



Approximately Efficient Two-Sided Combinatorial Auctions

RICCARDO COLINI-BALDESCHI, Facebook Core Data Science, UK

PAUL W. GOLDBERG, University of Oxford, UK

BART DE KEIJZER, University of Essex, Netherlands

STEFANO LEONARDI, Sapienza, University of Rome, Italy

TIM ROUGHGARDEN, Stanford University, USA

STEFANO TURCHETTA, NTT Data and University of Oxford, UK

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.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory**; **Algorithmic mechanism design**; **Computational pricing and auctions**; *Computational advertising theory*;

Additional Key Words and Phrases: Mechanism design, two-sided markets, posted price mechanisms, combinatorial auctions

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", by the ERC Advanced Grant 788893 AMDROMA "Algorithmic and Mechanism Design Research in Online Markets", and by the MIUR PRIN project ALGADIMAR "Algorithms, Games, and Digital Markets".

Authors' addresses: R. Colini-Baldeschi, Facebook Core Data Science, Facebook, 1 Rathbone Place, W1T 1FB, London, UK; email: rickuz@fb.com; P. W. Goldberg, University of Oxford, Room 255, Wolfson Building, Parks Road, Oxford OX1 3QD, United Kingdom; email: paul.goldberg@cs.ox.ac.uk; B. de Keijzer, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, Amsterdam, Netherlands; email: b.dekeijzer@essex.ac.uk; S. Leonardi, Department of Computer, Control and Management Engineering Antonio Ruberti, Sapienza University of Rome, Via Ariosto 25, 00185 Roma, Italy; email: leonardi@diag.uniroma1.it; T. Roughgarden, Department of Computer Science, Columbia University, 500 West 120th Street, Room 450, MC0401, New York, NY 10027; email: tim@cs.stanford.edu; S. Turchetta, Wolfson Building, Parks Road, Oxford OX1 3QD, United Kingdom; email: stefano.turchetta@cs.ox.ac.uk.

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.

© 2020 Association for Computing Machinery.

2167-8375/2020/03-ART4 \$15.00

<https://doi.org/10.1145/3381523>

ACM Reference format:

Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, and Stefano Turchetta. 2020. Approximately Efficient Two-Sided Combinatorial Auctions. *ACM Trans. Econ. Comput.* 8, 1, Article 4 (March 2020), 29 pages.
<https://doi.org/10.1145/3381523>

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, and a property that is often desirable in this setting is incentive-compatibility, i.e., 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 [10, 21, 34] 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 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 [27] 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.¹

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 article 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

¹The 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.

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 that 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 [24, 30, 31] 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 agents have no knowledge of the realized valuation function of the i th agent but only of her probability distribution.² 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 (as implied by the impossibility result of Myerson and Satterthwaite [27]).

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 [27] 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 article 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 [4] 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

²We note that these are standard assumptions, although knowledge of the distributions by the agents is not necessary for our DSIC results to hold.

Table 1. Summary of Our Results

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}_{add}	XOS	additive	6	interim IR	BIC	DSBB
\mathbb{M}_{add}	additive	additive	6	ex-post IR	DSIC	DSBB

current definition of SBB, which motivates the introduction of a strengthened version, that we call *direct-trade strong budget balance* (DSBB).

Let us briefly elaborate on why it is justified to ask for a (D)SBB mechanism over a WBB one. In many two-sided markets that one encounters in industry, it is true that the business running the mechanism is interested in making some profit. However, the way in which profit is made is usually not through running a mechanism with too-large budget imbalance between the buyers and sellers in the market. Rather, profit is made through entry fees or subscription fees (this happens in, e.g., stock exchanges) or through charging a relatively small commission for each transaction or volume of transactions that is made (which also happens in some stock exchanges and many two-sided market platforms in the sharing/access economy). Having an underlying (D)SBB pricing mechanism is then desirable over a WBB mechanism, as the former mechanism can be presented to its customers in a more transparent way, and the amount of profit made by the market platform is now easy to describe and control on the side, whereas the amount of profit made by a WBB mechanism can be difficult to describe and control and is often very instance dependent. This is not necessarily a problem in case the leftover budget is very small but gets problematic when the surplus is large. Therefore, imposing the (D)SBB requirement results in mechanisms that capture the intended primary goal in many practical mechanism design challenges: direct trades of goods for money between its customers, which is essentially the product that a commercial two-sided market platform primarily intends to offer to its customers. The way in which profit is made is then designed around such a mechanism as a secondary step.

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.

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}_{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}_{add} satisfies weaker IC and IR notions than $\mathbb{M}_{1\text{-supply}}$. More precisely, \mathbb{M}_{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}_{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.³ Furthermore, we note that our mechanisms

³The mechanism proposed in Reference [4] 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.

not only work for a more general setting than that of Reference [12] (see Section 1.3 for more details) but also improve the approximation ratio for double auctions from 16 to 6.

In the case of \mathbb{M}_{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 seller j , the items of seller j in the bundle were to be received with probability 1/2*. 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.⁴

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 [27].

A first feature all our mechanisms share to guarantee DSBB is being a generalised version of two-sided *sequential posted price mechanisms (SPMs)* [12] 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 lies in the choice of which market side is processed first. Otherwise, all mechanisms proposed in this article are oblivious to the order in which sellers and buyers are presented.

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* [19] 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}_{add} make use of an algorithm \mathbb{A} that, given a buyers' valuation profile and a set of items,

⁴This question also applies to standard demand queries [18], 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.

allocates the items to the buyers, without considering the sellers and their valuations. If one is not interested in achieving a low runtime, then one can take \mathbb{A} to be an exact algorithm that outputs an optimal allocation. Alternatively, by using a technique of Feldman et al. [19], one may take for \mathbb{A} a polynomial time approximation algorithm and combine this with sampling a sufficiently large number of valuation profiles from the distribution, 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 the case where \mathbb{A} is a polynomial time α -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 ϵ additive term, where ϵ is a parameter that results from the sampling procedure. This technique is described in further detail in Reference [19] 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.

To summarize, we utilize in the present article an implicit generalization of two-sided sequential posted pricing mechanisms of Reference [12]. We furthermore make use of the posted pricing techniques and concepts introduced in Reference [19] as one important component for achieving the correct item prices, and, moreover, we (indirectly) utilize the basic idea of using the median of a probability distribution for pricing, a concept that also, e.g., appears in Reference [25] and has been used in other works, such as References [4, 12, 25]. We then combine this with various novel insights on the sequence of offers, the prices, the probabilities, and the item bundling, to preserve incentive compatibility. The technical contribution in our present article thus lies in the insights and techniques used in the definition of the pricing procedure, in the bundles' selection procedure, and in the proofs that combine and connect these insights to establish the claimed approximation bounds. Last, on the conceptual side we introduce the notions of probabilistic demand queries and direct-trade SBB, which we believe are interesting concepts for future study.

1.4 Related Work

Due to the impossibility result of Myerson and Satterthwaite [27], no two-sided mechanism can simultaneously achieve optimal social welfare and satisfy the BIC, IR, and WBB constraints, even in the simple bilateral trade setting. Follow-up work thus had to focus on designing mechanisms that tradeoff among these properties.

The following articles 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 [20] showed that duplicating the number of agents by τ results in a market where the optimal IR, IC, and WBB mechanism's expected social welfare approximation factor approaches 1 at a rate of $O(\log \tau / \tau^2)$. Rustichini et al. [28] and Satterthwaite and Williams [31] 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 [24], an IC, WBB, and IR double auction is proposed that extracts at least a $(1 - 1/\ell)$ fraction of the maximum gain from trade,⁵ where ℓ is the number of traders in the optimal solution.

⁵The gain from trade (sometimes called incremental social-welfare) represents how much the total utility of all agents increases with respect to the situation where no trade happens. This is a meaningful objective in a two-sided market

The same type of approximation guarantee is obtained in Reference [3] for a more general model modeling spatially distributed markets, where there are multiple two-sided markets for the same good, and shipping an item between markets is possible at a certain cost.

Optimal revenue-maximising Bayesian auctions were characterized in Myerson [26], 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. [13], which studied maximising the auctioneer's revenue in Bayesian double auctions. The same objective was studied by Deshmukh et al. [14] yet in the *prior-free* model. In Reference [32], 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. Many other articles studied the problem of approximating the gain from trade when the mechanisms are IC, IR, and SBB or WBB [1, 7, 8, 11, 25, 33]. In Reference [8], a mechanism is presented that guarantees 2-approximation to the gain from trade, under DSIC, IR, and WBB-in-expectation. This approximation guarantee is with respect to an alternative benchmark from ours, called the *second-best* benchmark, where the mechanism is compared to the best possible BIC, IR, WBB-in-expectation mechanism. Moreover, Reference [2], proposes a mechanism that is shown to achieve approximation guarantees to both benchmarks simultaneously: a constant guarantee for the second-best benchmark, while their mechanism also approximates full efficiency as the market gets larger.

For the special case of *bilateral trade*, i.e., with one seller buyer and item, it was shown in Reference [5] that the gain from trade cannot be approximated within any constant under the IC, IR, and SBB constraints, and the asymptotic approximability factor in this setting was further characterized in Reference [11]. Note that this implies that the constant approximation factors that we establish in the present article rely crucially on taking the easier-to-approximate social welfare as our objective function, instead of gain from trade.

From a social-welfare perspective, mechanisms that are IC, IR, and SBB have been given for bilateral trade in Reference [4]. 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 [29] 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 [6]. They were later studied by Chawla et al. [9] for the objective of revenue maximisation. Additionally, Kleinberg and Weinberg [22] and Dütting and Kleinberg [16] strengthen these results further. Very relevant to our work is the article of Feldman et al. [19], 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. [17] 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 [24]. They obtain IR, DSIC, and WBB mechanisms with a good approximation of the social welfare, for knapsack, matching, and matroid allocation constraints. More recently, Colini-Baldeschi et al. [12] presented the first

because the sellers may have a positive utility even when no trade happens, simply because they possess some items. For this reason, approximating the gain from trade is harder than approximating the social welfare.

double auction⁶ 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. The setting we study in the present article is more general than the double auction setting by considering buyers and sellers with multiple distinct items and complex valuation functions.

Last, there is an alternative line of research where, among the desiderata of satisfying IC, IR, SBB, and efficiency (i.e., achieving optimal social welfare), the IC notion is weakened instead of efficiency: In Reference [23], for example, an efficient iterative mechanism is proposed for exchange markets that yields an outcome minimizing *regret* among the set of all budget-balanced payment rules, when agents bid truthfully.

2 PRELIMINARIES

As a general convention, we use bold 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 $\mathbb{R}_{\geq 0}$.⁷ 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 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 $(n, k, \mathbf{I}, \mathbf{F})$.

Throughout the article, 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.

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.

2.2 Mechanism Design Goals

The following discussion is specific to two-sided markets (the main focus of this article), 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 $(n, m, k, \mathbf{I}, \mathbf{G}, \mathbf{F})$ 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

⁶Recall that in double auctions (following the definition used in the classical body of literature in economic theory), items are identical and buyers and sellers are respectively assumed to have one unit of demand and supply, respectively.

⁷By 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 $v(\emptyset) = 0$.

$[k]$, and $X_1, \dots, X_n, Y_1, \dots, Y_m$ are mutually disjoint. Note that Y_j represents the bundle of items that remains in the seller's possession after execution of the mechanism. 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 ρ^B refers to the buyers' vector of payments and ρ^S to the sellers' one. An outcome is therefore a tuple (X, Y, ρ^B, ρ^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 article.

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 $u_j^S(w_j, (X, Y, \rho^B, \rho^S)) = w_j(Y_j) - \rho_j^S$.

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)**⁸

- *Dominant strategy incentive compatibility (DSIC)*: It is a dominant strategy for every agent to report her true valuation sincerely. That is, 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. That is, 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

⁸Technically, 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.

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} α -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 α -approximate the optimal social welfare for a low α , 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 [4] present a mechanism for exchange markets with subadditive valuation functions. It works by partitioning the set of agents into two sets A and B uniformly at random. For the agents in A , an agent i is offered a payment equal to his median valuation MED_i for the whole bundle. All the agents of A that accept this payment hand in their items, and the total set of collected items is then sold through a VCG auction to the agents in B . However, in running this VCG auction, an objective function is maximized that is a modified version of the regular social welfare, where an amount of $H(t_i)MED_i$ is subtracted for each agent $i \in A$, where t_i is the number of agents in B holding an item from i , and $H(\cdot)$ denotes the harmonic numbers. The auctioneer potentially takes items away from agents in A that subsequently remain unallocated after running the VCG mechanism. They prove the following for this mechanism, which we name \mathbb{M}_{bd} .

THEOREM 3.1 (BLUMROSEN AND DOBZINSKI [4]). *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}, 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.*

This mechanism gives us a constant approximation factor if the number of starting items of the agents is bounded by a constant, in particular in the unit-supply case.

The budget imbalance that \mathbb{M}_{bd} may generate can be arbitrarily high and is highly instance-dependent. We show now how we can use this mechanism as a black box to obtain an SBB

mechanism with only a slightly worse approximation ratio. Define mechanism \mathbb{M}_{sbb} as follows. When given as input a combinatorial exchange market $C = (n, k, I, F)$,

- (1) Select an agent in $i \in [n]$ uniformly at random.
- (2) Run Mechanism \mathbb{M}_{bd} on the combinatorial exchange market $C_{-i} = ([n] \setminus \{i\}, I_{-i} = (I_1, \dots, I_{i-1}, I_{i+1}, \dots, I_n), F_{-i} = (F_1, \dots, F_{i-1}, F_{i+1}, \dots, F_n))$.
Let (X_{-i}, ρ_{-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, X_{-i}) and output payment vector (p_i, ρ_{-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. Note that if we would like to ensure that all items are allocated to the agents (hence do not end up at the auctioneer) after running the mechanism, then we may additionally allocate any leftover items to the agent that we removed *a priori*. The following is a direct corollary of the DSIC, WBB, and ex-post IR properties of mechanism \mathbb{M}_{bd} .

THEOREM 3.2. *Mechanism \mathbb{M}_{sbb} is a DSIC, SBB, and ex-post IR mechanism for exchange markets with subadditive valuation functions.*

Second, the following theorem shows that the mechanism loses only a factor $2n/(n-1) \leq 3$ in the approximation ratio for $n \geq 3$.⁹

THEOREM 3.3. *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 three agents.*

PROOF. Fix a valuation vector \mathbf{v} of the agents, let $X_{\mathbf{v}}^{**} \subseteq \mathcal{A}$ be the social welfare maximising allocation when the agents have valuations \mathbf{v} . For an agent $i \in [n]$, denote by $X_{\mathbf{v}, -i}^{**}$ 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 $X_{\mathbf{v}}^{**}$ when i is removed, and all items of i are removed. Moreover, let $X_{\mathbf{v}, -i}^*$ 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 3.1, the expected social welfare of Mechanism \mathbb{M}_{sbb} is at least

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

⁹For $n = 2$ there exist alternative mechanisms that achieve a good approximation ratio. In this case, one may for example restrict to a mechanism that either considers reallocating the whole bundle of Agent 1 to Agent 2 or vice versa. The mechanism selects one of the two alternatives uniformly at random, and then runs a 2-approximate DSIC, SBB, IR mechanism for the bilateral trade setting (i.e., the setting with one buyer, one seller, one item) [5] where the bundle of the selling agent is treated as a single item. This yields a DSIC, SBB, IR mechanism where the random choice in the first step introduces an additional factor of 2 to the approximation factor, so that we get an upper bound of 4 on the total approximation factor.

$$\begin{aligned}
&= \frac{1}{4nH(s)} \sum_{i \in [n]} \sum_{\{j, j'\} : j, j' \in [n] \setminus \{i\} \wedge j \neq j'} \frac{1}{n-2} v_i(X_{\mathbf{v}}^{**}) \\
&= \frac{1}{4nH(s)} \sum_{i \in [n]} \frac{n-1}{2} v_i(X_{\mathbf{v}}^{**}) \\
&= \frac{n-1}{8nH(s)} \sum_{i \in [n]} v_i(X_{\mathbf{v}}^{**}),
\end{aligned}$$

where the second inequality follows from subadditivity, and the second-to-last equality is obtained through the insight that the number of terms in the inner summation is $\binom{n-1}{2} = (n-1)(n-2)/2$, which yields the factor $(n-1)/2$ when multiplied by the coefficient $1/(n-2)$. This proves the claim, since the above holds for every valuation vector \mathbf{v} . \square

This yields an ex-post IR, SBB, and 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. 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 3.4. 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 article, we will proceed to design mechanisms for two-sided markets that do satisfy DSBB.¹⁰

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 $j \in [k]$). When clear from context, in this section we will use w_j to denote the valuation of seller j for his initial bundle $\{j\}$. 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 SW^B, SW^S to respectively denote the buyers' and the sellers' contribution to the social welfare, i.e.,

¹⁰We note that the double auctions given in Reference [12] also satisfy the DSBB property. Moreover, one may define analogously "direct-trade" version of weak budget balance, say DWBB, which turns out to be satisfied by the WBB mechanisms of References [2, 3, 8, 24].

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

Our mechanism requires fixing a price for every item in the market. For a bundle of available items Λ 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 Λ that maximise i 's utility under the given item prices.

For a buyer i with an XOS 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 $S \subseteq [k]$. An 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. Regarding the computational aspects of our mechanisms: They assume the ability to compute an additive representative function of an XOS function for a given bundle of items. Regarding the computational complexity of the latter task, given value query access to an XOS-function, computing an additive representative function for a bundle of items can be done efficiently for some special subclasses of XOS-functions, such as submodular functions.¹¹ To the best of our knowledge for general XOS-functions, it is not known whether one can compute an additive representative function in polynomial time for a given bundle. Various previous works (e.g., References [15, 19]) that deal with XOS functions simply assume the existence of an *XOS-oracle* that performs this task without spending computational resources.

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

We define for each item $j \in [k]$ its *contribution* $SW_j^B(\mathbf{v})$ to the social welfare $SW(X^{\text{all}}(\mathbf{v}))$ 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 $X^{\text{all}}(\mathbf{v})$, then $SW_j^B(\mathbf{v}) = 0$.

This notion allows us to make a distinction between *high welfare items* and *low welfare items*, where these notions are with respect to the ratio of the social welfare of the item on the buyer's side, and the seller's social welfare. 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

¹¹For submodular functions one can order the items in the given bundle arbitrarily, set the coefficients outside of the bundle to 0, and set the coefficient of each item inside of the bundle to the marginal increase in valuation with respect to the set of predecessors in the ordering. It is then straightforward to prove that this yields the coefficients of an additive representative function for the given bundle and requires at most k value queries to compute.

j if we would allocate j according to $X^{\text{all}}(\mathbf{v})$ 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., $H := \{\ell \in [k] : \mathbb{E}[\text{SW}_\ell^B(\mathbf{v})] \geq 4\mathbb{E}[w_j]\}$, 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 :

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

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 mechanism makes an offer to the corresponding seller for this item, and 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 favourite bundle of items (given the price vector \mathbf{p}) among those items that are still in the market. If an item j is among the items 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 $\boldsymbol{\rho} = (\boldsymbol{\rho}^B, \boldsymbol{\rho}^S)$, 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 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 α 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, and 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, and DSBB, and 6-approximates the optimal social welfare. As mentioned in Section 1.3, one may alternatively take for \mathbb{A} a polynomial time α -approximation algorithm and use the technique in Reference [19], 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 ϵ 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., $OPT = OPT^S + OPT^B$ and $ALG = ALG^S + ALG^B$, 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*

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

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 $j \in H$, 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 ALG^B and ALG^S to OPT^B . The proof of this lemma is lengthy and therefore deferred to Appendix A.1.

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

$$4\alpha ALG = 4\alpha ALG^B + 4\alpha ALG^S \geq OPT^B.$$

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., $\sum_{j=1}^k \mathbb{E}[\text{SW}_\ell^B(\mathbf{v})]$, is an upper-bound on OPT^B/α . 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., $ALG^S > \frac{1}{4} \sum_{j \in L} \mathbb{E}[\text{SW}_j^B(\mathbf{v})]$,

- the only items that the agents can trade are those that have a high welfare w.r.t. $SW(X_i^{\text{all}}(\mathbf{v}))$. From that, we can (not trivially, see Appendix A.1) 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., $ALG^B > \frac{1}{4} \sum_{j \in H} \mathbb{E}[SW_j^B(\mathbf{v})]$.

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 when truthful for both buyers and sellers, and therefore the mechanism is also ex-post IR. Last, 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.

Note that if the coefficients of the additive valuation functions of the sellers would be drawn from independent distributions, the present setting can be handled satisfactorily by our previous mechanism (i.e., by treating the individual items of a single seller as if they were all items held by individual sellers) while preserving all desirable properties and the approximation ratio. However, in the present setting there generally exists a correlation among the coefficients of a seller's valuation function, which causes the incentive compatibility property of the first mechanism to not carry through to this case where we have additive sellers with multiple items. An example of an instance where the $\mathbb{M}_{1\text{-supply}}$ mechanism would fail is when there is one buyer and one seller of two items. The buyer holds (deterministically) a XOS valuation function consisting of two additive functions in its support, with respectively the coefficients (42, 0) and (0, 40). The seller holds (deterministically) an additive valuation function with coefficients (10, 0). Both items would be classified by the mechanism as high welfare items, and prices of 21 and 20 would be offered for the two respective items. If the seller behaves truthfully and accepts the offered price for both items, then the buyer would request only the first item, leaving the seller with a utility of 21. However, if the seller only accepts the price for the second item, then the buyer would request that item, and the seller obtains a utility of 30, which demonstrates that $\mathbb{M}_{1\text{-supply}}$ is not DSIC in this setting.

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. Here, the number of items and sellers is in general different, 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 algorithm \mathbb{A} on input \mathbf{v} is a mapping 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 an additive representative function of v_i for T . Also, we use the buyers' social welfare contribution $SW_\ell^B(\mathbf{v})$ for item $\ell \in [k]$ and buyers' valuation profile \mathbf{v} , as defined in Section 4.

Furthermore, we define the *sellers' social welfare contribution* $SW_\ell^S(\mathbf{w})$ for item $\ell \in I_j$ and sellers' valuation profile \mathbf{w} as $SW_\ell^S(\mathbf{w}) := w_j(\{\ell\})$. 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}_{add} . We let $H_j := \{\ell \in I_j : \mathbb{E}[SW_\ell^B(\mathbf{v})] \geq 4\mathbb{E}[SW_\ell^S(\mathbf{w})]\}$ 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

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

similarly 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 (given the price vector \mathbf{p}) from those items in H 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 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 (given the price vector \mathbf{p}), knowing that the item sets requested from each seller will be assigned to them with probability $1/2$.¹² 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.

- (1) For $\ell \in [k]$, compute $\mathbb{E}[SW_\ell^B(\mathbf{v})]$ and $\mathbb{E}[SW_\ell^S(\mathbf{w})]$.
- (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 Λ_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 $S_j := B \cap I_j$ 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 item 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 $j \in [m]$.

¹²Buyers may need to make complex calculations to establish which bundle maximises her expected utility.

Notice the mechanism \mathbb{M}_{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 seller j , the items of seller j in that bundle were to be received with probability $1/2$. Note that in answering such a query, a buyer has to take into account the positive correlation that exists between the presence of items in the market, as two items coming from the same seller are always either both in the market or both not in the market. This type of demand query places a heavier computational and cognitive burden on the agent than with standard demand queries. Answering such a demand query gets more difficult with increasing expressiveness of the class of valuation functions considered. In particular, the structure of the XOS valuation functions considered here will make answering such value queries easier than more expressive classes. This said, we will not address in the present article the complexity aspects of the buyer's task in answering such queries, though we believe that it is an important open question to investigate. This might be interesting even for a simpler type of probabilistic demand query where each item (rather than seller's bundles) are independently received with probability $1/2$.

The following example illustrates some important aspects of \mathbb{M}_{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 α_1 and α_2 .

Function	item 1 (α_1)	item 2 (α_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}_{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 , then 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.

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 $\{2\}$, in which case the mechanism will let the buyer pay seller 2 a price of 2 in exchange for item 2.

5.2 Results

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

THEOREM 5.2. *The mechanism \mathbb{M}_{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.

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

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

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

$$\begin{aligned}
&= \sum_{\ell \in H} \mathbb{E}[\text{SW}_{\ell}^B(\mathbf{v})] + \sum_{\ell \in L} \mathbb{E}[\text{SW}_{\ell}^B(\mathbf{v})] \\
&< \sum_{\ell \in H} \mathbb{E}[\text{SW}_{\ell}^B(\mathbf{v})] + 4 \sum_{\ell \in L} \mathbb{E}[\text{SW}_{\ell}^S(\mathbf{w})].
\end{aligned}$$

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

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

Now, since buyers can obtain only high-welfare items, their contribution to the expected social welfare of \mathbb{M}_{add} is greater than a constant fraction of the expected contribution of the high-welfare items to the allocation computed by \mathbb{A} . The proof of Proposition 5.5 is deferred to Appendix A.2.

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}_{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 α to be the approximation factor of algorithm \mathbb{A} , which is the algorithm that we assumed to generate allocation $X^{\text{all}}(\mathbf{v})$. \square

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}_{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}_{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, then 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_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 $v_i(\{\ell\}) - p_\ell \geq 0$. 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 naïve 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.

Additionally, we might consider asking every seller for her favorite 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. However, 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 a corresponding additive representative function.

We also comment on the reason for why we choose, in our two mechanisms, to scale down the probability of an item being "in the market" so that it is exactly $1/2$. The reason for scaling down these probabilities is twofold: First, it makes the definition of the probabilistic demand queries cleaner and easier to comprehend. It might be easier for a buyer to answer queries where the probability of an item being in the market is exactly $1/2$ instead of any arbitrary probability (although we do not yet have formal evidence for that). Second, our current proof technique is to decompose the social welfare into the buyer's and seller's contribution and show that both sides contribute a substantial amount. If our mechanisms would not ensure the scaling of the probability parameters, then this proof technique does not generalize appropriately, because the sellers' contribution to the social welfare can no longer be shown to be high enough. This loss can possibly be compensated for by the gain in social welfare on the buyer's side, but we expect that such a proof would differ significantly from the one provided, and in particular we would no longer be able to split the proof into two separate lemmas we do presently. We leave it as an open research direction to investigate whether such a modification can yield an improved approximation factor.

Finally, we think it is an interesting open question to design a mechanism for the case of additive sellers and XOS buyers, where buyers do not need to answer the probabilistic-type demand queries currently used, i.e., where a buyer has to specify a utility-maximizing bundle from a set of presented items, given item prices, under the assumption that each seller's bundle is allocated with probability $1/2$. Answering such a query places an additional burden on the buyer of which the complexity is as of yet unclear. Can we remove this burden by removing entirely the need for these demand queries, or can we replace those queries by a less demanding variant where, e.g., a buyer only needs to assume that each individual item is independently allocated with probability $1/2$?

APPENDIX

A PROOFS

A.1 Proof of Lemma 4.3

To do prove Lemma 4.3, 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 H .

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 Λ_1 at step 5 of $\mathbb{M}_{1\text{-supply}}$. For $i \in [n]$ let $\Lambda_i(\mathbf{v}_{<i}, Z)$ be the set Λ_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 $X_i \subseteq \Lambda_i(\mathbf{v}_{<i}, Z) \subseteq Z$. 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 A.1. *The total expected utility of the buyers for the allocation returned by $\mathbb{M}_{1\text{-supply}}$ is bounded from below by*

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

(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 H .)

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

$$p_j = \mathbb{E}_{\mathbf{v}}[\text{SW}_j^B(\mathbf{v}) - p_j] = \sum_{i=1}^n \mathbb{E}_{\mathbf{v}}[(\text{SW}_j^B(\mathbf{v}) - p_j) \mathbb{I}[j \in X_i^{\text{all}}(\mathbf{v})]]. \quad (2)$$

Fix $i \in [n]$, buyers' valuation profile \mathbf{v} , and set $Z \subseteq H$ of sellers who accepted the mechanism's offer, and now consider the set $\Lambda_i(\mathbf{v}_{<i}, Z) \subseteq H$ of available items that i can choose from. Notice that these do not depend on v_i , since the buyers in $\{1, \dots, i-1\}$ select their favorite bundle regardless of v_i . Moreover, notice that since the sellers are processed before the buyers, $\Lambda_1(\mathbf{v}_{<1}, \mathbf{z})$ only contains items whose sellers are willing to trade at the posted prices. Buyer i selects a bundle that maximises her utility, i.e., that is in $\mathcal{D}(v_i, \mathbf{p}, \Lambda_i(\mathbf{v}_{<i}, Z))$.

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} . Let $X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i})$ be the allocation of buyer i returned by $\mathbb{A}(v_i, \tilde{\mathbf{v}}_{-i})$. For $i \in [n]$, consider a 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

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

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 $B_i(\mathbf{v}, Z) \in \mathcal{D}(v_i, \mathbf{p}, \Lambda_i(\mathbf{v}_{<i}, Z))$ that maximises her utility, and $S_i(v_i, \mathbf{v}_{-i}, \tilde{\mathbf{v}}_{-i}, Z)$ is in $\Lambda_i(\mathbf{v}_{<i}, Z)$, 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 $Z \subseteq H$

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

The second-to-last inequality follows from the definition of the corresponding additive function $a(v_i, X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}, \cdot)$; that is, $v_i(S) \geq a(v_i, X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}, S)$ for all $S \subseteq [k]$.

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

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

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

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

For the first inequality, we simply renamed $\tilde{\mathbf{v}}_{-i} = \mathbf{v}_{-i}$, since they are equally distributed (and so equal under expectation) and then used Equation (2). For the last one, we used the fact that for any $i \in [n]$ it holds that $\Pr_{\mathbf{v}}[j \in \Lambda_i(\mathbf{v}_{< i}, Z)] \geq \Pr_{\mathbf{v}}[j \in \Lambda_{n+1}(\mathbf{v}, Z)]$. \square

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

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

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

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

We now prove Lemma 4.3 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 4.3. As just observed above, from Proposition A.1 and Proposition A.2, we have that

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

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

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

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,

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

A.2 Proof of Proposition 5.5

PROPOSITION A.3. *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 $X'_{i,j} := X'_i \cap H_j$ be the set of items in H 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 $X''_i(z) := \bigcup_{j \in [m]: z_j=1} X'_{i,j}$ for all $i \in [n]$. Then, for all $i \in [n]$ it holds that*

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

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

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

Observe that

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

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:

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

PROPOSITION A.4. 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 $\ell \in H_j$, it holds by definition of p_ℓ and H_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 is defined such that

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

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 A.5. *The expected total utility of the buyers is at least*

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

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 $X_i^H(\mathbf{v}, \tilde{\mathbf{v}}_{-i}) = X_i^{\text{all}}(v_i, \tilde{\mathbf{v}}_{-i}) \cap H \cap \Lambda_i(\mathbf{v}_{<i})$. 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 $S(\mathbf{v}, \mathbf{z}) = \bigcup_{j \in [m]: z_j=1} (S \cap H_j)$. Particularly, from this definition we obtain the random variable $X_i(\mathbf{v}, \tilde{\mathbf{v}}_{-i}, \mathbf{z}) = \bigcup_{j \in [m]: z_j=1} (X_i^H(\mathbf{v}, \tilde{\mathbf{v}}_{-i}) \cap H_j)$. 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 maximises her expected utility and will therefore request the bundle $B_i(\mathbf{v})$ that maximises her expected utility, i.e.,

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

By Proposition A.4 each seller's requested items will be allocated with probability 1/2, as reflected by the Bernoulli random variables \mathbf{z} .

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

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

where the second inequality follows from Proposition A.3, and the last inequality follows from the definition of an 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 , then we obtain the following bound on the total expected utility of the buyers:

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

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

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

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

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

LEMMA A.6. *The expected sum of payments made by the buyers is equal to*

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

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

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

The second equality holds by the independence of the two events. \square

PROOF OF PROPOSITION 5.5. 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

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

by definition of p_{ℓ} . □

ACKNOWLEDGMENTS

The authors thank anonymous referees for their helpful comments. Part of this work was done while the author was visiting the Economics and Computation semester at the Simons Institute for the Theory of Computing. This work was completed while the first author was working at LUISS, University of Rome.

REFERENCES

- [1] Moshe Babaioff, Yang Cai, Yannai A. Gonczarowski, and Mingfei Zhao. 2018. The best of both worlds: Asymptotically efficient mechanisms with a guarantee on the expected gains-from-trade. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. 373.
- [2] Moshe Babaioff, Yang Cai, Yannai A. Gonczarowski, and Mingfei Zhao. 2018. The best of both worlds: Asymptotically efficient mechanisms with a guarantee on the expected gains-from-trade. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM, 373–373.
- [3] Moshe Babaioff, Noam Nisan, and Elan Pavlov. 2009. Mechanisms for a spatially distributed market. *Games Econ. Behav.* 66, 2 (2009), 660–684.
- [4] Liad Blumrosen and Shahar Dobzinski. 2014. Reallocation mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC'14)*. ACM, 617–617.
- [5] Liad Blumrosen and Shahar Dobzinski. 2016. (Almost) efficient mechanisms for bilateral trading. *ArXiv/CoRR* abs/1604.04876 (2016).
- [6] Liad Blumrosen and Thomas Holenstein. 2008. Posted prices vs. negotiations: An asymptotic analysis. In *Proceedings of the 9th ACM conference on Electronic Commerce (EC'08)*. ACM, 49–49.
- [7] Liad Blumrosen and Yenohatan Mizrahi. 2016. Approximating gains-from-trade in bilateral trading. In *Proceedings of the 12th International Conference on Web and Internet Economics (WINE)*, Yang Cai and Adrian Vetta (Eds.). Springer, 400–413.
- [8] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. 2017. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC'17)*. 589–590.
- [9] 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'10)*. ACM, 311–320.
- [10] Edward H. Clarke. 1971. Multipart pricing of public goods. *Publ. Choice* 11, 1 (1971), 17–33.
- [11] Riccardo Colini-Baldeschi, Paul Goldberg, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. 2017. Fixed price approximability of the optimal gain from trade. In *International Conference on Web and Internet Economics*. Springer, Cham.
- [12] 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.
- [13] Xiaotie Deng, Paul W. Goldberg, Bo Tang, and Jinshan Zhang. 2014. Revenue maximization in a Bayesian double auction market. *Theor. Comput. Sci.* 539, 1 (2014), 1–12.
- [14] 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'02)*. Springer, 361–373.
- [15] Shahar Dobzinski, Noam Nisan, and Michael Schapira. 2010. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Operat. Res.* 35, 1 (2010), 1–13.
- [16] Paul Dütting and Robert Kleinberg. 2015. Polymatroid prophet inequalities. In *Proceedings of the 23rd Annual European Symposium on Algorithms (ESA'15)*. 437–449.
- [17] 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'14)*. ACM, 241–258.

- [18] Uriel Feige and Shlomo Joseph. 2014. Demand queries with preprocessing. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP'14)*. LNCS 8572. Springer, Berlin, 477–488.
- [19] 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'15)*. ACM-SIAM, 123–135.
- [20] 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. *J. Econ. Theory* 48, 1 (1989), 304–332.
- [21] Theodore Groves. 1973. Incentives in teams. *Econometrica* 41, 4 (1973), 617–631.
- [22] 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.
- [23] Benjamin Lubin, Adam I. Juda, Ruggiero Cavallo, Sébastien Lahaie, Jeffrey Shneidman, and David C. Parkes. 2008. ICE: An expressive iterative combinatorial exchange. *J. Artif. Intell. Res.* 33 (2008), 33–77.
- [24] R. Preston McAfee. 1992. A dominant strategy double auction. *J. Econ. Theory* 56, 2 (1992), 434–450.
- [25] R. Preston McAfee. 2008. The gains from trade under fixed price mechanisms. *Appl. Econ. Res. Bull.* 1, 1 (2008), 1–10.
- [26] Roger B. Myerson. 1981. Optimal auction design. *Math. Operat. Res.* 6, 1 (1981), 58–73.
- [27] Roger B. Myerson and Mark A. Satterthwaite. 1983. Efficient mechanisms for bilateral trading. *J. Econ. Theory* 29 (1983), 265–281.
- [28] 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.
- [29] Tuomas Sandholm and Andrew Gilpin. 2004. *Sequences of Take-It-or-Leave-It Offers: Near-Optimal Auctions Without Full Valuation Revelation*. Springer, 73–91.
- [30] 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. *Rev. Econ. Stud.* 56, 4 (1989), 477–498.
- [31] Mark A. Satterthwaite and Steven R. Williams. 2002. The optimality of a simple market mechanism. *Econometrica* 70, 5 (2002), 1841–1863.
- [32] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. 2016. A random-sampling double-auction mechanism. *CoRR abs/1604.06210* (2016).
- [33] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. 2016. SBBA: A strongly-budget-balanced double-auction mechanism. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory (SAGT'16)*. Springer, 260–272.
- [34] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finan.* 16, 1 (1961), 8–37.

Received November 2017; accepted July 2019