Approximately Efficient Double Auctions with Strong Budget Balance*

Riccardo Colini-Baldeschi† Bart de Keijzer‡ Stefano Leonardi§ Stefano Turchetta¶

Abstract

Mechanism design for one-sided markets is an area of extensive research in economics and, since more than a decade, in computer science as well. Two-sided markets, on the other hand, have not received the same attention despite the numerous applications to web advertisement, stock exchange, and frequency spectrum allocation. This work studies double auctions, in which unit-demand buyers and unit-supply sellers act strategically.

An ideal goal in double auction design is to maximize the social welfare of buyers and sellers with *individually rational (IR)*, *incentive compatible (IC)* and *strongly budget-balanced (SBB)* mechanisms. The first two properties are standard. SBB requires that the payments charged to the buyers are entirely handed to the sellers. This property is crucial in all the contexts that do not allow the auctioneer retaining a share of buyers' payments or subsidizing the market.

Unfortunately, this goal is known to be unachievable even for the special case of *bilateral trade*, where there is only one buyer and one seller. Therefore, in subsequent papers, meaningful trade-offs between these requirements have been investigated.

Our main contribution is the first IR, IC and SBB mechanism that provides an $O(1)$-approximation to the optimal social welfare. This result holds for any number of buyers and sellers with arbitrary, independent distributions. Moreover, our result continues to hold when there is an additional matroid constraint on the sets of buyers who may get allocated an item. To prove our main result, we devise an extension of sequential posted price mechanisms to two-sided markets. In addition to this, we improve the best-known approximation bounds for the bilateral trade problem.

1 Introduction

In the last decade, algorithmic mechanism design [16] has provided a rich body of methods for designing algorithms that solve computational problems while ensuring that truthful report of input data is the best strategy for the agents who participate in the mechanism. The cornerstone method of mechanism design is the Vickrey-Clarke-Groves (VCG) mechanism [20, 4, 10], which computes an optimal solution to the problem of maximizing the social welfare, i.e., the total valuation of the agents. The VCG mechanism optimizes social welfare and provides the right incentives that make truth-telling a dominant strategy in both one-sided markets (the restricted setting of one non-strategic seller – the auctioneer – and many strategic buyers) and two-sided markets (a setting in which both many buyers and many sellers act strategically and a non-strategic auctioneer lets them trade). However, in two-sided markets it may require the auctioneer subsidizing the trade, and this is highly undesirable in several practical applications.

As opposed to one-sided markets, which were extensively studied by numerous economists (and since more than a decade by computer scientists as well), two-sided markets did not have the same spread. Two-sided markets naturally arise in selling display-ads in web advertisement, the New York Stock Exchange (NYSE), the US FCC spectrum license reallocation, and many other settings with multiple buyers and sellers. For example, ad exchange platforms for selling display-ads face asymmetric information regarding both the valuations of buyers – the value per advertiser's impression shown – as well as about the reservation prices of sellers – the profit that the publisher could obtain by sending the pageviews to competing ad exchanges.

An ideal goal in market design is to devise *individually rational (IR)*, *incentive compatible (IC)* mechanisms that maximize the social welfare of all agents. IR requires that participating in the mechanism is beneficial to all agents. IC requires that truthfully reporting one's preferences to the mechanism is the best strategy for each agent, independently from what the other agents report.

In two-sided markets, a further important requirement is *strong budget-balance (SBB)*, which states the payments of the buyers must entirely and exclusively be transferred to the sellers, i.e., the buyers and the sellers are allowed to trade without leaving to the mechanism any share of the payments, and without the mechanism adding money into the market. For example, in ad ex-

*This work is supported by the EU FET project MULTIPLEX no. 317532, the ERC StG Project PAAl 259515, and the Google Research Award for Economics and Market Algorithms.

†Sapienza University of Rome; colini@dis.uniroma1.it

‡Sapienza University of Rome; dekeijzer@dis.uniroma1.it

§Sapienza University of Rome; leonardi@dis.uniroma.it

¶University of Oxford; stefano.turchetta@cs.ox.ac.uk. Work conducted while visiting Sapienza University of Rome.

change auctions, the intermediation profit of the broker must be limited to a fixed percentage of the revenue of the publishers. This is an important feature of the market, as otherwise it will be perceived as unfair if the mechanism keeps an additional arbitrary share of the payments charged to the advertisers.

Though SBB is a most desired requirement in many applications, it is hard to achieve in practice since it imposes a constraint on the payments that is difficult to satisfy. A weaker version of SBB often considered in the literature is *weak budget-balance (WBB)*, which only requires the mechanism not injecting money into the market.

In this paper we study a standard and simple model of a two-sided market. In its basic form, there is a single type of items for sale. The buyers each want to acquire a single unit of this item, and the sellers each have a single unit to sell. So, to each buyer and seller there is a single number associated, called her *valuation*, which describes how much a buyer or seller values having an item in posession. The valuations of the buyers and sellers are drawn from independent, possibly distinct distributions. We refer to mechanisms that work in such a setting as *double auctions*.[1]

Unfortunately, Myerson and Satterthwaite [15] proved that it is impossible for an IR, Bayesian IC (BIC),[2] and WBB mechanism to maximize social welfare in such a market. This result implies that an IR, BIC, social welfare maximizing double auction necessarily subsidizes the market.

Since then, much of the literature on double auctions [12, 18, 19] has focused on trading off social welfare, incentive compatibility and budget-balance. The complexity of general double auction problems clearly grows if there are restrictions on the sets of buyers and sellers that can trade with one another. Such restrictions may arise from practical constraints. For instance, a restriction on the maximum amount of page views from a publisher that can be allocated to a set of advertisers results into a matroid constraint. A restriction on those advertisers that can appear on the web pages of a publisher translates into a matching constraint. In this work we consider the restriction where there is an arbitrary matroid constraint on the set of buyers who can trade.

Two recent works addressed the problem of approximating social welfare for double auctions and related problems under the WBB requirement. Dütting, Talgam-Cohen, and Roughgarden [8] 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 [12], to obtain IR, *dominant strategy IC (DSIC)*, WBB mechanisms for knapsack, matching and matroid constraints that provide a good approximation of the social. Secondly, Blumrosen and Dobzinski [1] proposed an IR, DSIC, WBB mechanism that 8-approximates the social welfare for *combinatorial exchange markets* – a more general setting where agents can simultaneously be buyers and sellers and there are different kinds of items.

The negative result of [15] was proved for a paradigmatic special case of double auctions, the so-called *bilateral trade* problem. In bilateral trade there is only one unit-demand buyer and one unit-supply seller. McAfee [13] proposed a mechanism for this problem that sets the price equal to the median of the seller's distribution. The mechanism, which is IR, DSIC, and SBB, provides a 2-approximation to the optimal gain from trade (a measure related to social welfare), only if the median of the buyer's distribution is higher than the median of the seller's distribution. Moreover, Blumrosen and Dobzinski [1] showed that this mechanism 2-approximates the optimal social welfare no matter what the relation between the medians is. This is accompanied by a lower bound of 1.1231 and a further improved upper bound of $55/28 \approx 1.9643$ for the problem.

1.1 Overview of the results. This paper studies the problem of designing double auctions that achieve a good approximation to the social welfare. That is, we study mechanism design for two-sided markets, in which both buyers and sellers can exhibit strategic behavior. We work in the Bayesian setting: i.e., buyers' and sellers' valuations are private information and drawn from independent, arbitrary distributions. This is a standard assumption which is motivated by the fact that purchase history is generally available in many practical applications. We want our mechanisms to satisfy *individual rationality (IR)*, *dominant strategy incentive compatibility (DSIC)*, and *strong budget-balance (SBB)*.

Our goal is in finding mechanisms that approximate social welfare of buyers and sellers up to a constant factor, for any number of buyers and sellers. We additionally aim to design, when possible, *fixed-price* mechanisms that do not discriminate between buyers or

[1] Some economics literature uses a stricter definition of a double auction where the mechanism may only charge a single common price at which each trading seller/buyer-pair trades. In the present paper however, we refer to this stricter notion as a *fixed-price double auction*, and instead use the term *double auction* for any mechanism that works in the two-sided market setting.

[2] Bayesian incentive compatibility is a less restrictive form of incentive compatibility. Informally, it only requires that reporting truthfully one's valuation to the mechanism gives an agent the best *expected* utility *conditioned* on the assumption that all other agents also report their valuations truthfully (taking into account the valuation distributions of the other agents).

sellers when pricing identical items. Fixed-price mechanisms are highly desirable because of their fairness and simplicity.

All the mechanism presented in this work satisfy the IR, DSIC,[3] and SBB requirements. Our main results are the following:

- A 16-approximate (non-fixed-price) double auction for when there is an additional general matroid constraint on the set of buyers that may receive an item.
- A 16-approximate fixed-price double auction for k sellers and n buyers. We show that this approximation factor reduces to 8 in the special case the sellers have identical distributions.
- A 4-approximate fixed-price double auction for 1 seller and n buyers.

To our knowledge, these are the first mechanisms that achieve all three major design requirements for double auctions and approximate the optimal social welfare to within a constant factor.

A first ingredient needed to obtain our results is the extension of *sequential posted price mechanisms (SPMs)* [2, 3] to two-sided markets. SPMs are a particularly elegant and well-studied class of mechanisms for one-sided markets. In our paper, we formulate *two-sided sequential posted price mechanisms (2SPMs)* as a tool for designing simple mechanisms in two-sided markets. This tool will hopefully find more applications in future research on double auction design.

A second ingredient of our results is the use of prophet inequalities [11, 7] in combination with the medians of the sellers' distributions [13, 1] in order to set threshold prices that lead to mechanisms that well-approximate the optimal social welfare for general matroid settings.

Another relevant contribution of this paper is the improvement of upper and lower bounds for the bilateral trade problem. We show that there exists a 25/13-approximate IR, DSIC, SBB mechanism, which improves over the best-known approximation factor of 55/28, shown in [1]. Furthermore, we increase the best-known lower bound of [1] from 1.1231 to 1.3360.

1.2 Related work. The impossibility result of [15] states that no two-sided mechanism can simultaneously satisfy BIC, IR, WBB and be socially efficient, even in the simple bilateral trade setting. Follow-up work thus had to focus on designing mechanisms that trade off among these properties.

The following papers studied the convergence rate to social efficiency as a function of the number of agents when all i.i.d. sellers' and buyers' valuations are respectively drawn from regular distributions F and G while obeying to IR and WBB. In a two-sided market with multiple buyers and multiple sellers, Gresik and Satterwhite [9] showed that duplicating the number of agents by τ results in a market where the optimal IR, IC, WBB mechanism's inefficiency drops down as a function of $O(\log \tau/\tau^2)$. They also compared this mechanism with a fixed-price one in the case of a continuum of buyers and sellers. The latter mechanism's inefficiency is $O(1/\sqrt{\tau})$ and therefore performs worse. Further, [17, 19] investigated in detail the family of *c-double auction mechanisms*, where each choice of the index $c \in [0, 1]$ determines a different mechanism in the family, whose inefficiency approaches zero at the quickest possible rate. This mechanism is not truthful despite the incentive to lie being bounded by $O(1/m)$, where m is the number of buyers (assumed to be equal to the number of sellers). We note that the results mentioned above only hold for i.i.d. agents, are solely asymptotic and that the hidden constants must depend on the distributions of the sellers and buyers. Our interest is different as it is in finding *universal* constant approximation guarantees.

In McAfee [12], an IC, WBB, IR double auction is proposed that extracts at least a $(1 - 1/\ell)$ fraction of the maximum social welfare, where ℓ is the number of traders in the optimal solution. In the same paper it is also shown that when the valuations are i.i.d. and the buyers have the same distribution, as well as the sellers, then the inefficiency drops linearly as a function of the minimum between the number of buyers and sellers.

Optimal revenue-maximizing Bayesian auctions were characterized in [14], which provides an elegant tool applicable to single-parameter, one-sided auctions. Numerous subsequent articles dealt with extending these results. Related to our work is [5], which studied maximizing the auctioneer's revenue in Bayesian double auctions. The same objective was studied in [6] yet in the *prior-free* model.

Recently, Dütting et al. [8] provided black-box reductions from double auctions to one-sided mechanisms. They are for a prior-free setting and have applications to matroid, knapsack, and matching feasibility constraints. These mechanisms satisfy only WBB, though.

Also relevant to our work are the IC, IR, SBB mechanisms for bilateral trade of [1], which considers fixed-price mechanisms where the price is set to the median of the seller's distribution. In addition to this, the authors

[3] They actually also satisfy the stronger (but perhaps less well-known) incentive constraint of *weak group strategy proofness (WGSP)*.

proposed a mechanism for combinatorial exchange markets – a broader setting than double auctions in which there are multiple types of items and agents have subadditive valuation functions. This mechanism is IC, IR, WBB and 8-approximates the optimal social welfare. Our main result loses a factor of two against this mechanism, but ensures SBB and additionally allows for a matroid constraint to be specified the set of buyers that may trade.

Sequential posted price mechanisms (SPMs) in one-sided markets have recently gained particular attention from the computer science community 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 with i.i.d. agents [2]. They were then studied for the objective of revenue maximization in [3]. Additionally, the results of [11, 7], which extended prophet inequalities respectively to matroid and polymatroid constrains, found interesting applications to the design of multi-parameter SPMs. We present a two-sided version of SPMs and show that in this setting constant approximations to the optimal social welfare are attainable. To achieve this, we also make use of the prophet inequalities.

2 Preliminaries

2.1 Two-sided mechanisms. We use n to denote the number of buyers and k to denote the number of sellers. We assume that every seller has one item. The items are all identical, and each buyer wants at most one item. To a buyer $i \in [n]$,[4] and a seller $j \in [k]$, there are associated numbers $v_i \in \mathbb{R}_{\geq 0}$, $w_j \in \mathbb{R}_{\geq 0}$ that respectively represent buyer i's valuation and seller j's valuation for an item. These numbers are not publicly known: v_i is known only to buyer i and w_j is known only to seller j. However, for each of the buyers and sellers there is a probability distribution from which her valuation is drawn. This probability distribution is assumed to be public knowledge.

We use g_i and f_j for the probability density function of the valuation of buyer i and seller j, respectively. Let $G_i(x) := \int_{-\infty}^{x} g_i(x)\,dx$ and $F_j(x) := \int_{-\infty}^{x} f_i(x)\,dx$ be the corresponding cumulative distribution functions. We also define the inverse cumulative distribution functions of F_j and G_i (e.g., for seller j we have $F_j^{-1}(y) = \inf\{x \in \mathbb{R} | F(x) \geq y\}$). Then, the median of seller j's distribution is $m_j := F_j^{-1}(\frac{1}{2})$.

Let $\mathbf{v} = (v_1, \ldots, v_n)$ ($\mathbf{w} = (w_1, \ldots, w_k)$) be a valuation profile of the buyers (sellers). Further, let $\mathbf{V}$ ($\mathbf{W}$) be the set of all possible buyers' (sellers') valuation profiles. We use the notation $c_{(i)}$ to denote the i-th largest element of a vector $\mathbf{c}$.[5]

A *(direct revelation) mechanism* $\mathbb{M}$ is a function that takes as input a valuation vector reported by the buyers, and a valuation vector reported by the sellers, and outputs an *allocation* (that specifies which buyers receive an item, and which sellers sell their item) together with a specification of the prices to be paid by the sellers and buyers. Obviously, the output allocation has to satisfy the constraint that the number of buyers getting an item is equal to the number of sellers selling an item. When valuations $(\mathbf{v}, \mathbf{w}) \in \mathbf{V} \times \mathbf{W}$ are reported to the mechanism, we denote the output allocation by $\mathbf{a}(\mathbf{v}, \mathbf{w})$ and we denote the output prices by $\pi(\mathbf{v}, \mathbf{w})$. Additionally, we use the following notation. For every buyer i, $a_i^B(\mathbf{v}, \mathbf{w}) \in \{0, 1\}$ indicates if a buyer has an item after the mechanism is run and $\pi_i^B(\mathbf{v}, \mathbf{w})$ indicates how much the buyer i has to pay. The values $a_j^S(\mathbf{v}, \mathbf{w}) \in \{0, 1\}$ and $\pi_j^S(\mathbf{v}, \mathbf{w})$ are defined analogously for seller j. Since sellers usually do not pay but receive money, π_j^S is always non-positive. We let $\mathbf{a}^B(\mathbf{v}, \mathbf{w})$ denote the vector $(a_1^B(\mathbf{v}, \mathbf{w}), \ldots, a_n^B(\mathbf{v}, \mathbf{w}))$ and analogously define $\mathbf{a}^S(\mathbf{v}, \mathbf{w})$ for the sellers, so that $\mathbf{a}(\mathbf{v}, \mathbf{w}) = (\mathbf{a}^B(\mathbf{v}, \mathbf{w}), \mathbf{a}^S(\mathbf{v}, \mathbf{w}))$.

After the mechanism is run on the reported valuation profile $(\mathbf{v}, \mathbf{w})$, each buyer and seller experiences a certain utility. The utility of a buyer $i \in [n]$ is a function of her true valuation $v_i^* \in \mathbf{V}$ (which is distributed according to G_i), the allocation $a_i^B(\mathbf{v}, \mathbf{w})$ output by the mechanism, and the price π_i^B output by the mechanism. This utility is then given by $u_i^B(v_i^*, a_i^B(\mathbf{v}, \mathbf{w}), \pi_i^B(\mathbf{v}, \mathbf{w})) = a_i^B(\mathbf{v}, \mathbf{w}) \cdot v_i^* - \pi_i^B(\mathbf{v}, \mathbf{w})$ and the utility of a seller $j \in [k]$ is defined similarly as $u_j^S(w_j^*, a_j^S(\mathbf{v}, \mathbf{w}), \pi_j^S(\mathbf{v}, \mathbf{w})) = a_j^S(\mathbf{v}, \mathbf{w}) \cdot w_j^* - \pi_j^S(\mathbf{v}, \mathbf{w})$. Sellers and buyers act strategically so as to maximize their utility.

We refer to the following as desirable properties that a two-sided mechanism $\mathbb{M}$ should satisfy:

- individual rationality (IR) states that all participants do not decrease their utility by participating in the mechanism: So, for all $i \in [n]$ and $j \in [k]$ and for all valuation profiles $(\mathbf{v}, \mathbf{w}) \in \mathbf{V} \times \mathbf{W}$, it holds that $u_i^B(v_i, a_i^B(\mathbf{v}, \mathbf{w}), \pi_i^B(\mathbf{v}, \mathbf{w})) \geq 0$ and $u_j^S(w_j, a_i^S(\mathbf{v}, \mathbf{w}), \pi_j^S(\mathbf{v}, \mathbf{w})) \geq w_j$.
- dominant strategy incentive compatibility (DSIC): it is a dominant strategy for all buyers and sellers to reveal their true valuation. Formally,

[4] For a number $a \in \mathbb{N}$, we use $[a]$ for the set $\{1, \ldots, a\}$.

[5] Subsequently, $v_{(i)}$, $w_{(i)}$, and $m_{(i)}$ respectively are the i-th largest valuation among buyers, the i-th largest valuation among sellers, and the i-th largest value among the medians of the sellers.

for every reported valuation profile $(\mathbf{v}, \mathbf{w})$, for every buyer i, true valuation v_i^* it must hold that $u_i^B(v_i^*, a_i^B((v_i^*, \mathbf{v}_{-i}), \mathbf{w}), \pi_i^B((v_i^*, \mathbf{v}_{-i}), \mathbf{w})) \geq u_i^B(v_i^*, a_i^B(\mathbf{v}, \mathbf{w}), \pi_i^B(\mathbf{v}, \mathbf{w}))$, where $(v_i^*, \mathbf{v}_{-1})$ is the vector obtained from $\mathbf{v}$ by replacing v_i with v_i^*. The sellers' formulation is symmetric.

- strong budget-balance (SBB): the amount of money paid by the buyers is totally and exclusively transferred to the sellers. Formally, for every valuation profile $(\mathbf{v}, \mathbf{w})$, it must hold that $\sum_{i \in [n]} \pi_i^B(\mathbf{v}, \mathbf{w}) + \sum_{j \in [k]} \pi_j^S(\mathbf{v}, \mathbf{w}) = 0$.

Additionally, we desire our mechanism to produce an outcome where the social welfare (i.e., the sum of everyone's utility) is close to optimal. We will measure this as follows. First, let A be the set of feasible allocations of items to the sellers and buyers. For an allocation $\mathbf{a} = (\mathbf{a}^B, \mathbf{a}^S) \in A$ we denote the *social welfare* of $(\mathbf{a}^B, \mathbf{a}^S)$, for valuation profile $(\mathbf{v}, \mathbf{w})$, by $\mathbf{SW}(\mathbf{v}, \mathbf{w}, \mathbf{a}) = \sum_{i \in [n]} a_i^B v_i + \sum_{j \in [k]} a_j^S w_j$. We can then define the *expected social welfare of mechanism* $\mathbb{M}$ as $\mathbf{ALG} = \mathbf{E}[\mathbf{SW}(\mathbf{v}, \mathbf{w}, \mathbf{a}(\mathbf{v}, \mathbf{w}))]$, where the expected value is taken over the randomness of the participants' valuations and the randomness of the mechanism.[6] For a valuation profile $(\mathbf{v}, \mathbf{w})$, define the *optimal allocation* as the allocation $\bar{\mathbf{a}}(\mathbf{v}, \mathbf{w}) \in A$ that maximizes $\mathbf{SW}(\mathbf{v}, \mathbf{w}, \cdot)$. The expected optimal social welfare is then given by $\mathbf{OPT} = \mathbf{E}[\mathbf{SW}(\mathbf{v}, \mathbf{w}, \bar{\mathbf{a}}(\mathbf{v}, \mathbf{w}))]$.

Our goal is to design a mechanism $\mathbb{M}$ that extracts at least a constant fraction α of $\mathbf{OPT}$, i.e., $\mathbf{ALG} \geq \mathbf{OPT}/\alpha$. In such a case, we say that the mechanism *$1/\alpha$-approximates the optimal social welfare.*

2.2 Matroid constraints. A finite matroid $\mathcal{M}$ is a pair $(\mathcal{U}, \mathcal{I})$ where $\mathcal{U}$ is a finite set (called the *ground set*), $\mathcal{I}$ is a family of subsets of $\mathcal{U}$ (called *independent sets*), and $\mathcal{I}$ has the following properties: (i.) $\emptyset \in \mathcal{I}$, (ii.) $\forall I' \subset I \subseteq \mathcal{U}$ if $I \in \mathcal{I}$ then $I' \in \mathcal{I}$, (iii.) if $T, V \in \mathcal{I}$ and $|T| > |V|$ then $\exists t \in T$ such that $V \cup \{t\} \in \mathcal{I}$. A maximum-cardinality independent set of a matroid is called a *basis*.

2.3 One-sided Sequential Posted Price Mechanisms. A *one-sided sequential posted-price mechanism* (SPM) is a mechanism for a one-sided market with only buyers who are interested in getting at most one item. The buyers' valuations are drawn from independent distributions, and there is a matroid constraint $\mathcal{M} = ([n], \mathcal{I})$ on the set of buyers that may get an item. An SPM is parametrized by an ordering of buyers $\boldsymbol{\sigma}$ and a collection of prices $(p_i)_{i \in [n]}$. The price paid by buyer i is p_i if i gets an item, and 0 otherwise. The following simple iterative process determines who purchases an item. At iteration i, it is decided whether buyer σ_i is given an item. Let $T_{<i}$ be the buyers who have received an item during the first $i-1$ iterations. Buyer σ_i is allocated an item iff her reported valuation is at least p_i, and $T_{<i} \cup \{\sigma_i\} \in \mathcal{I}$.

SPMs are conceptually simple and satisfy several other desirable properties such as DSIC and IR. In [3], it is shown that the following SPM, which we call $\mathbb{M}_{\text{one}}$, achieves a 2-approximation to the optimal social welfare: Let q_i be the (a priori) probability that buyer $i \in [n]$ gets an item in the optimal allocation, and let $p_i = G_i^{-1}(1-q_i)$. Let $\boldsymbol{\sigma}$ be the sequence of buyers sorted by non-increasing p_i. The crucial properties satisfied by these prices and this order are (i) the probability that i accepts the offer is exactly q_i,[7] and (ii) the greedy order, which prioritizes agents who buy at higher prices.

Theorem 2.1. ([3]) *The SPM $\mathbb{M}_{one}$ 2-approximates the optimal social welfare in one-sided markets with an arbitrary matroid constraint on the set of buyers that may get an item.*[8]

3 Extending Sequential Posted Price Mechanisms to Two-Sided Markets

We describe how we extend one-sided sequential posted price mechanisms to two-sided markets where there is a matroid constraint $\mathcal{M} = ([n], \mathcal{I})$ on the set of buyers that may receive an item.[9]

A *two-sided sequential posted-price mechanism (2SPM)* is parametrized by an ordering of buyers $\boldsymbol{\sigma}$, an ordering of the sellers $\boldsymbol{\lambda}$, and a collection of prices $(p_{ij})_{i \in [n], j \in [k]}$. Given a $(0,1)$-array A, we use ones(A) to denote the set of all indexes i such that $A[i] = 1$. Moreover, we write $A[a, \dots, b]$ to denote the subarray corresponding to the subset of indexes $\{a, \dots, b\}$. A 2SPM has the following structure. In the description

[6] In the subsequent sections we use the subscripts s and b (e.g., $\mathbf{ALG}_s$) to denote the sum of all sellers' or all buyers' utilities.

[7] To ensure this holds also in case i's distribution contains point masses, the mechanism is slightly modified to offer p_i with probability $q_i/(1 - G_i(p_i))$.

[8] Chawla et al. [3] state this result for revenue maximization, but show in the appendix that it also extends to social welfare maximization.

[9] In this paper, SPMs and its extension to two-sided markets are defined as direct revelation mechanisms. This is solely done for ease of exposition. Under its original definition, an SPM is not a direct revelation mechanism: instead, a buyer solely interacts with the mechanism by replying whether she accepts the take-it-or-leave-it offer made by the mechanism. The two-sided SPM mechanisms and all other mechanisms defined in this paper can also be defined in such a way, which comes with the advantage that there is no need for a buyer or seller to reveal her valuation to the mechanism.

below, we say that buyer i and seller j *accept to trade at price p_{ij}* iff $v_i \geq p_{ij} \geq w_j$. Otherwise, we say that *buyer i refuses the offer* (if $v_i < p_{ij}$) or *seller j refuses the offer* (if $w_j > p_{ij}$).

1. Let A, π be two arrays of length $n+k$
2. Set $A[i] := 0$ for all $i \in [n]$ and $A[i] := 1$ for all $i \geq n+1$, and set $\pi[i] := 0$ for all $i \in [n+k]$
3. Set $i := 1, j := 1$
4. **While** $i \leq n$ and $j \leq k$ **do**:
   (a) If $\text{ones}(A[1, \ldots, n]) \cup \{\sigma_i\} \notin \mathcal{I}$, set $p_{\sigma_i \lambda_j} := +\infty$.
   (b) Offer σ_i and λ_j to trade at price $p_{\sigma_i \lambda_j}$.
   (c) If they both accept, set $A[\sigma_i] := 1, A[n+\lambda_j] := 0, \pi[\sigma_i] := p_{\sigma_i \lambda_j}, \pi[n+\lambda_j] := -p_{\sigma_i \lambda_j}$, and increment i and j.
   (d) If σ_i refuses, increment i. If λ_j refuses, increment j.
5. Every buyer $i \in [n]$ s.t. $A[i] = 1$ gets an item and pays $\pi[i]$ to the auctioneer, who in turn pays $-\pi[n+j]$ to every seller $j \in [k]$ s.t. $A[n+j] = 0$.

Observe that although 2SPMs satisfy SBB and IR, they are not necessarily incentive compatible. As an example, consider a buyer i with valuation v_i who is offered to trade with a seller j at price p_{ij}, which i and j would normally accept. However, if i knows that $p_{i(j+1)} < p_{ij}$, and $w_{j+1} < p_{i(j+1)}$ with a high probability, then she could profit by submitting a lower valuation, so that the mechanism lets i trade at a lower price. The following conditions on the set of prices suffice to ensure a 2SPM is DSIC.[10]

PROPOSITION 3.1. *Let $(p_{ij})_{i\in[n], j\in[k]}$ be a collection of prices of a 2SPM such that for all $i \in [n]$ and all $j, j' \in [k]$, it holds that $p_{ij} = p_{ij'}$. Then, (i) no buyer has an incentive to lie about her preferences. Moreover,(ii) if the prices are posted in a non-increasing order, then also no seller has an incentive to lie about her preferences.*

Proof. First, observe that by assumption prices only depend on i. Thus, a different seller cannot decrease i's price. The price is personal and is independent on i's reported valuation. The order in which the offers are made is also fixed and independent of i's reported valuation. Regarding the second claim, a seller will not report a valuation lower than her true valuation, as it might cause a seller to accept a price below her valuation. A seller will also not report a higher valuation, because it will cause the seller to accept a lower priced offer, by non-increasingness of the posted offers.

[10] We remark that a symmetric version of these conditions (w.r.t. buyers and sellers) would also be sufficient.

The *two-sided fixed-price mechanism (2FPM) with fixed price p* is the mechanism that picks two permutations $\boldsymbol{\sigma}$ on $[n]$ and $\boldsymbol{\lambda}$ on $[k]$ uniformly at random, and runs the 2SPM with parameters $(\boldsymbol{\sigma}, \boldsymbol{\lambda}, (p_{ij})_{i\in[n], j\in[k]})$ where $p_{i,j} = p$ for all $i \in [n], j \in [k]$. The 2FPM is thus a fixed-price double auction which satisfies IR, DSIC, and SBB.[11]

4 Two-sided Sequential Posted Price Mechanisms with Matroid Constraints

In this section we devise a 2SPM for double auctions with general matroid constraints. First, to get a 4α-approximate mechanism under the matroid $\mathcal{M}$, it is enough to get an α-approximation under the matroid $\mathcal{M}'$ defined below.

REMARK 4.1. *Let $\mathcal{M} = ([n], \mathcal{I})$ be a matroid of rank c. Define a new matroid $\mathcal{M}' = ([n], \mathcal{I}')$, where $\mathcal{I}' := \{S \in \mathcal{I} : |S| \leq c/4\}$. Suppose every $i \in [n]$ has a weight $v_i \geq 0$, and let B be a maximum-weight basis of $\mathcal{M}$. Then, there exists a $B' \in \mathcal{I}'$ such that $4 \cdot \sum_{i\in B'} v_i \geq \sum_{i\in B} v_i$.*

Using the above notation, let $(\bar{p}_i)_{i\in[n]}$ and $\boldsymbol{\sigma}$ be the prices and ordering defined as in the SPM of [3] (Theorem 2.1) under the matroid constraint $\mathcal{M}'$. Then, for all $i \in [n]$, set $p_{ij} := p_i := \max\{\bar{p}_i, m_{(k-|B|/2)}\}$, where $m_{(\ell)}$ denotes the ℓ-th largest median of the sellers' distributions. The remainder of this section is dedicated to proving the following result.

THEOREM 4.1. *The mechanism that takes a uniform random permutation $\boldsymbol{\lambda}$ of $[k]$ and runs the 2SPM $\mathbb{M}$ for matroid $\mathcal{M}'$, with parameters $\left(\boldsymbol{\sigma}, \boldsymbol{\lambda}, (p_{ij})_{i\in[n], j\in[k]}\right)$, is IR, DSIC, SBB, and 16-approximates the optimal social welfare under the matroid constraint $\mathcal{M}$.*

Proof. The incentive compatibility of the mechanism $\mathbb{M}$ follows from the definition of the prices and Proposition 3.1. Hence, we focus on showing $\mathbb{M}$'s approximation guarantee, which will follow from the subsequent lemmas, some of which are deferred to the full version of this paper due to space constraints.

LEMMA 4.1. *Let $S := \{j \in [k] : w_j \leq m_{(k-|B|/2)}\}$ and $\bar{S} := \{j \in [k] : w_j \geq m_{(k-|B|/2)}\}$. With probability at least $1/2$ it holds that both $|S| \geq |B|/4$ and $|\bar{S}| \geq (1/2) \cdot (k - |B|/2)$.*

[11] SBB and IR trivially hold. DSIC holds because every 2SPM mechanism in the support of the 2FPM mechanism satisfies the conditions of Proposition 3.1.

Let **ALG** denote the social welfare extracted by $\mathbb{M}$ and **OPT** be the optimal social welfare. Furthermore, let the subscripts s, b respectively denote the sellers' and buyers' contributions to the social welfare in both **ALG** and **OPT** (e.g., $\mathbf{ALG} = \mathbf{ALG}_s + \mathbf{ALG}_b$).

Lemma 4.2. $4/3 \cdot \boldsymbol{ALG}_s \geq \mathbf{E}\left[\sum_{j=1}^{k} w_j\right] \geq \boldsymbol{OPT}_s$.

Proof. Note that $\mathbb{M}$ matches at most $|B'| = |B|/4$ sellers due to $\mathcal{M}'$. Hence, at least $k - |B|/4$ sellers keep their item. Moreover, due to prices being posted in a non-increasing order, a seller j who declines an offer certainly has a valuation w_j that is higher than $w_{j'}$ for any j' who accepts that offer or any subsequent one. So, we can lower-bound $\mathbb{M}$ by assuming that it lets the first $|B|/4$ sellers in the order trade. The mechanism gets the welfare of the last $(k - |B|/4)$ sellers. Due to the sellers being shuffled uniformly at random, every $j \in [k]$ is selected with probability $(k - |B|/4)/k$. Thus, taking the expectation w.r.t. the random order, we extract at least $(1 - |B|/4k) \cdot \sum_{j \in [k]} \mathbf{E}[w_j] \geq (3/4) \cdot \sum_{j \in [k]} \mathbf{E}[w_j]$, following from the fact that $|B| \leq k$. The claim follows by noting that $\mathbf{OPT}_s \leq \sum_{j \in [k]} \mathbf{E}[w_j]$.

Lemma 4.3. $16 \cdot \boldsymbol{ALG}_b + 8/3 \cdot \boldsymbol{ALG}_s \geq \boldsymbol{OPT}_b$.

Proof. First of all, observe that if it is the case that for all $i \in [n]$, $\bar{p}_i \geq m_{(k-|B|/2)}$, then we are using exactly the same prices of [3]'s SPM. Suppose there are at least $|B'|$ items, i.e., sellers willing to accept an offer of at least $m_{(k-|B|/2)}$. By Theorem 2.1, our mechanism gets at least $1/2 \cdot \mathbf{E}\left[\sum_{i \in B'} v_i\right]$, which as previously remarked (Remark 4.1) is at least $1/2 \cdot 1/4 \cdot \mathbf{E}\left[\sum_{i \in B} v_i\right]$. By Lemma 4.1, with probability at least $1/2$, there are indeed at least $|B|/4 = |B'|$ sellers who would accept such an offer. Thus, conditioning on this to hold, we lose at most a factor of two and get at least $(1/16) \cdot \mathbf{E}\left[\sum_{i \in B} v_i\right] = (1/16) \cdot \mathbf{OPT}_b$.

Let A' be the set of buyers who would receive an item by [3]'s SPM and A be the buyers selected by $\mathbb{M}$. Suppose $A' \setminus A \neq \emptyset$ as otherwise we are in the previous case. We show that the welfare loss $\mathbf{E}\left[\sum_{i \in A'} v_i - \sum_{i \in A} v_i\right]$ due to setting a higher price is covered by a fraction of the sellers' welfare. Note that for every $i \in A' \setminus A$, it must be that $v_i < m_{(k-|B|/2)}$; therefore, $\mathbf{E}\left[\sum_{i \in A' \setminus A} v_i\right] \leq |A' \setminus A| \cdot m_{(k-|B|/2)}$.

Let $\bar{S}$ be defined as in Lemma 4.1. By this lemma, $|\bar{S}| \geq 1/2 \cdot (k - |B|/2)$ with probability at least $1/2$. Recall that by assumption $|B| \leq k$ and observe that $(1/2)(k - |B|/2) \geq |B|/4$ is equivalent to $|B| \leq k$. Now observe that

$$2 \cdot \mathbf{E}\left[\sum_{j=1}^{k} w_j\right] \geq 2 \cdot \mathbf{E}\left[\sum_{j \in \bar{S}} w_j\right]$$

$$\geq 2 \cdot \mathbf{Pr}\left[|\bar{S}| \geq \frac{1}{2} \cdot \left(k - \frac{|B|}{2}\right)\right] \cdot \frac{1}{2} \cdot \left(k - \frac{|B|}{2}\right) \cdot m_{(k-|B|/2)}$$

$$\geq \frac{|B|}{4} \cdot m_{(k-|B|/2)} \geq |A' \setminus A| \cdot m_{(k-|B|/2)} \geq \mathbf{E}\left[\sum_{i \in A' \setminus A} v_i\right].$$

By Lemma 4.2, we know that $(8/3) \cdot \mathbf{ALG}_s \geq 2 \cdot \mathbf{E}\left[\sum_{j=1}^{k} w_j\right]$; thus, in total we have that $(8/3) \cdot \mathbf{ALG}_s + 16 \cdot \mathbf{ALG}_b \geq \mathbf{OPT}_b$.

5 Two-sided Fixed-Price Mechanisms

In this section we present two-sided fixed-price mechanisms for some important special cases of two-sided markets. First, we consider the setting where there is no matroid constraint on the buyers. We show that for this setting there exists a *fixed-price* double auction. We propose a 2FPM that achieves the same approximation factor of 16.

Again, let $m_{(\ell)}$ denote the ℓ-th largest median among all sellers' medians. Let $\mathbb{M}_k$ be the 2FPM where the fixed price p is the price computed by the procedure below.

1. Set $q := m_{(\lceil k/2 \rceil)}$.
2. Set $t_0(p) := 0$ and $v_{n+1} := 0$.
3. **For** $1 \leq i \leq \lceil k/4 \rceil$ **do**:
   Set $t_i(p) := \min\{n - \lceil k/4 \rceil + i,$
   $\min\{j \in [n] : j > t_{i-1}(p) \wedge v_j \geq p\}\}$.
4. **For** $i = \lceil k/4 \rceil + 1, \ldots, k$ **do**:
   Set $t_i(p) := \min\{n + 1,$
   $\min\{j \in [n] : j > t_{i-1}(p) \wedge v_j \geq p\}\}$.
5. Set r s.t. $\mathbf{Pr}\left[t_{\lceil k/4 \rceil}(r) = n\right] = \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(r) < n\right] = 1/2$.
6. Set $p := \max\{q, r\}$.

Theorem 5.1. *The 2FPM $\mathbb{M}_k$ is IR, DSIC, SBB, and 16-approximates the optimal social welfare.*[12]

The above theorem will follow from the subsequent remarks and lemmas.

Remark 5.1. *Let A, B be two events such that $A \implies B$, i.e., $\mathbf{Pr}[B|A] = 1$. Then, $\mathbf{Pr}[B] \geq \mathbf{Pr}[A]$.*

Remark 5.2. *Let $Y := \sum_{j=1}^{k} Y_j$ be the sum of k random indicator variables such that for all $j \in [k]$, $\mathbf{E}[Y_j] \geq 1/2$. Then, $\mathbf{Pr}[Y \geq \lceil k/2 \rceil] \geq 1/2$.*

[12] Additionally, in case of i.i.d. sellers, a slight variation of the above procedure (specifically, when we set r s.t. $\mathbf{Pr}\left[t_{\lceil k/2 \rceil}(r) = n\right] = \mathbf{Pr}\left[t_{\lceil k/2 \rceil}(r) < n\right] = 1/2$) can be shown to 8-approximate the optimal social welfare.

REMARK 5.3. *The optimal social welfare $\boldsymbol{OPT}$ is at most the sum of all sellers' valuations plus the sum of the k highest buyers' valuations, i.e.,*

$$\boldsymbol{OPT} \leq \sum_{j=1}^{k} \mathbf{E}\left[w_j\right] + \sum_{i=1}^{k} \mathbf{E}\left[v_{(i)}\right].$$

REMARK 5.4. *For any $p \geq 0$,*

$$\sum_{i=1}^{k} \mathbf{E}\left[v_{(i)}\right] \leq \sum_{i=1}^{k} \mathbf{E}\left[p + \left(v_{(i)} - p\right)^{+}\right] \leq kp + \sum_{i=1}^{n} \mathbf{E}\left[\left(v_i - p\right)^{+}\right].$$

LEMMA 5.1. *For any $p \in \{q, r\}$,*

$$\boldsymbol{ALG}_b \geq \frac{1}{2} \cdot \Big(p \cdot \lceil k/4 \rceil \cdot \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) < n\right] + \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) = n\right] \cdot \sum_{i=1}^{n} \mathbf{E}\left[\left(v_i - p\right)^{+}\right] \Big).$$

Proof. Let $Y_j \in \{0, 1\}$ be a random indicator variable being equal to one if and only if $w_j < p$. Further, let $Y := \sum_{j=1}^{k} Y_j$ be the number of sellers that accept the price. Let $S := \{j \in [k] : m_j \leq m_{(\lceil k/2 \rceil)}\}$. Clearly, $Y \geq \sum_{j \in S} Y_j$. For all $j \in S$, $\mathbf{E}\left[Y_j\right] \geq 1/2$; hence, by Remark 5.2, $\mathbf{Pr}\left[\sum_{j \in S} Y_j \geq |S|/2\right] \geq 1/2$, which in turn implies that $\mathbf{Pr}\left[Y \geq \lceil k/4 \rceil\right] \geq 1/2$.

Thus, there are at least $\lceil k/4 \rceil$ sellers who accept the price p. Observe that $\mathbf{ALG}_b$ is equal to:

$$\begin{aligned}
& \sum_{\ell=1}^{k} \sum_{i=1}^{\ell} \mathbf{E}\left[v_{t_i(p)} | v_{t_i(p)} \geq p \wedge Y = \ell\right] \\
& \cdot \mathbf{Pr}\left[v_{t_i(p)} \geq p \wedge Y = \ell\right] \\
= \quad & \sum_{i=1}^{k} \sum_{\ell=i}^{k} \mathbf{E}\left[v_{t_i(p)} | v_{t_i(p)} \geq p\right] \\
& \cdot \mathbf{Pr}\left[v_{t_i(p)} \geq p\right] \cdot \mathbf{Pr}\left[Y = \ell\right] \\
\geq \quad & \sum_{i=1}^{\lceil k/4 \rceil} \mathbf{E}\left[v_{t_i(p)} | v_{t_i(p)} \geq p\right] \\
& \cdot \mathbf{Pr}\left[v_{t_i(p)} \geq p\right] \cdot \sum_{\ell=i}^{k} \mathbf{Pr}\left[Y = \ell\right] \\
\geq \quad & \mathbf{Pr}\left[Y \geq \lceil k/4 \rceil\right] \cdot \sum_{i=1}^{\lceil k/4 \rceil} \mathbf{E}\left[v_{t_i(p)} | v_{t_i(p)} \geq p\right] \\
& \cdot \mathbf{Pr}\left[v_{t_i(p)} \geq p\right] \\
\geq \quad & \frac{1}{2} \cdot \sum_{i=1}^{\lceil k/4 \rceil} \mathbf{E}\left[v_{t_i(p)} | v_{t_i(p)} \geq p\right] \cdot \mathbf{Pr}\left[v_{t_i(p)} \geq p\right].
\end{aligned}$$

In the first equality we exploited the independence of $F_1, \ldots, F_k, G_1, \ldots, G_n$. In the first inequality we simply summed up to $\lceil k/4 \rceil$ instead of k. In the second one we used that $\mathbf{Pr}\left[Y \geq i\right] \geq \mathbf{Pr}\left[Y \geq \lceil k/4 \rceil\right]$ for any $i \leq \lceil k/4 \rceil$, and the last one follows from Remark 5.2.

Further, observe that by summing and removing p, we get that

$$\begin{aligned}
& \frac{1}{2} \cdot \sum_{i=1}^{\lceil k/4 \rceil} \left(p + \mathbf{E}\left[v_{t_i(p)} - p | v_{t_i(p)} \geq p\right]\right) \cdot \mathbf{Pr}\left[v_{t_i(p)} \geq p\right] \\
= \quad & \frac{1}{2} \cdot \Big(p \cdot \sum_{i=1}^{\lceil k/4 \rceil} \mathbf{Pr}\left[v_{t_i(p)} \geq p\right] + \sum_{i=1}^{\lceil k/4 \rceil} \mathbf{E}\left[\left(v_{t_i(p)} - p\right)^{+}\right] \Big).
\end{aligned}$$

Now, observe that for any $i = 1, \ldots, \lceil k/4 \rceil$, the event $(t_{\lceil k/4 \rceil}(p) < n)$ implies the event $(v_{t_i(p)} \geq p)$; thus, by Remark 5.1, $\mathbf{Pr}\left[v_{t_i(p)} \geq p\right] \geq \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) < n\right]$. Therefore, the above expression is at least

$$\frac{1}{2} \cdot \Big(p \cdot \lceil k/4 \rceil \cdot \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) < n\right] + \sum_{i=1}^{\lceil k/4 \rceil} \mathbf{E}\left[\left(v_{t_i(p)} - p\right)^{+}\right] \Big).$$

Moreover,

$$\begin{aligned}
& \sum_{i=1}^{\lceil k/4 \rceil} \mathbf{E}\left[\left(v_{t_i(p)} - p\right)^{+}\right] \\
= \quad & \sum_{i=1}^{\lceil k/4 \rceil} \sum_{j=1}^{n} \mathbf{E}\left[\left(v_j - p\right)^{+} | t_i(p) = j\right] \cdot \mathbf{Pr}\left[t_i(p) = j\right] \\
= \quad & \sum_{j=1}^{n} \sum_{i=1}^{\lceil k/4 \rceil} \mathbf{E}\left[\left(v_j - p\right)^{+} | j = t_i(p)\right] \cdot \mathbf{Pr}\left[j = t_i(p)\right] \\
= \quad & \sum_{j=1}^{n} \mathbf{E}\left[\left(v_j - p\right)^{+} | j \leq t_{\lceil k/4 \rceil}(p)\right] \cdot \mathbf{Pr}\left[j \leq t_{\lceil k/4 \rceil}(p)\right] \\
\geq \quad & \sum_{j=1}^{n} \sum_{\ell=\lceil k/4 \rceil}^{n} \mathbf{E}\left[\left(v_j - p\right)^{+} | j \leq t_{\lceil k/4 \rceil}(p) \wedge \ell = t_{\lceil k/4 \rceil}(p)\right] \\
& \cdot \mathbf{Pr}\left[\ell = t_{\lceil k/4 \rceil}(p)\right] \cdot \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) = n\right] \\
= \quad & \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) = n\right] \cdot \sum_{j=1}^{n} \sum_{\ell=\lceil k/4 \rceil}^{n} \mathbf{E}\left[\left(v_j - p\right)^{+} | j \leq \ell\right] \\
& \cdot \mathbf{Pr}\left[\ell = t_{\lceil k/4 \rceil}(p)\right] \\
= \quad & \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) = n\right] \cdot \sum_{i=1}^{n} \mathbf{E}\left[\left(v_i - p\right)^{+}\right].
\end{aligned}$$

In the first inequality we exploit the fact that $\mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) \geq j\right] \geq \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) = n\right]$, and in the last equality that v_j is independent of $j \leq \ell$ and we simply rename j to i.

Combining these inequalities, we finally get

$$\mathbf{ALG}_b \geq \frac{1}{2} \cdot \Big(p \cdot \lceil k/4 \rceil \cdot \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) < n\right] + \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) = n\right] \cdot \sum_{i=1}^{n} \mathbf{E}\left[(v_i - p)^+\right] \Big).$$

LEMMA 5.2. *If $p = r$, then*

$$\boldsymbol{ALG}_s \geq \frac{1}{16} \cdot \boldsymbol{OPT}_s.$$

Proof. As in Lemma 5.1, let Y denote the number of sellers that accept p. Moreover, let X indicate the number of buyers who have valuation above p. Then, $\mathbf{ALG}_s$ is at least

$$\sum_{\ell=0}^{k} \mathbf{Pr}[Y = \ell] \cdot \Big(\sum_{j=1}^{k-\ell} \mathbf{E}\left[w_{(j)} | Y = \ell\right] + \sum_{m=0}^{\ell-1} \mathbf{Pr}[X = m] \cdot \frac{\ell - m}{\ell} \cdot \sum_{j=k-\ell+1}^{k} \mathbf{E}\left[w_{(j)} | Y = \ell \wedge X = m\right] \Big),$$

from the sellers since we certainly get the highest $k - \ell$ sellers, who do not accept the price, and in case there are more sellers than buyers accepting, we select them uniformly at random due to the initial shuffling.

Notice that the above expectations do not depend on the conditionings. Hence, we can lower bound $\mathbf{ALG}_s$ further by

$$\sum_{\ell=0}^{\lfloor \frac{3}{4}k \rfloor} \mathbf{Pr}[Y = \ell] \cdot \sum_{j=1}^{k-\ell} \mathbf{E}\left[w_{(j)}\right] + \sum_{\ell=\lfloor \frac{3}{4}k \rfloor + 1}^{k} \mathbf{Pr}[Y = \ell] \cdot \Big(\sum_{j=1}^{k-\ell} \mathbf{E}\left[w_{(j)}\right] + \sum_{m=0}^{\lceil k/4 \rceil} \mathbf{Pr}[X = m] \cdot \frac{\ell - m}{\ell} \cdot \sum_{j=k-\ell+1}^{k} \mathbf{E}\left[w_{(j)}\right] \Big).$$

Denote the first of these two summations by S_1, and the second by S_2.

Note that, as $\ell \leq \lfloor \frac{3}{4}k \rfloor$, the sum S_1 is at least $\mathbf{Pr}\left[Y \leq \lfloor \frac{3}{4}k \rfloor\right] \cdot \sum_{j=1}^{\lceil k/4 \rceil} \mathbf{E}\left[w_{(j)}\right] \geq \frac{1}{4} \cdot \mathbf{Pr}\left[Y \leq \lfloor \frac{3}{4}k \rfloor\right] \cdot \sum_{j=1}^{k} \mathbf{E}[w_j]$.

For S_2 observe that $(\ell - m)/\ell = 1 - m/\ell$ and that $\ell \geq \lfloor \frac{3}{4}k \rfloor + 1$ and $m \leq \lceil k/4 \rceil$. Since $\ell \geq \lfloor \frac{3}{4}k \rfloor + 1 \geq \frac{3}{4}k + 1/4$ and $m \leq \lceil k/4 \rceil \leq k/4 + 3/4$. We have that

$$\frac{m}{\ell} \leq \frac{k+3}{4} \cdot \frac{4}{3k+1} \leq \frac{7}{8}, \text{ for all integers } k > 1.$$

Thus, for any $k > 1$, $1 - m/\ell \geq 1/8$, and so S_2 is at least

$$\sum_{\ell=\lfloor \frac{3}{4}k \rfloor + 1}^{k} \mathbf{Pr}[Y = \ell] \cdot \Big(\sum_{j=1}^{k-\ell} \mathbf{E}\left[w_{(j)}\right] + \frac{1}{8} \cdot \sum_{j=k-\ell+1}^{k} \mathbf{E}\left[w_{(j)}\right] \cdot \sum_{m=0}^{\lceil k/4 \rceil} \mathbf{Pr}[X = m] \Big).$$

Now observe that the event $(t_{\lceil k/4 \rceil}(p) = n)$ implies $(X \leq \lceil k/4 \rceil)$; hence, by Remark 5.1 it holds that $\mathbf{Pr}[X \leq \lceil k/4 \rceil] \geq \mathbf{Pr}\left[t_{\lceil k/4 \rceil} = n\right] = 1/2$. Hence, we obtain that S_2 is at least

$$\begin{aligned} &\sum_{\ell=\lfloor \frac{3}{4}k \rfloor + 1}^{k} \mathbf{Pr}[Y = \ell] \cdot \Big(\sum_{j=1}^{k-\ell} \mathbf{E}\left[w_{(j)}\right] + \frac{1}{16} \cdot \sum_{j=k-\ell+1}^{k} \mathbf{E}\left[w_{(j)}\right] \Big) \\ \geq\ & \frac{1}{16} \cdot \mathbf{Pr}\left[Y > \left\lfloor \frac{3}{4}k \right\rfloor\right] \cdot \sum_{j=1}^{k} \mathbf{E}[w_j]. \end{aligned}$$

Combining the above inequalities, we have that

$$\begin{aligned} \mathbf{ALG}_s &\geq \frac{1}{4} \cdot \mathbf{Pr}\left[Y \leq \left\lfloor \frac{3}{4}k \right\rfloor\right] \cdot \sum_{j=1}^{k} \mathbf{E}[w_j] + \frac{1}{16} \cdot \mathbf{Pr}\left[Y > \left\lfloor \frac{3}{4}k \right\rfloor\right] \cdot \sum_{j=1}^{k} \mathbf{E}[w_j] \\ &\geq \frac{1}{16} \cdot \sum_{j=1}^{k} \mathbf{E}[w_j] \geq \frac{1}{16} \cdot \mathbf{OPT}_s \end{aligned}$$

LEMMA 5.3. *If $p = r$, then*

$$\boldsymbol{ALG}_b \geq \frac{1}{16} \cdot \boldsymbol{OPT}_b.$$

Proof. Recall Lemma 5.1 and observe that since $\mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) = n\right] = \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) < n\right] = 1/2$, we have that

$$\mathbf{ALG}_b \geq \frac{1}{4} \cdot \Big(\lceil k/4 \rceil \cdot p + \sum_{i=1}^{n} \mathbf{E}\left[(v_i - p)^+\right] \Big).$$

Finally, recalling Remark 5.3, we get that

$$\begin{aligned} \mathbf{ALG}_b &\geq \frac{1}{4} \cdot \Big(\lceil k/4 \rceil \cdot p + \sum_{i=1}^{n} \mathbf{E}\left[(v_i - p)^+\right] \Big) \\ &\geq \frac{1}{16} \cdot \sum_{i=1}^{k} \mathbf{E}\left[v_{(i)}\right] \geq \frac{1}{16} \cdot \mathbf{OPT}_b. \end{aligned}$$

LEMMA 5.4. *If $p = q$, then*

$$\boldsymbol{ALG} \geq \frac{1}{12} \cdot \boldsymbol{OPT}.$$

Proof. Let $S \subseteq [k]$ be the set of sellers whose median is at least p. Observe that $|S| \geq \lfloor k/2 \rfloor$ and that

$$\begin{aligned}\mathbf{ALG}_s \quad &\geq \quad \frac{1}{2} \cdot \sum_{j \in S} \mathbf{E}\left[w_j | w_j \geq p\right] \\ &\quad + \quad \sum_{j \in [k] \setminus S} \mathbf{Pr}\left[w_j \geq p\right] \cdot \mathbf{E}\left[w_j | w_j \geq p\right].\end{aligned}$$

Moreover, $\mathbf{OPT}_s$ is at most

$$\begin{aligned}&\sum_{j \in S} \mathbf{E}\left[w_j\right] + \sum_{j \in [k] \setminus S} \mathbf{E}\left[w_j\right] \\ = \quad &\sum_{j \in S} \mathbf{E}\left[w_j\right] + \sum_{j \in [k] \setminus S} \Big(\mathbf{Pr}\left[w_j \geq p\right] \cdot \mathbf{E}\left[w_j | w_j \geq p\right] \\ &+ \mathbf{Pr}\left[w_j < p\right] \cdot \mathbf{E}\left[w_j | w_j < p\right]\Big) \\ \leq \quad &\sum_{j \in S} \mathbf{E}\left[w_j\right] + \sum_{j \in [k] \setminus S} \mathbf{Pr}\left[w_j \geq p\right] \cdot \mathbf{E}\left[w_j | w_j \geq p\right] \\ &+ \left\lceil \frac{k}{2} \right\rceil \cdot p.\end{aligned}$$

Since $\mathbf{OPT}_b \leq kp + \sum_{i=1}^{n} \mathbf{E}\left[(v_i - p)^+\right]$, we obtain that $\mathbf{OPT}$ is at most

$$\begin{aligned}\leq \quad &\sum_{i=1}^{n} \mathbf{E}\left[(v_i - p)^+\right] + \left(k + \left\lceil \frac{k}{2} \right\rceil\right) \cdot p \\ + \quad &\sum_{j \in S} \mathbf{E}\left[w_j\right] + \sum_{j \in [k] \setminus S} \mathbf{Pr}\left[w_j \geq p\right] \cdot \mathbf{E}\left[w_j | w_j \geq p\right].\end{aligned}$$

Note that

$$\sum_{j \in S} \mathbf{E}\left[w_j | w_j \geq p\right] \geq \sum_{j \in S} \mathbf{E}\left[w_j\right], \text{ and}$$

$$5 \cdot \sum_{j \in S} \mathbf{E}\left[w_j | w_j \geq p\right] \geq \left(k + \left\lceil \frac{k}{2} \right\rceil\right) \cdot p,$$

where the latter inequality holds since

$$\sum_{j \in S} \mathbf{E}\left[w_j | w_j \geq p\right] \geq \lfloor k/2 \rfloor \cdot p \geq \frac{k-1}{2} \cdot p,$$

and $5 \cdot (k-1) \geq 3k+1$ for all $k \geq 3$. For $k = 2$, the above statement trivially holds too as $\lfloor k/2 \rfloor = \lceil k/2 \rceil = 1$.

Observe that $\mathbf{ALG}_b$ is at least

$$\begin{aligned}&\sum_{i=1}^{\lceil k/4 \rceil} \mathbf{E}\left[v_{t_i(p)} | v_{t_i(p)} \geq p \wedge Y \geq \lceil k/4 \rceil\right] \\ &\cdot \mathbf{Pr}\left[v_{t_i(p)} \geq p \wedge Y \geq \lceil k/4 \rceil\right] \\ \geq \quad &\frac{1}{2} \cdot \Big(\mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) < n\right] \cdot \lceil k/4 \rceil \cdot p \\ &+ \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) = n\right] \cdot \sum_{i=1}^{n} \mathbf{E}\left[(v_i - p)^+\right]\Big) \\ \geq \quad &\frac{1}{4} \cdot \sum_{i=1}^{n} \mathbf{E}\left[(v_i - p)^+\right],\end{aligned}$$

where the first and second inequalities were previously shown, and the last one follows from $\mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) < n\right] \cdot \lceil k/4 \rceil \cdot p \geq 0$ and $\mathbf{Pr}\left[t_{\lceil k/4 \rceil}(p) = n\right] \geq \mathbf{Pr}\left[t_{\lceil k/4 \rceil}(r) = n\right] = 1/2$ since $p \geq r$.

So, in total we get that

$$\begin{aligned}\mathbf{ALG} \quad \geq \quad &\frac{1}{4} \cdot \sum_{i=1}^{n} \mathbf{E}\left[(v_i - p)^+\right] + \frac{1}{2} \cdot \sum_{j \in S} \mathbf{E}\left[w_j | w_j \geq p\right] \\ &+ \sum_{j \in [k] \setminus S} \mathbf{Pr}\left[w_j \geq p\right] \cdot \mathbf{E}\left[w_j | w_j \geq p\right],\end{aligned}$$

and therefore $12 \cdot \mathbf{ALG} \geq \mathbf{OPT}$.

Next, we consider the special case in which $k = 1$, i.e., there is a single seller and n buyers. This is harder than the bilateral trade setting as it also involves choosing the buyer with highest valuation among all of them. Let m be the median of the seller's distribution. Moreover, let $t_1(r)$ be the first buyer that accepts price r, i.e., $\min\{i \in [n] : v_i \geq r\}$, or the last buyer n, in case the former does not exist. Define r such that it solves $\mathbf{Pr}\left[t_1(r) < n\right] = \mathbf{Pr}\left[t_1(r) = n\right] = 1/2$, and set the price $p := \max\{r, m\}$. Let $\mathbb{M}_1$ be the mechanism that sells the item to $t_1(p)$ if and only if $v_{t_1(p)} \geq p$ and $w_1 \leq p$.

THEOREM 5.2. *For $k = 1$, mechanism $\mathbb{M}_1$ is IR, DSIC, SBB, and 4-approximates the optimal social welfare.*

This theorem can be proven with similar techniques used in the previous proof; therefore, due to space constraints, we defer it to the full version of the paper.

6 Bilateral Trade

Bilateral trade is the most fundamental special case of a double sided market: There is only one buyer and one seller. The seller has one item in posession. The valuation for the seller and buyer are drawn independently from two (possibly distinct) distributions on $\mathbb{R}$. We denote the cumulative distribution function of the seller and buyer by F and G respectively. We assume without loss of generality that these probability distributions are finite and discrete, and we denote by f and g the probability mass functions of the seller and buyer respectively. Hence, a bilateral trade instance is a pair (f, g) of such probability mass functions. A strongly budget-balanced mechanism for (f, g) must decide whether the seller will give the item to the buyer, and what amount the buyer pays to the seller.

For $p \in \mathbb{R}_{\geq 0}$, a *p-fixed price mechanism* is a mechanism that lets the seller and buyer trade if and only if the buyer's reported valuation exceeds p and the seller's reported valuation is below p. When trade occurs, the price that the buyer pays to the seller is p. There is an additional tie-breaking rule that specifies three things:

- Whether or not to trade in the case the reported valuations are both equal to p.
- Whether or not to trade in case the reported valuation of the buyer is p, while that of the seller is less than p.
- Whether or not to trade in case the reported valuation of the seller is p, while that of the buyer is greater than p.

Thus, for a fixed $p \in \mathbb{R}_{\geq 0}$ there are eight p-fixed price mechanisms in total. It is easy to see that such mechanisms are SBB, DSIC, and IR. In fact, we can show that p-fixed price mechanisms characterize the entire set of SBB, IC, and IR bilateral trade mechanisms.

PROPOSITION 6.1. *Let $\mathbb{M}$ be a bilateral trade mechanism that is SBB, DSIC, and IR. Then $\mathbb{M}$ is a p-fixed price mechanism for some $p \in \mathbb{R}_{\geq 0}$.*

Proof. It is clear that if $\mathbb{M}$ does not let the buyer and seller trade, then $\mathbb{M}$ does not charge any price, as it would violate IR for the buyer. If trade does occur, then the price charged is at most the buyer's valuation (otherwise it would again violate IR for the buyer), and at least than the buyer's valuation (otherwise it would violate IR for the seller).

For a valuation profile (v, w) (where v is the valuation of the buyer), let $p(v, w)$ denote the price charged by $\mathbb{M}$ when (v, w) is reported. Fix a valuation w of the seller. If trade occurs for a valuation profile (v, w) then for each valuation profile (v', w) where $v' > v$, trade also occurs, and $p(v', w) \geq p(v, w)$, otherwise the buyer would be better off reporting v' if her valuation is v, which would violate DSIC. Likewise, it must be that $p(v', w) \leq p(v, w)$, otherwise the buyer would be better off reporting v if her valuation is v'. This implies that for any valuation w of the seller, there is a threshold point $\bar{v}(w)$ such that trade occurs for all profiles (v, w) where $v > \bar{v}(w)$. By symmetry, we conclude that for any valuation v of the buyer, there is a threshold point $\bar{w}(v)$ such that trade occurs at price $\bar{w}(v)$ for all profiles (v, w) where $w < \bar{w}(v)$. We will refer to this as *trade monotonicity*.

Let (v, w) and (v', w') be two distinct valuation profiles where trade occurs. We now prove that $p(v, w) = p(v', w')$.

- First, if $w < \bar{w}(v')$, then trade occurs at profile (v', w), so we derive that $p(v', w') = p(v', w) = p(v, w)$.
- If $w > \bar{w}(v') = \bar{v}(w')$, then from $v > w \geq \bar{v}(w')$ we see that trade occurs at profile (v, w'), so $p(v', w') = p(v, w') = p(v, w)$.
- If $w = \bar{w}(v')$ and trade does not occur at (w', v), we consider three cases: If $v > w'$, then trade does not occur at (w', v'), which is a contradiction. If $v = w'$ and $v > v'$, then trade occurs at profile (w', v), so $p(w', v') = p(w', v) = p(w, v)$. If $v = w'$ and $v = v'$, then $w < v$ (because otherwise (v, w) and (v', w') would not be distinct), and $\bar{w}(v) = \bar{w}(v') = w < v = w$, i.e., the value w lies above the threshold point $\bar{w}(v)$. This contradicts that trade occurs at (v, w).
- If $w = \bar{w}(v')$ and trade occurs at (w', v), then $p(w, v) = p(w', v) = p(w', v')$.

Therefore, there is a single threshold p such that if trade occurs, then it happens at price p. Moreover, when we include the trade monotonicity property and the IR constraints, we conclude that trade occurs if the buyer's valuation is above p and the seller's valuation is below p. When one of these valuations is exactly p, then trade may occur, although by DSIC it must be that for two reported valuation profiles (p, w) and (p, w') where $w' < p$ and $w < p$, trade occurs at (p, w) if and only if trade occurs at (p, w'). Likewise for reported valuation profiles (v, p) and (v', p) where $v > p$ and $v' > p$, trade occurs at (p, w) if and only if trade occurs at (p, w'). It is now straightforward to verify that this is precisely the definition of a p-fixed price mechanism.

It is clear that for a fixed $p \in \mathbb{R}_{\geq 0}$, among the eight p-fixed price mechanisms, the one with the highest expected social welfare uses the tie breaking rule where trade always occurs in case of ties. We focus on this subclass of fixed-price mechanisms for the remainder of this section.

Let M_s be the median of the seller. Dobzinski and Blumrosen show in [1] that the expected social welfare of the M_s-fixed price mechanism is a 2-approximation to the optimal social welfare. They also show that, in order to improve upon this approximation factor, the fixed price p necessarily needs to be a function of the buyer's distribution in addition to only the seller's distribution. They prove subsequently that for any bilateral trade instance (f, g) there exists a price $p(f, g)$ such that the expected social welfare of the $p(f, g)$-fixed price mechanism is a 55/28-approximation to the optimal social welfare. Additionally, they provide a lower bound bilateral trade instance for which no fixed price mechanism achieves an expected social welfare that 1.1231-approximates the expected optimal social welfare.

In this section we are interested in what is the best possible approximation factor to the optimal welfare that can be achieved by an SBB, DSIC, and IR mechanism on any bilateral trade intance. We improve the lower bound of 1.1231 to 1.3360 and we improve the upper bound of 55/28 to 25/13.

Theorem 6.1. *Let (f, g) be a bilateral trade instance. There exists a (SBB, DSIC, and IR) fixed price mechanism of which the expected social welfare* 25/13-*approximates the optimal social welfare.*

Theorem 6.2. *There exists a bilateral trade instance for which no SBB, DSIC, and IR mechanism gives an expected social welfare that* 1.3360-*approximates the optimal social welfare.*

We prove the latter theorem first, as it is the easiest of the two.

Proof. For $\epsilon \in \mathbb{R}_\epsilon$, let (f_ϵ, g_ϵ) be the following instance where f_ϵ and g_ϵ defined as follows:

- The support of f_ϵ is $\{0, \epsilon, 2\epsilon\}$. The support of g_ϵ is $\{\epsilon, 2\epsilon, 1\}$.
- The probabilities of the seller are: $f_\epsilon(0) = 4/(5+\sqrt{17})$, $f_\epsilon(\epsilon) = 2(\sqrt{17}-4)$, $f_\epsilon(2\epsilon) = (13-3\sqrt{17})/2$.
- The probabilities of the buyer are: $g_\epsilon(\epsilon) = (7-\sqrt{17})/4 - \epsilon$, $g_\epsilon(2\epsilon) = (\sqrt{17}-3)/4$, $g_\epsilon(1) = \epsilon$.

It is straightforwardly verified that this is a valid bilateral trade instance, i.e., the probabilities of both distributions indeed sum to 1. Because the support of the seller has three points, we know that the p-fixed price mechanism that attains the maximum social welfare is $\mathbb{M}_0$, $\mathbb{M}_\epsilon$ or $\mathbb{M}_{2\epsilon}$. By manually computing the expected social welfare of these three mechanisms, we conclude that $\mathbb{M}_\epsilon$ performs best and achieves an expected social welfare of

$$\epsilon(6-\sqrt{17}-\epsilon 2(\sqrt{17}-4)).$$

The expected optimal social welfare is

$$\epsilon\left(19-4\sqrt{17}-\epsilon\frac{3(5-\sqrt{17})}{2}\right).$$

Therefore, the factor by which the best SBB, DSIC, IR mechanism approximates the expected optimal social welfare is

$$\frac{19-4\sqrt{17}-\epsilon 3(5-\sqrt{17})/2}{6-\sqrt{17}-\epsilon 2(\sqrt{17}-4)}.$$

Taking ϵ to 0 then proves the claim.

$$\lim_{\epsilon\to 0}\frac{19-4\sqrt{17}-\epsilon 3(5-\sqrt{17})/2}{6-\sqrt{17}-\epsilon 2(\sqrt{17}-4)} = \frac{19-4\sqrt{17}}{6-\sqrt{17}} \approx 1.3360.$$

The remainder of this section is devoted to proving the upper bound of 25/13. We will denote the p-fixed price mechanism by $\mathbb{M}_p$. Let X be a random variable with distribution F and let Y be a random variable with distribution G. Throughout this section, for an event E we will write $\mathbf{1}(E)$ to denote the indicator function that maps to 1 if E is true, and to 0 otherwise. Given a bilateral trade instance (f, g), the expected social welfare of $\mathbb{M}_p$ can be written as $\mathbf{E}[X] + \mathrm{GFT}_{f,g}(p)$, where $\mathrm{GFT}_{f,g}(p) = \mathbf{E}[(Y-X)\mathbf{1}(Y \geq p \geq X)]$ is the *gain from trading at price* p. Moreover, the expected optimal social welfare can be written as $\mathbf{E}[X] + \mathrm{GFT}_{f,g}(\mathbf{OPT})$, where $\mathrm{GFT}_{f,g}(\mathbf{OPT}) = \mathbf{E}[(Y-X)\mathbf{1}(Y > X)]$ is the *optimal gain from trade*. The following lemma bounds $\mathrm{GFT}_{f,g}(\mathbf{OPT})$ in terms of $\mathbf{E}[Y]$ and $\mathrm{GFT}_{f,g}(p)$.

Lemma 6.1. *Let $c \in [0,1]$ and let f and g be two finite discrete probability distributions on $\mathbb{R}_{\geq 0}$. Define $F^{-1}(c)$ as the point p such that $F(p) \geq c$ and $F(p-\epsilon) < c$ for all $\epsilon \in \mathbb{R}_{>0}$. Then it holds that*

$$GFT_{f,g}(\boldsymbol{OPT}) \leq c\mathbf{E}[Y] + \frac{1}{c}GFT_{f,g}(F^{-1}(c)).$$

Proof. For convenience we assume that there exists a point p such that $F(p) = c$. (Otherwise one may "perturb" f slightly such that this is the case.) Let S be the support of f and B be the support of g. We split the expected optimal gain from trade into three summations $\mathrm{GFT}_{\leq p}$, $\mathrm{GFT}_{>p}$, and GFT_p, defined as follows.

$$\begin{aligned}\mathrm{GFT}_{\leq p} &= \sum_{x\in S, y\in B: x<y\leq p} f(x)g(y)(y-x),\\ \mathrm{GFT}_{>p} &= \sum_{x\in S, y\in B: p<x<y} f(x)g(y)(y-x),\\ \mathrm{GFT}_p &= \mathrm{GFT}_{f,g}(p).\end{aligned}$$

It can be checked that these three summations indeed sum up to $\mathrm{GFT}_{f,g}(\mathbf{OPT})$. We bound $GFT_{\leq p}$ in terms of $\mathbf{E}[Y]$:

$$\begin{aligned}
\mathrm{GFT}_{\leq p} &= \sum_{x\in S, y\in B: x<y\leq p} f(x)g(y)(y-x), \\
&\leq \sum_{x\in S, y\in B: x<y\leq p} f(x)g(y)y \\
&\leq \sum_{x\in S: x<p} \sum_{y\in B: y\leq p} f(x)g(y)y \\
&\leq c \sum_{y\in B: y\leq p} g(y)y \\
&\leq c\mathbf{E}[Y],
\end{aligned}$$

where the third inequality follows because $F(p) = c$. We bound $\mathrm{GFT}_{>p}$ in terms of GFT_p:

$$\begin{aligned}
\mathrm{GFT}_{>p} &= \sum_{y\in B} \sum_{x\in S: p<x<y} f(x)g(y)(y-x), \\
&\leq \sum_{y\in B} \sum_{x\in S: p<x<y} f(x)g(y)(y-p) \\
&\leq (1-c) \sum_{y\in B: p<y} g(y)(y-p) \\
&= \frac{(1-c)}{c} \sum_{x\in S: x<p} f(x) \cdot \sum_{y\in B: p<y} g(y)(y-p) \\
&= \frac{(1-c)}{c} \sum_{x\in S: x<p} \sum_{y\in B: p<y} f(x)g(y)(y-x) \\
&= \frac{(1-c)}{c} \mathrm{GFT}_p,
\end{aligned}$$

where the second inequality follows because the total probability mass of f above point p is at most $(1-c)$, i.e., $F(p) = c$. The third inequality follows for the same reason. Therefore,

$$\begin{aligned}
\mathrm{GFT}_{f,g}(\mathbf{OPT}) &= \mathrm{GFT}_{\leq p} + \mathrm{GFT}_{>p} + \mathrm{GFT}_p \\
&\leq c\mathbf{E}[Y] + \left(1 + \frac{1-c}{c}\right) \mathrm{GFT}_p \\
&\leq c\mathbf{E}[Y] + \frac{1}{c}\mathrm{GFT}_{f,g}(F^{-1}(c)).
\end{aligned}$$

For a bilateral trade instance (f, g), we will write $R(f, g)$ to denote

$$\begin{aligned}
&\max\left\{\frac{\mathbf{E}[X] + \mathrm{GFT}_{f,g}(p)}{\mathbf{E}[X] + \mathrm{GFT}_{f,g}(\mathbf{OPT})} : p \in \mathbb{R}_{\geq 0}\right\} \\
&= \max\left\{\frac{\mathbf{E}[X] + \mathrm{GFT}_{f,g}(p)}{\mathbf{E}[X] + \mathrm{GFT}_{f,g}(\mathbf{OPT})} : p \in S\right\},
\end{aligned}$$

where S is the support of f. Thus, $R(f, g)$ is the best (inverse) approximation factor to the social welfare that is achievable by a fixed-price mechanism on instance (f, g). For a fixed price $p \in \mathbb{R}_{\geq 0}$, we define the *missed gain from trade at* p as

$$\begin{aligned}
\mathrm{MGFT}_{f,g}(p) = &\sum_{x\in S, y\in B: x<y<p} f(x)g(y)(y-x) \\
&+ \sum_{x\in S, Y\in B: p<x<y} f(x)g(y)(y-x).
\end{aligned}$$

Using this notion, we may rewrite the ratio $R(f, g)$ in an alternative form $1 - R'(f, g)$, where

$$\begin{aligned}
R'(f, g) &= \frac{\min\{\mathrm{MGFT}_{f,g}(p) : p \in S\}}{\mathbf{E}[X] + \mathrm{GFT}_{f,g}(\mathbf{OPT})} \\
&= \frac{\min\{\mathrm{MGFT}_{f,g}(p) : p \in S\}}{\mathbf{E}[\max\{X, Y\}]}.
\end{aligned}$$

We will now proceed with reasoning about bilateral trade instances on which any fixed price mechanism generates the worst possible relative expected social welfare. We will establish various structural properties that such a worst-case instance satisfies. Let $\mathcal{I}$ be the set of all bilateral trade instances. Let $\mathcal{J}$ be the set of bilateral trade instances in which the support of the seller has the form $\{0\}\cup T$ and the support of the buyer's distribution is $T \cup \{1\}$, or $T \cup \{0, 1\}$ where T is a finite subset of $(0, 1)$. I.e, the buyer's and seller's support are subsets of $[0, 1]$, and these sets coincide on $(0, 1)$. Value 0 is in the support of the seller (and possibly the buyer), and value 1 is in the support of the buyer.

Lemma 6.2. *Let (f, g) be an instance in $\mathcal{I}$. There exists an instance $(f', g') \in \mathcal{J}$ such that $R(f, g) \geq R(f', g')$. Hence,*

$$\inf\{R(f, g) : (f, g) \in \mathcal{I}\} = \inf\{R(f, g) : (f, g) \in \mathcal{J}\}$$

Proof. Let $(f, g) \in \mathcal{I}$ be an instance that is not in $\mathcal{J}$. Let S and B denote the support of f and g respectively. We prove the claim by considering various operations on f and g, and we prove that for any of these operations, $R(f, g)$ does not increase (or equivalently: $R'(f, g)$ does not decrease). After having performed all of these operations, our instance (f, g) will have been transformed into an appropriate instance in $\mathcal{J}$ while $R(f, g)$ will not have increased.

Let $x = \min S$. If $\ell > 0$. Consider moving some amount of probability mass of f from x to 0. By performing this operation the denominator of $R'(f, g)$, (i.e., the expected optimal social welfare) decreases, as the denominator can be written as $\mathbf{E}[\max\{X, Y\}]$. Moreover, all of the $|S|$ values in the min-expression in $R'(f, g)$ do not decrease. We conclude that $R'(f, g)$ increases and thus that $R(f, g)$ increases.

Next, let $x = \max S$. If $x \geq \max B$, consider removing all probability mass from x and rescaling the remaining probabilities by an appropriate $\lambda > 1$ such that $\sum_{x' \in S} \lambda f(x') = 1$. It is easy to see that that $R'(f, g)$ increases, as a scaling operation leaves the ratio $R'(\cdot)$ unaffected and removing probability mass only decreases the denominator of $R'(f, g)$.

Next, we scale the points in the supports of f and g by an appropriate positive constant, and we obtain a probability distribution where $1 = \max B$. It can easily be checked that $R(f, g)$ does not change by performing such a scaling operation.

Now, suppose that there are two points $x, x' \in S$, $x < x'$ such that there is no point of B in between x and x'. We consider the operation of augmenting f by removing all probability mass from the points x and x', and putting it on the point $x'' = (f(x)x + f(x')x')/(f(x)+f(x'))$. Let f' be the resulting distribution and let X' be a random variable with distribution f'. It then holds that $\mathbf{E}[X'] = \mathbf{E}[X] - f(X)x - f(x')x' + (f(x)+f(x'))(f(x)x + f(x')x')/(f(x)+f(x')) = \mathbf{E}[X]$. Likewise, for any point $y \in B$ we see that

$$g(y)f(x)(y-x)+g(y)f(x')(y-x') = g(y)(f(x)+f(x'))x''$$

Thus $\mathrm{GFT}_{f,g}(x) \leq \mathrm{GFT}(x') = \mathrm{GFT}_{f',g}(x'')$, as every pair of terms $g(y)f(x)(y-x)$ and $g(y)f(x)(y-x)$ that occur in $\mathrm{GFT}_{f,g}(x)$ gets replaced by the term $g(y)(f(x)+f(x'))x''$, which is equal. For the same reason it holds that $\mathrm{GFT}_{f,g}(z) = \mathrm{GFT}_{f',g}(Z)$ for $z \in B \setminus \{x\}$ and $\mathrm{GFT}_{f,g}(\mathbf{OPT}) = \mathrm{GFT}_{f',g}(\mathbf{OPT})$.

Therefore, $R(f', g) = R(f, g)$. We can repeat this operation until there are no two points left in S that have no point in B in between them. By similar operations, we obtain an instance for which there are no two points $y, y' \in B$ with a point of S in between y and y'.

Lastly, suppose there are points $x \in S$ and $y \in B$, where $y < x$ and there is no point $z \in S \cup B$ with $y < z < x$. We consider the operation of augmenting f by moving point x (and all of its probability mass) to y. Clearly, the denominator of $R'(f, g)$ does not increase by performing this operation. Moreover, $\mathrm{MGFT}_{f,g}(x)$ increases for all $x \in S$, so $R'(f, g)$ increases.

From now on, we restrict our attention to instances in $\mathcal{J}$ exclusively, and we will use the following notation: m is the cardinality of the support of the buyer (and hence the cardinality of the support of the seller is either m or $m+1$). Let $x_1 = 0$ and let $x_2 < \ldots < x_m$ be such that $\{x_1, \ldots, x_m\}$ is the support of the seller. Let $x_{m+1} = 1$, so that by the above lemma we may assume that $\{x_2, \ldots, x_{m+1}\}$ or $\{x_1, \ldots, x_{m+1}\}$ is the support of the buyer. For $i \in [m]$, we abbreviate $f(x_i)$ to p_i and for $j \in [m+1]$ we abbreviate $g(x_j)$ to q_j (where possibly $q_1 = 0$). The values of these parameters depend on the instance, but whenever we use this notation it will be clear from context which instance is being discussed.

For $\epsilon \geq 0$, let $\mathcal{K}(\epsilon)$ be the subset of instances in $\mathcal{J}$ where for all $i, j \in [m-1]$, the gain from trade of $\mathbb{M}_{x_i}$ is equal to the gain from trade of $\mathbb{M}_{x_j}$ and the gain from trade of $\mathbb{M}_{x_m}$ is between 1 and $(1-\epsilon)$ times the gain from trade of $\mathbb{M}_{x_i}$, for all $i \in [m-1]$. The next lemma shows that we may safely restrict our attention to studying bilateral trade instances in the set $\mathcal{K}(\epsilon)$.

We define a *canonical instance* as an instance in $\mathcal{J}$ such that there is no other instance (f', g') in $\mathcal{J}$ for which it holds that both $R(f', g') \leq R(f, g)$ and the seller and buyer have in (f', g') a support of lower cardinality than in (f, g)

LEMMA 6.3. *Let (f, g) be a canonical instance in $\mathcal{J}$. For any $\epsilon > 0$, there exists a canonical instance $(f', g') \in \mathcal{K}(\epsilon)$ such that $R(f', g') \leq R(f, g)$. Hence,*

$$\inf\{R(f, g) : (f, g) \in \mathcal{J}\} = \inf\{R(f, g) : (f, g) \in \mathcal{K}(\epsilon)\}.$$

Moreover,

- *the support of g is the support of g',*
- *$g(x_m) > g'(x_m)$,*
- *and $g(x_m) + g(x_{m+1}) \geq g'(x_m) + g'(x_{m+1})$.*

(The last part of this lemma are technical properties that are needed subsequently.)

Proof. Pick any $\epsilon > 0$ and let (f, g) be a canonical instance. We prove the claim by considering various operations on f and g, and by showing that for any of these operations, $R(f, g)$ does not increase (or equivalently: $R'(f, g)$ does not decrease). We show that an appropriate combination of such operations turns instance (f, g) into an instance in $\mathcal{K}(\epsilon)$ such that $R(f, g)$ has not increased.

We first define an operation $\mathrm{reduce}(i, \varepsilon, (f, g))$, where $i \in [m-1]$ and $\varepsilon \in [0, p_{i+1}]$ as follows: Augment f by subtracting a probability mass of ε from p_{i+1}, and adding it to p_i. By doing this, the denominator of $R'(f, p)$ decreases. Moreover, we see that

- the value $\mathrm{MGFT}_{f,g}(x_i)$ decreases by $\varepsilon \sum_{i+1 \leq j \leq m+1} q_j(x_j - x_{i+1})$;
- for $k > i$ the value $\mathrm{MGFT}_{f,g}(x_k)$ increases by $\varepsilon(x_{i+1} - x_i) \sum_{i+1 \leq j \leq k} q_j$;
- and for $k < i$ the value $\mathrm{MGFT}_{f,g}(x_k)$ increases by $\varepsilon(x_{i+1} - x_i) \sum_{i+1 \leq j \leq m+1} q_j$.

We also define the operation $\text{reduce}(m, \varepsilon, (f, g))$ which has the effect that $\text{MGFT}_{f,g}(x_m)$ decreases and $\{\text{MGFT}_{f,g}(x_i) : i \in [m-1]\}$ all increase, while the denominator of $R'(f, g)$ does not change. This operation works as follows: Let $\delta(\varepsilon) < \varepsilon$ be the value such that $\delta(\varepsilon) \sum_{1 \le i \le m} p_i(x_{m+1} - x_i) - \varepsilon \sum_{1 \le i \le m} p_i(x_m - x_i) = 0$. The operation $\text{reduce}(m, \varepsilon, (f, g))$ consists of subtracting ε from q_m, adding $\delta(\varepsilon)$ to q_{m+1}, and adding the remaining probability mass to q_1.

The value $\delta(\varepsilon)$ has been chosen in such a way that the denominator of $R'(f, g)$ remains unaffected. (This is because $\text{GFT}_{f,g}(\mathbf{OPT}) = \text{GFT}_{f,g}(\mathbf{OPT}) + \delta(\varepsilon) \sum_{1 \le i \le m} p_i(x_{m+1} - x_i) - \varepsilon \sum_{1 \le i \le m} p_i(x_m - x_i) = \text{GFT}_{f,g'}(\mathbf{OPT})$, where (f, g') is the instance resulting from $\text{reduce}(m, \varepsilon, (f, g))$.) Also, we see that the value $\text{MGFT}_{f,g}(x_m)$ decreases by $\varepsilon \sum_{1 \le j \le m} p_j(x_m - x_j)$, while for $k \in [m-1]$ the value $\text{MGFT}_{f,g}(x_k)$ changes by $\Delta(m, k, \varepsilon, (f, g))$, defined as

$$\sum_{k+1 \le i \le m} p_i(\delta(\varepsilon)x_{m+1} - \varepsilon x_m + x_i(\varepsilon - \delta(\varepsilon))).$$

It can be seen that this change is positive, because the factor $(\delta(\varepsilon)x_{m+1} - \varepsilon x_m + x_i(\varepsilon - \delta(\varepsilon)))$ is increasing in i, therefore the terms in the summation of $\Delta(m, 0, \varepsilon(f, g)) = 0$ are first all negative and afterwards all positive. Thus, the partial sum $\Delta(m, k, \varepsilon(f, g))$ is positive, and $\text{MGFT}_{f,g}(x_k)$ increases.

We have thus shown that for all $i \in [m]$ and $\varepsilon \in \mathbb{R}_{\ge 0}$, the operation $\text{reduce}(i, \varepsilon, (f, g))$ decreases the denominator of $R'(f, g)$, decreases the value $\text{MGFT}_{f,g}(x_i)$, and increases the value $\text{MGFT}_{f,g}(x_j)$ for $j \in [m] \setminus \{i\}$. Therefore, for every i there is a threshold $t_i(f, g)$ such that the function $\text{reduce}(i, \cdot, (f, g))$ increases the value $\min\{\text{MGFT}_{f,g}(x_i) : i \in [m]\}$ up to the point $t_i(f, g)$ and decreases it after $t_i(f, g)$. In other words, applying operation $\text{reduce}(i, t_i(f, g), (f, g))$ results in an instance (f', g') in which $R'(f', g') \ge R'(f, g)$, and $\text{MGFT}_{f',g'}(x_i) = \min\{\text{MGFT}_{f',g'}(x_j) : j \in [m]\}$. We next consider what happens if we perform sequencences of operations of $\text{reduce}(i, t_i(\cdot), \cdot), i \in [m]$ such that $t_i > 0$, starting on the instance (f, g). We call these *reduction sequences on* (f, g). This definition implies that for any reduction sequence on instance $(f, g) \in \mathcal{J}$, the resulting instance (f', g') satisfies $R(f, g) \ge R(f', g')$.

Note first that the following facts hold.

Proposition 6.2. *There does not exist a reduction sequence on (f, g) that yields an instance that reduces the support of f or g. Every reduction sequence on (f, g) yields an instance (f', g') that lies in $\mathcal{J}$, such that f' and f have the same support, and g' and g have the same support.*

Proof. The claim follows by contradiction: suppose that a reduction sequence on (f, g) results in an instance (f', g') where the support of f' is a strict subset of the support of f, or the support of g' is a strict subset of the support of g. (Reduction sequences cannot influence the support in any other way, by definition.) Then, by Lemma 6.2, there is an instance $(f'', g'') \in \mathcal{J}$ such that $R(f'', g'') \le R(f, g)$ where the cardinality of the support of f'' is less than the cardinality of the support of f, or the cardinality of the support of g'' is strictly less than the cardinality of the support of g. This contradicts our assumption that (f, g) is canonical.

Proposition 6.3. *On any instance $(f', g') \in \mathcal{J}$, there exists a finite reduction sequence that results in an instance (f'', g'') where $MGFT_{f'',g''}(x_m) = \max\{MGFT_{f'',g''} : i \in [m]\}$ and $MGFT_{f'',g''}(x_i) = MGFT_{f'',g''}(x_j)$ for all $j \in [m-1]$.*

Proof. This fact holds because for $i \in [m-1]$, the value $\text{MGFT}_{f',g'}(x_m)$ increases the most among all values $\{\text{MGFT}_{f',g'}(x_j) : j \in [m], j > i\}$ when performing the operation $\text{reduce}(i, t_i(f', g'), (f', g'))$, all values $\{\text{MGFT}_{f',g'}(x_j) : j \in [m], j < i\}$ all increase by an even larger amount, and all values $\{\text{MGFT}_{f',g'}(x_j) : j \in [m], j < i\}$ all increase by the same amount. Therefore, the reduction sequence $\text{reduce}(m-1, t_{m-1}(\cdot), \cdot), \text{reduce}(m-2, t_{m-2}(\cdot), \cdot), \ldots, \text{reduce}(1, t_1(\cdot), \cdot)$ yields an instance (f''', g''') in which $\text{MGFT}_{f''',g'''}(x_i) < \text{MGFT}_{f''',g'''}(x_{i+1})$ for all $i \in [m-1]$. If after that, we apply the reduction sequence $\text{reduce}(2, t_2(\cdot), \cdot), \text{reduce}(3, t_3(\cdot), \cdot), \ldots \text{reduce}(m-1, t_{m-1}(\cdot), \cdot)$, we obtain an instance (f'', g'') where $\text{MGFT}_{f'',g''}(x_i) = \text{MGFT}_{f'',g''}(x_{i+1})$ for all $i \in [m-2]$ and $\text{MGFT}_{f'',g''}(x_{m-1}) \le \text{MGFT}_{f'',g''}(x_m)$, as desired.

Now we distinguish between two cases:

Case 1: There exists a finite reduction sequence on (f, g) that yields in an instance (f', g') such that $t_i(f', g') = 0$ for all $i \in [m]$. In such an instance, all the ratios $R'(f', g')$ are equal and greater than $R'(f, g)$, and $(f', g') \in \mathcal{K}(0)$, which proves the claim.

Case 2: There does not exist a finite reduction sequence that results in an instance (f', g') such that $t_i(f', g') = 0$ for all $i \in [m]$. In that case, consider any infinite reduction sequence

$$(\text{reduce}(i(k), t_{i(k)}(\cdot), \cdot))_{k \in \mathbb{N}}$$

on (f, g) such that

- $i(k) = m$ infinitely often, and

- for all k where $i(k) = m$ it holds that $\mathrm{MGFT}_{f_k,g_k}(x_m) = \max\{\mathrm{MGFT}_{f_k,g_k}(x_j) : j \in [m]\}$, and $\mathrm{MGFT}_{f_k,g_k}(x_j) = \mathrm{MGFT}_{f_k,g_k}(x_{j'})$ for all $j, j' \in [m-1]$, where (f_k, g_k) is the instance resulting from the first k reduce-operations in the sequence.

Such a sequence exists by Proposition 6.3. Let S be the set of indices k such that $i(k) = m$. Because an application of reduce$(m, \epsilon, \cdot)$ decreases the value of q_m by ϵ, we know that $\sum_{k \in S} t_m(f_k, g_k) < q_m$. Because S is infinite, for all $\epsilon' \in \mathbb{R}_{>0}$, there must be an index $k \in S$ such that $t_{i(k)} < \epsilon'$. Note moreover that by definition, for any $\delta > 0$ and any $(f', g') \in \mathcal{J}$, the operation reduce$(m, \delta, (f', g'))$ increases $\mathrm{GFT}_{f',g'}(x_m)$ by at most an amount of δ. Therefore, in the instance (f_k, g_k) it holds that

$$\begin{aligned}
&\min\{\mathrm{GFT}_{f_k,g_k}(x_j) : j \in [m]\}\\
&= \mathrm{GFT}_{f_k,g_k}(x_m)\\
&\geq \max\{\mathrm{GFT}_{f_k,g_k}(x_j) : j \in [m]\} - \epsilon'.
\end{aligned}$$

When we choose $\epsilon' = \epsilon p_1 q_{m+1}$ (where p_1 and q_{m+1} refer to instance (f, g), i.e., the values $f(0)$ and $g(1)$) we obtain

$$\begin{aligned}
&\min\{\mathrm{GFT}_{f_k,g_k}(x_j) : j \in [m]\}\\
&\geq \max\{\mathrm{GFT}_{f_k,g_k}(x_j) : j \in [m]\} - \epsilon p_1 q_{m+1}\\
&\geq (1-\epsilon)\max\{\mathrm{GFT}_{f_k,g_k}(x_j) : j \in [m]\},
\end{aligned}$$

where the last inequality follows from the property that reduction sequences can only increase p_1 and q_{m+1}, hence $p_1 q_{m+1} \leq f_k(x_1) g_k(x_{m+1}) \leq \mathrm{GFT}_{f_k,g_k}(x_i)$ for all $i \in [m]$.

So (f_k, g_k) is an instance in $\mathcal{K}(\epsilon)$ where $R'(f_k, g_k) \geq R'(f, g)$. This completes the proof.

Next, we define yet another class of instances: For $\epsilon \in \mathbb{R}_{>0}$, let $\mathcal{L}(\epsilon)$ be the bilateral trade instances in $\mathcal{J}$ where $q_{m+1} < \epsilon$ and $x_m \leq \epsilon$. The purpose of this class of instances (and the next lemma), will be clear later on in this proof.

Lemma 6.4. *Let (f, g) be an instance in $\mathcal{J}$. For any $\epsilon > 0$, there exists an instance $(f', g') \in \mathcal{L}(\epsilon)$ such that $R(f', g') \leq R(f, g)$. Hence, for all $\epsilon \in \mathbb{R}_{>0}$,*

$$\inf\{R(f,g) : (f,g) \in \mathcal{J}\} = \inf\{R(f,g) : (f,g) \in \mathcal{L}(\epsilon)\}.$$

Moreover, the probability mass on the second highest point in the support of g equals the probability mass on the second highest point in the support of g'.

(The last part of this Lemma is a technical property that is needed subsequently.)

Proof. For an instance $(f, g) \in \mathcal{J}$, we define $R(f, g, i)$ as

$$\frac{\mathbf{E}[X] + \mathrm{GFT}_{f,g}(x_i)}{\mathbf{E}[X] + \mathrm{GFT}_{f,g}(\mathbf{OPT})},$$

so that $R(f, g) = \max\{R(f, g, i) : i \in [m]\}$.

Let $(f, g) \in \mathcal{J}$. We prove that for all $\varepsilon \in [0, 1]$ there exists a $\delta(\varepsilon) \leq \varepsilon$ such that

- scaling down the points $x_1, \ldots x_m$ (but not x_{m+1}) by ϵ,
- scaling down the probability q_{m+1} by $\delta(\varepsilon)$,
- putting the remaining buyer's probability mass of $q_1 := (1 - \delta(\varepsilon)) q_{m+1}$ on the point 0,

yields an instance (f', g') in $\mathcal{J}$ such that $R(f', g') \leq R(f, g)$.

Thus, it remains to define $\delta(\epsilon)$ and to prove that the scaling operation just mentioned yields an instance (f', g') in which $R(f', g', i) \leq R(f, g, i)$ for all $i \in [m]$. Define $\delta(\epsilon) \in [0, \epsilon]$ such that

$$\epsilon q_{m+1} \sum_{i \in [m]} (x_{m+1} - x_i) p_i = \delta(\epsilon) q_{m+1} \sum_{i \in [m]} (x_{m+1} - \epsilon x_i) p_i.$$

Such a choice of $\delta(\epsilon)$ exists because the terms in the right hand side summation are larger than the terms in the left hand side summation.

Observe that $R(f', g', i)$ can then be written as

$$\begin{aligned}
&= \frac{\begin{pmatrix} \epsilon \mathbf{E}_{X\sim f}[X] + \epsilon \mathrm{GFT}_{f,g}(x_i) \\ - \epsilon q_{m+1} \sum_{j \in [i]} (x_{m+1} - \epsilon x_j) p_j \\ + \delta(\epsilon) q_{m+1} \sum_{j \in [i]} (x_{m+1} - \epsilon x_j) p_j \end{pmatrix}}{\begin{pmatrix} \epsilon \mathbf{E}_{X\sim f}[X] + \epsilon \mathrm{GFT}_{f,g}(\mathbf{OPT}) \\ - \epsilon q_{m+1} \sum_{j \in [m]} (x_{m+1} - x_j) p_j \\ + \delta(\epsilon) q_{m+1} \sum_{j \in [m]} (x_{m+1} - \epsilon x_j) p_j \end{pmatrix}}\\
&= \frac{\begin{pmatrix} \epsilon \mathbf{E}_{X\sim f}[X] + \epsilon \mathrm{GFT}_{f,g}(x_i) \\ - \epsilon q_{m+1} \sum_{j \in [i]} (x_{m+1} - \epsilon x_j) p_j \\ + \delta(\epsilon) q_{m+1} \sum_{j \in [i]} (x_{m+1} - \epsilon x_j) p_j \end{pmatrix}}{\epsilon \mathbf{E}_{X\sim f}[X] + \epsilon \mathrm{GFT}_{f,g}(\mathbf{OPT})}\\
&= \frac{\begin{pmatrix} \epsilon \mathbf{E}_{X\sim f}[X] + \epsilon \mathrm{GFT}_{f,g}(x_i) \\ + q_{m+1} \sum_{j \in [i]} (x_{m+1}(\delta(\epsilon) - \epsilon) + x_j \epsilon (1 - \delta(\epsilon)) p_j) \end{pmatrix}}{\epsilon \mathbf{E}_{X\sim f}[X] + \epsilon \mathrm{GFT}_{f,g}(\mathbf{OPT})},
\end{aligned}$$

where the second equality follows from the definition of $\delta(\epsilon)$. Moreover, observe that the above is at most

$$\begin{aligned}&\leq \frac{\epsilon\mathbf{E}_{X\sim f}[X]+\epsilon\mathrm{GFT}_{f,g}(x_i)}{\epsilon\mathbf{E}_{X\sim f}[X]+\epsilon\mathrm{GFT}_{f,g}(\mathbf{OPT})}\\ &= R(f,g,i),\end{aligned}$$

since the factor $(x_{m+1}(\delta(\epsilon)-\epsilon)+x_j\epsilon(1-\delta(\epsilon))$ is increasing in i: The summation $S := \sum_{j\in[m]}(x_{m+1}(\delta(\epsilon)-\epsilon)+x_j\epsilon(1-\delta(\epsilon))p_j)$ therefore consists of first only negative terms, and after that only positive terms. By the definition of $\delta(\epsilon)$, we have that $S=0$, so the partial summation $\sum_{j\in[m]}(x_{m+1}(\delta(\epsilon)-\epsilon)+x_j\epsilon(1-\delta(\epsilon))p_j)$ is negative.

Summarizing the above two lemmas: Lemma 6.4 changes an instance in $\mathcal{J}$ such that the buyer's probability mass on the highest point is smaller than any constant ϵ. In doing this, the probability mass on the second highest point remains unaffected. Lemma 6.3 changes a canonical instance into another canonical instance such that all fixed price mechanisms perform equally well, except possibly when the fixed price is on the highest seller's point in which case it performs slightly worse (by factor $(1-\epsilon)$ for any ϵ). In doing so, probability mass may be moved from the second highest point of the seller to the highest point of the seller, and this is the only way in which the probability mass on those two points is modified. Therefore, by a repeated sequence of applications of Lemma 6.3 and Lemma 6.4, (starting with some canonical instance, and an arbitrarily small $\epsilon>0$,) we may obtain an instance that is in both $\mathcal{K}(\epsilon)$ and $\mathcal{L}(\epsilon)$. This implies the following corollary:

Corollary 6.1. *For all $\epsilon>0$,*

$$\begin{aligned}&\inf\{R(f,g):(f,g)\in\mathcal{I}\}\\ =\ &\inf\{R(f,g):(f,g)\in\mathcal{L}(\epsilon)\cap\mathcal{K}(\epsilon)\}.\end{aligned}$$

Note that the above corollary requires all the technical properties in the statements of Lemmas 6.3 and 6.4: The total amount by which $q(x_{m+1})$ increases throughout any sequence of applications of the above two lemmas is bounded by 1, thus there must be a finite sequence of such applications that yields an instance that is both in $\mathcal{K}(\epsilon)$ and $\mathcal{L}(\epsilon)$, for any $\epsilon>0$.

Now we are ready to state the key lemma by which we can prove Theorem 6.1. This lemma says that for any ϵ, we can safely restrict our attention to instances in $\mathcal{J}$ that are both in $\mathcal{L}(\epsilon)$ and $K(0)$, i.e., instances in $\mathcal{J}$ where every fixed-price achieves the same social welfare, and the buyer's probability mass on x_{m+1} is arbitrarily small.

Lemma 6.5. *For all $\epsilon>0$,*

$$\begin{aligned}&\inf\{R(f,g):(f,g)\in\mathcal{I}\}\\ =\ &\inf\{R(f,g):(f,g)\in\mathcal{L}(\epsilon)\cap\mathcal{K}(0)\}.\end{aligned}$$

Proof. Let (f,g) be an instance in $\mathcal{L}(\epsilon)\cap\mathcal{K}(\epsilon)$. We will transform this instance into an instance $R(f',g')\in\mathcal{L}(\epsilon)\cap\mathcal{K}(0)$, such that $R(f',g')\leq R(f,g)+h(\epsilon)$, where h is an increasing continuous function with $h(0)=0$. Because we can pick ϵ arbitrarily small, this would establish that

$$\begin{aligned}&\inf\{R(f,g):(f,g)\in\mathcal{L}(\epsilon)\cap\mathcal{K}(\epsilon)\}\\ =\ &\inf\{R(f,g):(f,g)\in\mathcal{L}(\epsilon)\cap\mathcal{K}(0)\},\end{aligned}$$

and subsequently applying Corollary 6.1 would prove the claim.

For $i\in[m]$, we define $R(f,g,i)$ as in the proof of the previous lemma, so that $\max\{R(f,g,i):i\in[m]\}=R(f,g)$. In instance (f,g) the values $R(f,g,1),\ldots,R(f,g,m-1)$ are all equal. Value $R(f,g,m)$ is lower, but at least $(1-\epsilon)R(f,g,m-1)$. Consider the operation where we add probability mass of δ to $f(x_m)$, (i.e., the highest point of the seller), and subsequently scale the seller's distribution by a constant so as to ensure that $\sum_{i\in[m]}f(x_i)=1$. For simplicity we only consider the first of these two steps, as the scaling step does not influence the ratios $R(f,g,i), i\in[m]$ in any way.

This operation adds to the denominator of $R(f,g)$ (and hence the denominator of $R(f,g,i), i\in[m]$) an amount of $\delta x_m+\delta g(x_{m+1})(1-x_m)=\delta(g(x_{m+1})+x_m(1-g(x_{m+1})))\leq\delta$. The numerator of $R(f,g,m)$ increases by the same amount. For $i\in[m-1]$, the numerator of $R(f,g,i)$ increases by only δx_m. Therefore, it is possible to choose δ such that the numerator of $R(f,g,m)$ becomes equal to the numerators of $R(f,g,i), i\in[m-1]$. Denote this choice of δ by $\delta(\epsilon)$, and let the resulting instance be (f',g'). In other words, instance (f',g') is obtained from (f,g) by adding $\delta(\epsilon)$ to $f(x_m)$, and $(f',g')\in\mathcal{K}(0)$. Denote the quantity $\mathbf{E}[f]+\mathrm{GFT}_{f,g}(x_i)$ by $N(f,g,i)$, i.e., $N(f,g,i)$ is the numerator of $R(f,g,i)$; and denote the quantity $\mathbf{E}[f]+\mathrm{GFT}_{f,g}(\mathbf{OPT})$ by $D(f,g)$, i.e., $D(f,g)$ is the denominator of $R(f,g)$ and $R(f,g,i)$. Then, for $i\in[m-1]$

$$\begin{aligned}R(f,g) &= \frac{N(f,g,i)+\delta(\epsilon)g(x_{m+1})}{D(f,g)+\delta(\epsilon)(g(x_{m+1})+x_m(1-g(x_{m+1})))}\\ &\leq \frac{N(f,g,i)+\delta(\epsilon)g(x_{m+1})}{D(f,g)}\\ &= R(f,g,i)+\frac{\delta(\epsilon)g(x_{m+1})}{D(f,g)}\\ &= R(f,g)+\frac{\delta(\epsilon)g(x_{m+1})}{D(f,g)}.\end{aligned}$$

It is easy to verify that we defined $\delta(\epsilon)$ as

$$\begin{aligned}\delta(\epsilon) &= \frac{\mathrm{GFT}_{f,g}(x_m) - \mathrm{GFT}_{f,g}(x_{m+1})}{g(x_{m+1})(1-x_m)} \\ &\leq \frac{\epsilon \mathrm{GFT}_{f,g}(x_m)}{g(x_{m+1})(1-x_m)}.\end{aligned}$$

Therefore,

$$\begin{aligned}R(f',g') &\leq R(f,g) + \frac{\epsilon \mathrm{GFT}_{f,g}(x_m)}{(1-x_m)D(f,g)} \\ &\leq R(f,g) + \frac{\epsilon}{1-x_m} \\ &\leq R(f,g) + \frac{\epsilon}{1-\epsilon}.\end{aligned}$$

The function $h(\epsilon) = \epsilon/(1-\epsilon)$ is increasing and continuous in ϵ, and satisfies $h(0) = 0$, which proves the claim.

The following lemma states an important relation between the buyer and seller distribution for instances in $\mathcal{K}(0)$.

Lemma 6.6. *For any instance in $\mathcal{K}(0)$,*

$$\mathbf{E}[X \mid X > 0] \geq \mathbf{E}[Y \mid Y < 1].$$

Proof. It holds that $\mathrm{MGFT}_{f,g}(x_i) = \mathrm{MGFT}_{f,g}(x_{i+1})$ for all $i \in [m-1]$. This means that

$$p_{i+1} \sum_{j \in [m]: j > i+1} q_j(x_j - x_{i+1}) = q_{i+1} \sum_{j < i+1} p_j(x_{i+1} - x_j).$$

Rewriting this yields

$$\frac{p_{i+1}}{q_{i+1}} = \frac{\sum_{j<i+1} p_j(x_{i+1} - x_j)}{\sum_{j \in [m]: j > i+1} q_j(x_j - x_{i+1})}.$$

The numerator of the right hand side of this equation is increasing in i, while the denominator is decreasing in i. Therefore, the sequence

$$\left(\frac{p_2}{q_2}, \frac{p_3}{q_3}, \ldots, \frac{p_m}{q_m}\right)$$

is increasing, so

$$\left(\frac{\mathbf{Pr}[X = x_2 \mid X > 0]}{\mathbf{Pr}[Y = x_2 \mid Y < 0]}, \frac{\mathbf{Pr}[X = x_3 \mid X > 0]}{\mathbf{Pr}[Y = x_3 \mid Y < 0]}, \ldots \right.$$
$$\left. \ldots, \frac{\mathbf{Pr}[X = x_m \mid X > 0]}{\mathbf{Pr}[Y = x_m \mid Y < 0]}\right)$$

is also increasing. Let $\bar{Y}$ denote the distribution of Y conditioned on $Y < 1$ and let $\bar{X}$ denote the distribution of X conditioned on $X > 0$. The increasingness of the above sequence implies that there is an index $i \in \{2, \ldots, m-1\}$ such that $\mathbf{Pr}\left[\bar{X} = x_j\right] < \mathbf{Pr}\left[\bar{Y} = x_j\right]$ for all $j < i$, while $\mathbf{Pr}\left[\bar{X} = x_j\right] \geq \mathbf{Pr}\left[\bar{Y} = x_j\right]$ for all $j \geq i$. Thus, the distribution $\bar{Y}$ can be obtained from $\bar{X}$ by a sequence of reallocations of probability mass from higher points to lower points. Each such a reallocation decreases the expected value of the distribution by a positive amount, so $\mathbf{E}\left[\bar{Y}\right] \leq \mathbf{E}\left[\bar{X}\right]$.

Lastly, we make use of the following well-known fact.

Proposition 6.4. *Let a/b be a rational number in $[0,1]$ and let $c \in \mathbb{R}_{\geq 0}$. Then $a/b \geq (a-c)/(b-c)$.*

Proof.

$$\frac{a-c}{b-c} \leq \frac{a - c(a/b)}{b-c} = \frac{ab - ac}{b(b-c)} = \frac{a(b-c)}{b(b-c)}.$$

The above proposition and lemmas enable us to prove the upper bound.

Proof. [of Theorem 6.1] By Lemma 6.5, it suffices to prove that $\lim_{\epsilon \to 0} \inf\{R(f,g) : (f,g) \in \mathcal{K}(0) \cap \mathcal{L}(\epsilon)\} \geq 13/25$. Let $\epsilon \in \mathbb{R}_{>0}$, $k \in [0,1)$ and let (f,g) be an arbitrary bilateral trade instance in $\mathcal{K}(0) \cap \mathcal{L}(\epsilon)$. The gain from trading at price x_m is $q_{m+1} \sum_{i \in m}(1-x_i)p_m = q_{m+1}(1 - \mathbf{E}[X])$, while the gain from trading at price x_1 is $p_1 \mathbf{E}[Y]$. Because our instance is a member of $\mathcal{K}(0)$, these quantities are equal. In fact, for all $p \in [0,1)$ it holds that

$$\mathrm{GFT}_{f,g}(p) = p_1 \mathbf{E}[Y] = q_{m+1} - q_{m+1}\mathbf{E}[X]. \tag{6.1}$$

Using this identity, we derive the following bounds:

$$\begin{aligned}q_{m+1} &= p_1 \mathbf{E}[Y] + q_{m+1}\mathbf{E}[X] \\ &\leq p_1 \mathbf{E}[Y] + \epsilon \mathbf{E}[X] \\ &\leq p_1 \mathbf{E}[Y \mid Y < 1] + p_1 q_{m+1} + \epsilon \mathbf{E}[X] \\ &\leq (\epsilon \mathbf{E}[X] + p_1 \mathbf{E}[Y \mid Y < 1]) \sum_{j=0}^{\infty} p_1^j \\ &= \frac{\epsilon}{1-p_1}\mathbf{E}[X] + \frac{p_1}{1-p_1}\mathbf{E}[Y \mid Y < 1],\end{aligned}$$

where we used that $(f,g) \in \mathcal{L}(\epsilon)$ for the first inequality, and the last equality follows by recursively substituting q_{m+1} by $p_1 \mathbf{E}[Y \mid Y < 1] + p_1 q_{m+1} + \epsilon \mathbf{E}[X]$. Similarly:

$$\begin{aligned}q_{m+1} &\geq p_1 \mathbf{E}[Y] \\ &= p_1(1 - q_{m+1})\mathbf{E}[Y \mid Y < 1] + p_1 q_{m+1} \\ &= \frac{p_1(1-q_{m+1})\mathbf{E}[Y \mid Y < 1]}{1-p_1} \\ &\geq \frac{p_1(1-\epsilon)}{1-p_1}\mathbf{E}[Y \mid Y < 1],\end{aligned}$$

where we used that $(f,g) \in \mathcal{L}(\epsilon)$ for the last inequality, and the last equality follows by recursively substituting q_{m+1} by $p_1(1-q_{m+1})\mathbf{E}[Y \mid Y<1]+p_1 q_{m+1}$. Using the above upper and lower bounds on q_{m+1}, we derive three useful bounds on $\mathbf{E}[Y]$:

$$\begin{aligned}
\mathbf{E}[Y] &= q_{m+1} + (1-q_{m+1})\mathbf{E}[Y \mid Y<1] \\
&\leq \frac{\epsilon}{1-p_1}\mathbf{E}[X] + \frac{p_1}{1-p_1}\mathbf{E}[Y \mid Y<1] \\
&\quad + \mathbf{E}[Y \mid Y<1] \\
&= \frac{\epsilon}{1-p_1}\mathbf{E}[X] + \frac{1}{1-p_1}\mathbf{E}[Y \mid Y<1]
\end{aligned}$$

Our second bound is as follows:

$$\begin{aligned}
\mathbf{E}[Y] &= q_{m+1} + (1-q_{m+1})\mathbf{E}[Y \mid Y<1] \\
&\geq \frac{p_1(1-\epsilon)}{1-p_1}\mathbf{E}[Y \mid Y<1] \\
&\quad + (1-\epsilon)\mathbf{E}[Y \mid Y<1] \\
&\geq \frac{1-\epsilon}{1-p_1}\mathbf{E}[Y \mid Y<1].
\end{aligned}$$

Our third bound follows from the first bound:

$$\begin{aligned}
\mathbf{E}[Y] &\leq \frac{\epsilon}{1-p_1}\mathbf{E}[X] + \frac{1}{1-p_1}\mathbf{E}[Y \mid Y<1] \\
&\leq \frac{\epsilon}{1-p_1}\mathbf{E}[X] + \frac{1}{1-p_1}\mathbf{E}[X \mid X>0] \\
&= \frac{\epsilon}{1-p_1}\mathbf{E}[X] + \frac{1}{(1-p_1)^2}\mathbf{E}[X] \\
&= \frac{1+\epsilon(1-p_1)^2}{(1-p_1)^2}\mathbf{E}[X],
\end{aligned}$$

where we used Lemma 6.6 for the second inequality.

Using the above three bounds, we will now derive two lower bounds $B_1(f,g)$ and $B_2(f,g)$ on $R(f,g)$, so that $R(f,g) \geq \max\{B_1(f,g), B_2(f,g)\}$.

$R(f,g)$ is equal to

$$\begin{aligned}
&= \frac{\mathbf{E}[X] + p_1\mathbf{E}[Y]}{\mathbf{E}[\max\{X,Y\}]} \\
&\geq \frac{\begin{pmatrix} \mathbf{E}[X \mid X>0](1-p_1) \\ + p_1(1-\epsilon)\mathbf{E}[Y \mid Y<1]/(1-p_1) \end{pmatrix}}{\mathbf{E}[X] + \mathbf{E}[Y]} \\
&\geq \frac{\begin{pmatrix} \mathbf{E}[X \mid X>0](1-p_1) \\ + p_1(1-\epsilon)\mathbf{E}[Y \mid Y<1]/(1-p_1) \end{pmatrix}}{\begin{pmatrix} \mathbf{E}[X \mid X>0](1-p_1) + \epsilon\mathbf{E}[X]/(1-p_1) \\ + \mathbf{E}[Y \mid Y<1]/(1-p_1) \end{pmatrix}} \\
&\geq \frac{\begin{pmatrix} \mathbf{E}[Y \mid Y<1](1-p_1) \\ + p_1(1-\epsilon)\mathbf{E}[Y \mid Y<1]/(1-p_1) \end{pmatrix}}{\begin{pmatrix} \mathbf{E}[Y \mid Y>0](1-p_1) + \epsilon\mathbf{E}[Y \mid Y<1] \\ + \mathbf{E}[Y \mid Y<1]/(1-p_1) \end{pmatrix}},
\end{aligned}$$

where we use Lemma 6.6 in combination with Proposition 6.4 for the third inequality. Further, it holds that the above expression is equal to

$$\begin{aligned}
&= \frac{\mathbf{E}[Y \mid Y<1](1-p_1+p_1(1-\epsilon)/(1-p_1))}{\mathbf{E}[Y \mid Y>0](1-p_1+\epsilon+1/(1-p_1))} \\
&= \frac{1-p_1+p_1(1-\epsilon)/(1-p_1)}{1-p_1+\epsilon+1/(1-p_1)} \\
&\geq \frac{1-p_1+p_1/(1-p_1)}{1-p_1+1/(1-p_1)} - 2\epsilon \\
&=: B_1(f,g).
\end{aligned}$$

The bound on $B_2(f,g)$ is derived as follows.

$$\begin{aligned}
R(f,g) &= \frac{\mathbf{E}[X] + p_1\mathbf{E}[Y]}{\mathbf{E}[X] + \mathrm{GFT}_{f,g}(\mathbf{OPT})} \\
&\geq \frac{\mathbf{E}[Y](1-p_1)^2/(1+\epsilon(1-p_1)^2) + p_1\mathbf{E}[Y]}{\begin{pmatrix} \mathbf{E}[Y](1-p_1)^2/(1+\epsilon(1-p_1)^2) \\ + \mathrm{GFT}_{f,g}(\mathbf{OPT}) \end{pmatrix}} \\
&\geq \frac{\mathbf{E}[Y](p_1 + (1-p_1)^2/(1+\epsilon(1-p_1)^2))}{\begin{pmatrix} \mathbf{E}[Y](1-p_1)^2/(1+\epsilon(1-p_1)^2) \\ + \sqrt{p_1}\mathbf{E}[Y] + \mathrm{GFT}_{f,g}(F^{-1}(\sqrt{p_1}))/\sqrt{p_1} \end{pmatrix}} \\
&= \frac{\mathbf{E}[Y](p_1 + (1-p_1)^2/(1+\epsilon(1-p_1)^2))}{\mathbf{E}[Y](1-p_1)^2/(1+\epsilon(1-p_1)^2) + 2\sqrt{p_1}\mathbf{E}[Y]} \\
&= \frac{p_1 + (1-p_1)^2/(1+\epsilon(1-p_1)^2)}{2\sqrt{p_1} + (1-p_1)^2/(1+\epsilon(1-p_1)^2)} \\
&= \frac{p_1 + (1-p_1)^2}{2\sqrt{p_1} + (1-p_1)^2} - \epsilon \\
&=: B_2(f,g),
\end{aligned}$$

where the first inequality follows from our third bound on $\mathbf{E}[Y]$ in combination with Proposition 6.4, and the second inequality follows from applying Lemma 6.1, taking $c = \sqrt{p_1}$. The second equality follows from (6.1). Using $B_1(f,g)$ and $B_2(f,g)$, we can prove the claim.

$$\begin{aligned}
&\lim_{\epsilon \to 0} \inf\{R(f,g) : (f,g) \in \mathcal{L}(\epsilon)\} \\
\geq\ &\lim_{\epsilon \to 0} \inf\{\max\{B_1(f,g), B_2(f,g)\} : (f,g) \in \mathcal{L}(\epsilon)\} \\
=\ &\inf\left\{ \max\left\{ \frac{1-p_1+p_1/(1-p_1)}{1-p_1+1/(1-p_1)}, \right.\right. \\
&\left.\left. \frac{p_1+(1-p_1)^2}{2\sqrt{p_1}+(1-p_1)^2} \right\} : p_1 \in [0,1) \right\}.
\end{aligned}$$

Basic calculus shows that the max-expression is minimized at $p_1 = 1/4$, at which it evaluates to $13/25$.

References

[1] Blumrosen, L., Dobzinski, S.: Reallocation mechanisms. In: Proceedings of the fifteenth ACM con-

ference on Economics and computation. pp. 617–617. EC '14, ACM (2014)

[2] Blumrosen, L., Holenstein, T.: Posted prices vs. negotiations: an asymptotic analysis. In: Proceedings of the 9th ACM conference on Electronic commerce. pp. 49–49. EC '08, ACM (2008)

[3] Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-parameter mechanism design and sequential posted pricing. In: Proceedings of the forty-second ACM symposium on Theory of computing. pp. 311–320. STOC '10, ACM (2010)

[4] Clarke, E.H.: Multipart pricing of public goods. Public Choice 11(1), 17–33 (1971)

[5] Deng, X., Goldberg, P.W., Tang, B., Zhang, J.: Revenue maximization in a bayesian double auction market. Theoretical Computer Science 539(1), 1–12 (2014)

[6] Deshmukh, K., Goldberg, A.V., Hartline, J.D., Karlin, A.R.: Truthful and competitive double auctions. In: Proceedings of the 10th Annual European Symposium on Algorithms. pp. 361–373. ESA '02, Springer-Verlag (2002)

[7] Dütting, P., Kleinberg, R.: Polymatroid prophet inequalities. In: Bansal, N., Finocchi, I. (eds.) Algorithms – ESA 2015, Lecture Notes in Computer Science, vol. 9294, pp. 437–449. Springer-Verlag (2015)

[8] Dütting, P., Talgam-Cohen, I., Roughgarden, T.: Modularity and greed in double auctions. In: Proceedings of the Fifteenth ACM Conference on Economics and Computation. pp. 241–258. EC '14, ACM (2014)

[9] Gresik, T.A., Satterwhite, M.A.: The rate at which a simple market converges to efficiency as the number of traders increases: An asymptotic result for optimal trading mechanisms. Journal of Economic Theory 48(1), 304–332 (1989)

[10] Groves, T.: Incentives in teams. Econometrica 41(4), 617–631 (1973)

[11] Kleinberg, R., Weinberg, S.M.: Matroid prophet inequalities. In: Proceedings of the forty-fourth Annual ACM Symposium on Theory of Computing. pp. 123–136. STOC '12, ACM (2012)

[12] McAfee, R.P.: A dominant strategy double auction. Journal of Economic Theory 56(2), 434–450 (1992)

[13] McAfee, R.P.: The gains from trade under fixed price mechanisms. Applied Economics Research Bulletin 1(1), 1–10 (2008)

[14] Myerson, R.B.: Optimal auction design. Mathematics of Operations Research 6(1), 58–73 (1981)

[15] Myerson, R.B., Satterwhite, M.A.: Efficient mechanisms for bilateral trading. Journal of Economic Theory 29, 265–281 (1983)

[16] Nisan, N., Ronen, A.: Algorithmic mechanism design (extended abstract). In: Proceedings of the Thirty-first Annual ACM Symposium on Theory of Computing. pp. 129–140. STOC '99, ACM (1999)

[17] Rustichini, A., Satterwhite, M.A., Williams, S.R.: Convergence to efficiency in a simple market with incomplete information. Econometrica 62(5), 1041–1063 (1994)

[18] Satterthwaite, M.A., Williams, S.R.: The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. The Review of Economic Studies 56(4), pp. 477–498 (1989)

[19] Satterwhite, M.A., Williams, S.R.: The optimality of a simple market mechanism. Econometrica 70(5), 1841–1863 (2002)

[20] Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. The Journal of Finance 16(1), 8–37 (1961)