

Clearing Markets via Bundles

Michal Feldman*

Brendan Lucier†

January 20, 2014

Abstract

We study algorithms for combinatorial market design problems, where a set of heterogeneous and indivisible objects are priced and sold to potential buyers subject to equilibrium constraints. Extending the CWE notion introduced by Feldman et al. [STOC 2013], we introduce the concept of a *Market-Clearing Combinatorial Walrasian Equilibrium* (MC-CWE) as a natural relaxation of the classical Walrasian equilibrium (WE) solution concept. The only difference between a MC-CWE and a WE is the ability for the seller to bundle the items prior to sale. This innocuous and natural bundling operation imposes a plethora of algorithmic and economic challenges and opportunities. Unlike WE, which is guaranteed to exist only for (gross) substitutes valuations, a MC-CWE always exists. The main algorithmic challenge, therefore, is to design computationally efficient mechanisms that generate MC-CWE outcomes that approximately maximize social welfare. For a variety of valuation classes encompassing substitutes and complements (including super-additive, single-minded and budget-additive valuations), we design polynomial-time MC-CWE mechanisms that provide tight welfare approximation results.

1 Introduction

The *resource allocation* problem lies at the heart of theoretical economics: how should scarce resources be allocated among individual agents with competing interests? Since the emergence of the Internet, which enables complex resource allocation on a grand scale, this has naturally become a central problem in computer science as well. Economists generally approach this problem by adopting the notion of *market equilibrium*. Broadly speaking, a market equilibrium is a set of resource prices that are stable in the sense that all agents are maximally happy with their allocations and no resources are left unallocated. A long line of work has been dedicated to addressing the existence of equilibrium prices, and it has been shown (see, e.g., [2]) that market equilibria exist very generally, as long as the market is *convex*.

While this result sounds appealing in its generality, the convexity assumption usually requires that resources be infinitely divisible. In many applications of interest, especially those with a computational aspect, resources are indivisible; in these cases the convexity assumption is inapplicable. Do the results from the convex environments carry over to non-convex environments? In general the answer is no: the existence of equilibrium prices is not guaranteed. As a result, the study of markets for indivisible goods tends to focus on specific cases for which such prices exists, such as when buyer values satisfy the *gross substitutes* condition.

To be more precise, consider the following concrete model. Suppose there are m indivisible and heterogeneous items that should be allocated among n agents. The agents have potentially arbitrary preferences over bundles of items. Formally, every agent $i \in [n]$ has a valuation function v_i that maps every subset S of items into the value $v_i(S)$ that agent i derives from the bundle S (in monetary terms). Given a price vector $\mathbf{p} = (p_1, \dots, p_m)$, a bundle S is said to be in agent i 's *demand set* if S maximizes i 's *utility* given \mathbf{p} , defined as the difference between $v_i(S)$ and $\sum_{j \in S} p_j$.

*mfeldman@tau.ac.il. Tel-Aviv University.

†blucier@microsoft.com. Microsoft Research New England.

A *Walrasian equilibrium* (WE) is an assignment of *item prices* to the m items, and an assignment of the objects to the agents, such that: (1) every agent is allocated a bundle in his demand set, and (2) the market *clears*; i.e., all items are allocated¹. Such a solution is truly appealing; every agent is maximally happy despite competing preferences, the market clears, and the pricing structure is natural, simple, and transparent. Unfortunately, WE do not exist in general. A WE is guaranteed to exist only for a rather narrow class of valuations, known as gross substitutes (GS) valuations (a strict subset of submodular functions) [23]. This eliminates any hope for the applicability of WE to environments with valuations that exhibit complementarities, and many forms of substitutes as well.

Recently, Feldman et al. [18] proposed a relaxation of WE, termed a *combinatorial Walrasian equilibrium* (CWE). In a CWE, the seller can choose to *bundle* objects prior to assigning prices. This is a natural power to afford the seller, since as the owner of the resources he has some inherent ability to define what is meant by an “item.” The generated bundles induce a reduced market — a market in which the items for sale are the bundles generated by the seller. In addition to the bundling operation, the CWE further relaxes the WE notion in that it allows for items to remain unallocated (even when they are priced above zero). It is easy to see that a CWE exists for any valuation profile, since the seller could bundle all objects into a single item. The important issue, then, is whether there exists a CWE that is (approximately) efficient with respect to social welfare. Indeed, Feldman et al. show there always exists a CWE with at least half of the optimal (unconstrained) social welfare [18].

The CWE notion relaxes the WE notion in two ways: (i) it allows bundling, (ii) it does not require market clearance. While the bundling relaxation is central to the notion of CWE, the second relaxation warrants some discussion. The relaxation of market clearance is somewhat at odds with the notion of a two-sided market equilibrium: prices might not be stable from the seller’s perspective. After all, if an object (i.e. bundle) does not sell, the seller may be tempted to decrease its price in order to to sell it and increase revenue. The concept of CWE therefore implicitly requires that the seller pre-commit to (sub-optimal) prices, in addition to committing to a bundling of the items. With this in mind, we consider whether the relaxation of market clearance is truly necessary. It is easy to see that the bundling relaxation alone is enough to guarantee existence of an equilibrium, so the question becomes one of welfare. Can we hope to achieve the welfare bound of [18] without relaxing market clearance?

To answer this question we define the notion of a *Market-clearing CWE* (MC-CWE), which allows the bundling operation, but requires market clearance. A MC-CWE is, therefore, precisely a WE over the reduced market; it differs from a WE only in the ability of the seller to pre-bundle the items, and in particular it is a stronger (more restrictive) concept than CWE.

For a number of valuation classes, encompassing both substitutes and complements, we provide two types of results. The first finds the fraction of the optimal social welfare that can be obtained in a MC-CWE outcome. The second addresses the same problem but under the additional requirement of operating in polynomial time. Note that the approximation result established in [18] is only semi-computational — given the optimal allocation, it finds in polynomial time a CWE outcome that gives at least a half of the optimal welfare. Here, we devise polynomial approximation algorithms that do not need any initial allocation. Moreover, all of our approximation results match the computational lower bounds for their corresponding valuation classes.

We note that while the focus of our paper is welfare maximization, our analysis and results have some immediate implications on revenue maximization. In particular, all of our approximation results carry over to revenue approximation, since for each of our mechanisms the bound we obtain on the social welfare is precisely the revenue extracted at equilibrium. We also observe that from a revenue maximization perspective, MC-CWE might be preferred over WE even for GS valuations, as it can extract *arbitrarily* higher revenue than any WE solution (see Appendix B).

¹More precisely, every unallocated item is priced at zero.

	Uniform BA identical budgets	Uniform BA	Single minded	Super additive
MC-CWE gap	1	$\geq 8/7$ ≤ 2	1	1
Poly-time MC-CWE approx.	$\leq 4/3$	$\leq 8/3$	$O(m/\sqrt{\log m})$ [value] $\theta(\sqrt{m})$ [demand]	$\theta(\sqrt{m})$

Figure 1: Summary of our approximation results. The columns correspond to valuation classes. The first column corresponds to uniform budget additive valuations with identical budgets, and the second column corresponds to uniform budget additive valuations with arbitrary budgets. The first row corresponds to the gap in social welfare due to MC-CWE, disregarding computational considerations. The second row corresponds to the approximation that can be achieved with a MC-CWE poly-time mechanism. All approximation results assume the value-query model, unless otherwise stated. Note that m is the number of items for sale.

1.1 Our Results and Techniques

We construct MC-CWE mechanisms with certain welfare guarantees for various valuation classes. These results are stated below and summarized in Table 1.1.

Super-additive valuations In the case where agent valuations are super-additive, we show that there always exists a MC-CWE that maximizes social welfare. Note that it is not always possible to maximize social welfare without bundling, even if the market clearance requirement is relaxed: there exist input instances in which all bidders are single-minded, but every outcome with item pricing obtains only an $O(\sqrt{m})$ fraction of the optimal social welfare [17]. The use of bundling is therefore crucial in generating a socially efficient equilibrium outcome.

We next turn to computational algorithms. We show how to construct a MC-CWE that obtains an $O(\sqrt{m})$ approximation to the optimal social welfare in a polynomial number of *demand queries*², matching known lower bounds [31]. This mechanism is new; as far as we are aware, existing $O(\sqrt{m})$ -approximations do not satisfy the conditions of MC-CWE. Our mechanism proceeds by first crafting an $O(\sqrt{m})$ -approximate allocation and prices, then applying local search to repeatedly satisfy agent demands (bundling objects and/or raising prices in the process) until every agent obtains a demanded set at the given prices. Our construction makes use of demand queries in a way similar to that of [18]: rather than querying demand sets over the original space of objects, we query demand over bundles of objects (under linear prices). With *value queries*, we show that the $O(m/\sqrt{\log m})$ -approximate mechanism due to [25] satisfies the MC-CWE property. We also show that in the case of *single-minded valuations*, our demand-query mechanism can be modified to achieve an $O(\sqrt{m})$ approximation in a polynomial number of value queries.

Sub-additive valuations: uniform budget additive We then turn to the space of sub-additive valuations. Since efficient WE exist for the class of GS valuations, efficient MC-CWE exist for this class as well. We therefore consider a class of non-GS valuations: those that are *uniform budget-additive*. In this auction problem, each item a has a common value v_a , and each agent values the item at either 0 or v_a . Furthermore, each bidder has a budget that limits his value for any set of items. For this class, we demonstrate that WE do not necessarily exist. Moreover, we provide an instance in which no MC-CWE can achieve more than a $7/8$ fraction of the optimal social welfare. On the other hand, we show any allocation can be converted (in polynomial time) into a MC-CWE outcome that achieves at least half of the original social welfare. Thus, at least half of the optimal welfare can always be achieved in a MC-CWE outcome.

Turning to computational consideration, the welfare-maximization problem for this valuation class is known to be APX-hard, and the best-known algorithm achieves an approximation of $4/3$ (see [1, 5, 11, 35]).

²A demand query returns the utility-maximizing set given a vector of item prices.

Combined with the aforementioned algorithm, which converts every outcome to a MC-CWE outcome with at least half of the welfare of the original outcome, this implies a MC-CWE mechanism that achieves a $8/3$ approximation. Our analysis is based on the observation that an outcome can be implemented at MC-CWE if and only if that outcome is an optimal solution (among all fractional solutions) to a certain linear program: the *configuration LP* for the assignment problem restricted to the bundles in the outcome allocation. Our construction is based upon local search, but of a different nature than our super-additive mechanism. Rather than attempting to improve social welfare, we repeatedly bundle objects to *reduce* the optimal *fractional* welfare, shrinking the gap between fractional and integral solutions to the configuration LP.

We further show that if agents have identical budgets, the factor-2 loss disappears: any allocation can be made MC-CWE without loss in social welfare. Yet, even within this restricted class, a Walrasian equilibrium may not exist. These results are driven by connections between MC-CWE and the configuration LP for the combinatorial assignment problem.

1.2 Relation to Prior Work

There is a long line of work studying pricing equilibria in theoretical economics. [21] and [36] initiated the study of envy-freeness outcomes. Market-clearing prices in the market assignment problem were studied by [34]. Subsequent works have studied conditions for the existence of Walrasian equilibria [3, 8, 23, 26, 27].

An alternative to Walrasian equilibrium is to allow a seller to set (non-linear) prices on arbitrary bundles. Such package auctions were formalized by [9]. Some notable examples of combinatorial auction mechanisms that make use of bundle prices can be found in [4], [37], and [33]. Our notion of MC-CWE differs from package auctions in that the seller commits to a partition of the objects, then sets linear prices over those bundles.

[19] study an extension of envy-freeness, multi-envy-freeness, in which no agent envies any subset of other agents. This concept is related to our notion of MC-CWE. However, crucially, they restrict their definition to agents with single-minded types, which dampens the distinction from envy-freeness.

Many of our mechanisms access agent valuations via *demand queries*, whereby agents provide feedback to sellers through their choice of demand set at a given price vector. See [6, 10, 15, 28] for discussions on the power of the demand query model.

The assignment problem with budget-additive bidders has received significant attention in purely algorithmic frameworks. [1] present an algorithm with approximation factor $(1 - 1/e)$, and also introduce the value-uniform variant of the problem. This approximation factor was subsequently improved [5, 35], leading eventually to a $4/3$ approximation algorithm due to [11].

In a related paper, [22] introduce the notion of *conditional equilibrium*, a relaxation of WE in which no buyer wishes to add additional items to his allocation under the given prices. They show that every conditional equilibrium achieves at least half of the optimal social welfare, and a conditional equilibrium always exists when buyers have submodular valuations. Our equilibrium concept MC-CWE differs in that it is based on reducing the space of objects via bundling, rather than relaxing the notion of envy-freeness directly.

2 Preliminaries

The auction setting considered in this work consists of 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; 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 *buyer stable* for auction $A = (M, \mathbf{v})$ if $x_i \in D_i(\mathbf{p})$ for every $i \in N$. A tuple (\mathbf{x}, \mathbf{p}) is said to be *seller stable* for auction $A = (M, \mathbf{v})$ if for every $j \in x_0$, $p_j = 0$. The seller stability condition is also known as *market clearance*.

A tuple (\mathbf{x}, \mathbf{p}) is said to be a *Walrasian equilibrium* (WE) for auction $A = (M, \mathbf{v})$ if it is both buyer stable and seller stable.

It is well known that a WE exists only under restricted valuation functions. In particular, the class of *gross substitutes* (GS) valuations is a maximal class that admits Walrasian equilibria [23]. This class is a strict subset of submodular valuations. It is also known that if a WE exists, it is economically efficient (i.e., maximizes social welfare — the sum of agents' valuations).

We next define the notion of a *Market-Clearing CWE* (MC-CWE). The crux of the concept is that items are pre-partitioned into indivisible bundles. The constructed bundles are treated as indivisible objects, and the MC-CWE notion reduces to WE over the bundles. Crucially, although prices are now associated with bundles (of the original market), unlike previous notions, prices are linear once bundles are fixed. A formal definition follows.

For a partition $\mathbf{A} = (\mathbf{A}_1, \dots, \mathbf{A}_k)$ of the item set M we slightly abuse notation and denote by $\mathbf{A} = \{\mathbf{A}_1, \dots, \mathbf{A}_k\}$ the reduced set of items, where the valuation of each agent i of a subset $S \subseteq \mathbf{A}$ is $v_i(\bigcup_{j: \mathbf{A}_j \in S} \mathbf{A}_j)$. We denote by $A_{\mathbf{A}}$ an auction over the reduced set of items \mathbf{A} with the induced valuation profile.

Every allocation \mathbf{x} induces a partition $\mathbf{A}(\mathbf{x}) = (x_0, \dots, x_n)$. 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 *Market-Clearing CWE* (MC-CWE) if (\mathbf{x}, \mathbf{p}) is a WE in the auction $A_{\mathbf{A}(\mathbf{x})}$. An allocation \mathbf{x} is said to be MC-CWE if it admits a price vector $\mathbf{p} \in \mathbb{R}_{\geq 0}^{n+1}$ such that (\mathbf{x}, \mathbf{p}) is MC-CWE. A mechanism is said to be MC-CWE if it maps every valuation profile \mathbf{v} to an outcome (\mathbf{x}, \mathbf{p}) that is MC-CWE.

Relation to Combinatorial Walrasian equilibrium (CWE)

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 CWE if (\mathbf{x}, \mathbf{p}) is buyer stable in the auction $A_{\mathbf{A}(\mathbf{x})}$. Note that a MC-CWE is weaker than a WE, since it allows for a pre-sale partition of the goods. On the other hand, MC-CWE is stronger than a CWE, since it requires seller stability on top of buyer stability.

The MC-CWE concept is also closely related to the notion of envy-freeness. See more on this relation in Appendix A.

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 [8]. A similar observation was already given in [18], but we state it here for completeness, as it is used in later sections.

For a given partition \mathbf{A} of the objects, the allocation of \mathbf{A} 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 $S \subseteq \mathbf{A}$ 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 $\sum_{i, S \supseteq \mathbf{A}_j} y_{i,S} \leq 1$ for every $\mathbf{A}_j \in \mathbf{A}$ (each element of the partition is allocated to at most one agent). A *fractional allocation* of \mathbf{A} 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 $A_{\mathbf{A}}$

is given by the following linear program, which computes the fractional allocation that maximizes social welfare.

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

The characterization given in [8] 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 MC-CWE allocation.

Recall that every allocation \mathbf{x} induces a partition $\mathbf{A}(\mathbf{x}) = (x_0, \dots, x_n)$. The WE characterization implies the following MC-CWE characterization.

Claim 2.1. *An allocation $\mathbf{x} = (x_0, x_1, \dots, x_n)$ is a MC-CWE for A iff the configuration LP for $A_{\mathbf{A}(\mathbf{x})}$ has an integral optimal solution that sets $y_{i,x_i} = 1$ for all $i \in N$.*

In light of Claim 2.1, the problem of finding a MC-CWE allocation is equivalent to the problem of finding bundles such that the optimal social welfare generated by fractional and integral allocations of *these bundles* are identical, then returning an efficient allocation of those bundles.

Gap in social welfare due to MC-CWE: a lower bound. In [18], a lower bound on the social welfare of any CWE allocation is established. Specifically, using the equivalence of Claim 2.1, it is shown that the welfare of any CWE allocation is at most $2/3$ of the welfare of an optimal (non-CWE) allocation. We note that the same example trivially carries over to the stronger notion of MC-CWE; i.e., the welfare of any MC-CWE allocation is at most $2/3$ of the optimal welfare. The established lower bound holds for general valuations, but disappears for some families of valuation functions (as will be shown in later sections of this paper).

3 Super-additive valuations

In this section we study MC-CWE outcomes when agent valuations are super-additive. We first show that there is no loss in efficiency due to MC-CWE. In particular, every efficient allocation is MC-CWE.

We say that allocation \mathbf{x} is *bundle-efficient* for \mathbf{v} if, for all functions $\beta : [n] \rightarrow [n]$, we have $\sum_i v_i(x_i) \geq \sum_i v_i\left(\bigcup_{j \in \beta^{-1}(i)} x_j\right)$. That is, a bundle-efficient allocation \mathbf{x} maximizes social welfare among all ways to allocate the (indivisible) bundles x_1, \dots, x_n to the agents.

Theorem 3.1. *If agents are super-additive and \mathbf{x} is a bundle-efficient allocation, then the price vector $p_i = v_i(x_i)$ is such that (\mathbf{x}, \mathbf{p}) is MC-CWE.*

Proof. Pick any agent i and set of agents S . If $i \in S$, then

$$v_i(x_i) - p_i = 0 \geq v_i(\bigcup_{j \in S} x_j) - \sum_{j \in S} v_j(x_j) = v_i(\bigcup_{j \in S} x_j) - \sum_{j \in S} p_j,$$

where the inequality follows from the efficiency of allocation \mathbf{x} . Consider next the case where $i \notin S$. It holds that

$$v_i(x_i) + \sum_{j \in S} v_j(x_j) \geq v_i(\bigcup_{j \in S \cup \{i\}} x_j) \geq v_i(x_i) + v_i(\bigcup_{j \in S} x_j),$$

where the first inequality follows from the efficiency of allocation \mathbf{x} , and the second inequality follows from the super additivity of v_i . It follows that

$$v_i(x_i) - p_i = 0 \geq v_i(\cup_{j \in S} x_j) - \sum_{j \in S} v_j(x_j) = v_i(\cup_{j \in S} x_j) - \sum_{j \in S} p_j.$$

The assertion of the theorem follows. □

A consequence of Theorem 3.1 is that the full surplus can always be extracted by the seller as revenue.

We note that the use of bundling is necessary for the statement of Theorem 3.1, even if we relax the market-clearing requirement of Walrasian equilibrium and even for single-minded bidders. There are known auction instances for which, given any set of item prices for which agent demand sets are disjoint, the social welfare of the resulting allocation is an $O(\sqrt{m})$ fraction of the optimal social welfare [17]. For completeness we describe such an example in Appendix C.

3.1 Polynomial-Time Mechanisms

We next study the power of poly-time approximation mechanisms for maximizing social welfare in MC-CWE outcomes (compared to the optimal welfare that can be achieved by any mechanism, poly-time or not, MC-CWE or not). Of particular interest is the question whether the MC-CWE requirement entails an additional loss on top of the loss incurred due to the poly-time requirement alone. In our analysis, we distinguish between mechanisms that operate in the value-query and demand-query models, as is standard in the literature. We find that in both models, it is possible to construct a MC-CWE mechanism with an approximation factor matching that of the best-known approximation algorithms for welfare maximization. In particular, there exists a poly-time MC-CWE mechanism that achieves an $O(\sqrt{m})$ approximation under the demand-query model, and a poly-time MC-CWE mechanism that achieves an $O(m/\sqrt{\log m})$ approximation under the value-query model.

We first present a MC-CWE approximation mechanism for superadditive valuations using demand queries. This mechanism, which we call SuperAdditiveMC-CWE is listed as Algorithm 1. proceeds in two phases. In the first phase, it builds a preliminary solution by repeatedly allocating the set that maximizes *value density*. That is, set S is allocated to agent i so that $v_i(S)/|S|$ is maximized, and this process is then iterated on the remaining items. A bidder can be allocated to multiple times in this phase, in which case she is allocated the union of the assigned sets. After all objects have been allocated in this manner, we check whether the welfare can be improved by allocating all objects to a single player; if so, we do so and the mechanism ends. Otherwise we proceed to phase 2, where we repeatedly apply local improvements to the allocation. Specifically, if we write (x_1, \dots, x_n) for the tentative allocation, we look for circumstances in which some player i has more value for a set of bundles from among $\{x_1, \dots, x_n\}$ than the players to whom those bundles were previously assigned; if such a case exists, we bundle all of these items together and reallocate them to player i , then repeat the process with this updated tentative allocation. Note that this step amounts to repeatedly satisfying the demand of a player in the market with items $\{x_1, \dots, x_n\}$ and prices $p_i = v_i(x_i)$, until no further demands are made (which must occur since these prices only increase). When this process terminates we return the resulting allocation. The following theorem establishes the $O(\sqrt{m})$ approximation and polynomial run time of the algorithm.

Theorem 3.2. *Algorithm SuperAdditiveMC-CWE returns a MC-CWE outcome that $O(\sqrt{m})$ -approximates the optimal social welfare over all assignments. Furthermore, it can be implemented in a polynomial number of demand queries.*

To prove Theorem 3.2, we first note that the loop on lines 11-14 of Algorithm 1 must halt after polynomially many iterations. The reason is that, on each iteration of the loop, the number of players j such that $x_j = \emptyset$ increases by at least one, except in the case where $|T| = 1$ and $x_i = \emptyset$. Consider a sequence of

Algorithm 1 SuperAdditiveMC-CWE

Require: Additive valuations v_1, \dots, v_n .

1: Initialize $R \leftarrow M$, $x_i \leftarrow \emptyset$ for each $i \in N$

 // Phase 1: Build initial allocation

2: **repeat**

3: $(i, S) \leftarrow \operatorname{argmax}_{i \in [n], S \subseteq R} \{v_i(S)/|S|\}$

4: $x_i \leftarrow x_i \cup S$, $R \leftarrow R \setminus S$;

5: **until** $R = \emptyset$

6: **if** $\exists i \in [n]$ such that $v_i(M) > \sum_j v_j(x_j)$ **then** $x_i \leftarrow M$ and $x_j \leftarrow \emptyset$ for all $j \neq i$;

 // Phase 2: Local search

7: **while** $\exists i \in [n], T \subseteq N$ such that $v_i(\cup_{j \in T} x_j) > \sum_{j \in T} v_j(x_j)$ **do**

8: Choose $i \in [n]$ and $T \subseteq N$ maximizing $v_i(\cup_{j \in T} x_j) - \sum_{j \in T} v_j(x_j)$.

9: $x_i \leftarrow x_i \cup (\cup_{j \in T} x_j)$, $x_j \leftarrow \emptyset$ for all $j \in T \setminus \{i\}$

10: **end while**

11: **return** \mathbf{x}

iterations of the loop during which the number of players with empty allocations does not increase. In each iteration, since we choose the i and $T = \{j\}$ that maximizes $v_i(x_j) - v_j(x_j)$, a certain player i can be chosen only once during this sequence. This sequence of iterations therefore has length at most n . Since the number of players with empty allocations can increase at most n times, we conclude that this loop must halt after at most n^2 iterations.

We next analyze the number of demand queries needed to implement Algorithm 1.

Lemma 3.3. *Algorithm 1 can be implemented in a polynomial number of demand queries.*

Proof. It is clear that all operations other than lines 3, 11, and 12 can be implemented in polynomial time, so we will focus on those. Line 3 can be implemented with n *Relative-demand query* in the sense of Blumrosen and Nisan [10]. A relative demand query involves generating a vector \mathbf{p} of non-zero object prices, and returns the set S maximizing $v_i(S)/\sum_{i \in S} p_i$. Line 3 can be implemented by posing a demand query to each bidder with the price vector that assigns a price of 1 to each item in R and an arbitrarily large price to each item in $M \setminus R$. Since a relative demand query can be implemented by a polynomial number of demand queries [10], Line 3 can be implemented in polynomially many demand queries as well.

Lines 11 and 12 can be implemented as follows. We focus on line 12; line 11 can be implemented in the same way. We will define a new auction instance, over the original player set N and a new set of objects $\{a'_1, \dots, a'_n\}$. Object a'_i is intended to represent the set x_i in the original auction. Given some $T \subseteq [n]$, the value of player i for some object set $x' = \{a'_j : j \in T\}$ will be defined as $v_i'(x') := v_i(\cup_{j \in T} x_j)$. Note that this new auction corresponds precisely to the original auction with the objects partitioned into indivisible bundles x_1, \dots, x_n . We can then implement line 12 by making a single demand query to each agent i , *with respect to this new bundled auction*, under prices given by $p'_j = v_j(x_j)$. If any player has positive value for the demanded set returned by this query, then this player and his demanded set will satisfy the condition on line 12. Whichever player has the highest positive value, plus his demanded set, will be one chosen on line 12. If every player has value 0 for his demanded set, then the condition on line 9 evaluates to false. \square

To show correctness, we first show that SuperAdditiveMC-CWE generates a MC-CWE outcome, then bound the social welfare of the outcome it returns.

Claim 3.4. *Algorithm SuperAdditiveMC-CWE generates a MC-CWE outcome.*

Proof. The outcome \mathbf{x} generated by SuperAdditiveMC-CWE is such that, for all $i \in [n]$ and all $T \subseteq N$, $v_i(\cup_{j \in T} x_j) \leq \sum_{j \in T} v_j(x_j)$. Thus, if we set prices $p_i = v_i(x_i)$, we have that $v_i(x_i) - p_i = 0 \geq v_i(\cup_{j \in T} x_j) - \sum_{j \in T} p_j$ for all i and all $T \subseteq N$. Thus outcome (\mathbf{x}, \mathbf{p}) is MC-CWE. \square

Claim 3.5. *Algorithm SuperAdditiveMC-CWE obtains an $O(\sqrt{m})$ approximation to the optimal social welfare.*

Proof. We will argue that, on any superadditive input profile, the social welfare generated by our mechanism is at least that of a mechanism proposed by Blumrosen and Nisan, which was shown to obtain an $O(\sqrt{m})$ approximation to the optimal social welfare [10]. The mechanism due to Blumrosen and Nisan, which we will refer to as the BN mechanism, is identical to the first phase of our mechanism, except that it only allows a single set to be allocated to any single player in the loop on lines 2-5.

Consider first the case in which all agents are single-minded. In this case, phase 1 of our mechanism is identical to the BN mechanism (since our mechanism can allocate to any agent at most once in lines 2-6, by single-mindedness). Thus, after phase 1, the preliminary allocation generated by our mechanism has social welfare equal to that of the BN mechanism. We then note that the social welfare of the allocation can only increase over each iteration of the loop at lines 11-15. This is because, whenever the condition on line 11 evaluates to true, it must be that $x_i = \emptyset$ (again, due to single-mindedness), so the increase in value to agent i is greater than the decrease in value to each agent $j \in T$. We conclude that the social welfare of the final allocation is at least that of the allocation returned by the BN mechanism.

We now consider the general case of superadditive valuations. We will compare our auction with a corresponding auction in which each original superadditive bidder is represented by 2^m single-minded bidders, one for each subset of M . That is, for each $i \in [n]$ and $S \subseteq M$, we will have a bidder (i, S) with a single-minded value of $v_i(S)$ for set S . Note first that the optimal social welfare for this new auction is identical to that of the original auction. Also, since this is an auction over single-minded bidders, our mechanism obtains at least as much social welfare as the BN auction for this new auction instance (as argued above). We will now argue that our mechanism obtains at least as much welfare in the original auction as in the new auction instance; this will complete the proof.

The only difference in the mechanism's behavior on these two input instances is that, on line 4 or 13, we may allocate multiple times to a single agent in the original input instance, whereas in the expanded auction instance these allocations are to separate bidders. Note, however, that if we allocate multiple sets S_1, \dots, S_k to a player i in the original instance, then we allocate precisely to players $(i, S_1), \dots, (i, S_k)$ in the expanded instance. Since bidder valuations are superadditive, the value of allocating multiple disjoint sets to a single bidder cannot be less than the value of allocating these sets separately to single-minded bidders with the same valuations. We conclude that the welfare generated by our mechanism in the original auction is at least that of the BN mechanism in the corresponding auction with single-minded types, and hence obtains an $O(\sqrt{m})$ approximation to social welfare as required. \square

We next move to the value-query model. We show that the $O(m/\sqrt{\log m})$ approximation mechanism due to [25] is guaranteed to generate CWE outcomes. We note that it nearly matches the lower bound of $O(m/\log m)$ on the approximability of CAs with superadditive bidders (using value queries) [28].

Theorem 3.6. *If agents are super-additive, then there exists a mechanism that makes a polynomial number of value queries and generates a CWE outcome that achieves a $O(m/\sqrt{\log m})$ approximation to the optimal social welfare.*

Proof. The mechanism groups the objects into $\log m$ bundles, each of size $m/\log m$, arbitrarily. It then returns the bundle-efficient allocation over those bundles. This is known to achieve a $O(m/\sqrt{\log m})$ approximation, and can be implemented in a polynomial number of value queries [25]. Since the allocation is bundle-efficient, it is CWE. \square

3.2 Single-minded Valuations and Value Queries

In the special case in which agents are single-minded, Algorithm 1 can be improved to run in a polynomial number of *value* queries, obtaining an $O(\sqrt{m})$ approximation to the optimal welfare. This mechanism is new; as far as we are aware, existing $O(\sqrt{m})$ -approximation mechanisms do not satisfy MC-CWE.

Algorithm 2 SingleMindedMC-CWE

Require: Single-minded valuations $(v_1, S_1), \dots, (v_n, S_n)$.

```
// Phase 1: Players with small desired sets
1:  $Q \leftarrow \{i \in N : |S_i| \leq \sqrt{m}\}$ 
2:  $\mathbf{x} \leftarrow$  outcome of greedily allocating sets to players in  $Q$ , by value.
3: for each object  $a_i$  not allocated in  $\mathbf{x}$  do
4:   add  $a_i$  to an arbitrary non-empty element of  $\mathbf{x}$ 
5: end for

// Phase 2: Players with large desired sets
6: for  $i \in N - Q$ , in decreasing order of  $v_i$  do
7:    $C_i \leftarrow \{j : x_j \cap S_i \neq \emptyset\}$ 
8:   if  $v_i > \sum_{j \in C_i} v_j$  then  $x_i \leftarrow \cup_{j \in C_i} x_j$ ,  $x_j \leftarrow \emptyset$  for each  $j \in C_i$ ;
9: end for
10: return  $\mathbf{x}$ 
```

Our algorithm, which we call SingleMindedMC-CWE and listed as Algorithm 2, proceeds as follows. We split the bidders into two groups: those with “small” desired sets (of size at most \sqrt{m}) and those with larger desired sets. We first generate a provisional allocation that only includes the bidders with small desired sets. We construct this preliminary allocation greedily: we order players from largest value to smallest, then allocate to players in this order if their desired set is available. Any object that is left unallocated is then added to an arbitrary *non-empty* allocation. Then, in the second phase of the algorithm, we consider those bidders with large desired sets. We order these large-set bidders from highest value to smallest, and for each bidder i in this order, say with value v_i for set S_i , we ask whether v_i is greater than the sum of values of all players whose allocations intersect set S_i . If so, we take all of those intersecting allocations from their respective bidders and allocate them all to player i . After this operation has been completed for every large-set bidder (in order from highest value to smallest), we return the resulting allocation. The following theorem establishes the $O(\sqrt{m})$ approximation and polynomial run time of the algorithm.

Theorem 3.7. *When agent valuations are single-minded, Algorithm SingleMindedMC-CWE returns a MC-CWE outcome that $O(\sqrt{m})$ -approximates the optimal social welfare over all assignments. Furthermore, it can be implemented in a polynomial number of value queries.*

The runtime requirement of Theorem 3.7 is immediate from the definition of Algorithm 2. We address the approximation and MC-CWE requirements separately.

Lemma 3.8. *Algorithm 2 obtains an $O(\sqrt{m})$ approximation when agents are single-minded.*

Proof. We first note that, upon each iteration of loop 6–12, the social welfare of allocation \mathbf{x} cannot decrease. This is because, on each iteration, the allocation is altered precisely when the value of player i for set $\cup_{j \in C_i} x_j$ is greater than the value being obtained from these objects in the pre-existing allocation.

After line 2, the allocation \mathbf{x} corresponds to the outcome of the greedy allocation among bids for sets of size at most \sqrt{m} . Thus, since social welfare does not decrease from this point onward, the final allocation generates at least the welfare of this greedy outcome.

Next, suppose that i is the player for which v_i is maximized. We will show that our algorithm generates social welfare at least v_i . If $i \in Q$, then we will have $S_i \subseteq x_i$ after line 2, so our algorithm’s welfare must be at least v_i . Otherwise, i will be the first agent considered in the loop starting at line 3. If v_i is larger than the value generated by our algorithm before this first iteration, then the condition on line 8 must evaluate true and we will have $S_i \subseteq x_i$ after the first iteration of the loop. Thus, again, our algorithm’s welfare must therefore be at least v_i . In either case, our mechanism obtains welfare at least $\max_i v_i$.

We therefore have that our mechanism generates at least as much welfare as an algorithm that chooses the welfare-maximizing outcome between (a) the greedy allocation among bids for sets of size at most \sqrt{m} , and (b) allocating all items to a single bidder. However, it is known that such an allocation obtains an $O(\sqrt{m})$ approximation to the optimal social welfare [30]. Since our mechanism generates at least this much welfare, it must be an $O(\sqrt{m})$ approximation as well. \square

Lemma 3.9. *Let \mathbf{x} be the allocation returned by Algorithm 2, and consider price vector $p_i = v_i(x_i)$. Then (\mathbf{x}, \mathbf{p}) is MC-CWE.*

Proof. For a given fixed input profile, let $Q = \{i \in [n] : |S_i| \leq \sqrt{m}\}$ be the set of agents with “small” demand sets and let \mathbf{x} be allocation chosen by the mechanism. For each player i such that $x_i = \emptyset$, let $C_i = \{j : x_j \cap S_i \neq \emptyset\}$. It suffices to show that, for each i such that $x_i = \emptyset$, we have $v_i \leq \sum_{j \in C_i} v_j$. (Since $S_i \subseteq x_i$ whenever $x_i \neq \emptyset$).

Observe that, after line 5, the allocation x_j forms a partition of the set of objects. We also note that, for agent i , the value of $\sum_{j \in C_i} v_j$ can only increase over each iteration of the loop on lines 4-8. This is because, for each object a_i , the value of the player j such that $a_i \in x_j$ can only increase (since the allocation only changes when a set of objects is transferred to a player with value at least that of all previous owners of those objects).

Let i be some player such that $x_i = \emptyset$. Suppose $i \in Q$, and that i did not receive his desired set from the greedy allocation on line 2. Then, since $x_i = \emptyset$, it must be that $v_i \leq \max\{j \in Q, j \neq i, S_j \cap S_i \neq \emptyset : v_j\}$ (from the definition of the greedy algorithm). Since the values of winners only increase over the course of the mechanism, we must have $v_i \leq \max\{j, x_j \cap S_i \neq \emptyset : v_j\}$ and hence $v_i \leq \sum_{j \in C_i} v_j$.

Next suppose that $i \notin Q$, and the condition on line 6 does not evaluate to true when agent i is considered on lines 4-8. Then, at this point in the algorithm we have $v_i \leq \sum_{j \in C_i} v_j$, and since winner values only increase over the course of the mechanism this will be true at the completion of the mechanism.

Finally, suppose that i is such that, on some iteration of the loop on lines 6-12, x_i is initially non-empty but is set to \emptyset over the course of the iteration. Then the condition on line 8 evaluated true for some j such that $x_i \cap S_j \neq \emptyset$, and hence after this iteration we have $v_i \leq v_j$ and $S_i \cap x_j \neq \emptyset$. Since winner values only increase over the course of the mechanism, we must have $v_i \leq \sum_{j \in C_i} v_j$ at the completion of the mechanism, as required. \square

4 Uniform Budget-Additive Valuations


A budget-additive valuation is specified by budget B and item values $v_j, j \in M$. The value of set S is then $v(S) = \min\{B, \sum_{j \in S} v_j\}$. Note that the uniform-valuation case implies $v_i(j) \leq B_i$ for all $i \in [n], j \in M$.

The problem of maximizing social welfare with budget-additive valuations has been extensively studied in recent years from a computational perspective [1, 5, 11, 16, 35]. It is known to be APX-hard [11], and the best known approximation ratio is $4/3$ (achieved via iterative rounding or primal-dual algorithms). This factor also matches the integrality gap of the corresponding linear program. In this section we consider problem instances in which agents have *uniform* budget-additive valuations [1], in the sense that for each object j there is a value v_j such that, for all $i, v_i(j) \in \{0, v_j\}$. In other words, each object has a fixed value v_j ; each player with non-zero value for j values it at v_j . Here too, the best approximation known is $4/3$.


4.1 Arbitrary budgets

We first give an example in which no MC-CWE allocation can achieve more than a $7/8$ fraction of the optimal welfare.

Claim 4.1. *There is a profile of uniform budget-additive valuations for which no MC-CWE allocation achieves more than $7/8$ of the optimal welfare.*



(a) An instance of uniform budget additive bidders that admits a gap in welfare due to MC-CWE.



(b) An instance of uniform budget additive bidders with identical budgets that admits no Walrasian equilibrium.

Figure 2: Instances of uniform budget-additive agents. Agents and items are represented by thin and thick nodes, respectively. An edge between an agent and item nodes means that the agent values the item. Values written next to agents and items correspond to budgets and (uniform) values, respectively.

Proof. Consider the instance given in Figure 2(a) with 5 players, $\{c_i\}_{i=1..5}$ (thin nodes), and 4 items, $\{a_i\}_{i=1..4}$ (thick nodes). The value written next to an agent node (resp., item node) corresponds to agent's budget (resp., item's value). If an agent values some item, we draw an edge between the agent and item.

In this example, the optimal fractional assignment has social welfare 8 (for example, the following achieves welfare 8: c_1 gets sets $\{a_1\}$ and $\{a_2\}$ with probability $1/2$ each, c_2 gets sets $\{a_2\}$ and $\{a_3, a_4\}$ with probability $1/2$ each, and c_4 gets sets $\{a_3\}$ and $\{a_4\}$ with probability $1/2$ each). However, the optimal integral assignment obtains value $8 - \epsilon$ (for example, in one optimal assignment c_1 gets $\{a_1\}$, c_2 gets $\{a_2\}$, c_4 gets $\{a_4\}$, and c_3 gets $\{a_3\}$). Therefore, in order to get a MC-CWE outcome, one must bundle some items together. We claim that, for any bundling choice, one cannot achieve a higher welfare than 7. The argument proceeds via a case analysis:

1. If the bundle is $\{a_1, a_2\}$, then if it is allocated to c_1 its contribution is 3, and the total value of other items is 4, for a total of 7. If $\{a_1, a_2\}$ is allocated to c_2 then one of a_3 or a_4 must be allocated to a player of budget $1 - \epsilon$, and hence the total value is at most $4 + 2 + 1 - \epsilon < 7$.
2. If the bundle is $\{a_2, a_3\}$ or $\{a_2, a_4\}$, then the bundle can contribute at most 4 to the total welfare, and the total value of other items is 3, for a total of 7.
3. If the bundle is $\{a_3, a_4\}$, then if it is allocated to c_4 its contribution is 2 and the total value of other items is 5 for a total of 7. If it is allocated to c_2 then the only other player who can obtain value is c_1 with a budget of 3, so the total value is at most $4 + 3 = 7$.
4. If the bundle is $\{a_1, a_3\}$ or $\{a_1, a_4\}$, say $\{a_1, a_3\}$, then if the bundle is allocated to c_1 or c_3 then its value is at most 1, which leads to total value at most 7 since the other items have value 6. If the bundle is allocated to c_2 (for a value of 2) then a_2 can contribute value at most 3 and a_4 can contribute value at most 2, for a total of 7. Finally, if the bundle is allocated to c_4 then whichever of the remaining items is not allocated to c_2 will contribute one less than its value to the total, and hence the total value is at most $2 + 2 + 4 - 1 = 7$.

We conclude that every possible bundling and allocation results in a welfare of at most 7, and hence the gap is established. \square

We next present an algorithm that converts any allocation into a MC-CWE allocation while preserving at least half of the original welfare. The algorithm, presented as Algorithm 3, proceeds as follows. Given an allocation profile $\mathbf{x} = (x_1, \dots, x_n)$, the algorithm checks whether there is an instance in which $v_j(x_i) > v_i(x_i)$. Note that this can occur only if $\sum_{k \in x_i} v_k > B_i$ and $B_j > B_i$. If there is no such instance, then the algorithm terminates and the current allocation is returned. Otherwise, if there is a pair of agents i and j with

Algorithm 3 UniformBudgetAdditiveMC-CWE

Require: Uniform budget additive valuations v_1, \dots, v_n , agent budgets B_1, \dots, B_n , and an allocation profile $\mathbf{x} = (x_1, \dots, x_n)$ (w.l.o.g., every item is allocated to a player with positive value for it).

```
1: Initialize  $S \leftarrow \emptyset$ 
2: if  $\sum_{k \in x_i} v_i(k) \leq B_i$  for every  $i$  then
3:   return  $\mathbf{x}$ 
4: else
5:   for  $i \in [n]$ , in a non-decreasing order of  $B_i$  do
6:      $S \leftarrow \{j \in x_i : \exists i' \text{ s.t. } v_{i'}(j) > 0 \text{ and } B_{i'} > B_i\}$ 
7:     while  $\sum_{k \in x_i} v_i(k) > B_i$  and  $S \neq \emptyset$  do
8:        $j \leftarrow \operatorname{argmin}_{j \in S} \{v_j\}$ 
9:        $i' \leftarrow \operatorname{argmax}_{i' \in [n]} \{B_{i'} : v_{i'}(j) > 0\}$ 
10:       $x_i \leftarrow x_i - j$ ,  $S \leftarrow S - j$ ,  $x_{i'} \leftarrow x_{i'} \cup j$  (i.e., move  $j$  from  $x_i$  to  $x_{i'}$ )
11:    end while
12:  end for
13:  return  $\mathbf{x}$ 
14: end if
```

$v_j(x_i) > v_i(x_i)$, then the item of lowest value in x_i is removed from x_i and added to x_j . The algorithm then repeats. Note that the algorithm must terminate, since every iteration results in an item being shifted from one agent to another agent with strictly larger budget.

The intuition behind UniformBudgetAdditiveMC-CWE is that we would like to reduce the social welfare of the optimal *fractional* allocation of the bundles in \mathbf{x} . Indeed, if there are no instances in which $v_j(x_i) > v_i(x_i)$, then it must be that the assignment \mathbf{x} is an optimal *fractional* allocation, and hence by Claim 2.1 is MC-CWE. Thus, by transforming the input allocation into an allocation that satisfies this property, we are reducing the integrality gap of our reduced market to 1. What remains to show is that this iterative procedure does not drastically reduce the welfare of the allocation.

Theorem 4.2. *Suppose \mathbf{x} is an arbitrary allocation with welfare $SW(\mathbf{x})$. Given \mathbf{x} , Algorithm UniformBudgetAdditiveMC-CWE returns, in polynomial time, a MC-CWE outcome \mathbf{x}' such that $SW(\mathbf{x}') \geq \frac{1}{2}SW(\mathbf{x})$.*

Proof. First, note that we can assume without loss of generality that \mathbf{x}' is such that $j \in x'_i$ implies $v_i(j) > 0$ (else, take any object j such that $j \in x'_i$ and $v_i(j) = 0$ and re-allocate it to an arbitrary agent with non-zero value for it; this cannot reduce welfare).

We next prove that $v_i(x'_i) = \max_{\ell \in [n]} v_\ell(x'_i)$ for every $i \in [n]$. If $\sum_{j \in x'_i} v_j$ does not exceed B_i , then $v_i(x'_i)$ is clearly the largest possible value of x'_i by any agent. If $\sum_{j \in x'_i} v_j$ exceeds B_i , then $v_i(x'_i) = B_i$. Assume toward contradiction that there exists an agent i' such that $v_{i'}(x'_i) > v_i(x'_i)$. This implies that $B_{i'} > B_i$ and there exists an item $j \in x'_i$ such that $v_{i'}(j) > 0$. But if this were the case, item j would have moved to an agent with higher budget in line 10 of the algorithm; reaching a contradiction. It follows that the optimal fractional allocation of bundles \mathbf{x}' achieves the same social welfare of allocation \mathbf{x}' , and hence \mathbf{x}' is CWE.

It remains to prove that $SW(\mathbf{x}') \geq \frac{1}{2}SW(\mathbf{x})$. Since the mechanism moves items in a non-decreasing order, and $v_i(j) \leq B_i$ for every i, j , it is guaranteed that every agent i that exceeds his budget in \mathbf{x} , is left with value of at least $B_i/2$ in \mathbf{x}' . The desired approximation follows. \square

As a corollary, the known $4/3$ approximation algorithm for budget-additive valuations can be turned into an $8/3$ approximation MC-CWE mechanism in the case of uniform values. Furthermore, the allocation returned by Algorithm UniformBudgetAdditiveMC-CWE has the property that the full surplus can be extracted from the MC-CWE outcome as revenue.

Lemma 4.3. *If agents are sub-additive and \mathbf{x} is an allocation such that $v_i(x_i) = \max_{\ell \in [n]} v_\ell(x_i)$ for every $i \in [n]$, then the price vector $p_i = v_i(x_i)$ is such that (\mathbf{x}, p) is MC-CWE.*

Proof. Pick any agent i and set of agents S . Since for every i , $v_i(x_i) = \max_{\ell \in [n]} v_\ell(x_i)$, it follows that $\sum_{j \in S} v_i(\mathbf{x}_j) - \sum_{j \in S} v_j(\mathbf{x}_j) \leq 0$. We get:

$$\begin{aligned} v_i(x_i) - p_i &= 0 \geq \sum_{j \in S} v_i(\mathbf{x}_j) - \sum_{j \in S} v_j(\mathbf{x}_j) = \sum_{j \in S} v_i(\mathbf{x}_j) - \sum_{j \in S} p_j \\ &\geq v_i(\cup_{j \in S} \mathbf{x}_j) - \sum_{j \in S} p_j, \end{aligned}$$

where the last inequality follows from sub-additivity. The assertion of the lemma follows. \square

As a corollary, we get the following approximation result.

Corollary 4.4. *For every instance of uniform budget-additive valuations, there exists a poly-time mechanism that generates a MC-CWE outcome that is within factor $8/3$ of the optimal revenue.*

Proof. We have shown that one can turn the $4/3$ SW-approximation algorithm for budget-additive valuations into an $8/3$ SW-approximation CWE mechanism that has the property $v_i(x_i) = \max_{\ell \in [n]} v_\ell(x_i)$ for every $i \in [n]$. The last lemma implies that the full surplus can be extracted by a CWE mechanism. Observing that the social welfare is an upper bound on the profit establishes the assertion of the corollary. \square

4.2 Identical Budgets

In this section we study the restricted case of identical budgets. We note that the welfare maximization problem is NP-hard, even under this restriction, as it includes PARTITION as a special case³. Moreover, we shall show that there are input instances in this class for which Walrasian equilibria do not exist⁴. Nevertheless, we shall show that any allocation can be transformed into a MC-CWE allocation with no loss to the social welfare.

We first show that there exist input instances for which no Walrasian equilibrium exists. Consider the instance given in Figure 2(b) with 4 players, c_i, d_i for $i \in \{1, 2\}$, and 7 items: a_i, b_i, α_i , and β for $i \in \{1, 2\}$. Each player has budget 2. Each item has value 1, except for item β which has value 2. For $i \in \{1, 2\}$, player c_i has value for objects a_i, b_i , and α_i , and player d_i has value for objects a_i, b_i , and β . In this example, the optimal fractional assignment has social welfare 8. For example, the fractional assignment in which c_i gets sets $\{a_i, \alpha_i\}$ and $\{b_i, \alpha_i\}$ with probability $1/2$, and d_i gets sets $\{a_i, b_i\}$ and $\{\beta\}$ with probability $1/2$, for $i \in \{1, 2\}$, achieves welfare 8. However, the optimal integral assignment obtains value at most 7. Thus the optimal fractional welfare does not occur at an integral solution, and hence a WE does not exist.

We now show that, for any allocation, there exists a MC-CWE allocation that obtains at least as much social welfare, and moreover this MC-CWE allocation can be found efficiently.

Theorem 4.5. *Suppose \mathbf{x} is an arbitrary allocation. There exists MC-CWE allocation \mathbf{x}' with $SW(\mathbf{x}') \geq SW(\mathbf{x})$, and \mathbf{x}' can be found in polynomial time given \mathbf{x} .*

Proof. Given \mathbf{x} , we construct \mathbf{x}' by taking any object j such that $j \in x_i$ with $v_i(j) = 0$ and re-allocating it to an arbitrary agent with non-zero value for it. This can be done in polytime and can only increase the social welfare of the resulting allocation. It remains to show that this allocation is MC-CWE.

³For a given instance of PARTITION with integers a_1, \dots, a_n such that $\sum_{j \in [n]} a_j = 2B$, construct an instance of our problem with two agents, each with budget B , and n items, with item j having value a_j .

⁴While this is well known for the case of arbitrary budget additive valuations, here we consider the further restricted class of uniform values and identical budgets.

Suppose \mathbf{x}' is such that $j \in x'_i$ implies $v_i(j) > 0$. In this case, $v_i(x'_i) = \min\{B, \sum_{j \in x'_i} v_j\} = \max_{\ell} v_{\ell}(x'_i)$ for each i . The social welfare of allocation \mathbf{x}' is therefore $\sum_i \max_{\ell} v_{\ell}(x'_i)$, which is an upper bound on the value of any fractional assignment of bundles x'_1, \dots, x'_n . Thus the optimal fractional allocation of these bundles obtains the same welfare as \mathbf{x}' , and hence \mathbf{x}' is MC-CWE as required. \square

As a corollary, the known $4/3$ approximation algorithms for budget-additive valuations can be made MC-CWE for the case of uniform values and identical budgets. Note that the complementary case, in which the budgets are identical but the item values are non-uniform, also has instances where no MC-CWE allocation can achieve the welfare of an optimal allocation. Consider, for example, an instance with 3 players, c_1, c_2, c_3 , each with budget 2, and 3 items, a_1, a_2, a_3 , where $v_{c_1}(a_1) = 2, v_{c_1}(a_2) = v_{c_1}(a_3) = 1, v_{c_2}(a_2) = v_{c_2}(a_3) = 2$, and $v_{c_3}(a_1) = 4\epsilon, v_{c_3}(a_3) = \epsilon$ (and all other values are zero). One can easily verify that there exists a fractional solution that achieves welfare $4 + 2\epsilon$, while the highest welfare that can be obtained by any integral solution is $4 + \epsilon$. However, by bundling together any two items, the highest achievable welfare is reduced to at most 4.

References

- [1] N. Andelman and Y. Mansour. Auctions with budget constraints. In *SWAT*, pages 26–38, 2004.
- [2] K. J. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22:265 – 290, 1954.
- [3] R. J. Aumann. Values of markets with a continuum of traders. *Econometrica*, 43(4):611–46, July 1975.
- [4] L. M. Ausubel and P. R. Milgrom. Ascending auctions with package bidding. *The B.E. Journal of Theoretical Economics*, 0(1):1, 2002.
- [5] Y. Azar, B. Birnbaum, A. R. Karlin, C. Mathieu, and C. T. Nguyen. Improved approximation algorithms for budgeted allocations. In *Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I, ICALP '08*, pages 186–197, Berlin, Heidelberg, 2008. Springer-Verlag.
- [6] A. Badanidiyuru, S. Dobzinski, and S. Oren. Optimization with demand oracles. In *Proceedings of the 13th ACM conference on Electronic commerce, EC '12*, 2012.
- [7] M.-F. Balcan, A. Blum, and Y. Mansour. Item pricing for revenue maximization. In *Proceedings of the 9th ACM conference on Electronic commerce, EC '08*, pages 50–59, New York, NY, USA, 2008. ACM.
- [8] S. Bikhchandani and J. W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74(2):385–413, June 1997.
- [9] S. Bikhchandani and J. M. Ostroy. The package assignment model. *Journal of Economic Theory*, 107(2):377–406, December 2002.
- [10] L. Blumrosen and N. Nisan. On the computational power of demand queries. *SIAM J. Comput.*, 39(4):1372–1391, Oct. 2009.
- [11] D. Chakrabarty and G. Goel. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and gap. In *Foundations of Computer Science, 2008. FOCS '08. IEEE 49th Annual IEEE Symposium on*, pages 687 –696, oct. 2008.
- [12] M. Cheung and C. Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science, FOCS '08*, pages 35–44, Washington, DC, USA, 2008. IEEE Computer Society.

- [13] E. Cohen, M. Feldman, A. Fiat, H. Kaplan, and S. Olonetsky. Envy-free makespan approximation: extended abstract. In *ACM Conference on Electronic Commerce*, pages 159–166, 2010.
- [14] E. Cohen, M. Feldman, A. Fiat, H. Kaplan, and S. Olonetsky. Truth, envy, and truthful market clearing bundle pricing. In *WINE*, pages 97–108, 2011.
- [15] S. Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In *Proceedings of the 43rd annual ACM symposium on Theory of computing*, STOC '11, pages 139–148, New York, NY, USA, 2011. ACM.
- [16] U. Feige and J. Vondrak. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. In *Foundations of Computer Science, 2006. FOCS '06. 47th Annual IEEE Symposium on*, pages 667–676, oct. 2006.
- [17] M. Feldman, N. Gravin, and B. Lucier. Combinatorial walrasian equilibrium. In *STOC*, pages 61–70, 2013.
- [18] M. Feldman, N. Gravin, and B. Lucier. Combinatorial walrasian equilibrium. In *STOC*, 2013.
- [19] A. Fiat and A. Wingarten. Envy, multi envy, and revenue maximization. In *Proceedings of the 5th International Workshop on Internet and Network Economics*, WINE '09, pages 498–504, Berlin, Heidelberg, 2009. Springer-Verlag.
- [20] L. Fleischer and Z. Wang. Lower bound for envy-free and truthful makespan approximation on related machines. In *Proceedings of the 4th international conference on Algorithmic game theory*, SAGT'11, pages 166–177, Berlin, Heidelberg, 2011. Springer-Verlag.
- [21] D. Foley. Resource allocation and the public sector. 1967.
- [22] H. Fu, R. Kleinberg, and R. Lavi. Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In *ACM Conference on Electronic Commerce*, page 586, 2012.
- [23] F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *J. of Economic Theory*, 87(1):95–124, 1999.
- [24] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '05, pages 1164–1173, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.
- [25] R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47(1):104 – 123, 2004.
- [26] J. Kelso, Alexander S and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, November 1982.
- [27] H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–79, June 1983.
- [28] V. Mirrokni, M. Schapira, and J. Vondrak. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *Proceedings of the 9th ACM conference on Electronic commerce*, EC '08, pages 70–77, New York, NY, USA, 2008. ACM.

- [29] A. Mu’Alem. On multi-dimensional envy-free mechanisms. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory*, ADT ’09, pages 120–131, Berlin, Heidelberg, 2009. Springer-Verlag.
- [30] A. Mu’alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64(2):612–631, 2008.
- [31] N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, July 2006.
- [32] S. Pápai. Groves sealed bid auctions of heterogeneous objects with fair prices. *Social Choice and Welfare*, 20(3):371–385, 2003.
- [33] D. C. Parkes and L. H. Ungar. Iterative combinatorial auctions: Theory and practice. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, pages 74–81. AAAI Press, 2000.
- [34] L. S. Shapley and M. Shubik. The assignment game i: The core. *International Journal of Game Theory*, 1:111–130, 1971. 10.1007/BF01753437.
- [35] A. Srinivasan. Budgeted allocations in the full-information setting. In *Proceedings of the 11th international workshop, APPROX 2008, and 12th international workshop, RANDOM 2008 on Approximation, Randomization and Combinatorial Optimization: Algorithms and Techniques*, APPROX ’08 / RANDOM ’08, pages 247–253, Berlin, Heidelberg, 2008. Springer-Verlag.
- [36] H. R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, September 1974.
- [37] P. R. Wurman and M. P. Wellman. Akba: a progressive, anonymous-price combinatorial auction. In *Proceedings of the 2nd ACM conference on Electronic commerce*, EC ’00, pages 21–29, New York, NY, USA, 2000. ACM.

A Relation to envy-free (EF) notions

The MC-CWE notion is closely related to several notions of *envy freeness* that have been extensively studied in recent years. In what follows, we shall describe the difference and similarities between MC-CWE and these notions.

Envy free 1 (EF1) (as in, e.g., [13, 14, 20, 29, 32]).

Let $\mathbf{x} = (x_1, \dots, x_n)$ be an allocation, and let $\mathbf{p} = (p_1, \dots, p_n)$ be prices, such that p_i is the price of x_i . An outcome (\mathbf{x}, \mathbf{p}) is said to be EF1 if no agent wishes to switch her outcome with another; i.e., for every $i, j \in [n]$, $v_i(x_i) - p(x_i) \geq v_i(x_j) - p(x_j)$.

An obvious way to view the difference between MC-CWE and EF1 is that EF1 requires that no agent envies a single other agent, while MC-CWE precludes envy in any set of other agents. However, a more subtle examination reveals that any EF1 outcome can be extended to an outcome that exhibits the latter (stronger) requirement, but only if one allows for arbitrary bundle prices. This is in stark contrast to the linear prices required by MC-CWE. Moreover, MC-CWE, unlike EF1, requires market clearing.

Envy free 2 (EF2) (as in, e.g., [7, 12, 24]).

Let $\mathbf{x} = (x_1, \dots, x_n)$ be an allocation, and let $\mathbf{p} = (p_1, \dots, p_m)$ be item prices, assigning a price p_j for every $j \in M$. An outcome (\mathbf{x}, \mathbf{p}) is said to be EF2 if $x_i \in D_i(\mathbf{p})$ for every $i \in [n]$. This notion entails the first requirement of a Walrasian equilibrium, but does not require market clearing.

MC-CWE is, on the one hand, a relaxation of this notion, as it allows for a pre-sale bundling phase. On the other hand, it requires market clearing, whereas EF2 does not.

B Revenue enhancement

We now point out that revenue can increase substantially when we optimize revenue with respect to equilibria with bundles, as opposed to the individual item prices of standard Walrasian equilibrium. For example, consider a case with two bidders and three objects $\{a_1, a_2, a_3\}$. Player 1's valuation is given by $v_1(S) = 1$ if $a_1 \in S$, otherwise $v_1(S) = 0$. Choose an arbitrarily large constant R . Then player 2's valuation is additive up to a capacity of 2 objects, with values $R - 1, R, R$ for items a_1, a_2, a_3 respectively. In this example, both players satisfy gross substitutes, and the efficient allocation is $(\{a_1\}, \{a_2, a_3\})$. This allocation can be supported by a Walrasian equilibrium, and the maximum supporting prices are $p_1 = 1, p_2 = 2, p_3 = 2$, for a total revenue of 5. However, the revenue-maximizing MC-CWE prices set the price of $\{a_1\}$ to 1 and the price of $\{a_2, a_3\}$ to $R + 2$, for a total revenue of $R + 3$. This gap can be made arbitrarily large by increasing R .

C Necessity of Bundling

We show by way of example that the use of bundling is necessary for the statement of Theorem 3.1, even if we relax the market-clearing requirement of Walrasian equilibrium and even for single-minded bidders. We present an auction instance such that, for any set of item prices for which agent demand sets are disjoint, the social welfare of the resulting allocation is an $O(\sqrt{m})$ fraction of the optimal social welfare.

Example 1. Consider a combinatorial auction with $n = \sqrt{m} + 1$ single-minded bidders, as follows. One bidder desires set M for a value of m . Each of the remaining bidders i desires a set S_i of size \sqrt{m} , for a value of $1 + 2\sqrt{m}$. These sets have the property that $|S_i \cap S_j| = 1$ for all $i \neq j$, and for every object $a \in M$, a is contained in at most 2 sets; the existence of such sets follows easily by induction on their size⁵.

The optimal outcome allocates all objects to the first bidder, for a total value of m . Suppose this optimal outcome can be supported by item prices \mathbf{p} . Then $\sum_{a \in M} p_a \leq m$, but a counting argument demonstrates that for some set S_i , it must be that $\sum_{a \in S_i} p_a \leq 2\sqrt{m} < 1 + 2\sqrt{m}$. Player i would demand set S_i at these prices, a contradiction. Thus, at any pricing equilibrium, some set S_i must be allocated to a buyer i ; but this implies that no other set S_j can be allocated, as $S_j \cap S_i \neq \emptyset$ for all $j \neq i$. Thus the total social welfare at any equilibrium outcome with item prices is at most $1 + 2\sqrt{m}$.

⁵That is, that there exist t sets of size t with this property, for all $t \leq \sqrt{m}$.