

Combinatorial Walrasian Equilibrium

Michal Feldman
The Hebrew University of
Jerusalem & Harvard
University
michal.feldman@huji.ac.il

Nick Gravin
Nanyang Technological
University
ngravin@pmail.ntu.edu.sg

Brendan Lucier
Microsoft Research New
England
brlucier@microsoft.com

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 poly-time algorithm that, given an arbitrary allocation X , computes a CWE that achieves at least half of the welfare of X . 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, these results are complemented by strong lower bounds when the seller is restricted to using item prices only, which motivates the use of bundles. 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.

Categories and Subject Descriptors

F.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

Keywords

Combinatorial auctions; Walrasian Equilibrium; Envy-free

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'13, June 1–4, 2013, Palo Alto, California, USA.

Copyright 2013 ACM 978-1-4503-2029-0/13/06 ...\$15.00.

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, 13, 15, 26, 28, 32].

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 demanded 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 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 $\arg\max_S v_i(S) - \sum_{j \in S} p_j$. If no item is desired by more than

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

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 [6]. 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* [21]. 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., [7, 9, 33]). 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 [16] — 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 equals the sum of the bundles’ prices². The outcome of such a process is a CWE if every bidder obtains a utility-maximizing set of

²This essential feature — linearity of prices in the reduced market — distinguishes the proposed solution concept from previous notions in the spirit of bundle pricing [7].

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.

An observation that facilitates our analysis is a characterization of the set of CWE allocations (i.e., allocations that admit supporting CWE prices). Specifically, an outcome 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.

Characterization in hand, we study the problem of finding CWE outcomes that maximize social welfare and/or revenue in general market settings. 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 complementarities. 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 approximately efficient Bayesian incentive compatible mechanisms [24, 25].

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

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 poly-time algorithm that preserves the same approximation ratio. 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 the definition of CWE implicitly involves agents answering demand queries, hence our assumption that we have access to demand oracles is effectively driven by the notion of CWE itself. 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: there are instances 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 2 ($O(\log n)$ -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 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 [10] 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 aforementioned configuration LP provides a useful framework to study efficiency of stable pricing and existence of Walrasian equilibria. Generally speaking, the techniques of LP relaxations and rounding and especially the configuration LP are prevailing instruments in the study of combinatorial auctions. However, there is a new distinguishing feature in the context of Combinatorial Walrasian Equilibrium, as we have to determine a proper partition of the market into bundles. One may think of every bundle in a partition as imposing a set of additional linear constraints on the configuration LP of the initial (unpartitioned) market. In order to meet the goal of achieving a nearly efficient and stable state, we need to cast a carefully chosen set of those new constraints which push the solution of a derived LP to be achieved in an integral point corresponding to a proper allocation.

The main combinatorial tool employed in our design and analysis 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 [20], 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 choice of every bidder becomes scarcer and at greater expenses. 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 distinguish 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. Some of the earliest work in this spirit of envy-freeness is due to Foley [18] and Varian [35]. 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 [34]. Characterizations of existence of Walrasian equilibria were studied in, for example, [2, 6, 21, 27, 29].

An alternative line of work considers markets with non-linear bundle prices. Such package auctions were formalized by Bikhchandani and Ostroy [7]. Applications to combi-

natorial auctions include mechanisms due to Ausubel and Milgrom [3], Wurman and Wellman [36], and Parkes and Ungar [33]. 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. [22] provide approximation algorithms for envy-free pricing in certain special cases, leading to a line of work improving on the attainable approximation factors [5, 11, 23] and a polylogarithmic lower bound [10]. Mu’alem [31] studies the revenue maximization question for agents with general types. The notion of envy-freeness has also been applied to problems in machine scheduling [12].

Fiat et al. [17] 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, 13, 15, 26, 28, 32] 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, 8, 14, 30].

Fu, Kleinberg and Lavi [19] 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 is quite weak from the agents’ happiness perspective. In particular, it violates basic envy-freeness conditions, even for submodular valuations.

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 defined 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 standard, we assume that each agent has a quasi-linear utility function. That is, if agent i is allocated bundle X_i under prices \mathbf{p} , then the utility of agent i 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 $v_i(\bigcup_{j: \Gamma_j \in S} \Gamma_j)$. We denote by A_Γ an auction over the reduced set of items Γ with the induced valuation profile.

Every allocation \mathbf{X} induces a partition of the objects, $\Gamma(\mathbf{X}) = (X_0, \dots, X_n)$, where X_0 denotes the unallocated objects. 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.

Relation to WE.

A tuple (\mathbf{X}, \mathbf{p}) , where $\mathbf{X} = (X_0, X_1, \dots, X_n)$ and $\mathbf{p} = (p_1, \dots, p_m)$, is 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*. We emphasize that 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, so items with positive price can be left unsold.

2.1 Characterization of WE

We will make use of the following characterization of an allocation that can be supported in a WE [6]. For a given partition Γ 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 $S \subseteq \Gamma$ 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 to 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). 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.

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

The characterization given in [6] states that a WE exists if and only if the optimal fractional solution to the allocation LP occurs at an integral solution.

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. 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 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 linear gap 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. We claim that there is no stable pricing that sells more than a single item. For every $m \geq 2$, the optimal integral solution obtains a value of $m/2$ (by giving all the items to the XOS agent). By the characterization given in [6] (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 one object, and thus the highest welfare that can be obtained in a stable allocation is 1, resulting in a linear gap of $m/2$.

2.3 Efficiency loss due to CWE: A lower bound

In this section we prove a lower bound on the efficiency 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 the valuation $v_i(\{i\}) = 1$ and $v_i(\{1, 2, 3\} \setminus \{i\}) = 2 + \epsilon$. The optimal allocation has a social welfare of 3 with each object $i \in \{1, 2, 3\}$ being allocated to agent i . Any CWE that supports this allocation must, in fact, be a WE (since no items are bundled and all items are allocated). By the characterization given in [6], such a WE exists only if 3 is the optimal fractional solution of the corresponding configuration LP, but this is not the case: the fractional solution $y_{1,\{2,3\}} = y_{2,\{1,3\}} = y_{3,\{1,2\}} = 1/2$ obtains a total value of $3 + 3\epsilon/2 > 3$. 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 $2 + \epsilon$. Thus, no CWE in this market can approximate the optimal social welfare within a factor better than $\frac{2+\epsilon}{3}$ (for an arbitrary small $\epsilon > 0$).

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.

Our algorithm is a natural extension of the tâtonnement process that is used to achieve a Walrasian equilibrium for (gross) substitutes valuations. In a traditional tâtonnement process, 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.

Our algorithm is defined formally as Algorithm 1. Informally, the algorithm begins by bundling objects according to an initial allocation (Y in the statement of Theorem 3.1) 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 contains 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.

Algorithm 1 Simple CWE Algorithm

Input: Valuations \mathbf{v} ; initial 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  $= N$   
2: while Pool  $\neq \emptyset$  do  
3:   Remove an arbitrary 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:         $\Gamma \leftarrow \text{Bundle}(\Gamma, S) \setminus * \Gamma := \{\Gamma_j : \Gamma_j \notin S\} \cup \{S\} *$   
13:     else  
14:       if  $\exists b \neq a$  such that  $X_b = S$  then  
15:         ResolveConflict $(S, a, b)$   
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:    $p(S) \leftarrow p(S) + \epsilon$   
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  $\cup \{b\}$ 
```

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, until it is not demanded by one of these agents. The algorithm terminates when all agents' demands are satisfied.

We now present a formal description of the process that is simple and intuitive, but may run in exponential time (in m and n). In Section 4 we present a refined algorithm that obtains the same result with a poly-time implementation.

THEOREM 3.1. *Given an initial allocation \mathbf{Y} , Algorithm 1 computes a CWE (\mathbf{X}, \mathbf{p}) such that $\text{SW}(\mathbf{X}) \geq \frac{1}{2}\text{SW}(\mathbf{Y})$.*

PROOF. First, it is easy to see that the procedure must terminate. Indeed, bundles monotonically merge and prices monotonically increase. Moreover, on every iteration of the procedure, either a bundle merges (lines 7-12), a price strictly increases (lines 14-15 and **ResolveConflict**), or the size of the pool strictly decreases (otherwise). Thus, assuming fixed price increments of ϵ , the algorithm is guaranteed to terminate.

Second, upon termination, the obtained allocation and prices is a CWE. Indeed, every agent that is removed from the pool receives his most desired bundle given the current prices, and agents that are deallocated due to some other agent's demand go back to the pool. Thus, when the pool becomes empty, every agent receives his demand at the current prices, and the final allocation and pricing is CWE.

It remains to show that the returned CWE obtains at least half the social welfare of the original allocation \mathbf{Y} . 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

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

where the last inequality 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 (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 $j \in U$. Indeed, if $j \notin U$, then agent j would have gained 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 , contradicting the fact that $\emptyset \in D_j(\mathbf{p})$.

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. This means that agent j 's utility from X_j is at least $\frac{1}{2}v_j(Y_j)$.

Since all unallocated bundles Y_j were allocated in \mathbf{Y} to agents $i \in U$, summing over all these bundles, we get

$$\sum_{i \in U} u_i(X_i) \geq \sum_{Y_j \notin \bigcup_{i \in U} X_i} \frac{1}{2} v_j(Y_j). \quad (2)$$

By plugging the last inequality in (1) we obtain

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

and the assertion of the theorem follows. \square

Applying CWE algorithm to the optimal allocation \mathbf{Y} we derive the following corollary.

COROLLARY 1. *For every valuation profile \mathbf{v} , there exists a CWE that obtains at least half of the optimal social welfare.*

4. POLY-TIME IMPLEMENTATION

Here we discuss how one could implement our 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 in order just to verify whether a given outcome is stable. We also emphasize that our procedure takes as input an initial target allocation \mathbf{Y} ; the problem of finding an initial allocation \mathbf{Y} is outside the scope of our procedure. However, we can think of \mathbf{Y} as being generated from some approximation algorithm tailored to a particular class of valuations.

The polytime algorithm is given as Algorithm 2. Informally speaking, Algorithm 2 makes two main changes to the more straightforward Algorithm 1. First, after each main iteration of the algorithm, we invoke procedure **RaisePrices** to increase the prices of the allocated bundles; prices are raised as much as possible without changing the demand set of any bidder with an allocation. Second, whenever the demand set of a player from the pool is a singleton currently allocated to another bidder b , we immediately award the allocation to a . We then ensure that b is the next bidder to be considered, and allocate to b a particular set from his demand correspondence, which is determined during the previous call to **RaisePrices**.

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: $\Gamma \leftarrow$ **Bundle**(Γ, S) $\setminus * \Gamma := \{\Gamma_j : \Gamma_j \not\subseteq S\} \cup \{S\} * \setminus$
- 6: $p(S) \leftarrow \sum_{\Gamma_j \in S} p(\Gamma_j)$
- 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) $\setminus * \text{Attempt to allocate } T_b \text{ (from } \mathbf{RaisePrices} \text{) to } b. * \setminus$

RaisePrices():

- 1: Initialize: $N \leftarrow \{i : X_i \neq \emptyset\}$
- 2: **while** $N \neq \emptyset$ **do**
- 3: **for** $i \in N$ **do**
- 4: choose $S_i \in D_i(\mathbf{p}, \Gamma \setminus \{X_j : j \in N\})$ $\setminus * \text{Demand excluding allocations to agents in } N * \setminus$
- 5: $d_i \leftarrow u_i(X_i, \mathbf{p}) - u_i(S_i, \mathbf{p})$
- 6: $a \leftarrow \operatorname{argmin}_{i \in N} d_i$
- 7: **for** $i \in N$ **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 $N * \setminus$
- 10: $N \leftarrow N - \{a\}$

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 become unallocated.*

Also like Algorithm 1, after an invocation of **AllocateDemand** completes, each bidder that isn't in Pool is allocated a demanded set.

LEMMA 1. *After a call to **AllocateDemand** terminates, each bidder $i \notin$ Pool is allocated to his 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 \mathbf{p} to a maximal vector such that each bidder i with $X_i \neq \emptyset$ has $X_i \in D_i(\mathbf{p}, \Gamma)$.

LEMMA 2 (CORRECTNESS OF RAISEPRICES). *After each call to **RaisePrices**() , 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 N and an allocation \mathbf{X} , write Γ_N for $\{X_j : j \in N\}$ and Γ_{-N} for $\Gamma \setminus \Gamma_N$.

We first show that **RaisePrices** is equivalent to a different procedure which does not run in polynomial time. In this alternative procedure, the set N is defined as before, and the prices of elements of Γ_N are raised uniformly and continuously until the threshold at which the demand set of some $a \in N$ changes (note that this must occur eventually; the new demanded set may be \emptyset). When this occurs, a is removed from N , and the prices continue to increase for the elements remaining in Γ_N . This process continues until N 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 N . Suppose the demand set of some $a \in N$ changes to S , and that the price vector at the point of the change is \mathbf{p}' . We claim that $S \subseteq \Gamma_{-N}$. The reason is that any S that includes elements of Γ_N 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 $S \in D_a(\mathbf{p}', \Gamma_{-N}) = D_a(\mathbf{p}, \Gamma_{-N})$. 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 in N 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 N , choose the minimum, and raise the price of each object in Γ_N by this amount. This is precisely what is done by **RaisePrices**, and hence the procedures are equivalent as required.

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 N , 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_{-N} , 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 N . 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 $X_i = \emptyset$. 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 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 2, we know that the demanded set T_a of a is contained entirely in Γ_{-N} . Thus, since $T_a = X_b$, it must be that $b \notin N$ when a is removed from N , and hence $\pi(b) < \pi(a)$. \square

THEOREM 4.1. *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 n times. Since each invocation of **AllocateDemand** also runs in polynomial time by the above lemma, the result follows. \square

THEOREM 4.2. *Algorithm 2 returns an CWE with social welfare at least half of the optimal allocation.*

PROOF. The fact that Algorithm 2 returns an CWE follows immediately from Lemma 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 extract 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 log n separation. 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.

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 instead 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.

THEOREM 5.1. *Given an arbitrary allocation \mathbf{Y} , one can find a CWE that extracts revenue within factor $1/O(\log n)$ of $\text{SW}(\mathbf{Y})$ 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 $\text{SW}(\mathbf{X}) \geq \frac{1}{2}\text{SW}(\mathbf{Y})$. 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 1. *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 1

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

Due to Claim 1, 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 4. $\exists t \in \{0, 1, \dots, \ell + 1\}$ 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 $\text{SW}_t - \text{SW}_{t+1} \geq \frac{\text{SW}_0}{4\ell}$. We get:

$$\begin{aligned} \frac{\text{SW}_0}{4\ell} &\leq \text{SW}_t - \text{SW}_{t+1} \\ &= \sum_{i \in W_t \setminus W_{t+1}} v_i(X_i) \\ &\leq \sum_{i \in W_t \setminus W_{t+1}} \left(p_i + 2^t \frac{\text{SW}_0}{2k} \right) \\ &\leq \sum_{i \in W_t} \left(p_i + 2^t \frac{\text{SW}_0}{2k} \right) \\ &\leq 2 \sum_{i \in W_t} \left(p_i + 2^{t-1} \frac{\text{SW}_0}{2k} \right) = 2\text{REV}_t, \end{aligned}$$

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 $\text{REV}_t \geq \frac{\text{SW}_0}{8\ell}$, completing the proof of Lemma 4. \square

We are now ready to complete the proof of Theorem 5.1. Recall that SW_0 is within factor $\frac{1}{2}$ of $\text{SW}(\mathbf{Y})$. 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

$\text{SW}(\mathbf{Y})$. 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 2. 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, there is a gap between the 2-approximation result for social welfare and the lower bound of $3/2$. We conjecture that $3/2$ is the true bound, but closing this gap seems a challenging task. 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. An important question is whether our results extend to the stronger equilibrium notion of CWE with market clearance.

Third, it would be interesting to study how well item-pricing equilibria, without market clearance, can approximate social welfare. *A priori*, it seems intuitive that bundling should be most helpful when agent valuations exhibit complementarities. However, our negative result for fractionally sub-additive valuations, from Section 2.2, shows that item pricing can be insufficient for achieving a good welfare approximation even when valuations are complement-free. For the more restricted class of gross-substitutes valuations, a WE always exists, and thus optimal welfare can be achieved in equilibrium. Identifying the family of valuations for which a constant fraction of the optimal SW can be achieved via item-pricing equilibrium would be an interesting research direction. This question seems especially interesting for the class of submodular valuations.

Fourth, our algorithm receives an initial allocation as input and returns a CWE allocation that performs well with respect to the given allocation. This result leaves open the design of 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 poly-time process that would work for arbitrary inputs, but some families of valuations (e.g. submodular) seem particularly appealing in this context.

Fifth, 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?

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. That is, what is the best social welfare that can be obtained by an incentive-compatible CWE mechanism, poly-time or not?

7. 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] S. Bikhchandani and J. W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *J. of Economic Theory*, 74(2):385–413, 1997.
- [7] S. Bikhchandani and J. M. Ostroy. The package assignment model. *J. of Economic Theory*, 107(2):377–406, 2002.
- [8] L. Blumrosen and N. Nisan. On the computational power of demand queries. *SIAM J. Comput.*, 39(4):1372–1391, 2009.
- [9] L. Blumrosen and N. Nisan. Informational limitations of ascending combinatorial auctions. *J. Economic Theory*, 145(3):1203–1223, 2010.
- [10] P. Briest. Towards hardness of envy-free pricing. *ECCC*, 13(150), 2006.
- [11] M. Cheung and C. Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *FOCS*, pages 35–44, 2008.
- [12] E. Cohen, M. Feldman, A. Fiat, H. Kaplan, and S. Olonetsky. Envy-free makespan approximation. In *EC*, pages 159–166, 2010.
- [13] P. Dhangwatnotai, S. Dobzinski, S. Dughmi, and T. Roughgarden. Truthful approximation schemes for single-parameter agents. In *FOCS*, pages 15–24, 2008.
- [14] S. Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In *STOC*, pages 139–148, 2011.
- [15] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *STOC*, pages 644–652, 2006.
- [16] U. Feige and J. Vondrak. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. In *FOCS*, pages 667–676, 2006.
- [17] A. Fiat and A. Wingarten. Envy, multi envy, and revenue maximization. In *WINE*, pages 498–504, 2009.
- [18] D. Foley. Resource Allocation and the Public Sector. *Yale Economics Essays*, 7:45–98, 1967.
- [19] 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.
- [20] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [21] F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *J. of Economic Theory*, 87(1):95–124, 1999.
- [22] 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.
- [23] J. Hartline and Q. Yan. Envy, truth, and profit. In *EC*, pages 243–252, 2011.
- [24] J. D. Hartline, R. Kleinberg, and A. Malekian. Bayesian incentive compatibility via matchings. In *SODA*, pages 734–747, 2011.
- [25] J. D. Hartline and B. Lucier. Bayesian algorithmic mechanism design. In *STOC*, pages 301–310, 2010.
- [26] R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47(1):104–123, 2004.
- [27] J. Kelso, Alexander S and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.
- [28] D. Lehmann, L. I. O’callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *J. ACM*, 49(5):577–602, 2002.
- [29] H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *J. of Political Economy*, 91(3):461–79, 1983.
- [30] V. Mirrokni, M. Schapira, and J. Vondrak. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *EC*, pages 70–77, 2008.
- [31] A. Mu’Alem. On multi-dimensional envy-free mechanisms. In *ADT*, pages 120–131, 2009.
- [32] C. H. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *FOCS*, pages 250–259, 2008.
- [33] D. C. Parkes and L. H. Ungar. Iterative combinatorial auctions: Theory and practice. In *AAAI*, pages 74–81, 2000.
- [34] L. S. Shapley and M. Shubik. The assignment game i: The core. *International J. of Game Theory*, 1:111–130, 1971.
- [35] H. R. Varian. Equity, envy, and efficiency. *J. of Economic Theory*, 9(1):63–91, 1974.
- [36] P. R. Wurman and M. P. Wellman. Akba: a progressive, anonymous-price combinatorial auction. In *EC*, pages 21–29, 2000.