

Approximately Efficient Two-Sided Combinatorial Auctions

Riccardo Colini-Baldeschi, Paul Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden,
and Stefano Turchetta

¹ LUISS Rome, rcolini@luiss.it

² University of Oxford, paul.goldberg@cs.ox.ac.uk, stefano.turchetta@cs.ox.ac.uk

³ Centrum Wiskunde & Informatica (CWI), Amsterdam, keijzer@cwi.nl

⁴ Sapienza University of Rome, leonardi@diag.uniroma1.it

⁵ Stanford University, tim@cs.stanford.edu

Abstract. Mechanism design for one-sided markets has been investigated for several decades in economics and in computer science. More recently, there has been an increased attention on mechanisms for two-sided markets, in which buyers and sellers act strategically. For two-sided markets, an impossibility result of Myerson and Satterthwaite states that no mechanism can simultaneously satisfy *individual rationality (IR)*, *incentive compatibility (IC)*, *strong budget-balance (SBB)*, and be efficient. Follow-up work mostly focused on designing mechanisms that trade off among these properties, in settings where there exists one single type of items for sale, the buyers ask for one unit of the items, and the sellers initially hold one unit of the item.

On the other hand, important applications to web advertisement, stock exchange, and frequency spectrum allocation, require us to consider two-sided combinatorial auctions in which buyers have preferences on subsets of items, and sellers may offer multiple heterogeneous items. No efficient mechanism was known so far for such two-sided combinatorial markets. This work provides the first IR, IC and SBB mechanisms that provides an $O(1)$ -approximation to the optimal social welfare for two-sided markets. An initial construction yields such a mechanism, but exposes a conceptual problem in the traditional SBB notion. This leads us to define the stronger notion of *direct trade strong budget balance (DSBB)*. We then proceed to design mechanisms that are IR, IC, DSBB, and again provide an $O(1)$ -approximation to the optimal social welfare. Our mechanisms work for any number of buyers with XOS valuations – a class in between submodular and subadditive functions – and any number of sellers. We provide a mechanism that is dominant strategy incentive compatible (DSIC) if the sellers each have one item for sale, and one that is bayesian incentive compatible (BIC) if sellers hold multiple items and have additive valuations over them. Finally, we present a DSIC mechanism for the case that the valuation functions of all buyers and sellers are additive. We prove our main results by showing that there exists a variant of a sequential posted price mechanism, generalised to two-sided combinatorial markets, that achieves the desired goals.

1 Introduction

One-sided markets have been studied in economics for several decades and more recently in computer science. Mechanism Design in one-sided markets aims to find an efficient (high-welfare) allocation of a set of items to a set of agents, while ensuring that truthfully reporting the input data is the best strategy for the agents. The cornerstone method in mechanism design is the Vickrey-Clarke-Groves (VCG) mechanism [22, 4, 12] that optimizes the social welfare of the agents while providing the right incentives for truth-telling: VCG mechanisms are *dominant strategy incentive compatible (DSIC)*, and in many mechanism design settings, VCG is also *individually rational (IR)*. IR requires that participating in the mechanism is beneficial to each agent. DSIC requires that truthfully reporting one’s preferences to the mechanism is a dominant strategy for each agent, independently of what the other agents report.

Recently, increased attention has been on the problems that arise in two-sided markets, in which the set of agents is partitioned into *buyers* and *sellers*. As opposed to the one-sided setting (where one could say that the mechanism itself initially holds the items), in the two-sided setting the items are initially held by the set of sellers, who express valuations over the items they hold, and who are assumed to act rationally and strategically. The mechanism’s task is now to decide which buyers and sellers should trade, and at which prices. There is a growing interest in two-sided markets that can be attributed to various important applications. Examples range from selling display-ads on ad exchange platforms, the US FCC spectrum license reallocation, and stock exchanges. Two-sided markets are usually studied in a Bayesian setting: there is public knowledge of probability distributions, one for each buyer and one for each seller, from which the valuations of the buyers and sellers are drawn.

In two-sided markets, a further important requirement is *strong budget-balance (SBB)*, which states that monetary transfers happen only among the agents in the market, i.e., the buyers and the sellers are allowed to trade without leaving to the mechanism any share of the payments, and without the mechanism adding money into the market. A weaker version of SBB often considered in the literature is *weak budget-balance (WBB)*, which only requires the mechanism not to inject money into the market.

Unfortunately, Myerson and Satterthwaite [16] proved that it is impossible for an IR, Bayesian incentive compatible (BIC),⁶ and WBB mechanism to maximize social welfare in such a market, even in the *bilateral trade* setting, i.e., when there is just one seller and one buyer.⁷

Despite the numerous above mentioned practical contexts that need the application of *combinatorial* two-sided market mechanisms, we are not aware of any mechanism that approximates the social welfare while meeting the IR, IC and SBB requirements. The purpose of this paper is to provide mechanisms that satisfy these requirements and achieve an $O(1)$ -approximation to the social welfare for a broad class of agents’ valuation functions. We do, in fact, design mechanisms that work under the assumption of the valuations being *fractionally subadditive (XOS)*, a generalisation of submodular functions that are contained in the class of subadditive functions.

Our work builds upon previous work on an important special case of a two-sided market that is the one in which each seller holds a single item, items are identical, and each agent is only interested in holding a single item. In this setting, the valuations of the agents are thus given by a single number, representing the agent’s appreciation for holding an item. A mechanism for this

⁶ Bayesian incentive compatibility is a form of incentive compatibility that is less restrictive than DSIC. It only requires that reporting truthfully is in expectation the best strategy for a player when everyone else also does so, i.e., truthful reporting is a Bayes-Nash equilibrium.

⁷ The VCG mechanism can also be applied to two-sided markets; however, in this setting, VCG is either not IR or it does not satisfy WBB.

setting is known in the literature as a *double auction*. The goal of several works on double auctions [14, 19, 20] has been that of trading off the achievable social welfare with the strength of the incentive compatibility and budget balance constraints. A recent work addressed the problem of approximating social welfare in double auctions and related problems under the WBB requirement. Dütting, Talgam-Cohen, and Roughgarden [9] indeed proposed a greedy strategy that combines the one-sided VCG mechanism, independently applied to buyers and to sellers with the trade-reduction mechanism of McAfee [14]. They obtain IR, DSIC, WBB mechanisms with a good approximation of the social welfare, for knapsack, matching and matroid allocation constraints. More recently, Colini-Baldeschi et al. [5] presented the first double auction that satisfies IR, DSIC, and SBB, and approximates the optimal (expected) social welfare up to a constant factor. These results hold for any number of buyers and sellers with arbitrary, independent distributions on valuations. The mechanisms are also extended to the setting where there is an additional matroid constraint on the set of buyers who can purchase an item.

1.1 The Model

As stated above, the set of agents is partitioned into a set of *sellers*, each of which is initially endowed with a set of heterogeneous items, and a set of *buyers*, having no items initially. Buyers have money that can be used to pay for items. Every agent has its own, private valuation function, which maps subsets of the items to numbers, and agents are assumed to optimize their (quasi-linear) utility, which is given by the valuation of the set of items that the mechanism allocates to the agents, minus the payment that the mechanism asks from the agents. A seller will typically receive money instead of pay money, which we model by a negative payment.

For each agent we are given a (publicly known) probability distribution over a set of valuation functions, from which we assume her valuation function is drawn. The mechanism and the other agents have no knowledge of the actual valuation function of the i -th agent, but only of her probability distribution. The general aim of the mechanism is to reallocate the items so as to maximize the expected social welfare (the sum of the agents' valuations of the resulting allocation).

Let OPT be the expected social welfare of an optimal allocation of the items. Note that this is a well-defined quantity, even though computing an optimal allocation may be computationally hard, and even though there might not exist an appropriate mechanism that is guaranteed to always output an optimal allocation.

We are interested in mechanisms that satisfy *individual rationality (IR)*, *dominant strategy incentive compatibility (DSIC)* (or failing that, the weaker notion of *Bayesian incentive compatibility (BIC)*), and *strong budget-balance (SBB)*, and that reallocate the items in such a way that the expected social welfare is within some constant fraction of OPT , where expectation is taken over the given probability distributions of the agents' valuations, and over the randomness of the allocation that the mechanism outputs. The obstacle of interest to us is not the computational one, but the requirement of achieving mechanisms with such approximation guarantees, that are also DSIC, IR, and SBB. The main focus of this paper is the existence of mechanisms that obtain a constant approximation to the optimal social welfare and not the computational issues related to them. We remark, however, that our mechanisms can be implemented in polynomial time at the cost of an additional welfare loss of a factor c , where c is the best-known approximation factor for the problem of optimising buyers' social welfare, if both a poly-time approximation algorithm and an approximation for the query oracle exist (see e.g., [10]).

1.2 Overview of the Results

The present paper starts off by showing that there is a straightforward trick that one may apply to turn any WBB mechanism into an SBB one, with a small loss in approximation factor. The technique is to pre-select one random agent that is taken out of the market a priori whose role is to receive all leftover money, and it can be applied in order to obtain from [1] a $O(1)$ -approximate DSIC, IR, and SBB mechanism for a very broad class of markets. This therefore demonstrates a weakness in the current notion of SBB, which motivates the introduction of the strengthened notion *direct trade strong budget balance (DSBB)*, that requires that a monetary transfer between two agents in the market is only possible in case the agents trade items.

The goal of our proposed mechanisms is twofold: (i) achieve a constant approximation to the optimal social welfare, and (ii) design mechanisms that respect the stronger notion of DSBB. Note that with non-unit-supply sellers, a constant approximation is not known even in the context of WBB or standard SBB.⁸ We present the following mechanisms:

- A 6-approximate DSIC mechanism for buyers with XOS-valuations and sellers with one item at their disposal (i.e., *unit-supply sellers*);
- a 6-approximate BIC mechanism for buyers with XOS-valuations and non-unit supply sellers with additive valuations;
- a 6-approximate DSIC mechanism for buyers with additive valuations and sellers with additive valuations.

XOS functions lie in between the class of monotone submodular functions and subadditive functions in terms of generality. Additive and XOS functions are frequently used in the mechanism design literature. To our knowledge, these are the first mechanisms for these two-sided market settings that are simultaneously incentive compatible, (D)SBB, IR, and approximate the optimal social welfare to within a constant factor.

A first ingredient needed to obtain our results is the extension of *two-sided sequential posted price mechanisms (SPMs)* [5] for double auctions to two-sided markets. SPMs are a particularly elegant and well-studied class of mechanisms for one-sided markets. All the presented SPMs do not require any assumption on the arrival order of the agents. A second ingredient of our result is to use the *expected marginal contribution of an item to the social welfare* as the price of the item in a sequential posted price mechanism for buyers with XOS valuations [10].

1.3 Related Work

Due to the impossibility result of [16], no two-sided mechanism can simultaneously satisfy BIC, IR, WBB and be socially efficient, even in the simple bilateral trade setting. Follow-up work thus had to focus on designing mechanisms that trade off among these properties.

Some paper of the Economics literature studied the convergence rate to social efficiency as a function of the number of agents when all sellers' and buyers' valuations are independently respectively drawn from identical regular distributions, while satisfying IR and WBB. Gresik and Satterthwaite [11] showed that duplicating the number of agents by τ results in a market where the optimal IR, IC, WBB mechanism's inefficiency goes down as a function of $O(\log \tau / \tau^2)$. [17, 20] investigated a family of non-truthful double auction mechanisms, parameterized by a value $c \in [0, 1]$. We remark that the results mentioned above only hold for unit-demand buyers and unit-supply sellers, identical valuation distributions, and the hidden constants in these asymptotic result

⁸ The mechanism proposed in [1] achieves a constant approximation if the size of the initial endowment of each agent is bounded by a constant. Otherwise the mechanism achieves a logarithmic approximation to the optimal social welfare.

depend on the particular distribution. In contrast, our interest is in finding *universal* constant approximation guarantees for combinatorial settings.

In McAfee [14], an IC, WBB, IR double auction is proposed that extracts at least a $(1 - 1/\ell)$ fraction of the maximum social welfare, where ℓ is the number of traders in the optimal solution.

Optimal revenue-maximizing Bayesian auctions were characterized in [15], which provides an elegant tool applicable to single-parameter, one-sided auctions. Various subsequent articles dealt with extending these results. Related to our work is [6], which studied maximizing the auctioneer’s revenue in Bayesian double auctions. The same objective was studied in [7] yet in the *prior-free* model. Recently, Dütting et al. [9] provided black-box reductions from WBB double auctions to one-sided mechanisms. They are for a prior-free setting and can be applied with matroid, knapsack, and matching feasibility constraints on the allocations. More recently, [5] presented the first IR, IC and SBB mechanism for double auctions that $O(1)$ -approximates the optimal social welfare. In [21], mechanisms for some special cases of two-sided markets are presented that work by a combination of random sampling and random serial dictatorship. The mechanism is IR, SBB and DSIC and its *gain from trade* approaches the optimum when the market is sufficiently large. Mechanisms that are IC, IR, and SBB have been given for bilateral trade in [1]. In addition to this, the authors proposed a WBB mechanism for a general class of combinatorial exchange markets. We will use this result to construct our initial mechanism.

Sequential posted price mechanisms (SPMs) in one-sided markets have been introduced in [18] and have since gained attention due to their simplicity, robustness to collusion, and their easy implementability in practical applications. One of the first theoretical results concerning SPMs is an asymptotic comparison among three different types of single-parameter mechanisms [2]. They were later studied for the objective of revenue maximisation in [3]. Additionally, [13, 8] strengthen these results further. Very relevant to our work is the paper of Feldman et al. [10] showing that sequential posted price mechanisms can approximate social welfare up to a constant factor of 1/2 for XOS valuation functions if the published price for an item is equal to the expected additive contribution of the item to the social welfare.

2 Preliminaries

As a general convention, we use boldface notation for vectors and use $[a]$ to denote the set $\{1, \dots, a\}$. We will use $\mathbb{I}(X)$ to denote the indicator function that maps to 1 if and only if event/fact X holds.

Markets. A *two-sided market* comprises a set of two distinct types of agents: the *sellers*, who initially hold items for sale, and the *buyers*, who are interested in buying the sellers’ items. All agents possess a monotone and normalized valuation function, mapping subsets of items to $\mathbb{R}_{\geq 0}$.⁹ Formally, we represent a two-sided market as a tuple $(n, m, k, \mathbf{I}, \mathbf{G}, \mathbf{F})$, where $[n]$ denotes the set of buyers, $[m]$ denotes the set of sellers, $[k]$ denotes the set of all items for sale, $\mathbf{I} := (I_1, \dots, I_m)$ is a vector of (mutually disjoint) sets of items initially held by each seller, called the *initial endowment*, where it holds that $\bigcup_{j=1}^m I_j = [k]$. Vectors $\mathbf{G} = (G_1, \dots, G_n)$ and $\mathbf{F} = (F_1, \dots, F_m)$ are vectors of probability distributions, from which the buyers’ and sellers’ valuation functions are assumed to be drawn. A (*combinatorial*) *exchange market* is a more general version of the above defined two-sided market where an agent can act as both a buyer and a seller. Thus, everyone may initially own items and may both sell and buy items. As a result, in this setting, we override the notation and simply use n to denote the total number of agents. Formally, an exchange market is thus a tuple $(n, k, \mathbf{I}, \mathbf{F})$.

⁹ By *monotone* we mean that getting more items cannot decrease an agent’s overall valuation, and with normalized we mean that the empty set is mapped to 0.

In two-sided markets sellers are assumed to only value items in their initial bundle and are therefore not interested in buying from other sellers, i.e., $\forall j \in [m]$ and $\forall S \subseteq [k]$, $w_j(S) = w_j(S \cap I_j)$. On the contrary, in exchange markets, no such restriction on the valuation functions exists.

Throughout the paper, we reserve the usage of the letter i to denote a single buyer, the letter j to denote a single seller, and the letter ℓ to denote a single item. Moreover, we use v_i to denote buyer i 's valuation function and w_j to denote seller j 's valuation function.

Mechanism Design Goals. In order to avoid overloading this section, the following content will only be described for two-sided markets (which are the main focus of the paper). Nonetheless, any of the subsequent concepts can naturally be extended to combinatorial exchange markets.

Given a two-sided market, our aim is to redistribute the items among the agents so as to maximize the *social welfare* (the sum of the agents' valuations). An *allocation* for a given two-sided market $(n, m, k, \mathbf{I}, \mathbf{G}, \mathbf{F})$ is a pair of vectors $(\mathbf{X}, \mathbf{Y}) = ((X_1, \dots, X_n), (Y_1, \dots, Y_m))$ such that the union of $X_1, \dots, X_n, Y_1, \dots, Y_m$ is $[k]$, and $X_1, \dots, X_n, Y_1, \dots, Y_m$ are mutually non-intersecting.

Redistribution of the items is done by running a *mechanism* \mathbb{M} . A mechanism interacts with and receives input from the buyers, and outputs an *outcome*, consisting of an allocation (\mathbf{X}, \mathbf{Y}) and a payment vector $(\boldsymbol{\rho}^B, \boldsymbol{\rho}^S) \in \mathbb{R}^n \times \mathbb{R}^m$, where $\boldsymbol{\rho}^B$ refers to the buyers' vector of payments and $\boldsymbol{\rho}^S$ to the sellers' one. The outcome of a mechanism \mathbb{M} is represented by a tuple $(\mathbf{X}, \mathbf{Y}, \boldsymbol{\rho}^B, \boldsymbol{\rho}^S)$.

Agents are assumed to maximize their *utility*, which is defined as the valuation for the bundle of items that they get allocated, minus the payment charged by the mechanism. In particular, the utility $u_i^B(v_i, (\mathbf{X}, \mathbf{Y}, \boldsymbol{\rho}^B, \boldsymbol{\rho}^S))$ of a buyer $i \in [n]$ is $v_i(X_i) - \rho_i^B$, whereas for a seller $j \in [m]$ it is $u_j^S(w_j, (\mathbf{X}, \mathbf{Y}, \boldsymbol{\rho}^B, \boldsymbol{\rho}^S)) = w_j(Y_j) - \rho_j^S$.

Since agents are assumed to maximize their utility, they will strategically interact with the mechanism so as to achieve this. Our goal is to design mechanism such that there is a dominant strategy or Bayes-Nash equilibrium for the agents under which the mechanism returns an allocation with a high social welfare. For a allocation (\mathbf{X}, \mathbf{Y}) , the social welfare $\text{SW}(\mathbf{X}, \mathbf{Y})$ is defined as $\text{SW}(\mathbf{X}, \mathbf{Y}) = \sum_{i \in [n]} v_i(X_i) + \sum_{j \in [m]} w_j(Y_j)$.

The objective function we want to maximize is the above defined social welfare. We now describe three main constraints our mechanisms must adhere to. For each of these constraints we first introduce the strictest version and then a more relaxed one. Our mechanisms aim to satisfy the strictest versions, whenever possible.

- *Dominant strategy incentive compatibility (DSIC)*: It is a dominant strategy for every agent to report her true valuation sincerely, no matter what the others do, i.e., no agent can increase her utility by misreporting her true valuation.
- *Bayesian incentive compatibility (BIC)*: No agent can obtain a gain in her expected utility by declaring a valuation different from her true one, where the expectation is taken w.r.t. the others' valuations.¹⁰
- *Ex-post individual rationality (ex-post IR)*: It is not harmful for any agent to participate in the mechanism, i.e., there is guaranteed to be a strategy for an agent that yields the agent an increase in utility.
- *Interim individual rationality (interim IR)*: There is a strategy for each agent that yields her an expected increase in utility (where expectation is over the valuations of the other agents and the internal randomness of the mechanism).

¹⁰ Technically, as can be inferred, the DSIC properties are reserved for *direct revelation mechanisms*, i.e., where the buyer solely interacts with the mechanism by reporting his valuation function. It is well-known that mechanisms admitting a dominant strategy can be transformed into DSIC direct revelation mechanisms, and those with a Bayes-Nash equilibrium can be transformed into BIC direct revelation mechanisms. This way, we extend the DSIC and BIC definitions to non-direct revelation mechanisms.

- *Strong Budget Balance (SBB)*: The sum of all agents’ payments output by the mechanism is equal to zero. Conceptually, this means that no money is burnt or ends up at an external party, and no external party needs to subsidise the mechanism.
- *Weak Budget Balance (WBB)*: The sum of all payments is at least zero. In two sided-markets this translates in having the buyers’ payments being at least as large as the payments received by the sellers.

For valuation profiles (\mathbf{v}, \mathbf{w}) , $\text{OPT}(\mathbf{v}, \mathbf{w}) := \max_{\mathbf{X}, \mathbf{Y}} \{\text{SW}(\mathbf{X}, \mathbf{Y})\}$ denotes the *optimal social welfare*, where the max goes over all feasible allocations \mathbf{X}, \mathbf{Y} . The *expected optimal social welfare* is the value $\text{OPT} = \mathbb{E}_{\mathbf{v}, \mathbf{w}}[\text{OPT}(\mathbf{v}, \mathbf{w})]$. We say that a mechanism \mathbb{M} α -approximates the optimal social welfare for some $\alpha > 1$ iff $\text{OPT} \leq \alpha \mathbb{E}_{\mathbf{v}, \mathbf{w}}[\text{SW}(\mathbb{M}(\mathbf{v}, \mathbf{w}))]$. Our goal is to find mechanisms that α -approximates the optimal social welfare for a low α , are DSIC or BIC, SBB, and ex-post IR or interim IR.

Classes of Valuation Functions. We will consider probability distributions over the following classes of valuation functions. Let $v : 2^{[k]} \rightarrow \mathbb{R}_{\geq 0}$ be a valuation function. Then,

- v is *additive* iff there exists numbers $\alpha_1, \dots, \alpha_k \in \mathbb{R}_{\geq 0}$ such that $v(S) = \sum_{j \in S} \alpha_j$ for all $S \subseteq [k]$.
- v is *fractionally subadditive (or XOS)* iff there exists a collection of additive functions a_1, \dots, a_d such that for every bundle $S \subseteq [k]$ it holds that $v(S) = \max_{i \in [d]} a_i(S)$.

It is easy to see that every additive function is a XOS function. Further, it is well-known that the class of submodular functions are contained in the class of XOS functions. Due to space constraints, all proofs in the remainder of this paper have been deferred to the appendix.

3 An Initial Mechanism and Direct Trade Strong Budget Balance

In [1], Blumrosen and Dobzinski present a mechanism for exchange markets with subadditive valuation functions. They prove the following for this mechanism, which we name \mathbb{M}_{bd} .

Theorem 3.1 (Blumrosen and Dobzinski [1]). *Mechanism \mathbb{M}_{bd} is a DSIC, WBB, ex-post IR randomized direct revelation mechanism that $4H(s)$ -approximates the optimal social welfare for combinatorial exchange markets $(n, k, \mathbf{I}, \mathbf{F})$ with subadditive valuation functions, where $s = \min\{n, |I_i| : i \in [n]\}$ is the minimum of the number of agents and the number of items in an agent’s initial endowment, and $H(\cdot)$ denotes the harmonic numbers.*

In particular, this mechanism gives us a constant approximation factor if the number of starting items of the agents is bounded by a constant.

Now consider a mechanism \mathcal{M}_{sbb} that selects an agent $i \in [n]$ uniformly at random, runs \mathcal{M}_{bd} on the remaining agents, and allocates the surplus money of \mathbb{M}_{bd} to agent i . We then are able to proof the following.

Theorem 3.2. *Mechanism \mathbb{M}_{sbb} is DSIC, ex-post IR, SBB, and achieves an $8nH(s)/(n-1)$ -approximation to the optimal social welfare for exchange markets with subadditive valuations and at least 3 agents.¹¹*

The proof and a more precise description of \mathbb{M}_{sbb} are provided in Appendix A. This yields an ex-post IR, SBB, DSIC mechanism that $O(1)$ -approximates the social welfare if the number of items initially possessed by an agent is bounded by a constant. The principle that we used to construct Mechanism \mathbb{M}_{sbb} can more generally be used to turn any WBB mechanism into a SBB one, while

¹¹ For 2 agents, it is straightforward to come up with alternative mechanisms that have the desired properties.

preserving the DSIC and ex-post IR properties. It also reveals a problematic aspect of the notion of SBB: it allows for agents to receive money, while they are not involved in any trade. This motivates a strengthened notion of strong budget balance, which we call *direct trade strong budget balance*.

Definition 3.1. *A mechanism for an exchange market satisfies direct trade strong budget balance (DSBB) iff the outcome it generates can be achieved by a set of bilateral trades, where each trade consists of a reallocation of an item from an agent i to an agent j , and a monetary transfer from agent j to agent i . Moreover, each item may only be traded once.*

It can be seen that Mechanism \mathbb{M}_{sbb} does not satisfy DSBB. In the remainder of the paper we will proceed to design mechanisms for two-sided markets that do satisfy DSBB.¹² Moreover, two of our results provide an $O(1)$ -approximation even in settings in which the \mathbb{M}_{sbb} provides only a log-factor.

4 A Mechanism for Unit-Supply Two-Sided Markets with XOS Buyers

In this section we present a DSIC, ex-post IR, and DSBB mechanism for unit supply two-sided markets, when buyers have XOS valuation functions. This mechanism achieves constant approximation to the optimal social welfare. This result extends [10] to two-sided auctions and [5] to a setting in which buyers have combinatorial preferences over the items. In *unit-supply two-sided markets*, for simplicity, we use $[k]$ to denote to both the set of items and the set of sellers, where item j is owned by seller j (so $I_j = \{j\}$ for all $j \in [k]$). For each seller $j \in [k]$, we then treat F_j as a distribution over $\mathbb{R}_{\geq 0}$ instead of a distribution over functions.

We assume throughout this section that $(n, k, k, \mathbf{I}, \mathbf{G}, \mathbf{F})$ is a given unit-supply two-sided market, on which we run the mechanism to be defined. For an allocation $(\mathbf{X}, \mathbf{Y}) \in \mathcal{A}$, we shall use the superscripts SW^B, SW^S to respectively denote the buyers' and the sellers' contribution to the social welfare, i.e.,

$$\text{SW}^B(\mathbf{X}, \mathbf{Y}) := \sum_{i=1}^n v_i(X_i), \text{ and } \text{SW}^S(\mathbf{X}, \mathbf{Y}) := \sum_{j=1}^k w_j \mathbb{I}[j \in Y_j].$$

Our mechanism will work by specifying prices on the items in the market. For a bundle A and an item price vector $\mathbf{p} = (p_1, \dots, p_k)$, we define the *demand correspondence* of buyer $i \in [n]$ with valuation function v_i as

$$\mathcal{D}(v_i, \mathbf{p}, A) := \left\{ S \subseteq A : v_i(S) - \sum_{j \in S} p_j \geq v_i(T) - \sum_{j \in T} p_j \text{ for all } T \subseteq A \right\},$$

i.e., the set of bundles of items in A that maximize i 's utility under the given item prices, when she has valuation function v_i .

For a buyer i with valuation function v_i , we define the *additive representative function* for bundle $T \subseteq [k]$ as the additive function $a(v_i, T, \cdot) : 2^{[k]} \rightarrow \mathbb{R}_{\geq 0}$ such that $v_i(T) = a(v_i, T, T)$, and $v_i(S) \geq a(v_i, T, S)$ for all $S \subseteq [k]$. The additive representative function of a bundle is guaranteed to exist for each buyer i , for each valuation function in the support of F_i , by the definition of XOS functions.

Mechanism. Let \mathbb{A} be an algorithm that, given a valuation profile of the buyers \mathbf{v} , allocates the items $[k]$ to them, and does not take into account the sellers and their valuations. Our mechanism will use such an algorithm as a black-box, and it can be thought of as outputting either an allocation

¹² We note that the double auctions given in [5] also satisfy the DSBB property.

that is optimal for the buyers (in case one does not care about the runtime of the mechanism) or an approximately optimal one (in case one insists on the mechanism running in polynomial time).

Let $\mathbf{X}^{\text{all}}(\mathbf{v}) = (X_1^{\text{all}}(\mathbf{v}), \dots, X_n^{\text{all}}(\mathbf{v}))$ be the output allocation of $\mathbb{A}(\mathbf{v})$. Let $\text{SW}(\mathbf{X}^{\text{all}}(\mathbf{v}))$ be the total social welfare of the allocation $\mathbf{X}^{\text{all}}(\mathbf{v})$.

We define for each item $j \in [k]$ its *contribution* $\text{SW}_j^B(\mathbf{v})$ to the social welfare $\text{SW}(\mathbf{X}^{\text{all}}(\mathbf{v}))$ as follows: If there exists a player i that receives item j in allocation $X_i^{\text{all}}(\mathbf{v})$, then $\text{SW}_j^B(\mathbf{v}) = a(v_i, X_i^{\text{all}}(\mathbf{v}), \{j\})$. Otherwise, if j is not allocated to any buyer in $X_i^{\text{all}}(\mathbf{v})$, then $\text{SW}_j^B(\mathbf{v}) = 0$.

This notion allows us to make a distinction between *high welfare items* and *low welfare items*. An item $j \in [k]$ is said to have *high welfare with respect to* $\text{SW}(X_i^{\text{all}}(\mathbf{v}))$ iff $\mathbb{E}_{\mathbf{v}}[\text{SW}_j^B(\mathbf{v})] \geq 4\mathbb{E}[w_j]$, i.e., the expected social welfare contribution of j if we would allocate j according to $\mathbf{X}^{\text{all}}(\mathbf{v})$ is at least four times as high as the social welfare that results from leaving item j at its seller.

Formally, let L be the set of high welfare items, i.e., $L := \{j \in [k] : \mathbb{E}[\text{SW}_j^B(\mathbf{v})] \geq 4\mathbb{E}[w_j]\}$, and let \bar{L} be the set of low welfare items, i.e. $\bar{L} := [k] \setminus L$. For each high welfare item $j \in L$, the mechanism makes use of the following associated *item price* p_j :

$$p_j := \frac{1}{2}\mathbb{E}_{\mathbf{v}}[\text{SW}_j^B(\mathbf{v})].$$

Observe that $p_j \geq 2\mathbb{E}[w_j]$ for all $j \in L$, by our definition of high welfare items. The reason why L is chosen in such a way is twofold: first, the items in \bar{L} if kept by their sellers provide a welfare loss of at most a constant factor; second, every item in L is guaranteed to be sold (if sold) at a high price.

Our (randomized) mechanism does the following simple procedure. First, it goes to every seller in L (in any order) and asks each of them whether they would sell their item for a price of p_j . As mentioned above, by definition of the prices, every seller $j \in L$ accepts the price with probability at least $1/2$ because of Markov's inequality. To make sure that this probability is exactly $1/2$, the seller j is only given the opportunity to sell her item at the price p_j with probability q_j such that (in expectation) the offer is accepted with probability exactly $1/2$. Formally, the mechanism makes an offer to the seller j with probability

$$q_j := \frac{1}{2F_j(p_j)}, \text{ where } F_j(p_j) = \Pr[w_j \leq p_j].$$

Once we have gone through every seller and know which items are in the market, we move to the buyers and ask each (in any order) for their favorite bundle according to the items currently available for purchase.

We call the mechanism sketched above $\mathbb{M}_{1\text{-supply}}$, which we will now present more precisely:

1. Let $L := \{j \in [k] : \mathbb{E}_{\mathbf{v}}[\text{SW}_j^B(\mathbf{v})] \geq 4\mathbb{E}[w_j]\}$.
2. For all $j \in L$, set $p_j := \frac{1}{2}\mathbb{E}_{\mathbf{v}}[\text{SW}_j^B(\mathbf{v})]$.
3. Let $A_1 := \emptyset, X_i := \emptyset$ for all $i \in [n]$ and $Y_j := \{j\}$ for all $j \in [k]$.
4. For all $j \in L$:
 - (a) Set $q_j := 1/(2\Pr[w_j \leq p_j])$.
 - (b) With probability q_j , offer payment p_j in exchange for her item. Otherwise, skip this seller.
 - (c) If j accepts the offer, set $A_1 := A_1 \cup \{j\}$.
5. For all $i \in [n]$:
 - (a) Buyer i chooses a bundle $B_i \in \mathcal{D}(v_i, \mathbf{p}, A_i)$ that maximizes her utility.
 - (b) Allocate the accepted items to buyer i , i.e., $X_i := B_i$ and $Y_j := \emptyset$ for all $j \in B_i$.
 - (c) Remove the selected items from the available items, i.e., $A_{i+1} := A_i \setminus B_i$.
6. Return the outcome consisting of allocation $(\mathbf{X} = (X_1, \dots, X_n), \mathbf{Y} = (Y_1, \dots, Y_k))$ and payments $\boldsymbol{\rho} = (\boldsymbol{\rho}^B, \boldsymbol{\rho}^S)$, where $\rho_i^B = \sum_{j \in X_i} p_j$ for $i \in [n]$ and $\rho_j^S = -p_j \mathbb{I}[Y_j = \emptyset]$ for $j \in [k]$.

Note that Algorithm \mathbb{A} is used in the first steps of mechanism $\mathbb{M}_{1\text{-supply}}$, where $\mathbb{E}_{\mathbf{v}}[\text{SW}_j^B(\mathbf{v})]$ is computed. Let α be the factor by which \mathbb{A} is guaranteed to approximate the social welfare of the buyers.

Theorem 4.1. $\mathbb{M}_{1\text{-supply}}$ is ex-post IR, DSIC, DSBB, and $(2+4\alpha)$ -approximates the optimal social welfare.

In particular, taking for \mathbb{A} an optimal algorithm (i.e., $\alpha = 1$), we obtain that there exists a mechanism that is ex-post IR, DSIC, DSBB, and 6-approximates the optimal social welfare. Alternatively, one may take for \mathbb{A} an approximation algorithm in order to obtain polynomial time implementable mechanisms, as explained in further detail in [10].

5 A Mechanism for Two-Sided Markets with XOS Buyers and Additive Sellers

We now consider the setting in which sellers may own multiple distinct items and have an additive valuation function over them. We design a DSBB mechanism that is DSIC and ex-post IR on the sellers' side, and BIC and interim IR on the buyers' side. At the end of the section also proof that in the case of both buyers and sellers possessing additive valuation functions, the mechanism we present is DSIC and ex-post IR on both sides of the market.

We assume throughout this section that $(n, m, k, \mathbf{I}, \mathbf{G}, \mathbf{F})$ is a given two-sided market with XOS buyers and Additive Sellers, on which we run the mechanism to be defined. Like in the previous section, the buyers are still assumed to have XOS valuation functions over the items. Since now the number of items and sellers is different in general, we use m to denote the number of sellers and k for the number of items. The valuation w_j of a seller j is now an additive function. We reuse the following notation from Section 4: The allocation $(X_1^{\text{all}}(\mathbf{v}), \dots, X_n^{\text{all}}(\mathbf{v}))$ returned by an allocation algorithm \mathbb{A} on input \mathbf{v} returns an allocation of $[k]$ to $[n]$. We let $\alpha \geq 1$ again denote the approximation factor by which \mathbb{A} approximates the social welfare. For XOS valuation v_i and bundle $T \subseteq [k]$ we use $a(v_i, T, \cdot)$ to denote the corresponding additive function of v_i for T . Also we use the buyers' social welfare contribution $\text{SW}_\ell^B(\mathbf{v})$ for item $\ell \in [k]$ and buyers' valuation profile \mathbf{v} , as defined in Section 4. We define furthermore the *sellers' social welfare contribution* $\text{SW}_\ell^S(\mathbf{w})$ for item $\ell \in I_j$ and sellers' valuation profile \mathbf{w} as $\text{SW}_\ell^S(\mathbf{w}) := w_j(\{\ell\})$. Due to the fact that for $j \in [m]$, w_j is an additive function, there is no need for defining the notion of a *corresponding additive function* for a seller.

Mechanism. We aim to design a BIC, interim IR, and SBB mechanism that approximates the optimal social welfare within a constant. We propose the following mechanism, which we refer to as \mathbb{M}_{add} . We let $L_j := \{\ell \in I_j : \mathbb{E}[\text{SW}_\ell^B(\mathbf{v})] \geq 4\mathbb{E}[\text{SW}_\ell^S(\mathbf{w})]\}$ and $\bar{L}_j := I_j \setminus L_j$ for all $j \in [m]$, and we let $L := \bigcup_{j=1}^m L_j$ and $\bar{L} := [k] \setminus L$ denote the sets of high welfare items and low welfare items, respectively. Our mechanism will only allow trading items in L , and define for $\ell \in L$ the item price

$$p_\ell := \frac{1}{2}\mathbb{E}[\text{SW}_\ell^B(\mathbf{v})],$$

similar to what we did for $\mathbb{M}_{1\text{-supply}}$.

An essential difference between \mathbb{M}_{add} and $\mathbb{M}_{1\text{-supply}}$ is that the order in which buyers and sellers are processed is reversed. Mechanism \mathbb{M}_{add} roughly works as follows. It first asks every buyer which set of items it would like to receive from those items in L that have not been requested yet. Then \mathbb{M}_{add} offers every seller $j \in [m]$ a payment in exchange for the subset of all items in I_j that have been requested. This offer is made with a q_j ensures that the requested items of seller j are transferred to the buyers with probability $1/2$. The items of the sellers accepting the offer are transferred to

the buyers for the corresponding item prices. Buyers act strategically, and request a set of items that maximizes their expected utility, knowing that the item sets requested from each seller will be assigned to them with probability $1/2$ ¹³.

1. For $\ell \in [k]$, compute $\mathbb{E}[\text{SW}_\ell^B(\mathbf{v})]$ and $\mathbb{E}[\text{SW}_\ell^S(\mathbf{w})]$.
2. For all $j \in [m]$, compute L_j and q_j .
3. Compute L, \bar{L} .
4. Let $A_1 := L, X_i := \emptyset$ for all $i \in [n]$, and $Y_j := I_j$ for all $j \in [m]$.
5. For each buyer $i \in [n]$:
 - (a) Ask buyer i to select an expected-maximising bundle $B_i \subseteq A_i$ given the prices $\{p_\ell : \ell \in A_i\}$ from the set of available items (where the expectation is taken with respect to the randomness of both the mechanism and the valuations).
 - (b) Update the set of available items $A_{i+1} := A_i \setminus B_i$.
6. Let $B := \bigcup_{i=1}^n B_i$ be the set of all items demanded by the buyers.
7. For each seller $j \in [m]$:
 - (a) Let $S_j := B \cap L_j$ be the set of items owned by seller j that are demanded.
 - (b) Let $p(S_j) := \sum_{\ell \in S_j} p_\ell$ and let $q_j = 1/(2\Pr[w_j(S_j) \leq p(S_j)])$.
 - (c) With probability q_j , offer payment $p(S_j)$ in exchange for the bundle S_j . Otherwise, skip this seller.
 - (d) If the seller accepts the offer, allocate each items in S_j to the buyer that requested it (i.e., remove S_j from Y_j and add $S_j \cap B_i$ to X_i for all $i \in [n]$)
8. Return the outcome consisting of allocation $(\mathbf{X} = (X_1, \dots, X_n), \mathbf{Y} = (Y_1, \dots, Y_m))$ and payments $\boldsymbol{\rho} = (\boldsymbol{\rho}^B, \boldsymbol{\rho}^S)$, where $\rho_i^B = \sum_{\ell \in X_i} p_\ell$ for $i \in [n]$ and $\rho_j^S = \sum_{\ell \in I_j \setminus Y_j} -p_\ell$ for $j \in [m]$.

Example 5.1. There is 1 buyer and 2 unit-supply sellers. Each seller possesses 1 item. The buyer has two XOS valuation functions v_1 and v_2 , each one is chosen with probability 0.5. v_1 is composed by 3 additive functions a_1, a_2 , and a_3 , i.e., $v_1(S) = \max\{a_1(S), a_2(S), a_3(S)\}$. v_2 is composed by the only additive function a_4 . The additive functions are represented in the following table, recall that an additive function is represented by $a(S) = \sum_{j \in S} \alpha_j$ for all $S \subseteq [k]$. Each seller j has a valuation function $w_j = 0$, so we always want to reallocate items from sellers to the buyer.

Function	item 1 (α_1)	item 2 (α_2)
a_1	0	4
a_2	8	0
a_3	7	2
a_4	1	6

Now, we have to compute the prices that the mechanism has to post. Thus, we need the expected contribution to the optimal social welfare of every item. First, notice that the optimum allocates the items 1 and 2 to the buyer when her valuation is v_1 . In this case the contribution to the optimal social welfare of items 1 is 7, and the contribution of item 2 is 2. Similarly, if the buyer has valuation v_2 , the optimum still allocates items 1 and 2 to her, but in this case the contribution to the optimal social welfare of item 1 is 1, and the contribution of item 2 is 6. Thus, the expected contribution of every item to the optimal social welfare is 4, i.e., $\mathbb{E}[\text{SW}_j^B(\mathbf{v})] = 4$ for all $j = 1, 2$. Since the price p_j of each item is defined to be half of the expected contribution to the optimal social welfare, $p_j = 2$ for all the items.

¹³ The buyer may need to compute complicated expected values in order to establish which bundle is maximizing his expected utility. But this depends on the query oracle model we assume to have. However, most importantly, the focus of the paper is not on polynomial time implementability but on the existence of mechanisms that constantly approximate the optimal social-welfare.

Now the buyer has to compute the bundle that maximises her expected-utility given the posted prices, and the fact that each seller is available with probability 0.5. First, consider the case when the buyer has valuation v_1 . In this case the expected utility for the different bundles are:

$$u(\{1\}) = \frac{1}{2} \cdot (8 - 4) + \frac{1}{2} \cdot 0 = 2$$

$$u(\{2\}) = \frac{1}{2} \cdot (4 - 4) + \frac{1}{2} \cdot 0 = 0$$

$$u(\{1, 2\}) = \frac{1}{4} \cdot (8 - 4) + \frac{1}{4} \cdot (4 - 4) + \frac{1}{4} \cdot (9 - 8) + \frac{1}{4} \cdot 0 = \frac{5}{4}$$

The utility-maximising bundle that will be requested by the buyer in case of v_1 is $\{1\}$. Instead, if the valuation of the buyer is v_2 , then the requested bundle will be $\{1, 2\}$.

This example shows that the buyers are committed to difficult computations in order to understand which is the utility-maximising bundle. But, again, our goal is to understand the existence of mechanisms that achieve a constant approximation to the optimum social welfare, even if they are not computational efficient.

Theorem 5.1. *The mechanism \mathbb{M}_{add} is ex-interim IR, BIC, DSBB, and $(2 + 4\alpha)$ -approximates the optimal social welfare.*

By taking for \mathbb{A} an optimal algorithm (i.e., $\alpha = 1$), we obtain the existence of a mechanism that is ex-post IR, DSIC, DSBB, and 6-approximates the optimal social welfare. However, if both a polynomial time approximation algorithm \mathbb{A} and an approximation for the query oracle exist, then we get our mechanism to run in polynomial time.

Corollary 5.1. *For the special case that for all $i \in [n]$, distribution G_i is over additive valuation functions, \mathbb{M}_{add} is ex-post IR, DSIC, DSBB and $(2 + 4\alpha)$ -approximates the optimal social welfare.*

6 Discussion

An open problem is to extend or refine our mechanisms so that they satisfy the DSIC and ex-post IR properties for the case of XOS buyers and additive sellers. The first naive approach for doing so might be trying to consider every additive seller as a set of distinct unit-supply sellers and then run $\mathbb{M}_{1\text{-supply}}$. However, this is not guaranteed to work due to the fact that the items placed in the market by the seller influence the buyers' decisions on their bundle choices. Something we might additionally do is to ask every seller for her favourite bundle to place in the market, However, this may cause a seller to regret having chosen that particular bundle after seeing the realizations of the buyers' valuations. On the other hand, it also seems highly challenging to establish any sort of impossibility result for any reasonably defined class of posted price mechanisms for two-sided markets.

Another natural challenge is to extend the above mechanism to the setting in which both buyers and sellers possess an XOS valuation function over bundles of items. We suspect however that it is impossible to devise a suitable mechanism that uses item-pricing (i.e., mechanisms that fix a price vector offline and then request buyers or sellers one-by-one to specify their favorite item bundle). The main reason for why this seems to be the case is that if a buyer (or a seller) is asked to select a bundle of items B she desires the most, then it is impossible to guarantee that she receives the complete bundle. Instead, she may receive a subset of it, and in turn she may regret having chosen that bundle B and not another bundle B' after having observed the realizations of the sellers' valuations.

Bibliography

- [1] Blumrosen, L., Dobzinski, S.: Reallocation mechanisms. In: Proceedings of the 15th ACM Conference on Economics and Computation (EC). pp. 617–617. ACM (2014)
- [2] Blumrosen, L., Holenstein, T.: Posted prices vs. negotiations: an asymptotic analysis. In: Proceedings of the 9th ACM conference on Electronic Commerce (EC). pp. 49–49. ACM (2008)
- [3] Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-parameter mechanism design and sequential posted pricing. In: Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC). pp. 311–320. ACM (2010)
- [4] Clarke, E.H.: Multipart pricing of public goods. *Public Choice* 11(1), 17–33 (1971)
- [5] Colini-Baldeschi, R., Keijzer, B.D., Leonardi, S., Turchetta, S.: Approximately efficient double auctions with strong budget balance. In: Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). pp. 1424–1443. SODA '16, ACM-SIAM (2016)
- [6] Deng, X., Goldberg, P.W., Tang, B., Zhang, J.: Revenue maximization in a Bayesian double auction market. *Theoretical Computer Science* 539(1), 1–12 (2014)
- [7] Deshmukh, K., Goldberg, A.V., Hartline, J.D., Karlin, A.R.: Truthful and competitive double auctions. In: Proceedings of the 10th Annual European Symposium on Algorithms (ESA). pp. 361–373. Springer (2002)
- [8] Dütting, P., Kleinberg, R.: Polymatroid prophet inequalities. In: Proceedings of the 23rd Annual European Symposium on Algorithms (ESA). pp. 437–449 (2015)
- [9] Dütting, P., Talgam-Cohen, I., Roughgarden, T.: Modularity and greed in double auctions. In: Proceedings of the 15th ACM Conference on Economics and Computation (EC). pp. 241–258. ACM (2014)
- [10] Feldman, M., Gravin, N., Lucier, B.: Combinatorial auctions via posted prices. In: Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). pp. 123–135. ACM-SIAM (2015)
- [11] Gresik, T.A., Satterthwaite, M.A.: The rate at which a simple market converges to efficiency as the number of traders increases: An asymptotic result for optimal trading mechanisms. *Journal of Economic Theory* 48(1), 304–332 (1989)
- [12] Groves, T.: Incentives in teams. *Econometrica* 41(4), 617–631 (1973)
- [13] Kleinberg, R., Weinberg, S.M.: Matroid prophet inequalities. In: Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC). pp. 123–136. ACM (2012)
- [14] McAfee, R.P.: A dominant strategy double auction. *Journal of Economic Theory* 56(2), 434–450 (1992)
- [15] Myerson, R.B.: Optimal auction design. *Mathematics of Operations Research* 6(1), 58–73 (1981)
- [16] Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29, 265–281 (1983)
- [17] Rustichini, A., Satterthwaite, M.A., Williams, S.R.: Convergence to efficiency in a simple market with incomplete information. *Econometrica* 62(5), 1041–1063 (1994)
- [18] Sandholm, T., Gilpin, A.: Sequences of Take-It-or-Leave-It Offers: Near-Optimal Auctions Without Full Valuation Revelation, pp. 73–91. Springer (2004)
- [19] Satterthwaite, M.A., Williams, S.R.: The rate of convergence to efficiency in the buyer’s bid double auction as the market becomes large. *The Review of Economic Studies* 56(4), 477–498 (1989)
- [20] Satterthwaite, M.A., Williams, S.R.: The optimality of a simple market mechanism. *Econometrica* 70(5), 1841–1863 (2002)

- [21] Segal-Halevi, E., Hassidim, A., Aumann, Y.: A random-sampling double-auction mechanism. CoRR abs/1604.06210 (2016)
- [22] Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. The Journal of Finance 16(1), 8–37 (1961)

A An Initial Mechanism and Direct Trade Strong Budget Balance (Full Details)

In [1], Blumrosen and Dobzinski present a mechanism for exchange markets with subadditive valuation functions. They prove the following for this mechanism, which we name \mathbb{M}_{bd} .

Theorem A.1 (Blumrosen and Dobzinski [1]). *Mechanism \mathbb{M}_{bd} is a DSIC, WBB, ex-post IR randomized direct revelation mechanism that $4H(s)$ -approximates the optimal social welfare for combinatorial exchange markets $(n, k, \mathbf{I}, \mathbf{F})$ with subadditive valuation functions, where $s = \min\{n, |I_i| : i \in [n]\}$ is the minimum of the number of agents and the number of items in an agents initial endowment, and $H(\cdot)$ denotes the harmonic numbers.*

In particular, this mechanism gives us a constant approximation factor if the number of starting items of the agents is bounded by a constant.

We show now how we can use this mechanism as a black box in order to obtain an SBB mechanism with only a slightly worse approximation ratio. Define mechanism \mathcal{M}_{sbb} as follows. When given as input a combinatorial exchange market $C = (n, k, \mathbf{I}, \mathbf{F})$,

1. Select an agent in $i \in [n]$ uniformly at random.
2. Run Mechanism \mathbb{M}_{bd} on the combinatorial exchange market

$$C_{-i} = ([n] \setminus \{i\}, \mathbf{I}_{-i} = (I_1, \dots, I_{i-1}, I_{i+1}, \dots, I_n), \mathbf{F}_{-i} = (F_1, \dots, F_{i-1}, F_{i+1}, \dots, F_n)).$$

Let $(\mathbf{X}_{-i}, \boldsymbol{\rho}_{-i})$ be the outcome that Mechanism \mathbb{M}_{bd} outputs.

3. Set $X_i = I_i$ and set $p_i = -\sum_{j \in [n] \setminus \{i\}} p_j$. Output the allocation (X_i, \mathbf{X}_{-i}) and output payment vector $(p_i, \boldsymbol{\rho}_{-i})$.

So Mechanism \mathbb{M}_{sbb} essentially runs Mechanism \mathbb{M}_{bd} where one random agent is removed from the market. This agent receives the leftover money that Mechanism \mathbb{M}_{bd} generates, and does not receive or lose any items. The following is a direct corollary of the DSIC, WBB, and ex-post IR properties of mechanism \mathbb{M}_{bd} .

Theorem A.2. *Mechanism \mathbb{M}_{sbb} is a DSIC, SBB, and ex-post IR mechanism for exchange markets with subadditive valuation functions.*

Secondly, the following theorem shows that the mechanism loses only a factor $2n/(n-1) \leq 3$ in the approximation ratio for $n \geq 3$. (For $n = 2$ it is straightforward to come up with alternative mechanisms that achieve a good approximation ratio.)

Theorem A.3. *Mechanism \mathbb{M}_{sbb} achieves an $8nH(s)/(n-1)$ -approximation to the optimal social welfare for exchange markets with subadditive valuations and at least 3 agents.*

Proof. Fix a valuation vector \mathbf{v} of the agents, let $\mathbf{X}_{\mathbf{v}}^{**} \subseteq \mathcal{A}$ be the social welfare maximising allocation when the agents have valuations \mathbf{v} . For an agent $i \in [n]$, denote by $\mathbf{X}_{\mathbf{v}, -i}^{**}$ the allocation for C_{-i} where $(X_{\mathbf{v}, -i}^{**})_j = (X_{\mathbf{v}}^{**})_j \setminus I_i$ for $j \in [n] \setminus \{i\}$, i.e., the allocation obtained from $\mathbf{X}_{\mathbf{v}}^{**}$ when i is removed, and all items of i are removed. Moreover let $\mathbf{X}_{\mathbf{v}, -i}^*$ be the optimal allocation of the

combinatorial exchange market C_{-i} when the valuation function vector of the players $[n] \setminus \{i\}$ is fixed to \mathbf{v}_{-i} . Mechanism \mathbb{M}_{sbb} selects i uniformly at random, so by Theorem A.1, the expected social welfare of Mechanism \mathbb{M}_{sbb} is at least

$$\begin{aligned}
\frac{1}{4H(s)} \mathbb{E}_i \left[\sum_{j \in [n] \setminus \{i\}} v_j(\mathbf{X}_{\mathbf{v},-i}^*) \right] &\geq \frac{1}{4H(s)} \mathbb{E}_i \left[\sum_{j \in [n] \setminus \{i\}} v_j(\mathbf{X}_{\mathbf{v},-i}^{**}) \right] \\
&= \frac{1}{4nH(s)} \sum_{i \in [n]} \sum_{j \in [n] \setminus \{i\}} v_j((X_{\mathbf{v}}^{**})_j \setminus I_i) \\
&= \frac{1}{4nH(s)} \sum_{i \in [n]} \sum_{j \in [n] \setminus \{i\}} v_i((X_{\mathbf{v}}^{**})_i \setminus I_j) \\
&= \frac{1}{4nH(s)} \sum_{i \in [n]} \sum_{\substack{\{j,j'\}: j,j' \in [n] \setminus \{i\} \\ \wedge j \neq j'}} \frac{1}{n-2} (v_i((X_{\mathbf{v}}^{**})_i \setminus I_j) + v_i((X_{\mathbf{v}}^{**})_i \setminus I_{j'})) \\
&\geq \frac{1}{4nH(s)} \sum_{i \in [n]} \sum_{\substack{\{j,j'\}: j,j' \in [n] \setminus \{i\} \\ \wedge j \neq j'}} \frac{1}{n-2} v_i(\mathbf{X}_{\mathbf{v}}^{**}) \\
&= \frac{1}{4nH(s)} \sum_{i \in [n]} \frac{n-1}{2} v_i(\mathbf{X}_{\mathbf{v}}^{**}) \\
&= \frac{n-1}{8nH(s)} \sum_{i \in [n]} v_i(\mathbf{X}_{\mathbf{v}}^{**}),
\end{aligned}$$

where the second inequality follows from subadditivity. This proves the claim, since the above holds for every valuation vector \mathbf{v} . \square

This yields an ex-post IR, SBB, DSIC mechanism that $O(1)$ -approximates the social welfare if the number of items initially possessed by an agent is bounded by a constant.

The principle that we used to construct Mechanism \mathbb{M}_{sbb} can more generally be used to turn any WBB mechanism into a SBB one, while preserving the DSIC and ex-post IR properties. This principle also reveals a problematic aspect of the notion of SBB: it allows for agents to receive money, while they are not involved in any trade. This motivates a strengthened notion of strong budget balance, which we call *direct trade strong budget balance*.

Definition A.1. *A mechanism for an exchange market satisfies direct trade strong budget balance (DSBB) iff the outcome it generates can be achieved by a set of bilateral trades, where each trade consists of a reallocation of an item from an agent i to an agent j , and a monetary transfer from agent j to agent i . Moreover, each item may only be traded once.*

It can be seen that Mechanism \mathbb{M}_{sbb} does not satisfy DSBB. In the remainder of the paper we will proceed to design mechanisms for two-sided markets that do satisfy DSBB.¹⁴

B Proof of Theorem 4.1

We split the proof into two lemmas that bound the sellers' and the buyers' relative contributions to the social welfare. We use the notation OPT as defined in Section 2, and we use ALG to

¹⁴ We note that the double auctions given in [5] also satisfy the DSBB property.

denote the expected social welfare of the mechanism, i.e., $\mathbb{E}_{\mathbf{v}, \mathbf{w}}[\text{SW}(\mathbb{M}_{1\text{-supply}}(\mathbf{v}, \mathbf{w}))]$. Moreover, the superscripts S, B respectively denote the sellers' and buyers' contributions to the social welfare, e.g., $\text{OPT} = \text{OPT}^S + \text{OPT}^B$ and $\text{ALG} = \text{ALG}^S + \text{ALG}^B$. Recall that we use $\mathbf{X}^{\text{all}}(\mathbf{v})$ to denote the allocation resulting from Algorithm A on buyers' valuation vector \mathbf{v} .

The following lemma is a simple consequence of the fact that $\mathbb{M}_{1\text{-supply}}$ lets every seller in L gets an offer and accepts it with probability exactly $1/2$.

Lemma B.1. *If every seller $j \in L$ puts her item into the market with probability exactly $1/2$, then*

$$2\text{ALG}^S \geq \sum_{j=1}^k \mathbb{E}[w_j] \geq \text{OPT}^S.$$

Proof. The second inequality is trivial, so we focus on the first inequality. First, observe that

$$\Pr[w_j > p_j] \leq \Pr[w_j > 2\mathbb{E}[w_j]] < \frac{1}{2},$$

where the first inequality is because $j \in L$, and the second inequality is by Markov's inequality. Thus with probability at least $1/2$ a seller j is happy to sell his item at price p_j . But every seller receives an offer from the mechanism with probability $q_j := 1/(2\Pr[w_j \leq p_j])$, so every seller gets an offer and accepts it with probability exactly $1/2$. This implies that every $j \in L$ contributes in expectation at least a $\mathbb{E}[w_j]/2$ to the social welfare, since with non negative probability the buyers won't buy the item j . Moreover, a seller in \bar{L} who never trades, so their full expected valuations are contributed to the expected social welfare. \square

Next, we prove a more difficult bound that relates ALG^B and ALG^S to OPT^B .

Lemma B.2. *The buyers' contributions to the optimal social welfare is bounded by*

$$4\alpha\text{ALG}^B + 4\alpha\text{ALG}^S \geq \text{OPT}^B.$$

Before proving Lemma B.2, we point out that Theorem 4.1 follows straightforwardly from it.

Proof (of Theorem 4.1). The bound on the approximation ratio follows from the sum of the inequalities of Lemma B.1 and Lemma B.2. Moreover, it is a dominant strategy for a seller to accept if and only if the payment offered to her exceeds her valuation, and it is a dominant strategy for a buyer to choose a utility-maximising bundle for the items and item prices offered to her. Thus, when viewed as a direct revelation mechanism, $\mathbb{M}_{1\text{-supply}}$ is DSIC. It is clear that participating in the mechanism can never lead to a decrease in utility for both buyers and sellers, and therefore the mechanism is also ex-post IR. Lastly, it is straightforward to see that the mechanism is DSBB, as the definition of $\mathbb{M}_{1\text{-supply}}$ which we gave in terms of sequential posted pricing naturally yields us the required set of bilateral trades. \square

So it remains to prove Lemma B.2. In order to do this, we first prove two propositions: one of them bounds the expected sum of the buyers' utilities, and one of them bounds the expected sum of the buyers' payments. In both propositions we only consider items in L . Given a buyers' valuation profile \mathbf{v} , let $\mathbf{v}_{<i} = (v_1, \dots, v_{i-1})$. Further, let Z be a random variable that denotes the sellers that receive and accept an offer from the mechanism, i.e., the set A_1 at step 5 of $\mathbb{M}_{1\text{-supply}}$. For $i \in [n]$ let $A_i(\mathbf{v}_{<i}, Z)$ be the set A_i as given in the definition of $\mathbb{M}_{1\text{-supply}}$ when the valuation profile of the

buyers is \mathbf{v} and Z are the sellers in the market. Note that this implies that $X_i \subseteq A_i(\mathbf{v}_{<i}, Z) \subseteq Z$. Consequently, $\Lambda_{n+1}(\mathbf{v}, Z)$ is the subset of items for which the corresponding sellers have accepted the offer made to them by the mechanism, but remain allocated to the corresponding seller after execution.

Proposition B.1. *The total expected utility of the buyers for the allocation returned by $\mathbb{M}_{1\text{-supply}}$ is bounded from below by*

$$\mathbb{E} \left[\sum_{i \in [n]} u_i(\mathbb{M}_{1\text{-supply}}(\mathbf{v}, \mathbf{w})) \right] \geq \frac{1}{2} \sum_{j \in L} \Pr_{\mathbf{v}, Z} [j \in \Lambda_{n+1}(\mathbf{v}, Z) \mid j \in Z] p_j.$$

(Note that the random variables in this expression are \mathbf{v}, \mathbf{w} , and the decisions of the mechanism to make offers to the sellers in L .)

Proof. First, note that for each $j \in L$ it holds that $\Pr[j \in Z] = 1/2$. Recall that we defined $p_j := \frac{1}{2} \mathbb{E}_{\mathbf{v}} [\text{SW}_j^B(\mathbf{v})]$. Thus, observe that by definition of p_j , $\text{SW}_j^B(\mathbf{v})$, and the law of total probability, it holds for all $j \in L$ that

$$p_j = \mathbb{E}_{\mathbf{v}} [\text{SW}_j^B(\mathbf{v}) - p_j] = \sum_{i=1}^n \mathbb{E}_{\mathbf{v}} \left[(\text{SW}_j^B(\mathbf{v}) - p_j) \mathbb{I} [j \in X_i^{\text{all}}(\mathbf{v})] \right]. \quad (1)$$

Fix $i \in [n]$, buyers' valuation profile \mathbf{v} , and set $Z \subseteq L$ of sellers who accepted the mechanism's offer, and now consider the set $A_i(\mathbf{v}_{<i}, Z) \subseteq L$ of available items that i can choose from. Buyer i selects a bundle that maximizes her utility, i.e., that is in $\mathcal{D}(v_i, \mathbf{p}, A_i(\mathbf{v}_{<i}, Z))$.

Now consider an additional randomly drawn profile of valuation functions $\tilde{\mathbf{v}}_{-i}$ for all buyers except i , that is independent of \mathbf{v} . Consider the allocation $X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i})$ be the allocation of buyer i returned by $\mathbb{A}(v_i, \tilde{\mathbf{v}}_{-i})$. For $i \in [n]$, consider the corresponding additive representative function $a(v_i, X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}), \cdot)$, such that $a(v_i, X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}), \{j\}) = \text{SW}_j^B(v_i, \tilde{\mathbf{v}}_{-i})$. Let

$$S_i(v_i, \mathbf{v}_{-i}, \tilde{\mathbf{v}}_{-i}, Z) := X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}) \cap A_i(\mathbf{v}_{<i}, Z)$$

be the items in $X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i})$ that buyer i may choose from under valuation profile \mathbf{v} . As i chooses a bundle $B_i(\mathbf{v}, Z) \in \mathcal{D}(v_i, \mathbf{p}, A_i(\mathbf{v}_{<i}, Z))$ that maximizes her utility, and $S_i(v_i, \mathbf{v}_{-i}, \tilde{\mathbf{v}}_{-i}, Z)$ is in $\mathcal{D}_i(v_i, \mathbf{p}, A_i(\mathbf{v}_{<i}, Z))$, it follows that i 's utility for $B_i(\mathbf{v}, Z)$ is at least the utility she would get for choosing $S_i(v_i, \mathbf{v}_{-i}, \tilde{\mathbf{v}}_{-i}, Z)$. That is, for all \mathbf{v} and $Z \subseteq L$

$$\begin{aligned} v_i(B_i(\mathbf{v}, Z)) - \sum_{j \in B_i(\mathbf{v}, Z)} p_j &\geq \mathbb{E}_{\tilde{\mathbf{v}}_{-i}} \left[v_i(S_i(v_i, \mathbf{v}_{-i}, \tilde{\mathbf{v}}_{-i}, Z)) - \sum_{j \in S_i(v_i, \mathbf{v}_{-i}, \tilde{\mathbf{v}}_{-i}, Z)} p_j \right] \\ &\geq \mathbb{E}_{\tilde{\mathbf{v}}_{-i}} \left[\sum_{j \in S_i(v_i, \mathbf{v}_{-i}, \tilde{\mathbf{v}}_{-i}, Z)} (a(v_i, X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}), \{j\}) - p_j) \right] \\ &= \mathbb{E}_{\tilde{\mathbf{v}}_{-i}} \left[\sum_{j \in S_i(v_i, \mathbf{v}_{-i}, \tilde{\mathbf{v}}_{-i}, Z)} (\text{SW}_j^B(v_i, \tilde{\mathbf{v}}_{-i}) - p_j) \right], \end{aligned}$$

The second-to-last inequality follows from the definition of the corresponding additive function $a(v_i, X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}), \cdot)$.

Now summing the above expression over all $i \in [n]$ and taking the expectation over \mathbf{v} and Z , we get

$$\begin{aligned} \mathbb{E}_{\mathbf{v}, Z} \left[\sum_{i=1}^n \left(v_i(B_i(\mathbf{v}, Z)) - \sum_{j \in B_i(\mathbf{v}, Z)} p_j \right) \right] &\geq \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}_{-i}, Z} \left[\sum_{i=1}^n \sum_{j \in S_i(v_i, \mathbf{v}_{-i}, \tilde{\mathbf{v}}_{-i}, Z)} (\text{SW}_j^B(v_i, \tilde{\mathbf{v}}_{-i}) - p_j) \right] \\ &= \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}_{-i}, Z} \left[\sum_{i=1}^n \sum_{j \in L} (\text{SW}_j^B(v_i, \tilde{\mathbf{v}}_{-i}) - p_j) \right. \\ &\quad \left. \cdot \mathbb{I} [j \in X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i})] \mathbb{I} [j \in \Lambda_i(\mathbf{v}_{<i}, Z)] \right]. \end{aligned}$$

Note that we exploited the independence of the events $(j \in X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}))$ and $(j \in \Lambda_i(\mathbf{v}_{<i}, \mathbf{z}))$. Thus, switching the order of the sums and using linearity of expectation, we get that

$$\begin{aligned} \mathbb{E}_{\mathbf{v}, Z} \left[\sum_{i=1}^n \left(v_i(B_i(\mathbf{v}, Z)) - \sum_{j \in B_i(\mathbf{v}, Z)} p_j \right) \right] &\geq \sum_{j \in L} \sum_{i=1}^n \Pr_{\mathbf{v}, Z} [j \in \Lambda_i(\mathbf{v}_{<i}, Z)] \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}_{-i}} \left[(\text{SW}_j^B(v_i, \tilde{\mathbf{v}}_{-i}) - p_j) \mathbb{I} [j \in X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i})] \right] \\ &\geq \sum_{j \in L} \Pr_{\mathbf{v}, Z} [j \in \Lambda_{n+1}(\mathbf{v}, Z)] \sum_{i=1}^n \mathbb{E}_{\mathbf{v}} \left[(\text{SW}_j^B(\mathbf{v}) - p_j) \mathbb{I} [j \in X_i^{\text{all}}(\mathbf{v})] \right] \\ &= \sum_{j \in L} \Pr_{\mathbf{v}, Z} [j \in \Lambda_{n+1}(\mathbf{v}, Z)] p_j \\ &= \sum_{j \in L} \Pr_{\mathbf{v}, Z} [j \in \Lambda_{n+1}(\mathbf{v}, Z) \mid j \in Z] \Pr [j \in Z] p_j \\ &= \frac{1}{2} \sum_{j \in L} \Pr_{\mathbf{v}, Z} [j \in \Lambda_{n+1}(\mathbf{v}, Z) \mid j \in Z] p_j. \end{aligned}$$

For the last inequality, we used the fact that for any $i \in [n]$ it holds that $\Pr_{\mathbf{v}} [j \in \Lambda_i(\mathbf{v}_{<i}, Z)] \geq \Pr_{\mathbf{v}} [j \in \Lambda_{n+1}(\mathbf{v}, Z)]$. The first equality follows from (1). \square

Proposition B.2. *The expected sum of the payments charged by $\mathbb{M}_{1\text{-supply}}$ to the buyers is equal to*

$$\mathbb{E} \left[\sum_{i \in [n]} \rho_i^B \right] = \frac{1}{2} \sum_{j \in L} p_j \Pr_{\mathbf{v}, Z} [j \notin \Lambda_{n+1}(\mathbf{v}, Z) \mid j \in Z]$$

Proof. The revenue extracted by the mechanism, meaning the sum of the payments charged to the buyers, is equal to

$$\begin{aligned} \sum_{j \in L} p_j \Pr_{\mathbf{v}, Z} [j \notin \Lambda_{n+1}(\mathbf{v}, Z) \wedge j \in Z] &= \sum_{j \in L} p_j \Pr_{\mathbf{v}, Z} [j \notin \Lambda_{n+1}(\mathbf{v}, Z) \mid j \in Z] \Pr [j \in Z] \\ &= \frac{1}{2} \sum_{j \in L} p_j \Pr_{\mathbf{v}, Z} [j \notin \Lambda_{n+1}(\mathbf{v}, Z) \mid j \in Z]. \end{aligned}$$

\square

We now prove Lemma B.2 using the above two propositions. Observe that the buyers' contribution to the social welfare ALG^B extracted by $\mathbb{M}_{1\text{-supply}}$ is equal to the sum of all the buyers' utilities and all the buyers' payments.

Proof (of Lemma B.2). As just observed above, from Proposition B.1 and Proposition B.2, we have that

$$\begin{aligned} \text{ALG}^B &= \mathbb{E} \left[\sum_{i \in [n]} u_i(\mathbb{M}_{1\text{-supply}}(\mathbf{v}, \mathbf{w})) \right] + \sum_{j \in L} p_j \Pr_{\mathbf{v}, Z} [j \notin A_{n+1}(\mathbf{v}, Z) \wedge j \in Z] \\ &\geq \frac{1}{2} \sum_{j \in L} \Pr_{\mathbf{v}, Z} [j \in A_{n+1}(\mathbf{v}, Z) \mid j \in Z] p_j + \frac{1}{2} \sum_{j \in L} p_j \Pr_{\mathbf{v}, Z} [j \notin A_{n+1}(\mathbf{v}, Z) \mid j \in Z] \\ &= \frac{1}{2} \sum_{j \in L} p_j = \frac{1}{4} \sum_{j \in L} \mathbb{E} [\text{SW}_j^B(\mathbf{v})]. \end{aligned}$$

By definition of \bar{L} , for each $j \in \bar{L}$ it holds that $4\mathbb{E}[w_j] > \mathbb{E}[\text{SW}_j^B(\mathbf{v})]$. Every item in \bar{L} stays unsold so,

$$\text{ALG}^S \geq \sum_{j \in \bar{L}} \mathbb{E}[w_j] > \frac{1}{4} \sum_{j \in \bar{L}} \mathbb{E} [\text{SW}_j^B(\mathbf{v})].$$

Therefore,

$$\text{ALG}^B + \text{ALG}^S \geq \frac{1}{4} \sum_{j=1}^k \mathbb{E} [\text{SW}_j^B(\mathbf{v})].$$

Now recall that $\mathbb{E}[\text{SW}_j^B(\mathbf{v})]$ was defined by the allocation $\mathbf{X}^{\text{all}}(\mathbf{v})$, being the one returned by Algorithm A. So,

$$\frac{1}{4} \sum_{j=1}^k \mathbb{E} [\text{SW}_j^B(\mathbf{v})] = \frac{1}{4} \sum_{i=1}^n \mathbb{E}_{\mathbf{v}} [v_i(X_i^{\text{all}}(\mathbf{v}))] \geq \frac{1}{4\alpha} \text{OPT}^B.$$

□

C Proof of Theorem 5.1 and Corollary 5.1

Proposition C.1.

$$\sum_{\ell \in L} \mathbb{E} [\text{SW}_\ell^B(\mathbf{v})] + 4 \sum_{\ell \in \bar{L}} \mathbb{E} [\text{SW}_\ell^S(\mathbf{w})] > \sum_{i=1}^n \mathbb{E} [v_i(X_i^{\text{all}}(\mathbf{v}))].$$

Proof. Let $a(v_i, X_i^{\text{all}}(\mathbf{v}), \cdot)$ be the representative additive function of v_i for the bundle $X_i^{\text{all}}(\mathbf{v})$. Then,

$$\begin{aligned}
\sum_{i=1}^n \mathbb{E} \left[v_i(X_i^{\text{all}}(\mathbf{v})) \right] &= \sum_{i=1}^n \mathbb{E} \left[\sum_{\ell \in X_i^{\text{all}}(\mathbf{v})} a(v_i, X_i^{\text{all}}(\mathbf{v}), \{\ell\}) \right] \\
&= \sum_{i=1}^n \sum_{\ell=1}^k \mathbb{E} \left[a(v_i, X_i^{\text{all}}(\mathbf{v}), \{\ell\}) \mathbb{I} \left[\ell \in X_i^{\text{all}}(\mathbf{v}) \right] \right] \\
&= \sum_{\ell=1}^k \mathbb{E} \left[\text{SW}_{\ell}^B(\mathbf{v}) \right] \\
&= \sum_{\ell \in L} \mathbb{E} \left[\text{SW}_{\ell}^B(\mathbf{v}) \right] + \sum_{\ell \in \bar{L}} \mathbb{E} \left[\text{SW}_{\ell}^B(\mathbf{v}) \right] \\
&< \sum_{\ell \in L} \mathbb{E} \left[\text{SW}_{\ell}^B(\mathbf{v}) \right] + 4 \sum_{\ell \in \bar{L}} \mathbb{E} \left[\text{SW}_{\ell}^S(\mathbf{w}) \right].
\end{aligned}$$

The last inequality follows because

$$4 \sum_{\ell \in \bar{L}} \mathbb{E} \left[\text{SW}_{\ell}^S(\mathbf{w}) \right] > \sum_{\ell \in \bar{L}} \mathbb{E} \left[\text{SW}_{\ell}^B(\mathbf{v}) \right],$$

by definition of \bar{L} . □

Proposition C.2. *Let \mathbf{v} be a buyers' valuation function profile and let (X'_1, \dots, X'_n) be any allocation of items to the buyers, let $X'_{i,j} := X'_i \cap L_j$ be the set of items in L that are allocated to buyer $i \in [n]$ and belonged to seller $j \in [m]$. For each seller $j \in [m]$, let $z_j \in \{0, 1\}$ be a Bernoulli random variable such that $\mathbb{E}[z_j] = 1/2$. Let $X''_i(\mathbf{z}) := \bigcup_{j \in [m]: z_j=1} X'_{i,j}$ for all $i \in [n]$. Then, for all $i \in [n]$ it holds that*

$$\mathbb{E}_{\mathbf{z}} \left[v_i(X''_i(\mathbf{z})) \right] \geq \frac{1}{2} v_i(X'_i).$$

Moreover, given any vector $\mathbf{p} \in \mathbb{R}^k$ of item prices, the inequality also holds on the utilities of the buyers:

$$\mathbb{E}_{\mathbf{z}} \left[v_i(X''_i(\mathbf{z})) - \sum_{\ell \in X''_i(\mathbf{z})} p_{\ell} \right] \geq \frac{1}{2} \left(v_i(X'_i) - \sum_{\ell \in X'_i} p_{\ell} \right).$$

Proof. For the first claim, first note that due to subadditivity

$$\mathbb{E}_{\mathbf{z}} \left[v_i(X''_i(\mathbf{z})) \right] \geq v_i(X'_i) - \mathbb{E}_{\mathbf{z}} \left[v_i \left(\bigcup_{j \in [m]: z_j=0} X'_{i,j} \right) \right].$$

Observe that

$$\mathbb{E}_{\mathbf{z}} \left[v_i(X''_i(\mathbf{z})) \right] = \mathbb{E}_{\mathbf{z}} \left[v_i \left(\bigcup_{j \in [m]: z_j=1} X'_{i,j} \right) \right] = \mathbb{E}_{\mathbf{z}} \left[v_i \left(\bigcup_{j \in [m]: z_j=0} X'_{i,j} \right) \right],$$

because the events $z_j = 0$ and $z_j = 1$ are equiprobable for all $j \in [m]$. Combining this with the above inequality establishes the first claim.

The second claim follows from the following derivation.

$$\begin{aligned}
\mathbb{E}_{\mathbf{z}} \left[v_i(X_i''(\mathbf{z})) - \sum_{\ell \in X_i''(\mathbf{z})} p_\ell \right] &= \mathbb{E}_{\mathbf{z}} [v_i(X_i''(\mathbf{z}))] - \mathbb{E}_{\mathbf{z}} \left[\sum_{j \in [m]; z_j=1} \sum_{\ell \in X'_{i,j}} p_\ell \right] \\
&= \mathbb{E}_{\mathbf{z}} [v_i(X_i''(\mathbf{z}))] - \mathbb{E}_{\mathbf{z}} \left[\sum_{j=1}^m \left(\sum_{\ell \in X'_{i,j}} p_\ell \right) \mathbb{I}[z_j = 1] \right] \\
&= \mathbb{E}_{\mathbf{z}} [v_i(X_i''(\mathbf{z}))] - \sum_{j=1}^m \left(\sum_{\ell \in X'_{i,j}} p_\ell \right) \mathbb{E}_{\mathbf{z}} [\mathbb{I}[z_j = 1]] \\
&= \mathbb{E}_{\mathbf{z}} [v_i(X_i''(\mathbf{z}))] - \sum_{\ell \in X'_i} p_\ell \frac{1}{2} \\
&\geq \frac{1}{2} v_i(X'_i) - \frac{1}{2} \sum_{\ell \in X'_i} p_\ell
\end{aligned}$$

□

Proposition C.3. *Let $j \in [m]$ be a seller. The probability that the mechanism \mathbb{M}_{add} makes in Step 7c an offer to j that she accepts, is $1/2$.*

Proof. For every $j \in [m]$ and $\ell \in L_j$, it holds by definition of p_ℓ and L_j that $p_\ell \geq 2\mathbb{E}[w_j(\{\ell\})]$. From Markov's inequality it follows that

$$\Pr \left[w_j(S_j) > \sum_{\ell \in S_j} p_\ell \right] \leq \Pr [w_j(S_j) > 2\mathbb{E}[w_j(S_j)]] < \frac{1}{2}.$$

Thus, $\Pr [w_j(S_j) \leq \sum_{\ell \in S_j} p_\ell] \geq 1/2$, meaning that j accepts the offer with probability at least $1/2$, in case she is made an offer. The mechanism makes the offer with probability q_j , and

$$q_j \Pr \left[w_j(S_j) \leq \sum_{\ell \in S_j} p_\ell \right] = 1/2.$$

□

For $i \in [n+1]$ and valuation profile \mathbf{v} , let $\mathbf{v}_{<i} = (v_1, \dots, v_{i-1})$ and let $\Lambda_i(\mathbf{v}_{<i})$ be the set Λ_i defined in Step 5b, when \mathbb{M}_{add} is run when the buyers in $[i-1]$ have valuation profile $\mathbf{v}_{<i}$. Given this definition, the set $\Lambda_{n+1}(\mathbf{v})$ are the items not requested by any buyer at the end of Step 5, when the buyers' valuation profile is \mathbf{v} .

Lemma C.1. *The expected total utility of the buyers is at least*

$$\frac{1}{2} \sum_{\ell \in L} \Pr_{\mathbf{v}}[\ell \in \Lambda_{n+1}(\mathbf{v})] p_\ell.$$

Proof. First, let us consider a fixed buyer $i \in [n]$ and a fixed buyers' valuation profile \mathbf{v} . Let $\tilde{\mathbf{v}}_{-i}$ be an independently sampled valuation profile for the buyers in $[n] \setminus \{i\}$, and consider the bundle $X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i})$ that \mathbb{A} allocates to i when the valuation profile is $(v_i, \tilde{\mathbf{v}}_{-i})$. Let $X_i^L(\mathbf{v}, \tilde{\mathbf{v}}_{-i}) =$

$X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}) \cap L \cap \Lambda_i(\mathbf{v}_{<i})$. Moreover, let \mathbf{z} be a vector of m Bernoulli random variables with $\mathbb{E}[z_j] = 1/2$ and define for a subset $S(\mathbf{v}) \subseteq \Lambda_i(\mathbf{v})$ the random variable $S(\mathbf{v}, \mathbf{z}) = \bigcup_{j \in [m]: z_j=1} (S \cap L_j)$. Particularly, from this definition we obtain the random variable $X_i(\mathbf{v}, \tilde{\mathbf{v}}_{-i}, \mathbf{z}) = \bigcup_{j \in [m]: z_j=1} (X_i^L(\mathbf{v}, \tilde{\mathbf{v}}_{-i}) \cap L_j)$. Also, note that when the buyers' valuations are \mathbf{v} , the mechanism will let i choose to request a bundle from the set $\Lambda_i(\mathbf{v}_{<i})$ with item prices \mathbf{p} , the buyer maximizes her expected utility and will therefore request the bundle $B_i(\mathbf{v})$ that maximizes her expected utility which can be expressed as

$$\mathbb{E}_{\mathbf{z}} \left[v_i(B(\mathbf{v}, \mathbf{z})) - \sum_{\ell \in B(\mathbf{v}, \mathbf{z})} p_\ell \right],$$

as by Propostion C.3 each sellers' requested items will be allocated with probability 1/2, as reflected by the Bernoulli variables \mathbf{z} .

Since $B_i(\mathbf{v})$ is an expected-utility-maximising bundle $X_{i,j}(\mathbf{v}, \tilde{\mathbf{v}}_{-i}, \mathbf{z}) \subseteq \Lambda_i(\mathbf{v})$, it holds that

$$\begin{aligned} \mathbb{E}_{\mathbf{z}} \left[v_i(B(\mathbf{v}, \mathbf{z})) - \sum_{\ell \in B(\mathbf{v}, \mathbf{z})} p_\ell \right] &\geq \mathbb{E}_{\tilde{\mathbf{v}}_{-i}, \mathbf{z}} \left[v_i(X_i(\mathbf{v}, \tilde{\mathbf{v}}_{-i}, \mathbf{z})) - \sum_{\ell \in X_i(\mathbf{v}, \tilde{\mathbf{v}}_{-i}, \mathbf{z})} p_\ell \right] \\ &\geq \frac{1}{2} \mathbb{E}_{\tilde{\mathbf{v}}_{-i}} \left[v_i(X_i^L(\mathbf{v}, \tilde{\mathbf{v}}_{-i})) - \sum_{\ell \in X_i^L(\mathbf{v}, \tilde{\mathbf{v}}_{-i})} p_\ell \right] \\ &\geq \frac{1}{2} \mathbb{E}_{\tilde{\mathbf{v}}_{-i}} \left[a(v_i, X_i^{\text{all}}(v_i, \mathbf{v}_{-i}), X_i^L(\mathbf{v}, \mathbf{v}_{-i})) - \sum_{\ell \in X_i^L(\mathbf{v}, \tilde{\mathbf{v}}_{-i})} p_\ell \right] \\ &= \frac{1}{2} \mathbb{E}_{\tilde{\mathbf{v}}_{-i}} \left[\sum_{\ell \in X_i^L(\mathbf{v}, \tilde{\mathbf{v}}_{-i})} (\text{SW}_\ell^B(v_i, \tilde{\mathbf{v}}_{-i}) - p_\ell) \right]. \end{aligned}$$

where the second inequality follows from Proposition C.2, and the last inequality follows from the definition of the additive representative function $a(v_i, X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}, \cdot))$.

If we sum over all $i \in [n]$ and take the expectation w.r.t. every v_i , we obtain the following bound on the total expected utility of the buyers.

$$\begin{aligned} \mathbb{E}_{\mathbf{v}, \mathbf{z}} \left[\sum_{i=1}^n (v_i(B(\mathbf{v}, \mathbf{z})) - \sum_{\ell \in B(\mathbf{v}, \mathbf{z})} p_\ell) \right] &\geq \frac{1}{2} \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}_{-i}} \left[\sum_{i=1}^n \sum_{\ell \in X_i^L(\mathbf{v}, \tilde{\mathbf{v}}_{-i})} (\text{SW}_\ell^B(v_i, \tilde{\mathbf{v}}_{-i}) - p_\ell) \right] \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}_{-i}} \left[\sum_{i=1}^n \sum_{\ell \in L} (\text{SW}_\ell^B(v_i, \tilde{\mathbf{v}}_{-i}) - p_\ell) \mathbb{I}[\ell \in X_i^L(\mathbf{v}, \tilde{\mathbf{v}}_{-i})] \right] \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}_{-i}} \left[\sum_{i=1}^n \sum_{\ell \in L} (\text{SW}_\ell^B(v_i, \tilde{\mathbf{v}}_{-i}) - p_\ell) \mathbb{I}[\ell \in X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i})] \mathbb{I}[\ell \in \Lambda_i(\mathbf{v}_{<i})] \right] \\ &= \frac{1}{2} \sum_{\ell \in L} \sum_{i=1}^n \mathbb{E}_{v_i, \tilde{\mathbf{v}}_{-i}} \left[(\text{SW}_\ell^B(v_i, \tilde{\mathbf{v}}_{-i}) - p_\ell) \mathbb{I}[\ell \in X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i})] \right] \mathbb{E}_{\mathbf{v}_{-i}} [\mathbb{I}[\ell \in \Lambda_i(\mathbf{v}_{<i})]]. \end{aligned}$$

For the second-to-last equality, we exploited the independence of the events $(\ell \in X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}))$ and $(\ell \in \Lambda_i(\mathbf{v}_{<i}))$. Then, $\mathbb{E}_{\mathbf{v}_{-i}} [\mathbb{I}[\ell \in \Lambda_i(\mathbf{v}_{<i})]] = \Pr[\ell \in \Lambda_i(\mathbf{v}_{<i})]$ and since $L = \Lambda_1(\mathbf{v}_{<1}) \supseteq \dots \supseteq \Lambda_{n+1}(\mathbf{v})$, it holds that $\Pr[\ell \in \Lambda_i(\mathbf{v}_{<i})] \geq \Pr[\ell \in \Lambda_{n+1}(\mathbf{v})]$. So, we have that the above expression is

at least

$$\begin{aligned} & \frac{1}{2} \sum_{\ell \in L} \Pr_{\mathbf{v}} [\ell \in \Lambda_{n+1}(\mathbf{v})] \sum_{i=1}^n \mathbb{E}_{v_i, \tilde{v}_{-i}} \left[(\text{SW}_{\ell}^B(v_i, \tilde{v}_{-i}) - p_{\ell}) \mathbb{I} [\ell \in X_i^{\text{all}}(v_i, \tilde{v}_{-i})] \right] \\ &= \frac{1}{2} \sum_{\ell \in L} \Pr_{\mathbf{v}} [\ell \in \Lambda_{n+1}(\mathbf{v})] \sum_{i=1}^n \mathbb{E}_{\mathbf{v}} \left[(\text{SW}_{\ell}^B(\mathbf{v}) - p_{\ell}) \mathbb{I} [\ell \in X_i^{\text{all}}(\mathbf{v})] \right]. \end{aligned}$$

The equality follows from renaming the random variable $v_j := \tilde{v}_j$ for all $j \neq i$. Now observe that by definition of the prices, $p_{\ell} = \sum_{i=1}^n \mathbb{E}_{\mathbf{v}} [(\text{SW}_{\ell}^B(\mathbf{v}) - p_{\ell}) \mathbb{I} [\ell \in X_i^{\text{all}}(\mathbf{v})]]$. Combining these derivations, we obtain the desired bound on the expected utilities

$$\mathbb{E}_{\mathbf{v}, \mathbf{z}} \left[\sum_{i=1}^n (v_i(B(\mathbf{v}, \mathbf{z}))) - \sum_{\ell \in B(\mathbf{v}, \mathbf{z})} p_{\ell} \right] \geq \frac{1}{2} \sum_{\ell \in L} \Pr_{\mathbf{v}} [\ell \in \Lambda_{n+1}(\mathbf{v})] p_{\ell}.$$

□

Lemma C.2. *The expected sum of payments made by the buyers is equal to*

$$\frac{1}{2} \sum_{\ell \in L} \Pr_{\mathbf{v}} [\ell \notin \Lambda_{n+1}(\mathbf{v})] p_{\ell}.$$

Proof. For $j \in [m]$, let z_j be the random $(0, 1)$ -variable that indicates whether seller j has been made an offer and accepted it in Step 7c of Mechanism \mathbb{M}_{add} , so $z_j = 1$ is a Bernoulli variable with expected value $1/2$. The expected sum of payments made by the buyers is then

$$\begin{aligned} \sum_{j=1}^m \sum_{\ell \in L_j} \Pr [\ell \notin \Lambda_{n+1}(\mathbf{v}) \wedge z_j = 1] p_{\ell} &= \sum_{j=1}^m \sum_{\ell \in L_j} \Pr [\ell \notin \Lambda_{n+1}(\mathbf{v})] \Pr [z_j = 1] p_{\ell} \\ &= \frac{1}{2} \sum_{j=1}^m \sum_{\ell \in L_j} \Pr [\ell \notin \Lambda_{n+1}(\mathbf{v})] p_{\ell} \\ &= \frac{1}{2} \sum_{\ell \in L} \Pr [\ell \notin \Lambda_{n+1}(\mathbf{v})] p_{\ell} \end{aligned}$$

The second equality holds by the independence of the two events. □

Lemma C.3.

$$\text{ALG}^B \geq \frac{1}{4} \sum_{\ell \in L} \mathbb{E}_{\mathbf{v}} [\text{SW}_{\ell}^B(\mathbf{v})].$$

Proof. The expected social welfare contribution of the buyers is equal to the sum of the expected utilities and expected payments. By the above two lemmas, their sum is at least

$$\frac{1}{2} \sum_{\ell \in L} \Pr [\ell \in \Lambda_{n+1}(\mathbf{v})] p_{\ell} + \frac{1}{2} \sum_{\ell \in L} \Pr [\ell \notin \Lambda_{n+1}(\mathbf{v})] p_{\ell} = \frac{1}{2} \sum_{\ell \in L} p_{\ell} = \frac{1}{4} \sum_{\ell \in L} \mathbb{E}_{\mathbf{v}} [\text{SW}_{\ell}^B(\mathbf{v})],$$

by definition of p_{ℓ} . □

Lemma C.4.

$$4\alpha \text{ALG}^B + 4\alpha \text{ALG}^S \geq \text{OPT}^B.$$

Proof. By the above lemma, $4\text{ALG}^B \geq \sum_{\ell \in L} \mathbb{E}_{\mathbf{v}}[\text{SW}_{\ell}^B(\mathbf{v})]$. Moreover, our mechanism leaves every item $\ell \in \bar{L}$ with its seller, and so $4\text{ALG}^S \geq 4 \sum_{\ell \in \bar{L}} \mathbb{E}_{\mathbf{w}}[\text{SW}_{\ell}^S(\mathbf{v})]$. Therefore,

$$4\text{ALG}^B + 4\text{ALG}^S \geq \sum_{\ell \in L} \mathbb{E}_{\mathbf{v}} \text{SW}_{\ell}^B(\mathbf{v}) + 4 \sum_{\ell \in \bar{L}} \mathbb{E}_{\mathbf{w}}[\text{SW}_{\ell}^S(\mathbf{v})] \geq \sum_{i=1}^n \mathbb{E}_{\mathbf{v}} v_i(X_i^{\text{all}}(\mathbf{v})) \geq \frac{1}{\alpha} \text{OPT}^B,$$

The second inequality holds by Proposition C.1, and the last inequality follows because we defined α to be the approximation factor of algorithm \mathbb{A} , which is the algorithm that we assumed to generate allocation $X^{\text{all}}(\mathbf{v})$. □

Lemma C.5.

$$2\text{ALG}^S \geq \text{OPT}^S.$$

Proof. The only items that our mechanisms potentially reallocates are the ones belonging to L . Every item in \bar{L} stays with its seller. For the items in L , the mechanism ensures every seller sells her demanded bundle with probability exactly 1/2, so for each seller it holds that she retains her full initial endowment with probability at least 1/2, which implies the claim. □

Proof (of Theorem 5.1). Every buyer chooses a bundle that maximizes her expected utility, so the mechanism is ex-interim IR and BIC on the buyers' side. On the sellers' side, it is actually ex-post IR and DSIC: the sellers solely have to decide between accepting or rejecting a single offer to receive a proposed payment in exchange for a bundle of items, and it is clearly a dominant strategy to accept if and only if such an exchange leads to an improvement in the seller's utility. The fact that the mechanism is DSBB follows from its definition, which makes clear that payments are defined by the appropriate sequence of trades and payments from buyers to sellers. The approximation guarantee follows by the sum of the inequalities of the above two lemmas. □

By taking for \mathbb{A} an optimal algorithm (i.e., $\alpha = 1$), we obtain the existence of a mechanism that is ex-post IR, DSIC, DSBB, and 6-approximates the optimal social welfare. Again, one may also take for \mathbb{A} a polynomial time approximation algorithm in order to obtain a polynomial time mechanism.

Proof (of Corollary 5.1). If a buyer $i \in [n]$ has an additive valuation function, it is a dominant strategy to request the items in $A_i(\mathbf{v}_{<i})$ for which it holds that $v_i(\{\ell\}) > p_{\ell}$. This follows from the simple fact that by additivity, the utility that a player has for any bundle of items S can be written as $\sum_{\ell \in S} v_i(\{\ell\}) - p_{\ell}$. Thus, for every item $\ell \in [k]$ that a buyer requests (recall that this item is then allocated to her for price p_{ℓ} with probability 1/2), a term of $(1/2)(v_i(\{\ell\}) - p_{\ell})$ gets added to her expected utility. So including ℓ in her requested bundle is profitable if and only if $v_i(\{\ell\}) - p_{\ell} \geq 0$. Ex-post IR property is also satisfied by following this strategy. □