

# Approximately Efficient Two-Sided Combinatorial Auctions

RICCARDO COLINI-BALDESCHI, Dipartimento di Economia e Finanza, LUISS

PAUL W. GOLDBERG, University of Oxford

BART DE KEIJZER, Centrum Wiskunde & Informatica (CWI)

STEFANO LEONARDI, Sapienza, University of Rome

TIM ROUGHGARDEN, Stanford University

STEFANO TURCHETTA, KPMG Italy and University of Oxford

---

We develop and extend a line of recent work on the design of mechanisms for two-sided markets. The markets we consider consist of buyers and sellers of a number of items, and the aim of a mechanism is to improve the social welfare by arranging purchases and sales of the items. A mechanism is given prior distributions on the agents' valuations of the items, but not the actual valuations; thus the aim is to maximise the expected social welfare over these distributions. As in previous work, we are interested in the worst-case ratio between the social welfare achieved by a truthful mechanism, and the best social welfare possible.

Our main result is an incentive compatible and budget balanced constant-factor approximation mechanism in a setting where buyers have XOS valuations and sellers' valuations are additive. This is the first such approximation mechanism for a two-sided market setting where the agents have combinatorial valuation functions. To achieve this result, we introduce a more general kind of demand query that seems to be needed in this situation. In the simpler case that sellers have unit supply (each having just one item to sell), we give a new mechanism whose welfare guarantee improves on a recent one in the literature. We also introduce a more demanding version of the *strong budget balance* (SBB) criterion, aimed at ruling out certain "unnatural" transactions satisfied by SBB. We show that the stronger version is satisfied by our mechanisms.

Additional Key Words and Phrases: Auctions; mechanism design; market intermediation; two-sided markets; pricing; approximation algorithms

## 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 [4, 13, 23] that optimises the social welfare 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). The IR requirement demands that participating in the mechanism is not harmful to any agent. The DSIC requirement demands 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 turned to the problems that arise in two-sided markets, in which the set of agents is partitioned into *buyers* and *sellers*. In contrast with the one-sided setting (where one could say that the mechanism itself initially holds the items), in the two-sided setting the items

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

EC17, June 26–30, 2017, Cambridge, MA, USA. Copyright 2017 ACM ISBN 978-1-4503-4527-9/17/06...\$15.00

© 2017 ACM. ACM ISBN 978-1-4503-4527-9/17/06

<http://dx.doi.org/10.1145/3033274.3085128>

are initially held by the sellers, who have 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. The growing interest in two-sided markets can be attributed to various important applications. Relevant examples are 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 sellers are allowed to trade without leaving to the mechanism any share of the payments, and without the mechanism adding money to 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. However, it is known from the work of Myerson and Satterthwaite [17] that it is generally impossible for an IR, BIC, and WBB mechanism to maximise social welfare in such a market, even in the *bilateral trade* setting, i.e., when there is just one seller and one buyer.<sup>1</sup>

The practical contexts noted above need the application of two-sided market mechanisms that can work in a *combinatorial* setting, i.e., where there are multiple distinct items in the market and agents having possibly complex valuations over the subsets of items that they may receive. However, we are not aware of any such mechanism that approximates the social welfare while meeting the IR, DSIC 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 results extend and improve on previous work which targeted an important special case of a two-sided market: 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 utility for holding an item. A mechanism for this setting is known in the literature as a *double auction*. The goal of several works on double auctions [15, 20, 21] has been that of trading off the achievable social welfare with the strength of the incentive compatibility and budget balance constraints. In our present work, we investigate this question for the much more general class of combinatorial two-sided markets.

## 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 optimise their (quasi-linear) utility, which is given by the valuation of the set of items that the mechanism allocates to an agent, minus the payment that the mechanism collects from the agent. A seller will typically receive money (instead of pay money), which we treat as 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

---

<sup>1</sup>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.

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 maximise 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 (satisfying IR, SBB, and DSIC), that is guaranteed to always output an optimal allocation.

We are interested in mechanisms that satisfy IR, SBB, and DSIC (or failing that, the weaker notion of *Bayesian incentive compatibility (BIC)*), 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. In contrast with one-sided combinatorial auction design (where the main challenge is polynomial-time implementability), for the two-sided case our primary goal is to design (and thus show the existence of) IR, SBB, and DSIC/BIC mechanisms that  $O(1)$ -approximate OPT. Such mechanisms circumvent the aforementioned impossibility result of Myerson and Satterthwaite [17] by weakening the requirement of optimal social welfare to that of *approximately* optimal social welfare (while nonetheless strengthening the WBB constraint into SBB).

## 1.2 Our Results and their Significance

The present paper starts off by showing that there is a straightforward technical trick that one may apply to turn any WBB mechanism into an SBB one, with a small loss in the approximation factor. Technically, one could e.g. apply it to the WBB mechanism of Blumrosen and Dobzinski [1] for combinatorial exchange markets; however, the trick is unsatisfactory in practice as it essentially consists of giving the leftover money to a random agent. This demonstrates a weakness in the current definition of SBB, which motivates the introduction of a strengthened version, that we call *direct-trade strong budget balance (DSBB)*.

Our goal is the design of individually rational, incentive compatible, and direct-trade strongly budget balanced mechanisms for combinatorial two-sided markets, that achieve a constant approximation to the optimal social welfare. We present two mechanisms adhering to these constraints for general families of combinatorial two-sided markets, as summarized in the table below.

Mechanism	Buyers' valuations	Sellers' valuations	Approximation ratio	IR	IC	BB
$\mathbb{M}_{1\text{-supply}}$	XOS	unit-supply	6	ex-post IR	DSIC	DSBB
$\mathbb{M}_{\text{add}}$	XOS	additive	6	interim IR	BIC	DSBB
$\mathbb{M}_{\text{add}}$	additive	additive	6	ex-post IR	DSIC	DSBB

Table 1. Summary of our results.

Our  $\mathbb{M}_{1\text{-supply}}$  mechanism handles the setting where all sellers have a single item for sale, and buyers have *fractionally subadditive* (XOS) valuation functions over the set of items in the market. Our  $\mathbb{M}_{\text{add}}$  mechanism can handle the more general case where sellers have multiple items for sale and have additive valuation functions over the items they possess, though  $\mathbb{M}_{\text{add}}$  satisfies weaker IC and IR notions than  $\mathbb{M}_{1\text{-supply}}$ . More precisely,  $\mathbb{M}_{\text{add}}$  is DSIC and IR on the sellers' side and BIC and interim-IR on the buyers' side. However, for the special case where buyers have additive

valuation functions,  $\mathbb{M}_{\text{add}}$  does satisfy the stronger IC and IR notions for both buyers and sellers. In all three cases, DSBB is satisfied (a strengthened variant of SBB), and our mechanisms achieve an  $O(1)$ -approximation to the optimal social welfare.

To our knowledge, these are the first mechanisms for combinatorial two-sided markets that simultaneously are IC, (D)SBB, IR, and approximate the optimal social welfare to within a constant factor. Notice that with non-unit-supply sellers, a constant approximation was not previously known even in the context of WBB or standard SBB.<sup>2</sup> Furthermore, we note that our mechanisms not only work for a more general setting than that of [5], but also improve the approximation ratio for double auctions from 16 to 6.

In the case of  $\mathbb{M}_{\text{add}}$ , buyers are required to answer a generalised type of demand query, in which the mechanism gives prices for the items, and asks a buyer which bundle she would like *if, for each item in that bundle, she were to receive it with probability 1/2*. Our usage of these queries could be criticised for imposing an excessive cognitive burden on the agents. Although we are not concerned here with that issue (we model agents as computationally unbounded as well as rational), our apparent need for such queries highlights the general question of how agents' computational limitations affect what outcomes can be achieved.<sup>3</sup>

### 1.3 Overview of the Techniques

The main challenge in two-sided market design is to find prices that stimulate truthful behavior and are suitable for both buyers and sellers, which have contrasting interests. In fact, even in the simplest imaginable setting – the bilateral trade – it is impossible to design a socially efficient mechanism satisfying IR, BIC and WBB [17].

A first feature all our mechanisms share to guarantee DSBB is being a generalised version of two-sided *sequential posted price mechanisms (SPMs)* [5] for double auctions to combinatorial two-sided markets. These mechanisms assign fixed, pre-computed prices to each item so that these prices are the only ones for which the items can be traded. This yields a sequence of bilateral trades in which the amount paid by the buyer equals the amount received by the seller.

While one-sided SPMs provide IR and IC for free, two-sided SPMs require additional conditions to be met. In combinatorial two-sided markets, if prices are fixed for every single item, it cannot be guaranteed that a bundle of items chosen by a buyer will surely be allocated to her, in case at that point the corresponding seller has not been queried yet about her willingness to sell the item. Symmetrically, when a seller would communicate to an SPM mechanism which bundle of items she is willing to sell given the proposed item prices, then the mechanism cannot guarantee to the seller that this bundle will surely be traded in case it has not yet queried the buyers which item sets they demand. The situation is further complicated by the fact that there may exist strong interdependencies among items within an agent's valuation function, which implies that the choice of bundle that a buyer requests (or that a seller makes available) depends strongly on the set of items that the sellers are prepared to sell (or that the buyers request). Therefore, a mechanism designer needs to be careful in proposing prices that are suitable for both sides of the market, and needs to be particularly careful in selecting the side of the market to process first. The choice that the mechanism made here can depend crucially on the types of valuation functions of the agents. Indeed, one main difference between our two mechanisms is the order in which we process each side of the market. Anyway all the mechanisms proposed in this paper are oblivious to the order in which sellers and buyers are presented.

<sup>2</sup>The mechanism proposed in [1] achieves a constant approximation to the optimal social welfare if the size of the initial endowment of each agent is bounded by a constant; otherwise the approximation factor is of logarithmic order.

<sup>3</sup>This question also applies to standard demand queries [10], which may be computationally hard to answer or may involve a high communication complexity, depending on the computational model used and on the way in which the valuation functions are represented.

To additionally achieve a mechanism that results in a high social welfare, we exclude some items from trade and introduce randomness into the mechanism. The main idea is to suppose that all the items are available to the set of buyers as in a one-sided auction, and to compute the *expected marginal contribution of an item to the social welfare* [11] under this assumption. Then, the mechanism compares this contribution to the seller's value for the item: if the seller's value is much higher, then we exclude the item from trade and leave it with the seller. Thus, the mechanism only trades items that are of relatively high expected value to the buyers' side of the market.

To estimate the expected marginal contribution of an item to the social welfare,  $\mathbb{M}_{1\text{-supply}}$  and  $\mathbb{M}_{\text{add}}$  make use of an algorithm  $\mathbb{A}$  that, given a buyers' valuation profile and a set of items, allocates the items to the buyers, without considering the sellers and their valuations. If one is not interested in achieving a low runtime, one can take  $\mathbb{A}$  to be an exact algorithm that outputs an optimal allocation. Alternatively, by using a technique of Feldman et al. [11], one may take for  $\mathbb{A}$  a polynomial time approximation algorithm and combine this with sampling a sufficiently number of valuation profiles from the distribution, in order to estimate the expected marginal contribution of an item to the social welfare accurately in polynomial time. This yields polynomial-time implementable approximation mechanisms. In particular, in case  $\mathbb{A}$  is a polynomial time  $\alpha$ -approximation algorithm, it will run within time  $POLY(1/\epsilon, n, m)$ , and approximate the optimal social welfare within an  $O(\alpha)$  multiplicative factor and an  $\epsilon$  additive term, where  $\epsilon$  is a parameter that results from the sampling procedure. This technique is described in further detail in [11] and works for distributions with bounded support.

Randomness is added to make sure every seller independently sells her bundle of items with a fixed probability of  $1/2$ ; which is used to bound the social welfare loss on both sides of the market by no more than a constant factor.

#### 1.4 Related Work

Due to the impossibility result of Myerson and Satterthwaite [17], no two-sided mechanism can simultaneously achieve optimal social welfare and satisfy the BIC, IR, WBB constraints, even in the simple bilateral trade setting. Follow-up work thus had to focus on designing mechanisms that trade off among these properties.

The following papers of the Economics literature studied the convergence rate to the optimal social welfare 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 [12] showed that duplicating the number of agents by  $\tau$  results in a market where the optimal IR, IC, WBB mechanism's expected social welfare approximation factor approaches 1 at a rate of  $O(\log \tau / \tau^2)$ . Rustichini et al. [18] and Satterthwaite and Williams [21] investigated a family of non-IC double auctions, and study the inefficiency and the extent to which agents misreport their valuations in these double auctions. We remark that these results only hold for unit-demand buyers and unit-supply sellers, identical valuation distributions, and the hidden constants in these asymptotic results depend on the specific valuation distributions. In contrast, our interest is in finding *universal* constant approximation guarantees for combinatorial settings and not necessarily identical distributions.

In McAfee [15], an IC, WBB, IR double auction is proposed that extracts at least a  $(1 - 1/\ell)$  fraction of the maximum social welfare, where  $\ell$  is the number of traders in the optimal solution.

Optimal revenue-maximising Bayesian auctions were characterized in Myerson [16], 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 the work of Deng et al. [6], which studied maximising the auctioneer's revenue in Bayesian double auctions. The same objective was

studied by Deshmukh et al. [7] yet in the *prior-free* model. In [22], 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 by Sandholm and Gilpin [19] and have 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 by Chawla et al. [3] for the objective of revenue maximisation. Additionally, Kleinberg and Weinberg [14] and Dütting and Kleinberg [8] strengthen these results further. Very relevant to our work is the paper of Feldman et al. [11], 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.

A line of recent work addressed the problem of approximating social welfare in double auctions and related problems under the WBB requirement. Dütting et al. [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 [15]. 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.

## 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.

### 2.1 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  $\mathbb{R}_{\geq 0}$ .<sup>4</sup> 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 called the *initial endowment*, where  $I_j$  is the set of items that is initially held by seller  $j \in [m]$ . 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: The valuation function of buyer  $i \in [n]$  is drawn from distribution  $G_i$ , and similarly the valuation function of seller  $j \in [m]$  is drawn from distribution  $F_j$ .

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

<sup>4</sup>By a *monotone* valuation function  $v$  we mean that  $v(S) \geq v(T)$  for all sets of items  $T \subseteq S$ . That is, getting more items cannot decrease an agent's overall valuation. By *normalized* we mean that  $v(\emptyset) = 0$ .

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  $(n, k, I, F)$ .

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  $\forall S \subseteq [k]$ ,  $w_j(S) = w_j(S \cap I_j)$ . Conversely, 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  $\ell$  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.

## 2.2 Mechanism Design Goals

The following discussion is specific to two-sided markets (the main focus of this paper), but these concepts can be extended straightforwardly to combinatorial exchange markets. Given a two-sided market, our aim is to redistribute the items among the agents so as to maximise the *social welfare* (the sum of the agents' valuations). An *allocation* for a two-sided market  $(n, m, k, I, G, F)$  is a pair of vectors  $(X, Y) = ((X_1, \dots, X_n), (Y_1, \dots, Y_m))$  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 disjoint. When discussing a given two-sided market, we will denote by  $\mathcal{A}$  the set of all allocations for that market.

Redistribution of the items is done by running a *mechanism*  $\mathbb{M}$ . A mechanism interacts with and receives input from the agents, and outputs an *outcome*, consisting of an allocation  $(X, Y)$  and a payment vector  $(\rho^B, \rho^S) \in \mathbb{R}^n \times \mathbb{R}^m$ , where  $\rho^B$  refers to the buyers' vector of payments and  $\rho^S$  to the sellers' one. An outcome is therefore a tuple  $(X, Y, \rho^B, \rho^S)$ . Note that when an agent is charged a negative payment, this should be interpreted as an agent receiving money. The payment of a seller is usually negative in a reasonable two-sided market mechanism, and this is also the case for the mechanisms proposed in the present paper.

Agents are assumed to maximise their *utility*, which is defined as the valuation for the bundle of items that they possess with respect to the allocation vector, minus the payment charged by the mechanism. In particular, the utility  $u_i^B(v_i, (X, Y, \rho^B, \rho^S))$  of a buyer  $i \in [n]$  with valuation function  $v_i$  is  $v_i(X_i) - \rho_i^B$ , whereas for a seller  $j \in [m]$  with valuation function  $w_j$  it is  $u_j^S(w_j, (X, Y, \rho^B, \rho^S)) = w_j(Y_j) - \rho_j^S$ .<sup>5</sup>

Furthermore, agents are assumed to be fully rational, so that they will strategically interact with the mechanism to achieve their goal of maximising utility. Our goal is to design a 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 an allocation  $(X, Y)$ , the social welfare  $SW(X, Y)$  is defined as

$$SW(X, Y) = \sum_{i \in [n]} v_i(X_i) + \sum_{j \in [m]} w_j(Y_j).$$

We now describe three main economic properties our mechanisms must satisfy. For each of these constraints we first introduce the strictest version and then a more relaxed one. We aim to satisfy the strictest versions, whenever possible.

- **Incentive compatibility (IC)**<sup>6</sup>

<sup>5</sup>Note that  $Y_j$  represents the bundle of items that remains in the seller's possession after execution of the mechanism.

<sup>6</sup>Technically, as can be inferred, the DSIC properties are reserved for *direct revelation mechanisms*, i.e., where the buyer solely interacts with the mechanism reporting her 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, the DSIC and BIC definitions naturally extend to non-direct revelation mechanisms.

- *Dominant strategy incentive compatibility (DSIC)*: It is a dominant strategy for every agent to report her true valuation sincerely. I.e., for every agent  $i$  and for every vector of valuations of all other players, it is impossible for agent  $i$  to increase her expected utility by misreporting her valuation.
- *Bayesian incentive compatibility (BIC)*: It is a Bayes-Nash equilibrium (BNE) for the agents to truthfully report their valuations to the mechanism. I.e., each agent  $i$  maximizes her expected utility by truthfully reporting her valuation if all other players also truthfully report their valuations.
- **Individual rationality (IR)**
  - *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 a utility that is not less than her initial utility. (The initial utility of a seller with bundle  $I_j$  is  $w_j(I_j)$ , and the initial utility of a buyer is  $v_i(\emptyset) = 0$ .)
  - *Interim individual rationality (interim IR)*: There is a strategy for each agent that yields her an *expected* utility that is not less than her initial utility (where expectation is over the random outcome of the mechanism, resulting from internal randomness of the mechanism and randomized strategies adopted by the agents).
- **Budget balance (BB)**
  - *Strong Budget Balance (SBB)*: The sum of all agents' payments output by the mechanism is equal to zero. Conceptually, this means that no money 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 generally means that the buyers' payments are at least as large as the payments received by the sellers. No external party needs to subsidise the mechanism.

For valuation profiles  $(\mathbf{v}, \mathbf{w})$ ,  $\text{OPT}(\mathbf{v}, \mathbf{w}) := \max\{\text{SW}(X, Y) : (X, Y) \in \mathcal{A}\}$  denotes the *optimal social welfare*. 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  $\mathbb{M}$   $\alpha$ -*approximates the optimal social welfare* for some  $\alpha > 1$  if and only if  $\text{OPT} \leq \alpha \mathbb{E}_{\mathbf{v}, \mathbf{w}}[\text{SW}(\mathbb{M}(\mathbf{v}, \mathbf{w}))]$ . Our goal is to find mechanisms that  $\alpha$ -approximate the optimal social welfare for a low  $\alpha$ , are DSIC (or BIC), SBB, and ex-post IR (or interim IR).

### 2.3 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* if and only if there exist 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)* if and only if there exists a collection of additive functions  $a_1, \dots, a_d$  such that for every bundle  $S \subseteq [k]$  it holds that  $v(S) = \max_{i \in [d]} a_i(S)$ .
- $v$  is *subadditive* if and only if for all  $S, T \subseteq [k]$  it holds that  $v(S \cup T) \leq v(S) + v(T)$ .

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, and XOS functions are contained in the class of the subadditive functions.

## 3 AN INITIAL MECHANISM AND DIRECT TRADE STRONG BUDGET BALANCE

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

**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  $(n, k, I, F)$  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  $\mathbb{M}_{sbb}$  that selects an agent  $i \in [n]$  uniformly at random, runs  $\mathbb{M}_{bd}$  on the remaining agents, and allocates the surplus money of  $\mathbb{M}_{bd}$  to agent  $i$ . We are then able to prove 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.<sup>7</sup>*

Because of space constraints, the proof of this theorem and of various other results in the remainder of this paper, has been omitted. 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 an SBB one, while 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 3.3.* A mechanism for an exchange market satisfies *direct trade strong budget balance* (DSBB) if and only if 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.

DSBB strengthens the traditional SBB notion and seems to be a reasonable requirement in most two sided markets and exchange markets settings. Note that the way in which we strengthen SBB is rather mild: DSBB still allows an arbitrarily large amount of money to be transferred from one agent to another as long as at least one item is exchanged in the opposite direction. DSBB does not even require such a bilateral exchange to be profitable for both parties, but does nonetheless seem to rule out the rather unsatisfactory type construction such as the one used in  $\mathbb{M}_{sbb}$ .

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.<sup>8</sup> Moreover, two of our results provide an  $O(1)$ -approximation even in settings where  $\mathbb{M}_{sbb}$  would only provide an approximation factor of logarithmic order.

#### 4 A MECHANISM FOR UNIT-SUPPLY SELLERS AND XOS BUYERS

In this section we present a DSIC, ex-post IR, and DSBB mechanism for two-sided markets, when sellers initially possess a single item and buyers have XOS valuation functions. This mechanism achieves a constant approximation to the optimal social welfare. In this setting, we use  $[k]$  to denote 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  $(n, k, k, I, G, F)$  is a given two-sided market, on which we run the mechanism to be defined. For an allocation  $(X, Y) \in \mathcal{A}$ , we shall use the notation

<sup>7</sup>For two agents, it is straightforward to come up with alternative mechanisms that have the desired properties.

<sup>8</sup>We note that the double auctions given in [5] also satisfy the DSBB property.

$SW^B, SW^S$  to respectively denote the buyers' and the sellers' contribution to the social welfare, i.e.,

$$\begin{aligned} SW^B(X, Y) &:= \sum_{i=1}^n v_i(X_i), \\ SW^S(X, Y) &:= \sum_{i=1}^k w_j(Y_j) = \sum_{j=1}^k w_j \mathbb{I}[j \in Y_j]. \end{aligned} \tag{1}$$

Our mechanism requires fixing a price for every item in the market. For a bundle of available items  $\Lambda$  and an item price vector  $\mathbf{p} = (p_1, \dots, p_k) \in \mathbb{R}_{\geq 0}^k$ , we define the *demand correspondence* of buyer  $i \in [n]$  with valuation function  $v_i$  as

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

i.e.,  $\mathcal{D}(v_i, \mathbf{p}, \Lambda)$  is the set of bundles of items in  $\Lambda$  that maximise  $i$ 's utility under the given item prices.

For a buyer  $i$  with valuation function  $v_i$ , we define the *additive representative function* for bundle  $T \subseteq [k]$  as any 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  $i$  and for each valuation function in the support of  $G_i$ , by the definition of XOS functions.

#### 4.1 Mechanism

Let  $\mathbb{A}$  be an algorithm that, given a buyers' valuation profile  $\mathbf{v}$  and a set of items  $[k]$ , allocates the items to the buyers, without considering the sellers and their valuations.  $\mathbb{A}$  can either be an exact algorithm that outputs an optimal allocation of  $[k]$  to the buyers (if one is not interested in the runtime) or an approximately optimal one (in the case that one insists on polynomial-time implementability). Our mechanism uses  $\mathbb{A}$  as a black-box for the computation of item prices.

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

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

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*  $SW(X_i^{\text{all}}(\mathbf{v}))$  if and only if  $\mathbb{E}_{\mathbf{v}}[SW_j^B(\mathbf{v})] \geq 4\mathbb{E}[w_j]$ , 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  $j$  with its seller.

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

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

Observe that  $p_j \geq 2\mathbb{E}[w_j]$  for all  $j \in H$ , by our definition of high welfare items.

The reason why  $H$  is chosen in such a way is twofold: first, the items in  $L$  if kept by their sellers provide a welfare loss of at most a constant factor; second, every item in  $H$  is guaranteed to be sold (if sold) at a high price, to make sure that the buyer receiving the item has a high valuation for it.

Our (randomized) mechanism does the following. First, it considers the sellers with an item in  $H$  (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  $j \in H$  accepts the price with probability at least  $1/2$ , by Markov's inequality (recall that  $p_j \geq 2\mathbb{E}[w_j]$  for all  $j \in H$ ).

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 this means that the mechanism makes an offer to the seller  $j$  with probability

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

An item in  $H$  is considered to be "in the market" if the corresponding seller accepts the mechanism's offer. After the mechanism has made the offers to the sellers of  $H$ , it knows which items are in the market and then asks each buyer (sequentially, in any order) for her favorite bundle of items among those items that are still in the market. If an item  $j$  gets requested by a buyer, then  $j$  is transferred from its corresponding seller  $j$ , and the buyer pays  $p_j$  to seller  $j$ . Item  $j$  is then removed from the set of items in the market, and the mechanism proceeds to the next buyer.

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

- (1) Let  $H := \{j \in [k] : \mathbb{E}_{\mathbf{v}}[\text{SW}_j^B(\mathbf{v})] \geq 4\mathbb{E}[w_j]\}$ .
- (2) For all  $j \in H$ , set  $p_j := \frac{1}{2}\mathbb{E}_{\mathbf{v}}[\text{SW}_j^B(\mathbf{v})]$ .
- (3) Let  $\Lambda_1 := \emptyset, X_i := \emptyset$  for all  $i \in [n]$  and  $Y_j := \{j\}$  for all  $j \in [k]$ .
- (4) For all  $j \in H$ :
  - (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.
  - (c) If  $j$  accepts the offer, set  $\Lambda_1 := \Lambda_1 \cup \{j\}$ .
- (5) For all  $i \in [n]$ :
  - (a) Buyer  $i$  chooses a bundle  $B_i \in \mathcal{D}(v_i, \mathbf{p}, \Lambda_1)$  that maximises 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.,  $\Lambda_{i+1} := \Lambda_i \setminus B_i$ .
- (6) Return the outcome consisting of allocation  $(X = (X_1, \dots, X_n), Y = (Y_1, \dots, Y_k))$  and payments  $\rho = (\rho^B, \rho^S)$ , 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 only 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  $\alpha$  be the factor by which  $\mathbb{A}$  is guaranteed to approximate the social welfare of the buyers.

## 4.2 Results

Now, we are ready to present the main result of this section:

**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. As mentioned in Section 1.3, one may alternatively take for  $\mathbb{A}$  a polynomial time  $\alpha$ -approximation algorithm and use the technique of [11], to obtain a mechanism with runtime  $POLY(1/\epsilon, n, m)$  that approximates the optimal social welfare within a  $2 + 4\alpha$  multiplicative factor and an  $\epsilon$  additive term.

We split the proof of Theorem 4.1 into two lemmas that separately 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 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  $\text{ALG} = \text{ALG}^S + \text{ALG}^B$ , consistent with the notation of (1).

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

LEMMA 4.2. *If every seller  $j \in H$  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 holds because  $j \in L$ , and the second inequality follows by Markov's inequality. Thus, with probability at least  $1/2$  a seller  $j$  is happy to sell her 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 in  $H$  accepts to trade with probability exactly  $1/2$ . This implies that every seller  $j \in H$  contributes in expectation at least  $\mathbb{E}[w_j]/2$  to the social welfare. Moreover, every seller in  $L$  never trades, so that such a seller contributes her full expected valuation to the expected social welfare.  $\square$

Next, we provide a more difficult bound that relates  $\text{ALG}^B$  and  $\text{ALG}^S$  to  $\text{OPT}^B$ .

LEMMA 4.3. *The buyers' contributions to the optimal social welfare is bounded by*

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

Intuitively, Lemma 4.3 uses two main ingredients:

- the partition of the items between high-welfare ( $H$ ) items and low-welfare items ( $L$ ), and
- the definition of  $\text{SW}_j^B(\mathbf{v})$  w.r.t. a one-sided (approximation) algorithm  $\mathbb{A}$ .

The latter tells us that the sum of the expected contributions of all the items, i.e.  $\sum_{j=1}^k \mathbb{E}[\text{SW}_j^B(\mathbf{v})]$ , is an upper-bound on  $\text{OPT}^B/\alpha$ . From the former we know that:

- the sellers do not trade items in  $L$ , and this is enough to ensure that their contribution to the expected social welfare is greater than a constant fraction of the expected contribution of the items in  $L$ , i.e.  $\text{ALG}^S > \frac{1}{4} \sum_{j \in L} \mathbb{E}[\text{SW}_j^B(\mathbf{v})]$ . Moreover,
- the only items that the agents can trade are those that have a high welfare w.r.t.  $\text{SW}(X_i^{\text{all}}(\mathbf{v}))$ . From that we can infer that the contribution of the buyers to the expected social welfare is greater than a constant fraction of the expected contribution of the items in  $H$ , i.e.  $\text{ALG}^B > \frac{1}{4} \sum_{j \in H} \mathbb{E}[\text{SW}_j^B(\mathbf{v})]$ .

By combining these bounds, the claim of Lemma 4.3 follows. Theorem 4.1 is then obtained straightforwardly from Lemma 4.2 and Lemma 4.3.

PROOF OF THEOREM 4.1. The bound on the approximation ratio follows from the sum of the inequalities of Lemma 4.2 and Lemma 4.3. 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.  $\square$

## 5 A MECHANISM FOR ADDITIVE SELLERS AND XOS BUYERS

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 we show that, in the case that both buyers and sellers have 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  $(n, m, k, I, G, F)$  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 additive representative 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.

Furthermore, we define 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 an additive representative function for a seller.

### 5.1 Mechanism

We aim to design a BIC, interim IR, and DSBB mechanism that approximates the optimal social welfare within a constant. We propose the following mechanism, which we refer to as  $\mathbb{M}_{\text{add}}$ . We let  $H_j := \{\ell \in I_j : \mathbb{E}[\text{SW}_\ell^B(\mathbf{v})] \geq 4\mathbb{E}[\text{SW}_\ell^S(\mathbf{w})]\}$  and  $L_j := I_j \setminus H_j$  for all  $j \in [m]$ , and we let  $H := \bigcup_{j=1}^m H_j$  and  $L := [k] \setminus H$  denote the sets of high-welfare items and low-welfare items, respectively. Our mechanism will only allow trading items in  $H$ . We define for  $\ell \in H$  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  $\mathbb{M}_{\text{add}}$  and  $\mathbb{M}_{1\text{-supply}}$  is that the order in which buyers and sellers are processed is reversed. Mechanism  $\mathbb{M}_{\text{add}}$  roughly works as follows. It first asks every buyer which set of items it would like to receive from those items in  $H$  that have not been requested yet. Then  $\mathbb{M}_{\text{add}}$  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 probability  $q_j$ , chosen in such a way 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 will request a bundle of items that maximises their expected utility, knowing that the item sets requested from each seller will be assigned to them with probability  $1/2$ .<sup>9</sup> In our mechanism, the sellers will each have a dominant strategy, while the buyers' aforementioned behaviour relies on the sellers playing their dominant strategies. This reliance results in a BIC (rather than a DSIC) mechanism. Below we describe the mechanism in more detail and we subsequently provide an example of the mechanism's execution on a simple instance.

<sup>9</sup>Buyers may need to make complex calculations in order to establish which bundle maximises her expected utility.

- (1) For  $\ell \in [k]$ , compute  $\mathbb{E}[\text{SW}_\ell^B(\mathbf{w})]$  and  $\mathbb{E}[\text{SW}_\ell^S(\mathbf{w})]$ .
- (2) For all  $j \in [m]$ , compute  $H_j$ .
- (3) Compute  $H$  and  $L$ .
- (4) Let  $\Lambda_1 := H, X_i := \emptyset$  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-utility maximising bundle  $B_i \subseteq \Lambda_i$  given the prices  $\{p_\ell : \ell \in \Lambda_i\}$  from the set of available items  $\Lambda_i$  (where the expectation is taken w.r.t. the randomness of the valuations and the mechanism).
  - (b) Update the set of available items  $\Lambda_{i+1} := \Lambda_i \setminus B_i$ .
- (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  $S_j := B \cap H_j$  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  $(X = (X_1, \dots, X_n), Y = (Y_1, \dots, Y_k))$  and payments  $\rho = (\rho^B, \rho^S)$ , 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  $j \in [m]$ .

Notice the mechanism  $\mathbb{M}_{\text{add}}$  runs in polynomial time, but it makes use of a variant of a standard demand query in which the mechanism gives prices for the items, and asks a buyer which bundle she would like *if, for each item in that bundle, she were to receive it with probability 1/2*. This places a heavier computational and cognitive burden on the agent than with standard demand queries. We will not address in the present paper the complexity aspects of the buyer's task to answer such queries, though we believe that it is an interesting open question to investigate.

The following example illustrates some important aspects of  $\mathbb{M}_{\text{add}}$ , and the strategies of the buyers under a BNE.

*Example 5.1.* There is one buyer and two unit-supply sellers. Each seller has one item. The buyer has two XOS valuation functions  $v_1$  and  $v_2$ , each chosen with probability 1/2. Valuation  $v_1$  is composed of 3 additive functions  $a_1, a_2$ , and  $a_3$ , i.e.,  $v_1(S) = \max\{a_1(S), a_2(S), a_3(S)\}$ . Valuation  $v_2$  consists of a single additive function  $a_4$ . Each seller  $j$  has a valuation function  $w_j = 0$ . Recall that a function  $a$  is additive if there exists  $\alpha_1, \dots, \alpha_k$  such that  $a(S) = \sum_{j \in S} \alpha_j$  for all  $S \subseteq [k]$ . The functions  $a_1$  to  $a_4$  are described in the table below by listing the values  $\alpha_1$  and  $\alpha_2$ .

Function	item 1 ( $\alpha_1$ )	item 2 ( $\alpha_2$ )
$a_1$	0	2
$a_2$	8	0
$a_3$	7	2
$a_4$	1	6

Let us compute the prices offered by the mechanism  $\mathbb{M}_{\text{add}}$  when  $\mathbb{A}$  is an optimal algorithm. Thus, we need to compute 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 item 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.,

$\mathbb{E}[SW_j^B(\mathbf{v})] = 4$  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.

When the mechanism asks a buyer to select a bundle that maximizes her expected utility, the buyer has to answer by taking into account the fact that each item in her requested bundle will be allocated with probability  $1/2$ . First, consider the case when the buyer has valuation  $v_1$ . In this case the expected utility for the different bundles are:

$$\begin{aligned} u(\{1\}) &= \frac{1}{2} \cdot (8 - 2) + \frac{1}{2} \cdot 0 = 3, \\ u(\{2\}) &= \frac{1}{2} \cdot (2 - 2) + \frac{1}{2} \cdot 0 = 0, \\ u(\{1, 2\}) &= \frac{1}{4} \cdot (8 - 2) + \frac{1}{4} \cdot (2 - 2) + \frac{1}{4} \cdot (9 - 4) + \frac{1}{4} \cdot 0 = \frac{11}{4}. \end{aligned}$$

The utility-maximising bundle that will be requested by the buyer in case of  $v_1$  is  $\{1\}$ , in which case the mechanism will let the buyer pay Seller 1 a price of 2 in exchange for the item. Instead, if the valuation of the buyer is  $v_2$ , then the requested bundle will be  $\{1, 2\}$ , in which case the mechanism will let the buyer pay both sellers a price of 2 in exchange for their items.

## 5.2 Results

Our main result for  $\mathbb{M}_{\text{add}}$  is the following theorem.

**THEOREM 5.2.** *The mechanism  $\mathbb{M}_{\text{add}}$  is 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 a mechanism that is ex-post IR, DSIC, SBB, and 6-approximates the optimal social welfare. Again, we split the proof of this theorem 5.2 into two lemmas that separately bound the sellers' and the buyers' relative contributions to the social welfare. Like the previous section, we use the notation OPT as defined in Section 2, and we use ALG to denote the expected social welfare of the mechanism. Moreover, we use again the superscripts  $B$  and  $S$  to refer to the buyers' and sellers' expected contribution to the social welfare of a given allocation, as we did in Section 4.

Let us first discuss how we bound the sellers' expected contribution to the optimal allocation.

**LEMMA 5.3.**

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

**PROOF.** The only items that our mechanisms potentially reallocates are the ones belonging to  $H$ . Every item in  $L$  stays with its seller. For the items in  $H$ , 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.  $\square$

Similarly, we want to provide an upper bound on the buyers' expected contribution to the optimal allocation. To do that we need two auxiliary propositions.

The first proposition exploits the partition of the items among high-welfare items and low-welfare items. Since the low-welfare items are not traded, the sum of the expected contribution of the buyers on the high-welfare items and the expected contribution of the sellers on the low-welfare items gives us an upper bound on the buyers' expected contribution in the allocation computed by  $\mathbb{A}$ .

**PROPOSITION 5.4.**

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

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,

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

The last inequality follows because by definition of  $L$ ,

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

□

Now, since buyers can obtain only high-welfare items, their contribution to the expected social welfare of  $\mathbb{M}_{\text{add}}$  is greater than a constant fraction of the expected contribution of the high-welfare items to the allocation computed by  $\mathbb{A}$ .

PROPOSITION 5.5.

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

Thus, using Proposition 5.4 and Proposition 5.5 we can prove that the sum of the buyers' expected contribution and the sellers' expected contribution of  $\mathbb{M}_{\text{add}}$  provides a constant approximation to the buyers' expected contribution in the optimal allocation.

LEMMA 5.6.

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

PROOF. By Proposition 5.5,  $4\text{ALG}^B \geq \sum_{\ell \in H} \mathbb{E}_{\mathbf{v}}[\text{SW}_{\ell}^B(\mathbf{v})]$ . Moreover, our mechanism leaves every item  $\ell \in L$  with its seller, and so  $\text{ALG}^S \geq \sum_{\ell \in L} \mathbb{E}_{\mathbf{w}}[\text{SW}_{\ell}^S(\mathbf{w})]$ . Therefore,

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

The second inequality holds by Proposition 5.4, and the last inequality follows because we defined  $\alpha$  to be the approximation factor of algorithm  $\mathbb{A}$ , which is the algorithm that we assumed to generate allocation  $X^{\text{all}}(\mathbf{v})$ . □

Finally, we are ready to prove the main theorem of this section.

PROOF OF THEOREM 5.2. On the sellers' side, the mechanism is 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. Every buyer chooses a bundle that maximises her expected utility, and this choice depends solely on the choice of strategies of the sellers. Therefore, the mechanism has a BNE in which the sellers play a dominant strategy, and the mechanism is thus ex-interim IR and BIC. 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 Lemmas 5.3 and 5.6.  $\square$

It is important to notice that the mechanism  $\mathbb{M}_{\text{add}}$  turns into a DSIC and ex-post IR mechanism if the buyers have additive valuations instead of XOS valuations.

**COROLLARY 5.7.** *For the special case that for all  $i \in [n]$ , distribution  $G_i$  is over additive valuation functions,  $\mathbb{M}_{\text{add}}$  is ex-post IR, DSIC, DSBB and  $(2 + 4\alpha)$ -approximates the optimal social welfare.*

**PROOF.** If a buyer  $i \in [n]$  has an additive valuation function, it is a dominant strategy to request the items in  $\Lambda_i(\mathbf{v}_{<:i})$  for which it holds that  $v_i(\{\ell\}) > p_\ell$ . This follows from the simple fact that by additivity, the utility that a buyer 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  $p_\ell$  with probability  $1/2$ ), a term of  $(1/2)(v_i(\{\ell\}) - p_\ell)$  gets added to her expected utility. So including  $\ell$  in her requested bundle is profitable if and only if  $v_i(\{\ell\}) - p_\ell \geq 0$ . Using the same argument, the ex-post IR property is also satisfied by following this strategy.  $\square$

## 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 an additive valuation function may have intrinsic interdependencies among the items (e.g. if there are duplicates among the items) and so the independence of these distinct unit-supply sellers is not guaranteed. Something we might additionally consider to do is to ask every seller for her favorite bundle to place in the market, yet 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 direction is to extend the above mechanism to the setting in which both buyers and sellers possess an XOS valuation function over bundles of items. A first challenge consists in finding a suitable definition of the sellers' social welfare contribution of an item using the corresponding representative additive function.

## ACKNOWLEDGMENTS

The authors thank anonymous referees for their helpful comments. Tim Roughgarden is supported by the NSF Grant CCF-1524062. Stefano Leonardi is partly supported by the Google Focused Award on "Algorithms for Large-scale Data Analysis". Part of this work was done while Stefano Leonardi and Tim Roughgarden were visiting the Economics and Computation semester at the Simons Institute for the Theory of Computing.

## REFERENCES

- [1] Liad Blumrosen and Shahar Dobzinski. 2014. Reallocation mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*. ACM, 617–617.
- [2] Liad Blumrosen and Thomas Holenstein. 2008. Posted Prices vs. Negotiations: an Asymptotic Analysis. In *Proceedings of the 9th ACM conference on Electronic Commerce (EC)*. ACM, 49–49.

- [3] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*. ACM, 311–320.
- [4] Edward H. Clarke. 1971. Multipart Pricing of Public Goods. *Public Choice* 11, 1 (1971), 17–33.
- [5] Riccardo Colini-Baldeschi, Bart De Keijzer, Stefano Leonardi, and Stefano Turchetta. 2016. Approximately Efficient Double Auctions with Strong Budget Balance. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '16)*. ACM-SIAM, 1424–1443.
- [6] Xiaotie Deng, Paul W. Goldberg, Bo Tang, and Jinshan Zhang. 2014. Revenue Maximization in a Bayesian Double Auction Market. *Theoretical Computer Science* 539, 1 (2014), 1–12.
- [7] Kaustubh Deshmukh, Andrew V. Goldberg, Jason D. Hartline, and Anna R. Karlin. 2002. Truthful and Competitive Double Auctions. In *Proceedings of the 10th Annual European Symposium on Algorithms (ESA)*. Springer, 361–373.
- [8] Paul Dütting and Robert Kleinberg. 2015. Polymatroid Prophet Inequalities. In *Proceedings of the 23rd Annual European Symposium on Algorithms (ESA)*. 437–449.
- [9] Paul Dütting, Inbal Talgam-Cohen, and Tim Roughgarden. 2014. Modularity and Greed in Double Auctions. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*. ACM, 241–258.
- [10] Uriel Feige and Shlomo Joseph. 2014. Demand Queries with Preprocessing. In *Proceedings of the 41st Int'l Colloquium on Automata, Languages, and Programming (ICALP)*. Springer LNCS 8572, 477–488.
- [11] Michal Feldman, Nick Gravin, and Brendan Lucier. 2015. Combinatorial Auctions via Posted Prices. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. ACM-SIAM, 123–135.
- [12] Thomas A. Gresik and Mark A. Satterthwaite. 1989. 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 (1989), 304–332.
- [13] Theodore Groves. 1973. Incentives in Teams. *Econometrica* 41, 4 (1973), 617–631.
- [14] Robert Kleinberg and Seth M. Weinberg. 2012. Matroid Prophet Inequalities. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*. ACM, 123–136.
- [15] R. Preston McAfee. 1992. A Dominant Strategy Double Auction. *Journal of Economic Theory* 56, 2 (1992), 434–450.
- [16] Roger B. Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- [17] Roger B. Myerson and Mark A. Satterthwaite. 1983. Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory* 29 (1983), 265–281.
- [18] Aldo Rustichini, Mark A. Satterthwaite, and Steven R. Williams. 1994. Convergence to Efficiency in a Simple Market with Incomplete Information. *Econometrica* 62, 5 (1994), 1041–1063.
- [19] Tuomas Sandholm and Andrew Gilpin. 2004. *Sequences of Take-It-or-Leave-It Offers: Near-Optimal Auctions Without Full Valuation Revelation*. Springer, 73–91.
- [20] Mark A. Satterthwaite and Steven R. Williams. 1989. 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 (1989), 477–498.
- [21] Mark A. Satterthwaite and Steven R. Williams. 2002. The Optimality of a Simple Market Mechanism. *Econometrica* 70, 5 (2002), 1841–1863.
- [22] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. 2016. A Random-Sampling Double-Auction Mechanism. *CoRR* abs/1604.06210 (2016).
- [23] William Vickrey. 1961. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance* 16, 1 (1961), 8–37.