

Fixed Price Approximability of the Optimal Gain From Trade

Riccardo Colini-Baldeschi¹, Paul Goldberg², Bart de Keijzer³,
Stefano Leonardi⁴, and Stefano Turchetta⁵

¹ LUISS Rome, rcolini@luiss.it

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

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

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

⁵ KPMG Italy, stefano.turchetta@gmail.com

Abstract. *Bilateral trade* is a fundamental economic scenario comprising a strategically acting buyer and seller (holding an item), each holding valuations for the item, drawn from publicly known distributions. It was recently shown that the only mechanisms that are simultaneously dominant strategy incentive compatible, strongly budget balanced, and ex-post individually rational, are *fixed price* mechanisms, i.e., mechanisms that are parametrised by a price p , and trade occurs if and only if the valuation of the buyer is at least p and the valuation of the seller is at most p . The *gain from trade (GFT)* is the increase in welfare that results from applying a mechanism. We study the GFT achievable by fixed price mechanisms. We explore this question for both the bilateral trade setting and a *double auction* setting where there are multiple i.i.d. unit demand buyers and sellers. We first identify a fixed price mechanism that achieves a GFT of at least $2/r$ times the optimum, where r is the probability that the valuation of the seller valuation does not exceed that of the buyer. This extends a previous result by McAfee. Subsequently, we improve this approximation factor in an asymptotic sense, by showing that a more sophisticated rule for setting the fixed price results in a GFT within a factor $O(\log(1/r))$ of the optimum. This is asymptotically the best approximation factor possible. For the double auction setting, we present a fixed price mechanism that for all $\epsilon > 0$ a gain from trade of at least $(1 - \epsilon)$ times the optimum with probability $1 - 2/e^{\#T\epsilon^2/2}$, where $\#T$ is the expected number of trades of the mechanism. This can be interpreted as a “large market” result: Full efficiency is achieved in the limit, as the market gets thicker.

1 Introduction

Bilateral trade is a fundamental economic scenario comprising a buyer and a seller. The seller holds one item, and can possibly trade this item with the buyer for some price. The buyer and the seller each have a (non-negative real-valued) valuation for the item that is up for trade. The buyer’s valuation is only known by the buyer and the seller’s valuation is only known by the seller. The buyer and

seller both want to maximise their utility, which is assumed to be *quasi-linear*, i.e., of the form $x \cdot v - p$, where x is a 0/1-variable that is set to 1 if and only if the agent holds the item, v is the agent's value for the item, and p is the price paid/received by the agent. In the buyer's case p is non-negative and represents how much the buyer has to pay. In the seller's case, the price p is non-positive because the seller receives money to transfer her item.

The main problem studied for this bilateral trade setting is one in mechanism design: which mechanism maximises the *social welfare* (i.e., total utility of both players)? A direct revelation mechanism for this setting solicits the valuations of the buyer and the seller. Subsequently it determines whether the buyer and the seller should trade and which prices they have to pay or receive. We would like any mechanism to satisfy the following properties:

- Dominant strategy incentive compatibility (DSIC): It should be a dominant strategy for the buyer and seller to submit their true valuations to the mechanism.
- Ex-post individual rationality (ex-post IR): Neither agent should end up with a negative utility if the agent's true valuation is submitted to the mechanism.
- Strong budget balance (SBB): The price paid by the buyer is equal to the price received by the seller, i.e., the mechanism does not extract money from the market, nor does it inject money into the market.

While the valuations of the buyer and seller are known by the buyer and seller only, it is assumed that there is still distributional public knowledge about their valuations. More precisely, it is assumed that there are two publicly known distributions from which the buyer and seller independently draw their valuations. The mechanism may use this knowledge in order to determine the outcome.

Ideally, we would want the mechanism to have the seller trade with the buyer whenever the buyer's valuation exceeds the seller's valuation. The expected total utility that would result from trading as such is referred to as the *optimal* social welfare. Unfortunately the optimal social welfare is not achievable, as shown by Myerson and Satterthwaite [16]: No bilateral trade mechanism is simultaneously DSIC, IR, *weakly* budget balanced, and social-welfare optimizing. Weak budget balance (WBB) is less restrictive than strong budget balance, as WBB only requires that no money be injected into the market, while the mechanism is allowed to extract money from the market.

For the classic bilateral trade setting, it was recently shown [5] that the only direct revelation mechanisms that are simultaneously incentive compatible, strongly budget balanced, and ex-post individually rational, are *fixed price* mechanisms, i.e., mechanisms that are parametrised by a price p , and trade occurs if and only if the valuation of the buyer is at least p and the valuation of the seller is at most p .

An alternative —and more challenging to approximate— objective to the social welfare is the *gain from trade*, which measures the expected increase in total utility that is achievable by applying the mechanism, with respect to the initial allocation. For example, if a seller holds an item that she values \$4 and a buyer values the same item \$10, whenever a fixed price mechanism sets a price

$4 \leq p \leq 10$, the buyer and the seller trade producing a gain from trade of \$6. Whenever the price p is set lower than \$4 or greater than \$10 no trade occurs, and the gain from trade is 0.

McAfee [15] has shown that if the median of the distribution of the seller's valuation is less than the median of the distribution of the buyer's valuation, then there is a fixed price mechanism for which the expected gain from trade is at least half of the optimal gain from trade. In fact, it was shown for this special case that by setting the fixed price anywhere in between the two medians, half of the optimal gain from trade is guaranteed. We extend this result by showing that the optimal gain from trade is at least $2/r$ times the gain from trade achievable by a fixed price mechanism, where r is the probability that the seller's valuation does not exceed the buyer's valuation (which is the condition under which a gain from trade is possible in the first place).

Subsequently, we improve this approximation factor in an asymptotic sense, by showing that a more sophisticated rule for setting the fixed price results in an expected gain from trade within a factor $O(\log(1/r))$ of the optimal gain from trade. This is asymptotically the best approximation factor possible, which is shown by an appropriate example of a bilateral trade setting for which every fixed price achieves an expected gain from trade of $\Omega(\log(1/r))$ times the optimum.

It follows from our results that our mechanisms cannot approximate the gain from trade if r is small. Indeed, we prove a general negative result showing that the ratio between the gain from trade of a DSIC mechanism and the optimal gain from trade can be arbitrarily small as the support of the distribution grows. A similar result was proved independently in [2].

We finally extend our study to the *double auction* setting, where there are multiple buyers and sellers, each seller holding one item and each buyer having a demand for obtaining at most one item. The valuations of the n buyers are independently drawn from a common probability distribution, and the same holds for the m sellers, although the probability distribution of the sellers may be distinct from that of the buyers.

1.1 Our Results

The first results presented in this paper concern the bilateral trade problem. It is known that if a mechanism has to satisfy IR, DSIC, and SBB, then it must be a fixed price mechanism [5], i.e., the mechanism fixes a price p and posts it to the buyer and the seller. We want to understand p has to be chosen.

McAfee's result of [15] states that in case the seller's median is less than the buyer's median, then setting the price in between the medians of the buyer and the seller results in a 2-approximation to the optimal gain from trade. Our first result is a strict generalization of [15] where the approximation to the optimal gain from trade is given as a function of the probability that a trade is efficient, in other words: the probability that the valuation drawn from the buyer is greater than the valuation drawn from the seller. This parameter is referred to as $r = \Pr_{v \sim f, w \sim g}[v \geq w]$, where f is the buyer's distribution and g is the seller's distribution.

In particular, we show that setting the price p such that $\Pr_{v \sim f}[v \geq p] = \Pr_{w \sim g}[w \leq p]$ results in a $r/2$ -approximation to the optimal gain from trade.

Then, we show how that it is possible to improve the approximation factor of $2/r$ considerably in an asymptotic sense: We prove that by using a more complex rule for determining the fixed price p , the optimal gain from trade is at most a factor of $O(\log(1/r))$ times the gain from trade when trading at price p . When r is small, this results in a big improvement when compared to the approximation factor that we established in the previous section. Our mechanism works by showing that we can decompose “roughly” the entire probability space into at most $\log(1/r) + 1$ such events, so that choosing the best fixed price corresponding to each of these events results in a gain from trade that is an $O(\log(1/r))$ -approximation to the optimal gain from trade. Finally, we want to consider the double auction setting. In this setting, we extend the definition of a fixed price mechanism in a natural way: the mechanism computes a single price p , buyers with a valuation greater than p and sellers with a valuation lower than p will be allowed to trade. If the sets of allowed buyers and allowed sellers have different cardinalities, agents will be removed from the biggest set uniformly at random so that the cardinality of the two sets will be equal. The fixed price mechanism that we propose for the double auction setting achieves a gain from trade that is a $1 - \epsilon$ approximation to the optimal gain from trade with probability $(1 - 1/e^{\epsilon^2 \#T/2})$ where $\#T$ is the expected number of trades of the mechanism. This implies that if the double auction instance is such that a relatively small expected number of trades can happen at this price, then a reasonably good approximation factor is achieved by our mechanism (see Section 5 for a detailed discussion). One may also interpret our result as a “large market” result: the approximation factor approaches 1 as we let the number of buyers and sellers in the market grow proportionally, since in that case the number of trades grows arbitrarily large. This is, to the best of our knowledge, the first fixed price mechanism that is DSIC, SBB, and ex-post IR, and achieves a near-optimal gain from trade under mild conditions on the size of the market.

1.2 Related Literature

The impossibility result of [16] proved that no two-sided mechanism can be simultaneously BIC, IR, WBB, and optimise the social welfare even in the simple bilateral trade setting. Thus, many subsequent works studied how it is possible to relax some of the constraints to achieve positive results in the context of maximise the social welfare or the gain from trade.

In [3], a BIC mechanism is devised that approximates the expected gain from trade in bilateral trade up to a factor of $1/e$ when the buyer’s distribution function satisfies a property known as the monotone hazard rate condition. The mechanism of [3] is not DSIC since it achieves this approximation factor from a Bayes-Nash equilibrium by using the valuation of the seller in the price offered to the buyer. It is also shown in the same work that no BIC mechanism can achieve an approximation bound better than $2/e$. Mechanisms that are DSIC/BIC, IR, and SBB have been given for bilateral trade in [1]. In addition to this, the

authors proposed a WBB mechanism for a general class of markets known as combinatorial exchange markets.

Mechanisms for double auctions with near-optimal gain from trade have been previously proposed for the prior-free setting. McAfee [14] has shown a WBB, DSIC, IR mechanism which achieves a $1 - 1/k$ -approximation to the optimal gain from trade if the number of trades under the optimum allocation is k . This rate of convergence requires the prior distributions of traders to be bounded above zero, over an interval $[0, 1]$, but the mechanism is not a function of the priors. More recently, Segal-Halevi et al. [20] devised a SBB mechanism with the same performance guarantee. The mechanisms of [14,20] are direct revelation mechanisms where the price depends crucially on the reported valuations. In contrast, the goal in our present paper is to find out how much gain from trade can be generated by means of setting a single fixed price, independent of the valuations of the players, at which all agents trade. Such mechanisms have the advantage that they are conceptually simpler and have a pricing scheme that is extremely easy to understand. They also can be implemented as sequential posted price mechanisms.

In [19], the authors present a mechanism that combines random sampling and random serial dictatorship techniques which is IR, SBB and DSIC, and asymptotically approaches the optimum gain from trade. Recently, [4] provides an IR, SBB, and BIC mechanism that achieves a constant approximation to the best gain from trade achievable among the IR, WBB, and BIC mechanisms, which is an alternative (more permissive) benchmark. Two recent papers by Feldman and Gonen [10,9] study a multi-unit variant of double auctions for online advertising purposes. They design IR, WBB, and DSIC mechanisms that well-approximate the gain from trade under certain technical conditions, as a function of the number of trades under the optimum allocation.

Deng et al. [7] study revenue maximisation in a setting of multiple buyers and sellers with uncorrelated priors and a single type of item being traded. The same objective was studied by [8] yet in the *prior-free* model. Gerstgrasser et al. [11] also study the objective of maximising expected revenue, in a setting where there is a small number of buyers and seller, who have a prior distribution whose support size represents the complexity of instances of the problem. In [11], this distribution is otherwise unrestricted, and in particular may be correlated. Giannakopoulos et al. [12] study a similar double auction setting to the one studied here in Section 5: there are multiple buyers and sellers and one kind of item, with unit supply and demand. Buyers have a common prior distribution on their valuations, as do sellers. [12] study market intermediation from the perspective of welfare maximisation, and revenue maximisation. Colini-Baldeschi et al. [6] also study market intermediation, in the context of buyers and seller of a collection of heterogenous items, aiming to maximise welfare and achieve a strong notion of budget balance.

In the context of social welfare in the economics literature, [13] showed that duplicating the number of agents by τ results in a market where the optimal IR, IC, WBB mechanism's expected social welfare approximation factor approaches

1 at a rate of $O(\log \tau / \tau^2)$. The papers [17] and [18] investigated a family of non-IC double auctions, and study the inefficiency and the extent to which agents misreport their valuations in these double auctions.

2 Preliminaries

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

Double Auction Setting. In a double auction setting there are n buyers and m sellers. Initially, each seller $j \in [m]$ holds one item and has a valuation w_j for it. The sellers are not interested in possess more than one item. Each buyer $i \in [n]$ is interested in obtaining no more than one item and has a valuation v_i for it. Moreover, they are indifferent among the different items.

The valuations of the buyers and the sellers are private knowledge, but they are independently drawn from publicly known distributions f and g , where f is the probability distribution for the valuation of a buyer and g is the probability distribution for the valuation of a seller. We treat f and g as probability density functions. All the buyers share the same distribution f and all the sellers share the same mass probability distribution g , but f and g may be distinct. Moreover, let G be the corresponding cumulative distribution functions of g and let \bar{F} be the corresponding complementary cumulative distribution function (or survival function) of f .

Given a double auction setting (n, m, f, g) , our goal is to redistribute the items from the sellers to the buyers. An *allocation* for a double auction setting (n, m, f, g) is a pair of vectors $(\mathbf{X}, \mathbf{Y}) = ((X_1, \dots, X_n), (Y_1, \dots, Y_m))$ such that all the elements $X_1, \dots, X_n, Y_1, \dots, Y_m \in \{0, 1\}$, and $\sum_{i \in [n]} X_i + \sum_{j \in [m]} Y_j = m$. The set \mathcal{A} represents the set of all allocations for the double auction setting.

The redistribution of the items from sellers to buyers is done by running a *mechanism* \mathbb{M} . A mechanism receives input from the agents, and outputs an allocation (\mathbf{X}, \mathbf{Y}) and a price p . The allocation (\mathbf{X}, \mathbf{Y}) and the price p represents the *outcome* of the mechanism \mathbb{M} . Thus, an outcome is a tuple $(\mathbf{X}, \mathbf{Y}, p)$. The price p represents how much a buyer has to pay to obtain an item and how much a seller has to receive to sell her item.⁶

Agents are assumed to be utility maximisers. The *utility* is defined as the valuation for the items that they possess with respect to the allocation vector, minus the payment charged by the mechanism. Specifically, the utility of a buyer i will be $u_i^B(\mathbf{X}, \mathbf{Y}, p) = (v_i - p) \cdot X_i$. Similarly, the utility of a seller j will be $u_j^S(\mathbf{X}, \mathbf{Y}, p) = w_j Y_j + p \cdot (1 - Y_j)$.

Furthermore, agents are assumed to be fully rational, so that they will strategically interact with the mechanism to achieve their goal of maximising utility.

⁶ More generally, we may define the notion of a mechanism such that more complex pricing schemes are possible, but our definition suffices for the mechanisms that we will define later in this paper.

Our goal is to design a mechanism that is DSIC, IR, SBB (as defined in the introduction) such that the the *gain from trade* is high. For an outcome $(\mathbf{X}, \mathbf{Y}, p)$, the gain from trade $\text{GFT}(\mathbf{X}, \mathbf{Y}, p)$ is defined as the increase in total utility as a result of running the mechanism. It can be expressed as follows.

$$\text{GFT}(\mathbf{X}, \mathbf{Y}, p) = \sum_{i=1}^n v_i X_i + \sum_{j=1}^m w_j (Y_j - 1)$$

For a double auction setting (n, m, f, g) , the *expected optimal gain from trade* is defined as

$$\text{OPT}_{n,m,f,g} = \mathbf{E}_{v \sim f^n, w \sim g^m} \left[\max \left\{ \sum_{i=1}^n v_i X_i + \sum_{i=1}^m w_j (Y_j - 1) \mid (\mathbf{X}, \mathbf{Y}) \in \mathcal{A} \right\} \right].$$

We will sometimes omit the subscript, as in those cases the instance being discussed will be clear from context.

We say that a mechanism \mathbb{M} α -approximates the optimal gain from trade for some $\alpha > 1$ if and only if $\text{OPT} \leq \alpha \mathbf{E}[\text{GFT}(\mathbf{X}, \mathbf{Y}, p)]$, where $(\mathbf{X}, \mathbf{Y}, p)$ is the random allocation that the mechanism generates, when valuations v and w are drawn from f^n and g^m respectively. Our goal is to find a DSIC, ex-post IR, and SBB mechanism that α -approximates the optimal gain from trade for a low α .

Bilateral Trade Setting. The bilateral trade setting is a special case of the double auction setting where there is only one unit-demand buyer and one unit-supply seller. Thus, we can represent a bilateral trade setting as a pair of valuation distribution function, one for the buyer f and one for the seller g , i.e., (f, g) . It is known that if a mechanism has to satisfy IR, DSIC, and SBB, then it must be a fixed price mechanism [5], i.e., the mechanism fixes a price p a priori, and trade happens if and only if both the buyer's valuation is at least p and the seller's valuation is at most p .

For a bilateral trade instance, the gain from trade of a fixed price mechanism with fixed price p will be denoted by $\text{GFT}_{f,g}(p)$. That is,

$$\text{GFT}_{f,g}(p) = \mathbf{E}_{v \sim f, w \sim g} [\max\{0, v - w\} \mathbf{1}(w \leq p \leq v)].$$

Moreover, note that for the bilateral trade setting we can express $\text{OPT}_{f,g}$ as $\mathbf{E}_{v \sim f, w \sim g} [\max\{0, v - w\}]$.

For the bilateral trade setting, the goal of this paper is to study how to set the price p such that the gain from trade achieved by the fixed price mechanism with price p is as close as possible to the optimal gain from trade. We will design fixed price mechanisms where the ratio between $\text{OPT}_{f,g}$ and $\text{GFT}_{f,g}(p)$ is a function of the probability that the buyer has a value greater than the seller, i.e., the provability that a trade is efficient. This probability will be represented by the parameter r . Thus, $r = \mathbf{Pr}_{v \sim f, w \sim g} [v \geq w]$.

Due to space constraints, proofs of the theorems and lemmas have been omitted, and will be provided in a full version of this paper.

3 An $O(1/r)$ -Approximation Mechanism for Bilateral Trade

In the bilateral trade setting there is only one unit-demand buyer and one unit-supply seller. It can be proven that if a mechanism has to satisfy IR, DSIC, and SBB, then it must be a fixed price mechanism [5], i.e., the mechanism fixes a price p and posts it to the buyer and the seller. Trade happens if and only if both the buyer's valuation is at least p and the seller's valuation is at most p .

We will show in this section that there exists a fixed price mechanism that achieves an expected gain from trade that is at least $r/2$ times the expected optimal gain from trade. In the fixed price mechanism that we propose for this, the fixed price p is set such that $\Pr_{v \sim f}[v \leq p] = \Pr_{w \sim g}[w \leq p]$. The main theorem that we prove is thus as follows.

Theorem 1. *Let (f, g) be a bilateral trade instance, let $p \in \mathbb{R}_{\geq 0}$ be any fixed price, and let q be the minimum of $\Pr_{v \sim f}[v \geq p]$ and $\Pr_{w \sim g}[w \leq p]$. Then,*

$$\frac{1}{q} GFT_{f,g}(p) \geq OPT_{f,g}. \quad (1)$$

Moreover, if p is chosen such that q is maximised (i.e., p is such that $\Pr_{w \sim g}[w \leq p] = \Pr_{v \sim f}[v \geq p]$), it holds that

$$\frac{2}{r} GFT_{f,g}(p) \geq OPT_{f,g}. \quad (2)$$

Note that this theorem strictly generalises McAfee's result of [15], which states that in case the seller's median is less than the buyer's median, then setting the price in between the medians of the buyer and the seller results in a 2-approximation to the optimal gain from trade: If we take p to be any price in between the median of the seller and the buyer, then q is at least $1/2$, and (1) then states that the gain from trade at fixed price p is at least half the optimal gain from trade.

4 Improving the Asymptotic Dependence on r

In this section, we show how it is possible to improve the approximation factor implementing a more involved rule to determine the fixed price p . When the trading price p will be set with the new rule the approximation factor will improve from $2/r$ to $O(\log(1/r))$. Notice that when r is small, this is a big improvement with respect to the approximation shown in the previous section. All logarithms used in this section are to base 2.

Let us first give a high level description of how we determine the fixed price of the mechanism. Let us consider any two points z and z' such that $\Pr[v \geq z] = 2\Pr[v \geq z']$. Let E be the event that the buyer's valuation exceeds z , and that the seller's valuation lies in between z and z' . Let \bar{F}_E be the complementary cumulative distribution function of the buyer conditioned on E and let G_E be

the cumulative distribution function of the seller conditioned on E . We now see that on the interval $[z, z']$, the function \bar{F}_E decreases from 1 to $1/2$ and the function G_E increases from 0 to 1. Thus, the functions cross each other in $[z, z']$ at a value of at least $1/2$, which means that the median of the buyer exceeds the median of the seller when conditioning on E . Using Theorem 1, we thus obtain that when conditioning on E there exists a fixed price that achieves a 2-approximation to the optimal gain from trade.

Our mechanism works by showing that we can decompose “roughly” the entire probability space into at most $\log(1/r)+1$ such events, so that choosing the best fixed price corresponding to each of these events results into an $O(\log(1/r))$ approximation to the optimal gain from trade. More precisely, we show that there are two sets of roughly $\log(1/r)+1$ such events, and we prove that in case one of these sets does not cover a fraction of the probability space that accounts for at least $1/2$ of the optimal gain from trade, then the other set of events does. To determine the desired fixed price, we can thus

1. first determine which of the two event sets “covers” a large part of the optimal gain from trade,
2. and subsequently select the best fixed price among the $\log(1/r)+1$ prices corresponding to the event set.

The two event sets have the following properties: one of them excludes the part of the probability space where the buyer’s complementary CDF is below the threshold $r/2$. The other one switches the roles of the seller and buyer, and excludes the part of the probability space where the seller’s CDF is below a the threshold $r/2$. From this property of the event sets (i.e., having these particular thresholds on the tails of the two distributions), we are able to show that one of the event sets covers a large part of the optimal gain from trade. We now proceed by making these ideas precise.

We first describe how we determine the price, which we denote by p^* , for a given instance (f, g) . In contrast with the last section, we assume (for convenience of exposition) without loss of generality that f and g are continuous distributions without point masses, where we treat f and g as probability density functions, and we let F and G be the corresponding cumulative distribution functions. We write \bar{F} to denote the buyer’s complementary cumulative distribution function $1 - F$. Let r be the probability $\Pr_{v \sim f, w \sim g}[v \geq w]$ of a trade being possible (as before). Let x be the value such that $F(x) = r/2$ and let y be the value such that $\bar{G}(y) = r/2$. We distinguish between two cases.

- If $\mathbf{E}_{v \sim f, w \sim g}[(v - w)\mathbf{1}(w \leq v \wedge w > y)] \geq \text{OPT}_{f,g}/2$, then let p^* be the price that achieves the maximum gain from trade among the prices $p_1, \dots, p_{\lceil \log(2/r) \rceil}$, where for $i \in [\lceil \log(2/r) \rceil]$, price p_i is such that

$$\begin{aligned} \Pr_{w \sim g}[w \leq p_i \mid \bar{F}^{-1}(1/2^{i-1}) \leq w \leq \bar{F}^{-1}(1/2^i)] \\ = \Pr_{v \sim f}[v > p_i \mid \bar{F}^{-1}(1/2^{i-1}) \leq v]. \end{aligned}$$

- Otherwise, let p^* be the price that achieves the maximum gain from trade among the prices $p'_1, \dots, p'_{\lceil \log(2/r) \rceil}$, where for $i \in [\lceil \log(2/r) \rceil]$, price p'_i is such that

$$\begin{aligned} \Pr_{v \sim f}[v > p_i \mid G^{-1}(1/2^i) \leq v \leq G^{-1}(1/2^{i-1})] \\ = \Pr_{w \sim g}[w \leq p_i \mid G^{-1}(1/2^i) \leq w], \end{aligned}$$

where we define $G^{-1}(1) = \infty$ if there exists no point $t \in \mathbb{R}_{\geq 0}$ such that $G(t) = 1$.

This completes the definition of the fixed price p^* .

First we can show that if the first of the two cases does not apply (i.e., if the inequality $\mathbf{E}_{v \sim f, w \sim g}[(v-w)\mathbf{1}(w \leq v \wedge w > y)] \leq \text{OPT}_{f,g}/2$ does not hold), then the symmetric inequality $\mathbf{E}_{v \sim f, w \sim g}[(v-w)\mathbf{1}(w \leq v \wedge w < x)] \leq \text{OPT}_{f,g}/2$ holds for the second case.

Lemma 1. *If $\mathbf{E}_{v \sim f, w \sim g}[(v-w)\mathbf{1}(w \leq v \wedge w > y)] > \text{OPT}_{f,g}/2$, then $\mathbf{E}_{w \sim f, v \sim g}[(v-w)\mathbf{1}(w \leq v \wedge v < x)] \leq \text{OPT}_{f,g}/2$.*

Using the above lemma, it is possible to prove the intended approximation factor for price p^* .

Theorem 2. *Let (f, g) be any bilateral trade instance, and let p^* be the price for (f, g) , as defined above. It holds that*

$$\text{OPT}_{f,g} \leq 4 \log \left(\left\lceil \frac{2}{r} \right\rceil \right) \text{GFT}_{f,g}(p^*).$$

Note that the approximation bound of $2/r$ that we established in the first section is better than the approximation bound of $4\lceil \log(2/r) \rceil$ when r is roughly greater than 0.05. At $r = 0.05$, the approximation factor $4\lceil \log(2/r) \rceil$ already takes a value around 20. Hence, the result of this section is intended to provide theoretical insight into how the approximability of the gain from trade depends on r asymptotically. An (asymptotically) matching lower bound is given in the appendix of [2], which shows that $\Theta(\log(1/r))$ is asymptotically the best possible factor by which the optimal gain from trade that can always be approximated.

5 A Fixed Price Double Auction

We now turn to the double auction setting. Recall that in this setting there are $n \geq 1$ buyers and $m \geq 1$ sellers. The sellers each hold one item, and neither the buyers or the sellers are interested in holding more than one item. As before, we refer to f for the probability distribution function from which the buyers' valuations are independently drawn, and to g for the probability distribution from which the sellers' valuations are independently drawn. We denote the (random) valuation of buyer $i \in [n]$ by v_i and the (random) valuation of seller $j \in [m]$ by w_j . See Section 2 for the definition.

In order to present the definition of a *fixed price mechanism* for the double auction setting, let us first introduce the concept of *feasible pair*.

Definition 1. Let (n, m, f, g) be an instance of a double auction setting, let $(v, w) \in \mathbb{R}^n \times \mathbb{R}^m$ be a valuation profile for the buyers and sellers, and let $p \in \mathbb{R}_{\geq 0}$. We call $(i, j) \in [n] \times [m]$ a feasible pair with respect to profile (v, w) and fixed price p iff $v_i \geq p \geq w_j$.

Now, we can define a fixed price mechanism as follows.

Definition 2. We define a fixed price mechanism \mathbb{M} for a double auction setting (n, m, f, g) as a direct revelation mechanism for which there is a price p such that the mechanism selects a uniform random maximal subset of feasible pairs with respect to reported profile (v, w) and p , and makes these pairs trade with each other. Moreover, for every selected trading pair (i, j) , the mechanism makes buyer i pay an amount of p to seller j . We refer to p as the price of \mathbb{M} .

This is perhaps the most natural generalization of the notion of a fixed price mechanism that one may think of. Please note that in a fixed price mechanism with price p , given a reported valuation profile (v, w) , the number of pairs that trade is always the minimum of $|\{v_i : v_i \geq p\}|$ and $|\{w_i : w_i \leq p\}|$.

It is easy to show that fixed price mechanisms clearly satisfy the three basic properties that we want:

Theorem 3. For every double auction setting, every fixed price mechanism is ex-post IR, SBB, and DSIC.

Fixed price mechanisms have some additional advantages.

- First, a fixed price mechanism is entirely symmetric: Each seller has the same expected utility, and each buyer has the same expected utility. The mechanism treats buyers with the same valuation entirely symmetrically and does not break ties in favour of one over the other. This symmetricity is desirable from the point of view of fairness.
- Secondly, the mechanism does not require the agents to fully reveal their entire valuation, since it can be implemented as a *two-sided sequential posted price mechanism* [5]. Under such an implementation, the mechanism goes over the buyers and sellers one by one. It proposes a take-it-or-leave-it price (equal to p , in this case) to each buyer and seller, which the buyers and sellers can choose to accept or reject. As soon as an accepting (buyer,seller)-pair is found, the mechanism lets this pair trade at price p . Taking a uniform random order of buyers and sellers will result in a random subset of feasible pairs who trade at price p , i.e., it will result in an implementation of the fixed price mechanism with price p . Under such an implementation, each buyer and seller has to reveal only one bit of information, which indicates whether her valuation is above or below p .

We aim to design a simple fixed price mechanism for which the gain from trade is a good approximation to the optimal gain from trade. The mechanism we use is as follows.

Definition 3. Given an instance (n, m, f, g) of a double auction setting, let \bar{p} be the price such that $n\bar{F}(\bar{p}) = mG(\bar{p})$. We refer to the fixed price mechanism with price \bar{p} as the balanced fixed price double auction. For ease of presentation, we refer to $\bar{F}(\bar{p})$ as \bar{q}^B and we refer to $G(\bar{p})$ as \bar{q}^S . We denote by $GFT(\bar{p})$ the expected gain from trade achieved by the balanced fixed price double auction, and we denote by $\#T$ the expected number of trades that the balanced fixed price double auction generates.

Observe that the balanced fixed price double auction is a generalization of the mechanism of Theorem 1 that achieves for the bilateral trade setting a $2/r$ -approximation of the optimal gain from trade. We note that the value $\#T$ is by definition equal to $n\bar{q}^B = m\bar{q}^S$.

The main result we prove in this section is as follows.

Theorem 4. For all $\epsilon \in [0, 1]$, with probability at least $1 - 2/e^{\#T\epsilon^2/2}$, the balanced fixed price double auction achieves a gain from trade that is at least $(1 - \epsilon)$ times the expected optimal gain from trade.

Note that $\#T$, the expected number of trades of the balanced fixed price double auction, needs to exceed $2\ln(2)/\epsilon^2$ by any constant for the above theorem to yield a constant approximation guarantee. The value $\#T$ can be regarded as a property of the instance (n, m, f, g) on which the mechanism is run, and is equal to the value where the functions $n\bar{F}$ and mG cross each other. The requirement on $\#T$ is reasonably mild: For example, the above theorem says that when $\#T$ is at least 10, the balanced fixed price double auction yields an expected gain from trade that is a (< 4)-approximation to the optimal gain from trade, by taking $\epsilon \approx 0.61$ (since $(1 - 2/e^{0.61^2 \cdot 5}) \cdot (1 - 0.61) > 0.25$). The theorem provides a constant approximation ratio for all instances where $\#T > 2\ln(2) \approx 1.38$, but grows unbounded as $\#T$ approaches $2\ln 2$ from above.

There is an interesting interpretation of this theorem in terms of large markets: Observe that increasing the number of buyers or sellers in the market also increases $\#T$. In particular, by increasing both the number of buyers and the number of sellers simultaneously, $\#T$ grows unboundedly. From our theorem we may therefore infer that the balanced fixed price double auction approximation approximates the gain from trade by a factor that goes to 1 as the market grows.

To prove the desired approximation property of the balanced fixed price double auction, we note that due to symmetry, we may assume that under the optimum allocation every buyer has the same a priori probability of trading with a seller, and every seller has the same a priori probability of trading with a buyer. This motivates the following definition.

Definition 4. For a double auction setting (n, m, f, g) we define the values q^S as the probability that any buyer receives an item under the optimum allocation, and we define q^B as the probability that any seller loses her item under the optimum allocation. We define the prices p^B and p^S as the prices closest to \bar{p} such that $\bar{F}(p^B) = q^B$ and $G(p^S) = q^S$. That is: p^B is such that a buyers' probability of her valuation exceeding p^B is equal to the probability of obtaining an item under

the optimum allocation, and if there multiple such prices then p^B is defined as the unique one closest to \bar{p} . Likewise, p^S is such that a sellers' probability of her valuation being at most p^S is equal to the probability of losing her item under the optimum allocation. Lastly, we let OPT denote the expected gain from trade achieved by the optimum allocation.

The values $GFT(\bar{p})$, OPT , $\#T$, \bar{p} , \bar{q}^B , \bar{q}^S , p , q^B , and q^S all depend (like r) on the instance (n, m, f, g) . We will leave this dependence implicit.

Theorem 4 can be proved by means of the following sequence of lemmas.

Lemma 2. *For every instance (n, m, f, g) of a double auction setting, the following property of the optimal allocation is satisfied.*

$$nq^B = mq^S.$$

The following lemma states that our price \bar{p} always lies in between p^B and p^S .

Lemma 3. *For every instance (n, m, f, g) of a double auction setting, $p^B \geq \bar{p} \geq p^S$ or $p^S \geq \bar{p} \geq p^B$.*

The following lemma provides a useful bound on OPT

Lemma 4. *For every instance (n, m, f, g) of a double auction setting, it holds that*

$$OPT \leq nq^B \mathbf{E}[v_1 \mid v_1 \geq p^B] - mq^S \mathbf{E}[w_1 \mid w_1 \leq p^S].$$

We then use the following technical lemma to bound OPT further.

Lemma 5. *For every instance (n, m, f, g) of a double auction setting, it holds that*

$$\begin{aligned} nq^B \mathbf{E}[v_1 \mid v_1 \geq p^B] - mq^S \mathbf{E}[w_1 \mid w_1 \leq p^S] \\ \leq n\bar{q}^B \mathbf{E}[v_1 \mid v_1 \geq \bar{p}] - m\bar{q}^S \mathbf{E}[w_1 \mid w_1 \leq \bar{p}]. \end{aligned}$$

Using the above lemmas, Theorem 4 can be proved by showing an appropriate bound on the gain from trade of the balanced fixed price double auction, along with applying a Chernoff bound.

Acknowledgements. We thank Tim Roughgarden for helpful discussions at the early stages of this work.

References

1. Liad Blumrosen and Shahar Dobzinski. Reallocation mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, pages 617–617. ACM, 2014.
2. Liad Blumrosen and Shahar Dobzinski. (Almost) Efficient Mechanisms for Bilateral Trading. *ArXiv/CoRR*, abs/1604.04876, 2016.

3. Liad Blumrosen and Yenohatan Mizrahi. Approximating Gains-from-Trade in Bilateral Trading. In Yang Cai and Adrian Vetta, editors, *Proceedings of the 12th International Conference on Web and Internet Economics (WINE)*, pages 400–413. Springer, 2016.
4. Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, pages 589–590, 2017.
5. Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately Efficient Double Auctions with Strong Budget Balance. In *Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1424–1443. SIAM, 2016.
6. Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, Stefano Turchetta, and Tim Roughgarden. Approximately Efficient Two-Sided Combinatorial Auctions. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 591–608. ACM, 2017.
7. Xiaotie Deng, Paul W. Goldberg, Bo Tang, and Jinshan Zhang. Revenue maximization in a Bayesian double auction market. *Theoretical Computer Science*, 539:1–12, 2014.
8. Kaustubh Deshmukh, Andrew V. Goldberg, Jason D. Hartline, and Anna R. Karlin. Truthful and competitive double auctions. In *Proceedings of the 10th Annual European Symposium on Algorithms (ESA)*, pages 361–373. Springer, 2002.
9. Moran Feldman and Rica Gonen. Markets with strategic multi-minded mediators. *ArXiv/CoRR*, abs/1603.08717, 2016.
10. Moran Feldman and Rica Gonen. Online truthful mechanisms for multi-sided markets. *CoRR*, abs/1604.04859, 2016.
11. Matthias Gerstgrasser, Paul W. Goldberg, and E. Koutsoupias. Revenue maximization for market intermediation with correlated priors. In Martin Gairing and Rahul Savani, editors, *Proceedings of the 9th International Symposium on Algorithmic Game Theory (SAGT)*, pages 273–285. Springer, 2016.
12. Yiannis Giannakopoulos, Elias Koutsoupias, and Philip Lazos. Online Market Intermediation. *ArXiv/CoRR*, abs/1703.09279, 2017.
13. Thomas A. Gresik and Mark A. Satterthwaite. 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.
14. R. Preston McAfee. A Dominant Strategy Double Auction. *Journal of Economic Theory*, 56(2):434–450, 1992.
15. R. Preston McAfee. The Gains from Trade under Fixed Price Mechanisms. *Applied Economics Research Bulletin*, 1, 2008.
16. Roger B. Myerson and Mark A. Satterthwaite. Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory*, 29(2):265–281, 1983.
17. Aldo Rustichini, Mark A. Satterthwaite, and Steven R. Williams. Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 62(5):1041–1063, 1994.
18. Mark A. Satterthwaite and Steven R. Williams. The optimality of a simple market mechanism. *Econometrica*, 70(5):1841–1863, 2002.
19. Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. A random-sampling double-auction mechanism. *ArXiv/CoRR*, abs/1604.06210, 2016.
20. Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. SBBA: A Strongly-Budget-Balanced Double-Auction Mechanism. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory (SAGT)*, pages 260–272. Springer, 2016.