

COMBINATORIAL WALRASIAN EQUILIBRIUM*

MICHAL FELDMAN[†], NICK GRAVIN[‡], AND BRENDAN LUCIER[‡]

Abstract. We study a combinatorial market design problem, where a collection of indivisible objects is to be priced and sold to potential buyers subject to equilibrium constraints. The classic solution concept for such problems is Walrasian Equilibrium (WE), which provides a simple and transparent pricing structure that achieves optimal social welfare. The main weakness of the WE notion is that it exists only in very restrictive cases. To overcome this limitation, we introduce the notion of a Combinatorial Walrasian equilibrium (CWE), a natural relaxation of WE. The difference between a CWE and a (non-combinatorial) WE is that the seller can package the items into indivisible bundles prior to sale, and the market does not necessarily clear.

We show that every valuation profile admits a CWE that obtains at least half of the optimal (unconstrained) social welfare. Moreover, we devise a polynomial time algorithm that, given an arbitrary allocation, computes a CWE that achieves at least half of its welfare. Thus, the economic problem of finding a CWE with high social welfare reduces to the algorithmic problem of social-welfare approximation. In addition, we show that every valuation profile admits a CWE that extracts a logarithmic fraction of the optimal welfare as revenue. Finally, to motivate the use of bundles, we establish strong lower bounds when the seller is restricted to using item prices only. The strength of our results derives partly from their generality — our results hold for arbitrary valuations that may exhibit complex combinations of substitutes and complements.

Key words. Combinatorial auctions, Walrasian Equilibrium, Envy-free, Approximation

AMS subject classifications. 91B26, 68W40, 68W25

1. Introduction. Recent years have been marked by an explosion of interest in the role of computer science theory in market design. Large-scale, computer-aided combinatorial markets are becoming a reality, with the FCC spectrum auctions emerging as a front-running example [3]. The potential outcome of this line of work is a system in which many bidders, each having complex preferences over combinations of items for sale, can express these preferences to an auction resolution algorithm that decides an appropriate outcome and payments. Spurred forward by this vision, the computer science community has generated an entire subfield of work on developing efficient algorithms for combinatorial allocation problems [1, 18, 21, 35, 37, 41].

Much of the existing work on combinatorial auctions has focused on the desideratum of incentive compatibility, where bidders are incentivized to report their preferences truthfully to an auction resolution mechanism. It is our view, however, that the connection between computational requirements and combinatorial market design is much broader than the design of incentive compatible mechanisms, and combinatorial extensions of complex markets are fundamental in a wider context. In this paper we study a classic market design problem: setting prices so that socially efficient outcomes arise when buyers select their most desired sets. We propose a natural combinatorial extension of this problem, whereby the seller can choose to bundle objects prior to assigning prices. We demonstrate that providing this basic operation to the seller leads to the existence of (and algorithms to find) near-optimal outcomes in settings that were previously known to suffer from severe limitations.

Background: Walrasian Equilibrium. A vast literature in economic theory is dedicated to methods of assigning prices to outcomes so that a market clears in equilibrium

*An extended abstract of this work has been presented at the 45'th ACM Symposium on the Theory of Computing (STOC) 2013.

[†]Tel-Aviv University, mfeldman@tau.ac.il

[‡]Microsoft Research, ngravin@microsoft.com brlucier@microsoft.com

and a socially efficient outcome arises. Suppose that we have a single seller¹ with a set M of m items for sale, and there is a set N of n buyers who have possibly complex preferences over the items, represented by a valuation function $v_i(\cdot) : 2^M \rightarrow \mathbb{R}_{\geq 0}$ that maps every subset of M to a real value. A strong notion of a pricing equilibrium for such a market is as follows. First, the seller sets prices $\{p_j\}_{j \in M}$ on the items for sale. Second, each buyer selects his most-desired set of items at those prices, i.e., a set S in $\operatorname{argmax}_S v_i(S) - \sum_{j \in S} p_j$. If no item is desired by more than one bidder (i.e., there is no over-demand), and all items are sold (i.e., there is no over-supply), this outcome is known as a Walrasian equilibrium (WE).

The WE solution concept is appealing: despite competition among the agents, every buyer is maximally happy with her allocation, the market clears, and the pricing structure is natural, simple, and transparent. Moreover, it is known that a WE, when it exists, is socially efficient; i.e., it maximizes *social welfare* — the sum of buyers’ valuations [7]. The main disadvantage of WE is that the concept is “too good to be true” — it is known to exist only for an extremely restrictive subset of sub-modular valuations, known as *gross substitutes* [30]. Since a motivating feature of combinatorial auctions is the ability to capture complementarities in the buyers’ preferences (i.e., super-additive valuation functions), this restriction limits the applicability of WE in many algorithmic mechanism design settings.

Circumventing the existence problem requires relaxing the WE notion, and several approaches can be taken with respect to this relaxation. One approach is to allow the seller to set arbitrary bundle prices instead of item prices; i.e., set a price p_S for every bundle S (see, e.g., [8, 10, 42]). This approach does lead to strong existence and efficiency results, but loses much of the simplicity and transparency that is offered by item pricing.

Another approach is to relax the requirement for market clearance (while still insisting that every buyer maximizes his utility). This approach is natural in settings where the seller wishes to maximize some well-defined objective function, such as social welfare or revenue, and might be able to credibly leave some unsold items in the market. This relaxation completely solves the existence problem; indeed, an outcome in which all items are priced prohibitively high would trivially adhere to the proposed equilibrium notion. It might, however, come at a huge social expense. This begs the question: can a good social welfare be supported in a pricing equilibrium that relaxes the market clearance condition? As we show in Section 2.2, the answer is surprisingly discouraging. In particular, even for the class of *fractionally subadditive* valuations [22] — a strict subset of subadditive functions that exhibits strong substitutability among items — the loss in social welfare can be linear in the number of items. Relaxing market clearance, therefore, is not sufficient, and a new approach is needed. In what follows we introduce a new equilibrium concept that captures a novel approach to the problem.

A New Concept: Combinatorial Walrasian Equilibrium (CWE). We propose to pair the notion of Walrasian equilibrium with a simple combinatorial operation, as follows. The seller first *partitions* the items for sale into indivisible *bundles*. This partition induces a *reduced market*, where individual items are no longer available, rather only the specified bundles. This operation can be perceived as redefining the items. Each individual bundle — now an indivisible item — is then assigned a price, overall pricing over the reduced market is *linear*, i.e., the price of a set of bundles

¹We note that the seller could be either a government agency wishing to maximize market efficiency, or a monopolist wishing to maximize revenue.

equals the sum of the bundles' prices². An *allocation* in the reduced market is then an assignment of bundles to bidders.

The outcome of such a process is a CWE if every bidder obtains a utility-maximizing set of bundles in the reduced market, and no conflict arises. Thus, the essential feature of a CWE is the ability of the seller to redefine his items by pre-bundling them prior to sale. This is an innocuous and natural power to afford the seller; after all, as the owner of the objects to be sold, it seems reasonable that he may choose to repackage them as he sees fit.

Clearly, a CWE is guaranteed to exist for every valuation profile, even without relaxing market clearance. Indeed, the seller could simply collect all objects together into a single grand bundle, and then sell that one bundle to the bidder who values it most. However, this may be a very inefficient outcome for the market. The natural combinatorial question, then, is whether there exists a partition of the objects (and associated prices) so that a CWE exists and has high social welfare. An additional question is how best to partition the objects and set prices in order to maximize the seller's revenue. We are interested in both the existential and computational aspects of these problems.

Our Results. We begin by providing characterization of the set of CWE allocations (i.e., allocations that admit supporting CWE prices). Specifically, an allocation can be implemented at CWE if and only if that outcome is an optimal solution to a certain linear program: the configuration LP for the assignment problem, restricted to the bundles in the outcome allocation. This implies that every CWE generates an efficient allocation of the bundles that are sold, though we note that this necessary condition is not always sufficient. In particular, the optimal allocation cannot necessarily be implemented at CWE; in Section 2.3 we exhibit an example where the welfare-optimal CWE attains only $2/3$ of the unconstrained optimal social welfare.

We next study the problem of finding CWE outcomes that maximize social welfare or revenue in general market settings. Recall that an *allocation* refers to an assignment of items to bidders. Our main result is the following:

Result 1 (2-approximation for social welfare): Given an allocation \mathbf{Y} , we provide an algorithm that computes a CWE \mathbf{X} that guarantees a social welfare of at least $\frac{1}{2}\text{SW}(\mathbf{Y})$, and runs in polynomial time, given an access to each bidder's *demand oracle*.

A direct corollary of the above result is that every instance admits a CWE that obtains at least half of the optimal (unconstrained) social welfare. Note that our result does not restrict the preferences of the bidders; it holds for arbitrary valuation functions, including those with complements. Moreover, since the result holds for arbitrary \mathbf{Y} , every social-welfare approximation can be converted into a CWE allocation that achieves the same approximation (up to factor 2). In other words, our algorithm can be interpreted as a *black-box reduction* that reduces the economic problem of finding a CWE with good social welfare to the algorithmic problem of social-welfare approximation for a given class of valuation functions. The fact that our method proceeds in a black-box manner is significant, as it allows a separation of the algorithmic and economic aspects of our pricing problem. Such reductions have been developed only rarely, such as for Bayesian incentive compatible mechanisms [6, 12, 32, 33].

The presentation of our algorithm is given in two stages. We first describe an algorithm that provides the desired approximation result, albeit it (a) might run

²This essential feature — linearity of prices in the reduced market — distinguishes the proposed solution concept from previous notions in the spirit of bundle pricing [8].

in exponential time, and (b) generates only an “approximate CWE” (in a sense we make explicit in Section 3). The advantage of this algorithm is its simplicity and its natural interpretation as an ascending price auction (see Section 3). A more challenging task is to modify the proposed algorithm into a polynomial time algorithm that preserves the same approximation ratio and generates an exact CWE. This is the content of Section 4. The algorithm is polynomial, given an access to a *demand oracle* of each agent (in the reduced market) — where agents get a set of item prices and respond with their most desired bundle given these prices. We note that in our context demand queries are unavoidable, since agents must be able to determine their demand sets. Additionally, since we consider reduced markets defined by the seller’s choice of partition, we will allow demand queries over any reduced market (not just the original instance without bundles).

We also provide a negative result illustrating the need for bundles:

Result 2 (linear lower bound on approximation factor, under item pricing and XOS valuations): There exists an instance in which no equilibrium with *item* prices gives a sublinear approximation to social welfare, even if valuations are fractionally subadditive.

We next consider the problem of revenue maximization and provide the following results.

Result 3 ($O(\log n)$ -approximation for revenue): Given an allocation \mathbf{Y} to n buyers, we provide an algorithm that computes a CWE \mathbf{X} that extracts revenue of $O(\log n)$ fraction of $\text{SW}(\mathbf{Y})$, and runs in polynomial time, given an access to agents’ *demand oracles*.

Moreover, this result is tight in terms of the trade-off between social welfare and revenue objectives for the outcomes that might be supported at CWE: there are instances in which no CWE extracts revenue more than a logarithmic fraction of the social welfare. A corollary of our result is that, for any class of valuations functions that admits a polytime constant approximation to social welfare, one can find (in polytime) a CWE that obtains an $O(\log n)$ approximation to the revenue-optimal CWE. Furthermore, a computational hardness result due to Briest [11] shows that one cannot hope for better than a polylogarithmic approximation: even in the special case of unit-demand bidders, where CWE reduces to envy-free pricing, there is a lower bound of $\Omega(\log^\epsilon(n))$ for the problem of approximating optimal revenue (subject to natural hardness assumptions).

Our techniques. The main combinatorial tool employed in the design and analysis of our algorithms resembles techniques taken from the theory of stable matching. In particular, our scheme proceeds in a fashion that is similar to the Gale-Shapley algorithm [29], with bidders and items residing in the two sides of the market, and bidders “making proposals” to the items. During the procedure, the price of each item reflects the item’s preferences over the bidders and it keeps growing monotonically. Meanwhile, the choices available to each bidder become scarcer and more expensive. Finally, the resulting allocation of buyers to bundles may be viewed as a matching, since every allocated set of items and/or bundles may be further treated as a single big bundle.

Despite the similarities to the Gale-Shapley algorithm, there are several important aspects that distinguishes our setup from the standard setting of stable matching. Firstly, our combinatorial auction model allows for bidders to demand *sets* of items. As a result, bundles demanded by the bidders may overlap in a complex way, which makes our task of resolving conflicts on the over-demanded items incomparably more

difficult than for the unit demand valuations in any matching setup. Secondly, the stable matching framework assumes no money in the market, while in our setting prices play a crucial role to guarantee stability. Finally, our routine begins with an initial allocation which is provided as part of the input and serves as a benchmark against which to compare the obtained social welfare. The initial allocation is indeed necessary if one is looking for an efficient implementation, due to the strong NP-hardness results on social-welfare approximation in combinatorial auctions.

The ascending-price nature of our basic algorithm leads to a potentially exponential runtime, as prices may climb slowly toward a stable profile. To address this problem, we must aggressively raise prices to “interesting” breakpoints. We then analyze the structure of the agents’ demands at these maximal price profiles, and find that by resolving the demands of agents in a particular order we can ensure that steady progress is made toward a final solution, leading to a polynomial runtime.

To construct a CWE with high revenue, a natural approach is to impose *reserves*: lower bounds on bundle prices. However, manipulating prices in this way can affect the structure of a final equilibrium in non-trivial ways, so that it is not clear that revenue will ultimately increase. To circumvent this issue, we begin with a CWE with high welfare, then modify prices by adding a constant amount to the price of *each* bundle. This operation is conceptually similar to imposing a reserve, but does not fundamentally change the structure of a stable allocation. Our approach to maximizing revenue then reduces to tuning the extent of this flat price increase.

Related Work. The study of pricing equilibria in markets and related concepts of outcome fairness have a rich history in theoretical economics, beginning with the introduction of competitive equilibria by Walras [45]. Some of the earliest modern work in the spirit of envy-free market outcomes is due to Foley [27] and Varian [44]. An envy-free outcome is one where no agent wishes to exchange outcomes with another. The line of work on market-clearing prices in our market assignment problem was initiated by Shapley and Shubik [43]. Characterizations of existence of Walrasian equilibria were studied in, for example, [2, 7, 30, 36, 38]. There is also a significant line of work devoted to computation of Walrasian equilibrium and the analysis of tâtonnement processes; see, for example, [14, 15, 17].

An alternative line of work considers markets with non-linear bundle prices. Such package auctions were formalized by Bikhchandani and Ostroy [8]. Applications to combinatorial auctions include mechanisms due to Ausubel and Milgrom [3], Wurman and Wellman [46], and Parkes and Ungar [42]. Our notion of CWE differs in that the seller commits to a partition of the objects, then sets linear prices over those bundles.

The problem of computing revenue-optimal envy-free prices has received recent attention in the computer science literature. Guruswami et al. [31] provide approximation algorithms for envy-free pricing in certain special cases, leading to a line of work improving on the attainable approximation factors [5, 13, 34] and a polylogarithmic lower bound [11]. Mu’alem [40] studies the revenue maximization question for agents with general types. The notion of envy-freeness has also been applied to problems in machine scheduling [16].

Fiat et al. [26] considered an extension of envy-freeness in which no agent envies any subset of other agents. This concept is related to our notion of CWE. However, crucially, they restrict their definition to agents with single-minded types, which dampens the distinction between multi-envy freeness and envy-freeness.

A significant line of work in the algorithmic mechanism design literature is concerned with the development of truthful mechanisms for combinatorial markets. See,

for example, [1, 18, 21, 24, 35, 37, 41] and references therein. The goal in this work is to develop algorithms that elicit truthful value revelation from the bidders. In contrast, we assume a full-information model and our goal is to develop an algorithmic pricing structure that satisfies certain transparency and fairness conditions.

Some of our algorithms make use of *demand queries*, a manner of eliciting preference information from bidders with complex valuations. For representative works on the power of demand queries, see [4, 9, 19, 39].

Fu, Kleinberg and Lavi [28] introduced the notion of *conditional equilibrium* as a WE relaxation, where no buyer wishes to add additional items to his allocation under the given prices, but may wish to drop ones. They show that, when buyers have submodular valuations, a conditional equilibrium always exists and every conditional equilibrium achieves at least half of the optimal social welfare. While their work is similar in spirit to the results herein, our equilibrium concept differs fundamentally in that it does not relax the requirement that every agent receives a bundle in his demand set. In particular, the conditional equilibrium notion violates basic envy-freeness conditions, in the sense that one bidder might prefer another agent's items, at their current prices, to his own allocation.

2. Model and preliminaries. We consider an auction framework with a set M of m indivisible objects and a set of n agents. Each agent has a valuation function $v_i(\cdot) : 2^M \rightarrow \mathbb{R}_{\geq 0}$ that indicates his value for every set of objects, is non-decreasing (i.e., $v_i(S) \leq v_i(T)$ for every $S \subseteq T \subseteq M$) and is normalized so that $v_i(\emptyset) = 0$. The profile of agent valuations is denoted by $\mathbf{v} = (v_1, \dots, v_n)$, and an auction setting is denoted by a tuple $A = (M, \mathbf{v})$.

A price vector $\mathbf{p} = (p_1, \dots, p_m)$ consists of a price p_j for each object $j \in M$. An *allocation* is a vector of sets $\mathbf{X} = (X_0, X_1, \dots, X_n)$, where $X_i \cap X_k = \emptyset$ for every $i \neq k$, and $\bigcup_{i=0}^n X_i = M$. In the allocation \mathbf{X} , for every $i \in N$, X_i is the bundle assigned to agent i , and X_0 is the set of unallocated objects; i.e., $X_0 = M \setminus \bigcup_{i=1}^n X_i$.

As is standard, we assume that each agent has a quasi-linear utility function; i.e., the utility of agent i being allocated bundle X_i under prices \mathbf{p} is $u_i(X_i, \mathbf{p}) = v_i(X_i) - \sum_{j \in X_i} p_j$. Given prices \mathbf{p} , the *demand correspondence* $D_i(\mathbf{p})$ of agent i contains the sets of objects that maximize agent i 's utility:

$$D_i(\mathbf{p}) = \left\{ S^* : S^* \in \operatorname{argmax}_{S \subseteq M} \{u_i(S, \mathbf{p})\} \right\}.$$

A tuple (\mathbf{X}, \mathbf{p}) is said to be *stable* for auction $A = (M, \mathbf{v})$ if $X_i \in D_i(\mathbf{p})$ for every $i \in N$. A price vector \mathbf{p} is *stable* if there exists an allocation \mathbf{X} such that (\mathbf{X}, \mathbf{p}) is stable. An allocation \mathbf{X} is *stable* if there exists a pricing \mathbf{p} such that (\mathbf{X}, \mathbf{p}) is stable.

For a partition $\Gamma = (\Gamma_1, \dots, \Gamma_k)$ of the item set M we slightly abuse notation and denote by $\Gamma = \{\Gamma_1, \dots, \Gamma_k\}$ the reduced set of items, where the valuation of each agent i of a subset $S \subseteq \Gamma$ is $v_i(\bigcup_{j: \Gamma_j \in S} \Gamma_j)$. That is, v_i can be naturally interpreted as a valuation over the set of bundles in Γ . Throughout the paper we will typically use j to index over partition elements. We denote by A_Γ an auction over the reduced set of items Γ with the induced valuation profile. We note that k could be either smaller or larger than n (e.g., by setting some elements of the partition to \emptyset).

Every allocation \mathbf{X} induces a partition $\Gamma(\mathbf{X}) = (X_0, \dots, X_n)$. A tuple (\mathbf{X}, \mathbf{p}) , where $\mathbf{X} = (X_0, \dots, X_n)$, and p_i is the price of X_i for every $X_i \neq \emptyset$, is a *Combinatorial Walrasian Equilibrium* (CWE) if (\mathbf{X}, \mathbf{p}) is stable in the auction $A_{\Gamma(\mathbf{X})}$. Clearly, X_0 may be an empty set, in which case no item remains unallocated (i.e., the market clears). Allowing for X_0 to be non-empty is essentially the relaxation of the market

clearance condition. An allocation \mathbf{X} is said to be CWE if it admits a price vector $\mathbf{p} \in \mathbb{R}_{\geq 0}^{n+1}$ such that (\mathbf{X}, \mathbf{p}) is CWE. A mechanism is said to be CWE if it maps every valuation profile \mathbf{v} to an outcome (\mathbf{X}, \mathbf{p}) that is CWE.

Relation to WE. A tuple (\mathbf{X}, \mathbf{p}) , where $\mathbf{X} = (X_0, X_1, \dots, X_n)$ and $\mathbf{p} = (p_1, \dots, p_m)$, forms a *Walrasian equilibrium* (WE) if (\mathbf{X}, \mathbf{p}) is stable in A and $p_j = 0$ for every item $j \in X_0$. When the latter condition is satisfied, we also say that (\mathbf{X}, \mathbf{p}) *clears the market*. A CWE is weaker than a WE in two ways: first, it allows for market reduction by means of bundling; second, it does not require market clearance.

2.1. Characterization. The characterization of a CWE allocation is closely related to the characterization of an allocation that can be supported in a WE [7]. For a given partition $\Gamma = \{\Gamma_1, \dots, \Gamma_k\}$ of the objects, the allocation of Γ to N can be specified by a set of integral variables $y_{i,S} \in \{0, 1\}$, where $y_{i,S} = 1$ if the set $S \subseteq \Gamma$ of partition elements is allocated to agent $i \in N$ and $y_{i,S} = 0$ otherwise. These variables should satisfy the following conditions: $\sum_S y_{i,S} \leq 1$ for every $i \in N$ (each agent is allocated at most one bundle) and $\sum_{i, S \supseteq \Gamma_j} y_{i,S} \leq 1$ for every $\Gamma_j \in \Gamma$ (each element of the partition is allocated to at most one agent), where we are using j to index over the elements of partition Γ . A *fractional allocation* of Γ is given by variables $y_{i,S} \in [0, 1]$ that satisfy the same conditions and intuitively might be viewed as an allocation of divisible items. The configuration LP for A_Γ is given by the following linear program, which computes the fractional allocation that maximizes social welfare.

$$\begin{aligned} & \max \sum_{i,S} v_i(S) \cdot y_{i,S} \\ & \text{s.t.} \sum_S y_{i,S} \leq 1 \text{ for every } i \in N \\ & \quad \sum_{i, S \supseteq \Gamma_j} y_{i,S} \leq 1 \text{ for every } \Gamma_j \in \Gamma \\ & \quad y_{i,S} \in [0, 1] \text{ for every } i \in N, S \subseteq \Gamma \end{aligned}$$

The characterization given in [7] states that a WE exists if and only if the optimal fractional solution to the allocation LP occurs at an integral solution. This characterization of a WE allocation can be used to derive a characterization of a CWE allocation. Recall that every allocation \mathbf{X} induces a partition $\Gamma(\mathbf{X}) = (X_0, \dots, X_n)$. Let $\Gamma(\dot{\mathbf{X}}) = (X_1, \dots, X_n)$; i.e., $\Gamma(\dot{\mathbf{X}})$ denotes the reduced set of items $\Gamma(\mathbf{X})$ excluding element X_0 . The WE characterization now implies the following CWE characterization.

CLAIM 1. *An allocation $\mathbf{X} = (X_0, X_1, \dots, X_n)$ is a CWE for A iff the configuration LP for $A_{\Gamma(\dot{\mathbf{X}})}$ has an integral optimal solution that sets $y_{i,X_i} = 1$ for all $i \in N$.*

Therefore, the problem of finding a CWE allocation is equivalent to the problem of finding a subset of the items and a bundling over this subset, such that the optimal fractional allocation over these bundles occurs at an integral allocation.

2.2. Stable item pricing. The proposed concept of CWE is weaker than the concept of WE both in that it allows to restrict the item set by bundling and in that it does not require market clearance. Clearly, relaxing any one of these conditions is sufficient to guarantee existence of a stable allocation. This begs the question whether it is possible to achieve good guarantees on the market efficiency (social welfare) by

relaxing only one of these conditions. The relaxation that allows bundling, but still requires market clearance, is studied in follow-up work [20, 25]; in particular, a lower bound of $\Omega(\log m)$ has been established for subadditive buyers [20]. In this section, we show that relaxing the market clearance condition alone is not sufficient, in general. In particular, for several families of valuation functions, we establish strong lower bounds on the social-welfare approximation that can be achieved in a stable item-priced allocation. These results reinforce the need for bundling, as captured by the CWE notion.

Unit-demand & single-minded valuations. Consider the following auction: bidder 1 is a unit-demand agent, who values any non-empty subset of the items at $1 + \epsilon$; bidder 2 is a single-minded agent, who desires the set of all items for a value of m . In the optimal allocation all m items must go to agent 2 resulting in a social welfare (SW) of m . However, in every stable pricing \mathbf{p} that supports this allocation, there exists an item $j \in [m]$ such that $p_j \leq 1$ (otherwise the set $[m]$ cannot be a demand set of agent 2). But this in turn implies that j is demanded by agent 1. Therefore, every stable allocation assigns a single item to agent 1 for a SW of $1 + \epsilon$, compared to the optimal SW of m , and the multiplicative linear gap of $\Omega(m)$ follows.

It might not come as a surprise that item prices are not sufficient to obtain high welfare if valuations are super-additive. After all, if items are complementary to each other, then bundling is an intuitive operation. Surprisingly, the next example shows that a linear gap may exist even if all valuations are sub-additive.

XOS (fractionally-subadditive) valuations. Consider an auction with m items and two agents with the following symmetric XOS valuations. Agent 1 is unit-demand and values every subset at $1/2 - \delta$, for a sufficiently small δ (that will be determined soon). Agent 2 values any subset of size k at $\max(1, k/2)$; it is easy to verify that this is an XOS valuation. The socially optimal outcome allocates all objects to the second agent for a total value of $m/2$. We claim that there is no stable pricing that sells more than two items. For every $m \geq 3$, an optimal integral solution obtains a value of $m/2$ (by giving all the items to agent 2). By the characterization given in [7] (see also Section 2), this allocation admits a stable pricing if and only if $m/2$ is the optimal fractional solution of the corresponding configuration LP. We will now show a fractional solution that obtains value greater than $m/2$ for every $\delta < \frac{1}{2(m-1)}$. Consider the fractional solution in which the allocation of the first (unit demand) agent is given by $y_{1,\{j\}} = 1/m$ for every $j \in [m]$, and the allocation of the second (XOS) agent is given by $y_{2,\{j\}} = \frac{1}{m(m-1)}$ for every $j \in [m]$, and $y_{2,[m]} = \frac{m-2}{m-1}$. One can easily verify that this is a feasible solution, and the welfare obtained by $\{y_{i,S}\}$ is given by $SW(y) = \frac{m}{2} + \frac{1}{2(m-1)} - \delta$, which is greater than $\frac{m}{2}$ for every $\delta < \frac{1}{2(m-1)}$, as required. We conclude that a stable outcome can allocate at most two objects, and thus the highest welfare that can be obtained in a stable allocation is $3/2 - \delta$, resulting in a multiplicative gap of $m/3$.

2.3. Efficiency loss due to CWE: A lower bound. In this section we prove a lower bound on the efficiency loss of combinatorial Walrasian equilibria. In particular, we show that there are instances in which no CWE obtains more than a $(2/3 + \epsilon)$ fraction of the optimal social welfare, for every $\epsilon > 0$. Consider an auction with 3 items and 3 bidders. Each agent i has valuation v_i such that $v_i(\{1, 2, 3\}) = v_i(\{1, 2, 3\} \setminus \{i\}) = 2 + \epsilon$, and for any other set S we have $v_i(S) = 1$ if $i \in S$ and $v_i(S) = 0$ otherwise. The optimal integral allocation has a social welfare of 3 with each object $i \in \{1, 2, 3\}$ being allocated to agent i . On the other hand, the fractional solution

$y_{1,\{2,3\}} = y_{2,\{1,3\}} = y_{3,\{1,2\}} = 1/2$ to the configuration LP obtains a total value of $3 + 3\varepsilon/2$ that exceeds the value of optimal integral solution. Therefore, in any CWE either one item is unsold or at least two items are bundled together. In either case the social welfare cannot exceed $2 + \varepsilon$. Thus, no CWE in this market can approximate the optimal social welfare within a factor better than $\frac{2+\varepsilon}{3}$ (for an arbitrary small $\varepsilon > 0$).

3. Social welfare approximation. In Section 2.3 we established a lower bound of 1.5 on the social welfare approximation that can be achieved at a CWE. In this section we show that there always exists a CWE that gives a 2-approximation to the optimal social welfare for arbitrary valuations.

Moreover, we show existence by describing a natural ascending-price process that converges to a CWE. This will be a natural extension of the classic Tâtonnement process that finds a Walrasian equilibrium for (gross) substitutes valuations. In traditional Tâtonnement, prices are initially set to zero, agents iteratively respond to the current prices with their demand, and prices of over-demanded items increase monotonically. In our case, in addition to increasing prices of over-demanded items, the market orchestrator will make use of an additional tool — merging bundles. These operations will be *monotone*: prices increase monotonically, and bundles never break apart.

We can describe this modified Tâtonnement process as an algorithm, listed as Algorithm 1. Informally, the algorithm begins by bundling objects according to an initial allocation (Y in the statement of Theorem 3.2) and setting properly-designed initial prices (specifically, pricing every Y_j at $v_j(Y_j)/2$). It maintains a pool of buyers who are not allocated a demanded set (initially all buyers). The algorithm iteratively chooses a buyer from the pool and asks for his most-demanded set. Whenever a buyer's demand set S is comprised of more than one item, the items in S are bundled together (irrevocably), the bundle S is allocated to him, and any agents who were allocated subsets of S are deallocated and placed back in the pool. It should be noted that with our “aggressive” bundling, setting initial prices too low (as in standard tâtonnements) can lead to a big welfare loss. Therefore, the initial prices must be carefully chosen. If the demanded set is a singleton that is already allocated to another agent, the conflict is resolved (in subprocedure **ResolveConflict**) by gradually increasing the price of the item (in discrete jumps of some $\varepsilon > 0$), until it is not demanded by one of these agents. The algorithm terminates when all agents' demands are satisfied.

We note that the direct implementation of the ascending-price auction described in Algorithm 1 is simple and intuitive, but has the disadvantage that it increments prices slowly in small, discrete jumps. As a result, the algorithm runs in pseudo-polynomial time, and it finds only an ε -CWE (in a sense we make explicit in Definition 3.1 below). In Section 4 we will present a modified algorithm that increments prices differently and thereby finds an exact CWE in polynomial time.

DEFINITION 3.1. *A tuple (\mathbf{X}, \mathbf{p}) is ε -stable for a given auction if $u_i(X_i, \mathbf{p}) \geq u_i(X_i', \mathbf{p}) - \varepsilon$ for every i and every set X_i' of items in the auction. A tuple (\mathbf{X}, \mathbf{p}) is an ε -approximate Combinatorial Walrasian Equilibrium (ε -CWE) if (\mathbf{X}, \mathbf{p}) is ε -stable in the auction $A_{\Gamma(\mathbf{X})}$. That is, each agent can gain at most ε utility by switching to a different set of bundles at the given prices.*

We will show that Algorithm 1 finds an ε -CWE while degrading the social welfare of an initial allocation by at most half, plus an additive error term that arises from the discretization of prices.

THEOREM 3.2. *Given an initial allocation \mathbf{Y} , Algorithm 1 computes an ε -CWE (\mathbf{X}, \mathbf{p}) such that $SW(\mathbf{X}) \geq \frac{1}{2}SW(\mathbf{Y}) - n\varepsilon$.*

Algorithm 1 Simple CWE Algorithm**Input:** Valuations \mathbf{v} ; initial allocation $\mathbf{Y} = (Y_1, \dots, Y_n)$; discretization parameter ε **Output:** A CWE (\mathbf{X}, \mathbf{p})

```

1: Initialize:  $\Gamma = \{Y_i : Y_i \neq \emptyset\}$ ;  $p(Y_i) = \frac{1}{2}v_i(Y_i)$  for all  $i$ ;  $X_i = \emptyset$  for all  $i$ ; Pool =  $N$ 
2: while Pool  $\neq \emptyset$  do
3:   Remove any element  $a$  from Pool
4:   if  $D_a(\mathbf{p}, \Gamma) \neq \emptyset$  then
5:     Choose  $S \in D_a(\mathbf{p}, \Gamma)$ 
6:      $X_a \leftarrow S$ 
7:     if  $|S| > 1$  then
8:       Set  $p(S) := \sum_{\Gamma_j \in S} p(\Gamma_j)$ 
9:       for  $i$  such that  $X_i \in S$  do
10:         $X_i \leftarrow \emptyset$ 
11:        Pool  $\leftarrow$  Pool  $\cup \{i\}$ 
12:         $\Gamma \leftarrow \mathbf{Bundle}(\Gamma, S) \setminus * \Gamma := \{\Gamma_j : \Gamma_j \notin S\} \cup \{S\} * \setminus$ 
13:     else
14:       if  $\exists b \neq a$  such that  $X_b = S$  then
15:         ResolveConflict( $S, a, b$ )
16:      $\setminus * \text{Note: If } D_a(\mathbf{p}, \Gamma) = \emptyset \text{ then } a \text{ is not added back to Pool } * \setminus$ 
16: Return  $(\mathbf{X}, \mathbf{p})$ 

```

ResolveConflict(S, a, b):

```

1: while  $S \in D_a(\mathbf{p}, \Gamma) \cap D_b(\mathbf{p}, \Gamma)$  do
2:    $p(S) \leftarrow p(S) + \varepsilon$ 
3: if  $S \notin D_a(p, \Gamma)$  then  $X_a \leftarrow \emptyset$ ; Pool  $\leftarrow$  Pool  $\cup \{a\}$ 
4: else  $X_b \leftarrow \emptyset$ ; Pool  $\leftarrow$  Pool  $\cup \{b\}$ 

```

Proof. First, it is easy to see that the procedure must terminate. Indeed, bundles monotonically merge and prices monotonically increase. Thus, assuming fixed price increments of ε , the algorithm is guaranteed to terminate.

Second, we claim that upon termination, the obtained allocation and prices is an ε -CWE. More specifically, we claim that after every iteration of the while loop on line 2, any agent not in the pool is obtaining at most ε less than his optimal utility under the current bundles and prices. Thus, when the pool becomes empty, every agent receives an ε -approximate demand at the current prices, and the final allocation and pricing is an ε -CWE. To see why the claim is true, note that when an agent a is removed from the pool, he is allocated a most-demanded set S (line 6). Since prices never decrease, this set S must remain a most-demanded set for agent a unless its price increases, which occurs only in **ResolveConflict**. If **ResolveConflict** causes an allocated set S to no longer be demanded, then the corresponding agent is returned to the pool, with one exception: if the competing agents in **ResolveConflict** stop demanding the set S simultaneously, only one of them is returned to the pool. In this case, since the remaining agent demanded the set at the previous price increment, his shortfall in utility can be at most one price increment, which is ε .

It remains to bound the welfare generated by the ε -CWE. Let U be the set of agents that get non-empty allocations in \mathbf{X} ; i.e., $U = \{i : X_i \neq \emptyset\}$. Then, the social

welfare is given by

$$\begin{aligned}
 \sum_{i \in U} v_i(X_i) &= \sum_{i \in U} (v_i(X_i) - p(X_i)) + \sum_{i \in U} p(X_i) \\
 &= \sum_{i \in U} u_i(X_i) + \sum_{i \in U} p(X_i) \\
 (3.1) \qquad &\geq \sum_{i \in U} u_i(X_i) + \sum_{i \in U} \sum_{Y_j \subset X_i} \frac{1}{2} v_j(Y_j),
 \end{aligned}$$

where recall that $\mathbf{Y} = (Y_1, \dots, Y_n)$ is the given initial allocation. The last inequality of (3.1) follows directly from the following facts: (i) every bundle Y_j is originally priced at $\frac{1}{2}v_j(Y_j)$, (ii) the price of a bundle that is created by a merge of two bundles equals the sum of their prices, and (iii) a bundle's price can only increase.

The second term in the RHS of (3.1) captures the welfare that comes from bundles Y_j that are being allocated in \mathbf{X} . We next need to take care of the welfare that comes from Y_j 's that are not being allocated in \mathbf{X} . We first observe that every bundle that is allocated in our procedure, keeps being allocated until termination. This is because a bundle can be deallocated from one agent only if it is being allocated to another agent. Therefore, every bundle Y_j that is not allocated in \mathbf{X} has never been allocated, and therefore still has a price of $\frac{1}{2}v_j(Y_j)$ (as priced originally).

Given the last observation, we conclude that for every bundle Y_j that is not allocated in \mathbf{X} , it must be that agent j is being allocated some other non-empty bundle in \mathbf{X} (i.e., $j \in U$). Indeed, if $j \notin U$, then agent j would gain a utility of $v_j(Y_j) - p(Y_j) \geq v_j(Y_j) - \frac{1}{2}v_j(Y_j) \geq 0$ from the bundle Y_j . This contradicts the fact that $\emptyset \in D_j(\mathbf{p})$, unless $v_j(Y_j) = 0$.

However, the agent $j \in U$ who has been allocated the bundle Y_j in \mathbf{Y} is allocated in the CWE \mathbf{X} another bundle which is preferred by him, up to an additive ε . This means that agent j 's utility from X_j is at least $\frac{1}{2}v_j(Y_j) - \varepsilon$.

Since all unallocated bundles Y_j were either allocated in \mathbf{Y} to agents $i \in U$, of which there are at most n , or satisfy $v_j(Y_j) = 0$, summing over these bundles we get

$$(3.2) \qquad \sum_{i \in U} u_i(X_i) \geq \sum_{Y_j \notin \bigcup_{i \in U} X_i} \frac{1}{2} v_j(Y_j) - n\varepsilon.$$

By plugging the last inequality in (3.1) we obtain

$$\sum_{i \in U} v_i(X_i) \geq \frac{1}{2} \sum_{Y_j} v_j(Y_j) - n\varepsilon = \frac{1}{2} \text{SW}(Y) - n\varepsilon$$

and the assertion of the theorem follows. \square

Applying CWE algorithm to the optimal allocation \mathbf{Y} and taking the limit as ε tends to 0 we derive the following corollary.

COROLLARY 3.3. *For every valuation profile \mathbf{v} , there exists a CWE that obtains at least half of the optimal social welfare.*

4. Polynomial-time implementation. Here we discuss how one could implement the ascending-price procedure efficiently in a polynomial number of demand queries to the agents. We note that in our context demand queries are indeed unavoidable, as agents must know their demand sets even to verify whether a given outcome is stable. We also emphasize that our procedure takes as input an initial

target allocation Y ; the problem of finding an initial allocation Y is outside the scope of our procedure. However, we can think of Y as being generated from some approximation algorithm tailored to a particular class of valuations.

The polynomial-time algorithm is listed as Algorithm 2. Informally speaking, Algorithm 2 makes two main changes to the more straightforward Algorithm 1. First, rather than incrementing prices when necessary to resolve a conflict, the new algorithm will attempt to maintain prices as high as possible, given the current allocation. That is, prices are raised as much as possible³ without changing the demand correspondence of any bidder that is not in the pool (i.e., any bidder currently holding a demanded set). The advantage of this approach is that since these maximal prices are easily computed from a given allocation, the algorithm can “skip” many rounds of incremental price increases and simply jump to a maximal price vector. Moreover, maintaining maximal prices has another advantage: when prices are maximal, each agent must be indifferent between their current allocation and at least one other set in their demand correspondence. This indifference simplifies the process of resolving conflicts when serving an agent from the pool. This leads to the second change to the algorithm, which is the way in which conflicts are resolved. For each buyer b with an allocation X_b , the algorithm also maintains an “alternate” allocation T_b which is also utility-maximizing. Whenever the demand set of an agent a from the pool is a singleton currently allocated to another buyer b , we immediately award the allocation to a and instead switch bidder b ’s allocation to T_b . Any conflicts over T_b (i.e., because it was previously allocated to some other bidder) are then resolved recursively.

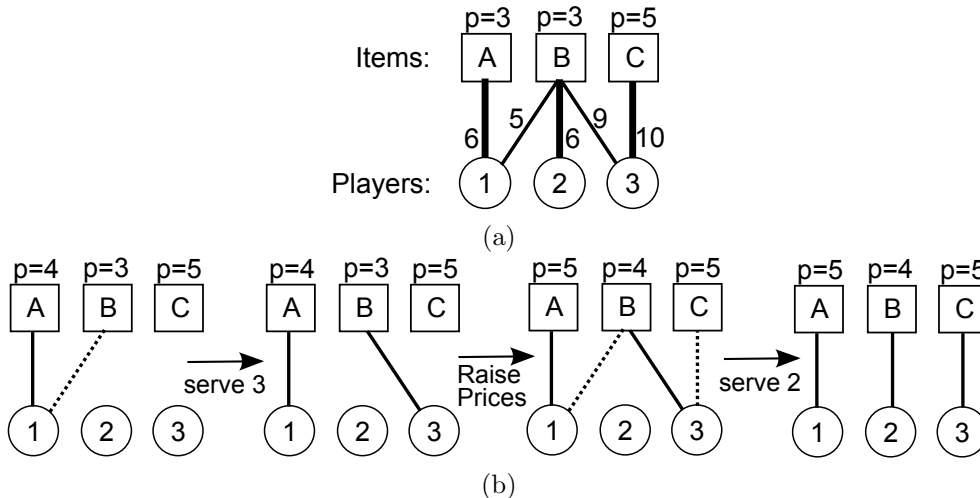


FIG. 4.1. An example execution of Algorithm 2. (a) Agent valuations and initial prices. All agents are unit demand. Bold lines denote the optimal allocation. (b) The execution of Algorithm 2, beginning after serving bidder 1 and applying **RaisePrices**. Solid lines indicate allocations X_i , dotted lines indicate sets T_i .

We note that the alternative allocations T_a , together with the provisional allocations X_a , can be interpreted as an *augmenting path* in the sense of matchings in the market. We will illustrate the augmenting-path nature of the algorithm with an

³In fact, following the literature on Walrasian prices [36], the set of price vectors that support a given allocation of non-pool buyers forms a lattice, since these are equilibrium prices in a reduced market. Our procedure picks the maximal element in this lattice.

example, depicted in Figure 4.1. In this example, there are three unit-demand agents $\{1, 2, 3\}$ and three items $\{A, B, C\}$. The valuations and initial prices are illustrated in Figure 4.1(a). If agent 1 is removed from the pool first, then she is allocated her most-demanded bundle $\{A\}$, and prices are raised so that she is indifferent between this and her next most-demanded bundle $\{B\}$. We then have $X_1 = \{A\}$ and $T_1 = \{B\}$, see Figure 4.1(b). Likewise, when agent 3 is served next from the pool, she is allocated item B . Prices on A and B are then raised until some agent becomes indifferent between their allocation and something outside $\{A, B\}$, which in this example is agent 3 and $T_3 = \{C\}$. The result is an augmenting path: agent 1 is allocated $\{A\}$ but is indifferent between $\{A\}$ and $\{B\}$; set $\{B\}$ is allocated to agent 3, who is indifferent between $\{B\}$ and $\{C\}$. When agent 2 is removed from the pool and demands item B , the allocation of agent 3 is immediately shifted to her alternative allocation $\{C\}$, yielding the final allocation.

We now turn to proving the correctness of Algorithm 2. We first note that, like Algorithm 1, Algorithm 2 is monotone in the following sense:

(MONOTONICITY). *Over the course of Algorithm 2, prices only increase and no bundle is ever split. Moreover, once a bundle is allocated it never becomes unallocated.*

Also, after an invocation of `AllocateDemand` completes, each bidder that is not in Pool is allocated a demanded set. Note that each allocated set is truly utility-maximizing, rather than only approximately utility-maximizing as in Algorithm 1.

LEMMA 4.1. *After a call to **AllocateDemand** terminates, each bidder $i \notin \text{Pool}$ is allocated a most demanded set.*

Proof. If **AllocateDemand** is called with bidder a , then we have two cases. If the demanded set for a is S with $|S| > 1$, then a is allocated S and other conflicting bidders are added to Pool, so the result holds inductively. If $|S| = 1$, then a is allocated S , so in particular a is allocated his most demanded set. The call to **AllocateDemand** then terminates only if there is no conflicting bidder; in this case the result holds. Note that we have not yet argued that **AllocateDemand** will, in fact, terminate. \square

We next show that **RaisePrices** increases the price vector to be maximal, within the set of \mathbf{p} such that $X_i \in D_i(\mathbf{p}, \Gamma)$ for each bidder i with $X_i \neq \emptyset$.

LEMMA 4.2 (Correctness of **RaisePrices**). *After each call to **RaisePrices**(\cdot), for each i with $X_i \neq \emptyset$, $p_i(X_i)$ is the maximal value such that $X_i \in D_i(\mathbf{p}, \Gamma)$.*

Proof. For a subset of players M and an allocation \mathbf{X} , write Γ_M for $\{X_j : j \in M\}$ and Γ_{-M} for $\Gamma \setminus \Gamma_M$. We first show that **RaisePrices** is equivalent to a different procedure which does not run in polynomial time. In this alternative procedure, the set M is defined as before, and the prices of elements of Γ_M are raised uniformly and continuously until the threshold at which the demand set of some $a \in M$ changes (note that this must occur eventually; the new demanded set may be \emptyset). When this occurs, a is removed from M , and the prices continue to increase for the elements remaining in Γ_M . This process continues until M is empty.

To see that this is equivalent to **RaisePrices**, consider some iteration of this new process, say with initial price vector \mathbf{p} and set M . Suppose the demand set of some $a \in M$ changes to S , and that the price vector at the point of the change is \mathbf{p}' . We claim that $S \subseteq \Gamma_{-M}$. The reason is that any S that includes elements of Γ_M has its price increase by at least as much as X_a , and hence a cannot prefer it to X_a at prices \mathbf{p}' . Thus when the demand set of a changes, it must be to some $S \in D_a(\mathbf{p}', \Gamma_{-M}) = D_a(\mathbf{p}, \Gamma_{-M})$. At the point at which the demand of a changes, it must be that $u_a(X_a, \mathbf{p}') = u_a(S, \mathbf{p}') = u_a(S, \mathbf{p})$. Thus the price increase between \mathbf{p} and \mathbf{p}' is precisely $u_a(X_a, \mathbf{p}) - u_a(S, \mathbf{p})$, and moreover player a is precisely the player

Algorithm 2 Polytime CWE Algorithm**Input:** Valuations \mathbf{v} ; target allocation $\mathbf{Y} = (Y_1, \dots, Y_n)$ **Output:** A CWE (\mathbf{X}, \mathbf{p})

- 1: Initialize: $\Gamma = \{Y_i : Y_i \neq \emptyset\}$; $p(Y_i) = \frac{1}{2}v_i(Y_i)$ for all i ; $X_i = \emptyset$ for all i ; Pool = N ;
Reject = \emptyset ; $T_i = \emptyset$ for all i
- 2: **while** Pool $\neq \emptyset$ **do**
- 3: Remove an arbitrary element a from Pool
- 4: **if** $u_a(S, \mathbf{p}) \leq 0$ for each $S \in D_a(\mathbf{p}, \Gamma)$ **then**
- 5: Reject \leftarrow Reject $\cup \{a\}$
- 6: **else**
- 7: Choose $S \in D_a(\mathbf{p}, \Gamma)$
- 8: **AllocateDemand**(a, S)
- 9: **RaisePrices**()

AllocateDemand(a, S):

- 1: **if** $|S| > 1$ **then**
- 2: **for** i such that $X_i \in S$ **do**
- 3: $X_i \leftarrow \emptyset$
- 4: Pool \leftarrow Pool $\cup \{i\}$
- 5: $p(S) \leftarrow \sum_{\Gamma_j \in S} p(\Gamma_j)$
- 6: $\Gamma \leftarrow$ **Bundle**(Γ, S) $\setminus * \Gamma := \{\Gamma_j : \Gamma_j \not\subseteq S\} \cup \{S\} * \setminus$
- 7: $X_a \leftarrow S$
- 8: **else**
- 9: $X_a \leftarrow S$
- 10: **if** $\exists b \neq a$ such that $X_b = S$ **then**
- 11: $X_b \leftarrow \emptyset$
- 12: **AllocateDemand**(b, T_b) $\setminus * \text{Change } b\text{'s allocation to } T_b \text{ (set in } \mathbf{X} \text{)}$
 RaisePrices from the previous iteration of the algorithm). $* \setminus$

RaisePrices():

- 1: Initialize: $M \leftarrow \{i : X_i \neq \emptyset\}$
- 2: **while** $M \neq \emptyset$ **do**
- 3: **for** $i \in M$ **do**
- 4: choose $S_i \in D_i(\mathbf{p}, \Gamma \setminus \{X_j : j \in M\})$ $\setminus * \text{Demand, excluding allocations to agents in } M * \setminus$
- 5: $d_i \leftarrow u_i(X_i, \mathbf{p}) - u_i(S_i, \mathbf{p})$ $\setminus * \text{Note } u_i(S_i, \mathbf{p}) \geq 0 * \setminus$
- 6: $a \leftarrow \operatorname{argmin}_{i \in M} d_i$
- 7: **for** $i \in M$ **do**
- 8: $p(X_i) \leftarrow p(X_i) + d_a$
- 9: $T_a \leftarrow S_a$ $\setminus * T_a$ is the set demanded by a , excluding allocations to agents in $M * \setminus$
- 10: $M \leftarrow M - \{a\}$

in M for which this quantity is minimal (since a was the first for whom the demanded set changed). It is therefore equivalent to directly compute this quantity for each player in M , choose the minimum, and raise the price of each object in Γ_M by this minimum amount. This is precisely what is done by **RaisePrices**, and hence the procedures are equivalent as claimed.

The lemma now follows easily from the definition of this equivalent process. For each i with $X_i \neq \emptyset$, we have that at the point when i is removed from M , an increase of $p(X_i)$ would cause $X_i \notin D_i(\mathbf{p}, \Gamma)$. Moreover, one element in $D_i(\mathbf{p}, \Gamma)$ is contained in Γ_{-M} , and the price of this element does not change between the point at which i is removed from M and the conclusion of the process. Thus, when the process concludes, it will still be that an increase of $p(X_i)$ would cause $X_i \notin D_i(\mathbf{p}, \Gamma)$. Thus $p(X_i)$ is maximal such that $X_i \in D_i(\mathbf{p}, \Gamma)$, as required. \square

Note that **RaisePrices** defines an ordering over the players with $X_i \neq \emptyset$: the order in which they are removed from M . Given an iteration of Algorithm 2, we will write π for this permutation defined by the invocation of **RaisePrices** on the previous iteration. That is, $\pi(i)$ denotes the order in which player i was removed; for notational convenience we will set $\pi(i) = \infty$ for all i with $X_i = \emptyset$. For instance, in the example illustrated in Figure 4.1, after **RaisePrices** is called following the agent 3 being served, we have $\pi(3) = 1$ (since agent 3 is removed first, with $T_3 = \{C\}$), $\pi(1) = 2$, and $\pi(2) = \infty$. Note that on the first iteration of Algorithm 2 we have $\pi(i) = \infty$ for all i .

We now bound the number of iterations that can occur on a single invocation of **AllocateDemand**.

LEMMA 4.3. *An invocation of **AllocateDemand** can recurse at most n times.*

Proof. Note that **AllocateDemand** concludes with a potential tail recursion, which can be thought of as an iteration of **AllocateDemand** with a different agent. We must show that this tail recursion cannot occur more than n times in a single invocation of **AllocateDemand**. To show this, we'll show that if **AllocateDemand** on input a results in a tail recursion with input b , then it must be that $\pi(b) < \pi(a)$. In particular, this means that no agent i can be passed as input to **AllocateDemand** more than once in a recursive chain, and hence the number of recursive calls is at most n .

To prove the claim, note that a recursive call occurs precisely when agent a demands a single object from Γ , and this bundle is currently assigned to a bidder b . In an initial (i.e. non-recursive) call to **AllocateDemand** we have $\pi(a) = \infty$ (since a was drawn from Pool) and $\pi(b) < \infty$, so the result holds trivially. In a recursive call we have $\pi(a) < \infty$, and $X_b \in D_a(\mathbf{p}, \Gamma)$. However, recalling our notation from the proof of Lemma 4.2, we know that the demanded set T_a of a is contained entirely in Γ_{-M} . Thus, since $T_a = X_b$, it must be that $b \notin M$ when a is removed from M , and hence $\pi(b) < \pi(a)$. \square

THEOREM 4.4. *Algorithm 2 runs in polynomial time.*

Proof. In each iteration of the main loop, either a set of objects is bundled or an agent is added to the rejection set R . Each of these can happen at most m and n times, respectively. Since each invocation of **AllocateDemand** also runs in polynomial time by the above lemma, the result follows. \square

THEOREM 4.5. *Algorithm 2 returns a CWE with social welfare at least half of the optimal allocation.*

Proof. The fact that Algorithm 2 returns a CWE follows immediately from Lemma 4.1. The argument for the approximation factor guarantee is the same as for Algorithm 1, as this depends only on the starting condition (which is unchanged) and the fact that no object becomes unallocated after it has been allocated. \square

5. Revenue approximation. In this section we consider the objective of the seller's revenue. Clearly, for any valuation profile, the seller's revenue can never exceed the social welfare of the optimal allocation. Therefore, given an allocation, its social

welfare serves as a natural benchmark for the revenue objective. We prove that given an allocation \mathbf{Y} , the seller can compute in polynomial time a CWE that extracts revenue of $1/O(\log n)$ of the social welfare of \mathbf{Y} .

We first prove a lower bound: there are instances in which no CWE extracts revenue greater than $1/\ln(n)$ times the optimal social welfare.

An example with a logarithmic separation between revenue and social welfare. Consider a market that consists of n items and n unit-demand buyers, where buyer i has value $v_i(\{j\}) = 1/i$ for every item j . In any optimal allocation every agent gets exactly one item, which results in a social welfare of $\sum_{i=1}^n 1/i \approx \ln n$. Any reduced set of items has the same structure of the agent's valuations as before; i.e. the reduced market contains $m \leq n$ items with n unit-demand buyers, where buyer i has value $v_i(\{j\}) = 1/i$ for every item j in the market. It is easy to verify that due to the structure of valuations, in any CWE all allocated items must have the same price. Suppose that k agents receive non-empty bundles; then one of these agents has index $i \geq k$. For this agent, $v_i(\cdot) = 1/i \leq 1/k$. Therefore, the price on every sold item is at most $1/k$, which generates revenue of at most $k \cdot 1/k = 1$.

We next show that given an allocation \mathbf{Y} , one can compute a CWE that extracts revenue within a factor $1/O(\log n)$ of the social welfare of \mathbf{Y} . As a corollary, there always exists a CWE in which the revenue is at least a $1/O(\log n)$ fraction of the optimal social welfare, which is an upper bound on the optimal revenue.

To see how to construct a CWE with high revenue, consider beginning with a CWE with high social welfare. A natural approach to increasing revenue is to impose reserve prices: a lower bound on the price of each bundle. However, manipulating prices in this way can affect demanded sets in non-trivial ways, and it is not clear that the final outcome will actually generate more revenue (or even be stable at all). Instead of imposing a reserve price, we will consider adding a constant amount to the price of *each* bundle. This operation is conceptually similar to imposing a reserve, but does not change the structure of a stable allocation (beyond compelling some agents to leave empty-handed). We prove that there exists at least one choice for this per-bundle price increase such that the corresponding revenue is least a logarithmic fraction of the initial social welfare.

We note that our approach is not guaranteed to maintain the high social welfare of the original CWE, and in fact seems likely to reduce welfare significantly when the per-bundle surcharge is large. We leave open the problem of finding a CWE that simultaneously achieves high welfare and revenue.

THEOREM 5.1. *Given an arbitrary allocation \mathbf{Y} , one can find a CWE that extracts revenue within factor $1/O(\log n)$ of $SW(\mathbf{Y})$ in a polynomial number of demand queries.*

Proof. Given an allocation \mathbf{Y} , we first run Algorithm 2 with \mathbf{Y} as an input, and obtain a CWE (\mathbf{X}, \mathbf{p}) such that $SW(\mathbf{X}) \geq \frac{1}{2}SW(\mathbf{Y})$. This step can be done in a polynomial number of demand queries, as established in Section 4. Let $\mathbf{X} = (X_0, X_1, \dots, X_k)$ and $\mathbf{p} = (p_0, \dots, p_k)$; that is, in the CWE (\mathbf{X}, \mathbf{p}) , for every $i = 1, \dots, k$, agent i receives the bundle X_i at a price of p_i . We next make the following important observation.

CLAIM 2. *Let (\mathbf{X}, \mathbf{p}) be a CWE, where $\mathbf{X} = (X_0, \dots, X_k)$ and $\mathbf{p} = (p_0, \dots, p_k)$. For any positive constant σ let \mathbf{p}^σ be the price vector $(p_0 + \sigma, \dots, p_k + \sigma)$, and \mathbf{X}^σ be the allocation*

$$\forall i = 1, \dots, k \quad X_i^\sigma = \begin{cases} X_i & \text{if } v_i(X_i) \geq p^\sigma \\ \emptyset & \text{otherwise.} \end{cases}$$

Then, $(\mathbf{X}^\sigma, \mathbf{p}^\sigma)$ is a CWE.

Proof. For any non-empty set S it holds that $v_i(S) - \sum_{j \in S} p_j^\sigma \leq v_i(S) - \sum_{j \in S} p_j - \sigma$. On the other hand, $v_i(X_i) - p_i^\sigma = v_i(X_i) - p_i - \sigma$. Since (\mathbf{X}, \mathbf{p}) is a CWE, it follows that $v_i(X_i) - p_i \geq v_i(S) - \sum_{j \in S} p_j$ for every S . Combining the above inequalities, we get that $u_i(X_i, p^\sigma) \geq u_i(S, p^\sigma)$. In addition, $u_i(X_i, p^\sigma) \geq 0$ if and only if $v_i(X_i) \geq p_i^\sigma$. The assertion follows. \square

Let $\text{SW}_0 = \sum_{i=1}^k v_i(X_i)$ denote the social welfare of CWE (\mathbf{X}, \mathbf{p}) . In addition, let $\ell = \lceil \log(2k) \rceil$, and for every integer $t \in \{1, \dots, \ell + 1\}$ define $\sigma^{(t)} = 2^{t-1} \frac{\text{SW}_0}{2k}$. Let $\mathbf{p}^{(t)}$ and $\mathbf{X}^{(t)}$ be the vectors with $X_i^{\sigma^{(t)}}$ defined as in Claim 2

$$\mathbf{p}^{(t)} = (p_0 + \sigma^{(t)}, \dots, p_k + \sigma^{(t)}), \quad \mathbf{X}^{(t)} = (X_1^{\sigma^{(t)}}, \dots, X_k^{\sigma^{(t)}}),$$

Due to Claim 2, for every $t \in \{1, \dots, \ell + 1\}$, $(\mathbf{X}^{(t)}, \mathbf{p}^{(t)})$ is a CWE. For every CWE $(\mathbf{X}^{(t)}, \mathbf{p}^{(t)})$, we let SW_t , REV_t and W_t denote its social welfare, revenue, and the set of indices of allocated bundles, respectively. Note that for every t , $W_{t+1} \subseteq W_t$. Finally, let REV_0 denote the revenue of CWE (\mathbf{X}, \mathbf{p}) . The following is the key lemma in the proof of the theorem.

LEMMA 5.2. $\exists t \in \{0, 1, \dots, \ell + 1\}$ s.t. $\text{REV}_t \geq \frac{\text{SW}_0}{8\ell}$.

Proof. We first observe that

$$\begin{aligned} \text{SW}_0 &= \sum_{i=1}^k v_i(X_i) = \sum_{i \in W_1} v_i(X_i) + \sum_{i \notin W_1} v_i(X_i) \\ &\leq \sum_{i \in W_1} v_i(X_i) + \sum_{i \notin W_1} (p_i + \frac{\text{SW}_0}{2k}) \\ &\leq \text{SW}_1 + \text{REV}_0 + k \cdot \frac{\text{SW}_0}{2k} \\ &= \text{SW}_1 + \text{REV}_0 + \frac{\text{SW}_0}{2}. \end{aligned}$$

The first inequality follows from the fact that for every $i \notin W_1$, $v_i(X_i) \leq p_i^{(1)}$, and the second inequality follows by substituting $\text{SW}_1 = \sum_{i \in W_1} v_i(X_i)$ and $\sum_{i \notin W_1} p_i \leq \text{REV}_0$.

Therefore, $\text{SW}_1 \geq \frac{1}{2} \text{SW}_0 - \text{REV}_0$. One may assume that $\text{REV}_0 \leq \frac{1}{4} \text{SW}_0$, since otherwise the assertion of the lemma follows directly. Thus, $\text{SW}_1 \geq \frac{1}{4} \text{SW}_0$.

We next show that $\text{SW}_{\ell+1} = 0$. For every $i \in \{1, \dots, k\}$, $p_i^{(\ell+1)} = p_i + 2^\ell \frac{\text{SW}_0}{2k} \geq p_i + \text{SW}_0 \geq p_i + v_i(X_i) \geq v_i(X_i)$; thus $W_{\ell+1} = \emptyset$ and $\text{SW}_{\ell+1} = 0$.

Given that $\text{SW}_1 \geq \frac{1}{4} \text{SW}_0$ and $\text{SW}_{\ell+1} = 0$, there must exist some $t \in \{1, \dots, \ell\}$

such that $SW_t - SW_{t+1} \geq \frac{SW_0}{4^\ell}$. We get:

$$\begin{aligned}
\frac{SW_0}{4^\ell} &\leq SW_t - SW_{t+1} \\
&= \sum_{i \in W_t \setminus W_{t+1}} v_i(X_i) \\
&\leq \sum_{i \in W_t \setminus W_{t+1}} \left(p_i + 2^t \frac{SW_0}{2k} \right) \\
&\leq \sum_{i \in W_t} \left(p_i + 2^t \frac{SW_0}{2k} \right) \\
&\leq 2 \sum_{i \in W_t} \left(p_i + 2^{t-1} \frac{SW_0}{2k} \right) = 2\text{REV}_t,
\end{aligned}$$

where the second inequality follows from the fact that for every $i \in W_t \setminus W_{t+1}$, $v_i(X_i) \leq p_i^{(t+1)}$. We get that $\text{REV}_t \geq \frac{SW_0}{8^\ell}$, as required. \square

Recall that SW_0 is within factor $\frac{1}{2}$ of $SW(\mathbf{Y})$. Combining this with the last lemma, and noting that $\ell = \lceil \log(2k) \rceil$, where $k \leq n$ implies that REV_t is within factor $\frac{1}{O(\log n)}$ of $SW(\mathbf{Y})$. To conclude the proof we observe that after the execution of Algorithm 2 all the quantities $v_i(X_i)$, SW_t , REV_t can be easily computed in polynomial time. \square

By applying theorem 5.1 with the initial allocation \mathbf{Y} being the optimal allocation, the following corollary follows.

COROLLARY 5.3. *For every valuation profile \mathbf{v} , there exists a CWE that extracts revenue within a factor $1/O(\log n)$ of the optimal social welfare.*

6. Open Problems. Our results leave many open questions and avenues for future research. First, our results leave a gap between the 2-approximation result for social welfare of a CWE and the lower bound of $3/2$. We conjecture that $3/2$ is the true bound, but closing this gap seems a challenging task. In particular, the integrality gap for the configuration LP approaches 2 when the integral solution is a matching, so this technique cannot be used to improve the gap.

Second, the equilibrium notion studied in this paper does not require market clearance. To what extent do our results extend to the stronger equilibrium notion of CWE with market clearance? Some progress has been made on this question since the conference version of this paper first appeared: market-clearing CWE can yield constant-factor approximations for certain valuation classes, including single-minded valuations [25], but a lower bound of $\Omega(\log m)$ has been established for subadditive buyers [20].

It would also be interesting to study how well item-pricing equilibria, without market clearance, can approximate social welfare. In this paper we established a lower bound of $\Omega(m)$ when valuations are fractionally subadditive. For the more restricted class of gross-substitutes valuations, a WE always exists, and thus optimal welfare can be achieved in equilibrium. Identifying more general families of valuations for which a constant fraction of the optimal SW can be achieved would be an interesting research direction. Again, some partial progress has been made: since the conference version of this paper first appeared, a lower bound of $\Omega(\sqrt{m})$ was established for submodular valuations [24].

Our algorithm receives an initial allocation as input and returns a CWE allocation that performs well with respect to the given allocation. Is there a natural process

that arrives at a good approximation, without receiving an initial allocation? The severe NP-hardness results for various valuation families preclude the possibility of a polynomial time process that would work for arbitrary inputs, but some families of valuations (e.g., submodular) seem particularly appealing in this context.

Our algorithms make use of demand queries, but of a special form: we allow the seller to define a partition of the objects into bundles, and can then ask demand queries with respect to the reduced market in which these bundles are the objects for sale. This interpretation of demand queries seems natural, since the assumption that agents can determine their demands is not tied to the particular items for sale. What is the additional computational power afforded by this (seemingly stronger) definition of demand queries?

We have shown two separate results: the existence of a CWE that obtains a good approximation to the optimal social welfare, as well as a CWE that obtains a logarithmic fraction of the optimal revenue. Is there a CWE that simultaneously obtains a good approximation to both social welfare and revenue?

Finally, this paper operates in the full information regime, where incentive compatibility is not a concern. An interesting question is to what extent our results can be extended to private-information settings. Since the appearance of the conference version of this paper, some progress has been made on item pricing mechanisms in the Bayesian setting, where valuations are private but drawn from known distributions. For this setting, there is a polynomial time, dominant-strategy incentive compatible mechanism that obtains a constant fraction of the optimal social welfare for submodular (in fact, fractionally subadditive) valuations [23]. However, despite the connection to item pricing, this mechanism does not necessarily generate stable outcomes. One could additionally ask: what is the best social welfare that can be obtained by an incentive-compatible CWE mechanism, whether in polynomial time or not?

REFERENCES

- [1] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *FOCS*, pages 482–491, 2001.
- [2] R. J. Aumann. Values of markets with a continuum of traders. *Econometrica*, 43(4):611–46, 1975.
- [3] L. M. Ausubel and P. R. Milgrom. Ascending auctions with package bidding. *B.E. Journal of Theoretical Economics*, 0(1):1, 2002.
- [4] A. Badanidiyuru, S. Dobzinski, and S. Oren. Optimization with demand oracles. In *EC*, pages 110–127, 2012.
- [5] M.-F. Balcan, A. Blum, and Y. Mansour. Item pricing for revenue maximization. In *EC*, pages 50–59, 2008.
- [6] X. Bei and Z. Huang. Bayesian incentive compatibility via fractional assignments. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '11*, pages 720–733, 2011.
- [7] S. Bikhchandani and J. W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *J. of Economic Theory*, 74(2):385–413, 1997.
- [8] S. Bikhchandani and J. M. Ostroy. The package assignment model. *J. of Economic Theory*, 107(2):377–406, 2002.
- [9] L. Blumrosen and N. Nisan. On the computational power of demand queries. *SIAM J. Comput.*, 39(4):1372–1391, 2009.
- [10] L. Blumrosen and N. Nisan. Informational limitations of ascending combinatorial auctions. *J. Economic Theory*, 145(3):1203–1223, 2010.
- [11] P. Briest. Towards hardness of envy-free pricing. *ECCC*, 13(150), 2006.
- [12] Y. Cai, C. Daskalakis, and S. M. Weinberg. Understanding incentives: Mechanism design becomes algorithm design. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 618–627, 2013.

- [13] M. Cheung and C. Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *FOCS*, pages 35–44, 2008.
- [14] Y. K. Cheung, R. Cole, and N. Devanur. Tatonnement beyond gross substitutes?: Gradient descent to the rescue. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC '13*, pages 191–200, New York, NY, USA, 2013. ACM.
- [15] B. Codenotti, B. McCune, and K. R. Varadarajan. Market equilibrium via the excess demand function. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 74–83, 2005.
- [16] E. Cohen, M. Feldman, A. Fiat, H. Kaplan, and S. Olonetsky. Envy-free makespan approximation. In *EC*, pages 159–166, 2010.
- [17] R. Cole and L. Fleischer. Fast-converging tatonnement algorithms for one-time and ongoing market problems. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 315–324, 2008.
- [18] P. Dhangwatnotai, S. Dobzinski, S. Dughmi, and T. Roughgarden. Truthful approximation schemes for single-parameter agents. In *FOCS*, pages 15–24, 2008.
- [19] S. Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In *STOC*, pages 139–148, 2011.
- [20] S. Dobzinski, M. Feldman, I. Talgam-Cohen, and O. Weinstein. Welfare and revenue guarantees for competitive bundling equilibrium. *CoRR*, abs/1406.0576, 2014.
- [21] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *STOC*, pages 644–652, 2006.
- [22] U. Feige and J. Vondrak. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. In *FOCS*, pages 667–676, 2006.
- [23] M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In *SODA*, pages 123–135, 2015.
- [24] M. Feldman, N. Gravin, and B. Lucier. Pricing without conflict in combinatorial auctions. In *Working paper*, 2015.
- [25] M. Feldman and B. Lucier. Clearing markets via bundles. In *SAGT*, pages 158–169, 2014.
- [26] A. Fiat and A. Wingarten. Envy, multi envy, and revenue maximization. In *WINE*, pages 498–504, 2009.
- [27] D. Foley. Resource Allocation and the Public Sector. *Yale Economics Essays*, 7:45–98, 1967.
- [28] H. Fu, R. Kleinberg, and R. Lavi. Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In *EC*, page 586, 2012.
- [29] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [30] F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *J. of Economic Theory*, 87(1):95–124, 1999.
- [31] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *SODA*, pages 1164–1173, 2005.
- [32] J. D. Hartline, R. Kleinberg, and A. Malekian. Bayesian incentive compatibility via matchings. In *SODA*, pages 734–747, 2011.
- [33] J. D. Hartline and B. Lucier. Bayesian algorithmic mechanism design. In *STOC*, pages 301–310, 2010.
- [34] J. D. Hartline and Q. Yan. Envy, truth, and profit. In *EC*, pages 243–252, 2011.
- [35] R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47(1):104–123, 2004.
- [36] A. S. Kelso, Jr. and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.
- [37] D. Lehmann, L. I. O’callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *J. ACM*, 49(5):577–602, 2002.
- [38] H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *J. of Political Economy*, 91(3):461–79, 1983.
- [39] V. Mirrokni, M. Schapira, and J. Vondrak. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *EC*, pages 70–77, 2008.
- [40] A. Mu’Alem. On multi-dimensional envy-free mechanisms. In *ADT*, pages 120–131, 2009.
- [41] C. H. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *FOCS*, pages 250–259, 2008.
- [42] D. C. Parkes and L. H. Ungar. Iterative combinatorial auctions: Theory and practice. In *AAAI*, pages 74–81, 2000.
- [43] L. S. Shapley and M. Shubik. The assignment game i: The core. *International J. of Game Theory*, 1:111–130, 1971.

- [44] H. R. Varian. Equity, envy, and efficiency. *J. of Economic Theory*, 9(1):63–91, 1974.
- [45] L. Walras. *Elements of Pure Economics: Or the Theory of Social Wealth*. Elements of Pure Economics, Or the Theory of Social Wealth. Routledge, 2003.
- [46] P. R. Wurman and M. P. Wellman. Akba: a progressive, anonymous-price combinatorial auction. In *EC*, pages 21–29, 2000.