending auction [5] (see the Appendix for details). **MaxWE** and **MaxWErApprox** refer to Guruswami et al.'s algorithms (see Sect. 4). **LP** refers to our revenue-maximizing EFER algorithm. The algorithm **LP** is qualified by the type of allocation given as input: **LP Optimal** refers to the case when an optimal allocation is given as input, and **LP Greedy** refers to the case when the greedy allocation is given as input. We report results where the greedy approximation orders goods by descending order of remaining supply. We also experimented with ordering goods by ascending order of remaining supply, but saw no qualitative differences in the results.

Results. We report on two sets of experiments: singleton CCMMs, and general. In both cases, we take as a baseline the CK auction, which, in the case of unit-demand markets, like MaxWE, is guaranteed to produce an efficient outcome, but at the expense of low revenue. For singleton CCMMs, we compare our algorithm to MaxWErApprox as well.

As expected, the efficiency of CK and MaxWE in singleton CCMMs is perfect. Also, these algorithms yield zero violations, as they are known to generate Walrasian equilibria. However, their revenues are low, particularly in under-demanded markets, because there are more unallocated goods in these markets,

each of which must be priced at zero, thereby limiting the price of allocated goods, in order to avoid envy. Our algorithm and MaxWErApprox both impose a reserve, enabling them to disregard some of these overly constraining goods, but ours obtains substantially more revenue than MaxWErApprox, with fewer (but not significantly fewer) violations. In over-demanded markets, the performance of our algorithm is comparable to that of MaxWErApprox, and both algorithms outperform MaxWE, which outperforms CK. Note that our algorithm produces *unrestricted* envy-free outcomes in singleton CCMMs.

In general CCMMs, CK is nearly efficient in under-demanded markets, but much less so in the over-demanded case. But in neither case does it strike a good balance between efficiency and revenue. In contrast, our algorithm is able to accrue upwards of 72% revenue in over-demanded markets, while retaining an efficiency of 89%, and it does so with hardly any EF violations, compared to CK, which obtains an efficiency of only 79% and revenue of only 41%. As in the case of singleton CCMMs, our algorithm does not fare as well in the harder case of under-demanded markets, where once again there are more unallocated goods, although it again produces nearly-*unrestricted* envy-free outcomes. These experiments suggest that while not perfect, finding revenue-maximizing restricted envy-free outcomes in polynomial time is a reasonable heuristic for maximizing revenue among nearly-*unrestricted* envy-free outcomes.

6 Conclusion and Future Directions

CCMMs with unit-demand valuations have some important properties, e.g., a WE always exists and there are polynomial-time algorithms to find such outcomes. A WE outcome guarantees that all bidders are envy-free and that the market clears, and thus, by the first welfare theorem of economics, yields an allocation that maximizes social welfare. Guruswami et al. [7] proposed an algorithm for the unit-demand case, sacrificing social welfare in an attempt to maximize seller revenue, while maintaining the envy-freeness property. In this paper, we proposed an algorithm that generalize this well-known algorithm for the unit-demand case to the case of single valued, size-interchangeable CCMMs. In future work, we plan to look more closely at algorithms [3, 4, 7, 9] that have been proposed for the more difficult case of single-minded bidders, and to perhaps generalize results about those algorithms to the single-valued, size-interchangeable bidder setting studied here. We will also explore alternative solution concepts where we combine the restricted EF condition with the objective of maximizing the number of allocated bidders.

Appendix

Mixed ILP to Find Optimal Allocations

Given a market $(U, C, E, \mathbf{N}, \mathbf{I}, \mathbf{R})$, Algorithm 5 is a mixed ILP that can be used to find an optimal allocation.

Algorithm 5. Mixed ILP Optimal Allocation

Input: Market $M = (U, C, E, \mathbf{N}, \mathbf{I}, \mathbf{R})$ **Output:** An optimal allocation \mathbf{X} maximize $\sum_{j=1}^m R_j y_j$
 subject to

- (1) $\forall i : \sum_{j=1}^m x_{ij} \leq N_i$
- (2) $\forall i, j : \text{If}(i, j) \notin E \text{ then } x_{ij} = 0$
- (3) $\forall j : y_j \leq \frac{1}{I_j} \sum_{i=1}^n x_{ij} \leq y_j$
- (4) $y_j \in \{0, 1\}$
- (5) $\forall i, j : x_{ij} \in \mathbb{Z}^+$

Constraints (1), (2) and (5) imply that a solution to the Mixed-ILP is a feasible allocation. Constraints (3) and (4) imply that a bidder attains reward R_j if and only if it is completely fulfilled, and together with constraint (5), imply that if $y_j = 0$ then $x_{ij} = 0$ for all i . The objective of the mixed ILP implies that the solution maximizes bidders' rewards over all feasible allocations and thus, it is an optimal allocation. To obtain an allocation that respects reserve price r , change the objective of the mixed ILP to $\sum_{j=1}^m (R_j - rI_j)y_j$, where $r \in \mathbb{R}^+$ is the reserve price parameter.

Proof of Theorem 2

Finding an optimal allocation is NP-hard. To prove this, we reduce from the following version of set packing: Given a universe $\mathcal{U} = \{u_1, u_2, \dots, u_n\}$ and a family of subsets $\mathcal{S} = S_1, S_2, \dots, S_k \subseteq \mathcal{U}$, find the maximum number of pairwise disjoint sets in \mathcal{S} .

Consider an input $(\mathcal{U}, \mathcal{S})$ to the set packing problem as described above. Let us construct a market $(U, C, E, \mathbf{N}, \mathbf{I}, \mathbf{R})$ from $(\mathcal{U}, \mathcal{S})$ as an input to the optimal allocation problem. At a high level, the input market consists of n goods each offered in exactly 1 copy and k bidders where each bidder corresponds to a member $S_j \in \mathcal{S}$ that demand as many goods as elements in S_j and attains a reward of exactly 1. A goods is connected to a bidder only if the index of the good is contained in the set S_j associated with the bidder.

Formally, given $(\mathcal{U}, \mathcal{S})$ where $\mathcal{U} = \{1, 2, \dots, n\}$ and $\mathcal{S} = \{S_1, S_2, \dots, S_k\}$, construct $f(\mathcal{U}, \mathcal{S}) = (U, C, E, \mathbf{N}, \mathbf{I}, \mathbf{R})$ as follow: (1) let $U = \mathcal{U}$ and $N_i = 1$ for all $i = 1, 2, \dots, n$. (2) let $C = \{1, 2, \dots, k\}$, and associate each bidder $j \in C$ to $S_j \in \mathcal{S}$ so that $I_j = |S_j|$. Also, $R_j = 1$ for all $j = 1, 2, \dots, k$. (3) add edge (i, j) to E only if $i \in S_j$. Clearly the transformation f is polynomial on the size of the input $(\mathcal{U}, \mathcal{S})$.

We now show that a set packing for $(\mathcal{U}, \mathcal{S})$ corresponds to an optimal allocation for $f(\mathcal{U}, \mathcal{S})$ and vice versa. Suppose that l is the maximum number of pairwise disjoint sets in \mathcal{S} and that $S_1, S_2, \dots, S_l \in \mathcal{S}$ are these sets. By our transformation f we know that each bidder j associated with a set S_j from the previous list is connected to as many goods as $|S_j|$. Since all these sets are pairwise disjoint, all bidders are connected to different goods. Therefore, each of these bidders can be fulfilled which means that the value of the optimal allocation is at least l . Moreover, we know that is not possible to fulfill more than l bidders since l is the

maximum number of pairwise disjoint sets, and selecting more than l bidders would imply, by our transformation f , that at least one good has a supply greater than 1. Therefore, l is the value of the optimal allocation of $f(\mathcal{U}, \mathcal{S})$.

Suppose that l is the value of the optimal allocation of $f(\mathcal{U}, \mathcal{S})$. This means that l is the maximum number of bidders that can be fulfilled. Bidder j is fulfilled only if its allocation is at least $|S_j|$. By construction we know that a bidder is connected to exactly $|S_j|$ many goods. Therefore, each allocated bidder j is fulfilled by exactly $|S_j|$ goods. Moreover, none of these goods are allocated to different bidders since there is exactly 1 copy of each good. Therefore, the sets associated with the selected bidders must not overlap in any element, i.e., they must be pairwise disjoint. This shows that there are at least l pairwise disjoint sets in $(\mathcal{U}, \mathcal{S})$. We also know that there must be at most l pairwise disjoint sets or otherwise the value of the optimal allocation would have been more than l . Therefore, l is the maximum number of pairwise disjoint sets in \mathcal{S} . \square

Crawford and Knoer Ascending Auction

In the unit-demand setting, it is well known that Walrasian Equilibria exist [6]. Furthermore, Crawford and Knoer [5] proposed an ascending auction mechanism which, for price increment ϵ , yields an ϵ WE.¹ We describe the workings of their mechanism in a unit-demand CCMM in Algorithm 6.²

Algorithm 6. Crawford-Knoer Ascending Auction (Unit-demand)

Input: Market $M = (U, C, E, N, I, R)$, where $\forall i : N_i \geq 1$ and $\forall j : I_j = 1$

Output: A pricing \mathbf{p} and an allocation \mathbf{X}

For every i , set $p_i = 0$

For every i, j , set $x_{ij} = 0$

while *TRUE* **do**

foreach *Unallocated bidder* j **do**

 Let $i^* \in \arg \max_{i \in U} \{R_j - (p_i + \epsilon)\}$.

 Add (i^*, j) to B .

if $B = \emptyset$ **then**

 Halt with current allocation \mathbf{X} and prices \mathbf{p} .

else

 Choose $(i, j) \in B$.

$x_{ij} = 1$.

if $\sum_l x_{il} > N_i$ **then**

$p_i = p_i + \epsilon$.

 Completely unallocate all bidders $j' \neq j$ such that $x_{ij'} > 0$.

Output the final pair (\mathbf{X}, \mathbf{p}) .

¹ In an ϵ WE, envy-free-ness is satisfied up to ϵ .

² The only difference between our presentation and the original one is that N_i may exceed 1, so in the final **if** statement, it may be necessary to unallocate goods from more than one bidder.

This algorithm, as stated, generalizes to size-interchangeable CCMMs, except that at each step of the algorithm we must query bidders for their favorite bundles at the current prices plus ϵ , rather than their favorite individual goods.

References

1. Briest, P.: Uniform budgets and the envy-free pricing problem. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP 2008. LNCS, vol. 5125, pp. 808–819. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-70575-8_66](https://doi.org/10.1007/978-3-540-70575-8_66)
2. Chen, N., Deng, X.: Envy-free pricing in multi-item markets. In: Abramsky, S., Gavaille, C., Kirchner, C., Meyer auf der Heide, F., Spirakis, P.G. (eds.) ICALP 2010. LNCS, vol. 6199, pp. 418–429. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-14162-1_35](https://doi.org/10.1007/978-3-642-14162-1_35)
3. Chen, N., Deng, X., Goldberg, P.W., Zhang, J.: On revenue maximization with sharp multi-unit demands. *J. Comb. Optim.* **31**(3), 1174–1205 (2014)
4. Cheung, M., Swamy, C.: Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In: IEEE 49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, pp. 35–44. IEEE (2008)
5. Crawford, V.P., Knoer, E.M.: Job matching with heterogeneous firms and workers. *Econometrica J. Econometric Soc.* **49**(2), 437–450 (1981)
6. Gul, F., Stacchetti, E.: Walrasian equilibrium with gross substitutes. *J. Econ. Theory* **87**(1), 95–124 (1999). <http://EconPapers.repec.org/RePEc:eee:jetheo:v:87:y:1999:i:1:p:95-124>
7. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On profit-maximizing envy-free pricing. In: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1164–1173. Society for Industrial and Applied Mathematics (2005)
8. Huang, L.S., Li, M., Zhang, B.: Approximation of Walrasian equilibrium in single-minded auctions. *Theor. comput. sci.* **337**(1), 390–398 (2005)
9. Monaco, G., Sankowski, P., Zhang, Q.: Revenue maximization envy-free pricing for homogeneous resources. In: Proceedings of the 24th International Conference on Artificial Intelligence, pp. 90–96. AAAI Press (2015)
10. Myerson, R.B.: Optimal auction design. *Math. Oper. Res.* **6**(1), 58–73 (1981)
11. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: *Algorithmic Game Theory*. Cambridge University Press, New York (2007)
12. Tan, K.T.K.: The first fundamental theorem of welfare economics (2008) (Unpublished)
13. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *J. Finan.* **16**(1), 8–37 (1961)
14. Walras, L.: *Elements of Pure Economics: Or the Theory of Social Wealth*. Elements of Pure Economics, Or the Theory of Social Wealth, Routledge (2003). <https://books.google.com/books?id=hwjRD3z0Qy4C>
15. Wellman, M., Wurman, P.R.: A trading agent competition for the research community. In: IJCAI-99 Workshop on Agent-Mediated Electronic Trading (1999)

Agent-Mediated Electronic Commerce. Designing
Trading Strategies and Mechanisms for Electronic
Markets

AMEC/TADA 2015, Istanbul, Turkey, May 4, 2015, and
AMEC/TADA 2016, New York, NY, USA, July 10, 2016,

Revised Selected Papers

Ceppi, S.; David, E.; Hajaj, C.; Robu, V.; Vetsikas, I.A.
(Eds.)

2017, IX, 157 p. 49 illus., Softcover

ISBN: 978-3-319-54228-7