Check for updates

Bilateral Trade with Correlated Values*

Shahar Dobzinski
Weizmann Institute of Science
Rehovot, Israel
shahar.dobzinski@weizmann.ac.il

Ariel Shaulker
Weizmann Institute of Science
Rehovot, Israel
ariel.shaulker@weizmann.ac.il

ABSTRACT

We study the bilateral trade problem where a seller owns a single indivisible item, and a potential buyer seeks to purchase it. Previous mechanisms for this problem only considered the case where the values of the buyer and the seller are drawn from independent distributions. In contrast, this paper studies bilateral trade mechanisms when the values are drawn from a joint distribution.

We prove that the buyer-offering mechanism guarantees an approximation ratio of $\frac{e}{e-1} \approx 1.582$ to the social welfare even if the values are drawn from a joint distribution. The buyer-offering mechanism is Bayesian incentive compatible, but the seller has a dominant strategy. We prove the buyer-offering mechanism is optimal in the sense that no Bayesian mechanism where one of the players has a dominant strategy can obtain an approximation ratio better than $\frac{e}{e-1}$. We also show that no mechanism in which both sides have a dominant strategy can provide any constant approximation to the social welfare when the values are drawn from a joint distribution.

Finally, we prove some impossibility results on the power of general Bayesian incentive compatible mechanisms. In particular, we show that no deterministic Bayesian incentive-compatible mechanism can provide an approximation ratio better than $1+\frac{\ln 2}{2} \approx 1.346$.

CCS CONCEPTS

• **Theory of computation → Algorithmic game theory and mechanism design**.

KEYWORDS

bilateral trade, incentive compatibility.

ACM Reference Format:
Shahar Dobzinski and Ariel Shaulker. 2024. Bilateral Trade with Correlated Values. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24), June 24–28, 2024, Vancouver, BC, Canada.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3618260.3649659

1 INTRODUCTION

This paper focuses on the bilateral trade problem where a seller owns a single indivisible item, and a potential buyer seeks to purchase it. The seller has a value of $s \geq 0$ associated with retaining the item (and 0 otherwise), whereas the buyer's value for obtaining it is $b \geq 0$ (if the buyer does not receive the item, then the buyer's value is 0). The values b and s are private, but the probability distributions from which they were derived are known.

The two most common objectives in bilateral trade scenarios are to maximize the social welfare and to maximize the gains from trade. The former goal is aimed at maximizing the total value generated by the transaction (that is, the social welfare is b in case of trade and s otherwise), while the latter is focused on increasing the difference between the buyer's and the seller's surplus (i.e., the gains from trade is $b - s$ in case of trade, and 0 otherwise). Our interest in this paper is in incentive-compatible mechanisms that are strongly budget balanced. That is, the buyer's payment is fully transferred to the seller. See Section 2 for a precise statement of the problem.

The problem was extensively studied, and here, we only mention some of the papers closest to our research. Myerson and Satterthwaite [18] prove that under very mild conditions, no Bayesian mechanism can exactly maximize the gains from trade or, equivalently, the social welfare. Blumrosen and Dobzinski [3] show that a fixed-price mechanism guarantees at least half of the optimal social welfare (equivalently, provides a 2-approximation) for every distribution. This approximation ratio was later improved to 1.99 [8], then to $\frac{e}{e-1}$ [4], then to $\frac{e}{e-1} - 0.0001$ [11], and then to an almost optimal ratio of approximately 1.38 [7, 12].

There has also been much work regarding approximating the gains from trade. McAfee [15] shows that for some distributions of the buyer and seller's values, there exists a fixed price mechanism that recovers half of the optimal gains-from-trade. However, for every fixed fraction, there are distributions for which any fixed price guarantees less than this fraction of the optimal gains-from-trade [3]. To overcome this, Blumrosen and Mizrahi [5] propose the seller-offering mechanism, in which the seller makes the buyer the profit-maximizing take-it-or-leave-it offer, given his value s. In the seller-offering mechanism, only the buyer has a dominant strategy. Still, they show that if the buyer's value is drawn from a distribution with a monotone hazard rate, then the seller-offering mechanism is Bayesian incentive compatible and provides an e approximation to the optimal gains-from-trade. Brustle et al. [6] show that the better of the seller offering mechanism and the buyer offering mechanism (in which the buyer makes the seller a profit-maximizing take-it-or-leave-it-offer) recovers at least half of the gains from trade of every incentive-compatible mechanism. In a breakthrough result, Deng et al. [9] show that the better of the offering mechanisms provides 8.23 approximation to the optimal gains-from-trade. This constant was later improved to 3.15 by Fei [10]. We also know that the better of the offering mechanisms sometimes recovers strictly less than half of the optimal gains from trade [2].

Despite this extensive work, existing research on bilateral trade, including the works cited above, largely assumes that the values

*Work supported by ISF grant 2185/19 and BSF-NSF grant (BSF number: 2021655, NSF number: 2127781).

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

ACM ISBN 979-8-4007-0383-6/24/06
https://doi.org/10.1145/3618260.3649659

of the seller and the buyer are drawn independently[1]. This paper investigates a more realistic and technically challenging scenario in which the values are derived from a joint probability distribution. Remarkably, we demonstrate that despite the correlation, the buyer-offering mechanism approximates the social welfare within a constant factor. Furthermore, this factor is quite close to the best approximation ratio possible for independent distributions.

Theorem: For every joint distribution of the seller and buyer's values, there is a mechanism that provides an $\frac{e}{e-1}$-approximation to the optimal social welfare[2].

Interestingly, whereas the power of the buyer-offering mechanism is equal to the power of the seller-offering mechanism for gains-from-trade approximation, the buyer-offering mechanism is much more powerful than the seller-offering mechanism in our setting: it is not hard to see that the seller-offering mechanism does not guarantee any constant fraction of the optimal social welfare[3].

The buyer-offering mechanism is natural and simple. It is also appealing from a theoretical point of view: it is not only Bayesian incentive compatible, but in fact, one player (the offered player) has a dominant strategy as it can only accept or reject a take-it-or-leave-it offer. We call Bayesian incentive compatible mechanisms in which one side has a dominant strategy *one-sided dominant strategy* mechanisms. We prove that the buyer-offering mechanism is optimal among all one-sided mechanisms:

Theorem: There exists a joint distribution of the seller and buyer's values such that the approximation ratio of every one-sided dominant strategy mechanism is no better than $\frac{e}{e-1}$.

This theorem demonstrates, in particular, that even taking, e.g., the better of the seller-offering and the buyer-offering mechanism, or the better of the buyer-offering mechanism and some fixed price mechanism, does not improve the approximation guarantee of the buyer-offering mechanism.

Moreover, we show that mechanisms in which both sides have a dominant strategy cannot guarantee any fixed fraction of the optimal social welfare[4].

Theorem: For every constant $c > 1$, there exists a joint distribution of the seller and buyer's values such that the approximation ratio of every dominant strategy incentive compatible mechanism is at least c.

Unlike the impossibility of one-sided mechanisms (which requires careful analysis and subtle construction), the proof that dominant-strategy mechanisms have no power is technically simpler. We then continue analyzing the power of Bayesian incentive-compatible mechanisms, now proving limits on their power:

Theorem: There exists a joint distribution of the seller and buyer's values such that no deterministic Bayesian incentive-compatible mechanism provides a better than $1 + \frac{\ln 2}{2} \approx 1.346$ approximation to the optimal social welfare.

Previously, Blumrosen and Mizrahi [5] proved a bound of 1.07 on the approximation ratio of Bayesian mechanisms for independent distributions. Their proof relies on characterizing the second-best mechanism of [18]. In our setting, proving impossibilities with this approach is less promising: not only is the Myerson-Satterthwaite characterization often hard to compute for independent distributions, but it also does not apply to joint distributions. Thus, we present a new technique for proving impossibilities for Bayesian incentive-compatible mechanisms. Our approach is based on introducing and analyzing a certain family of "L-shaped" distributions, for which the second-best mechanism has a nicer structure. A disadvantage of our approach is that our results only apply to deterministic Bayesian mechanisms that are ex-post individually rational, whereas the results of Blumrosen and Mizrahi apply to randomized mechanisms that are interim individually rational ones. We note that all major mechanisms considered in the literature (e.g., fixed price mechanisms, buyer and seller offering mechanisms) are deterministic and ex-post individually rational[5]. Nevertheless, our technique is robust enough that, as a by-product, we improve the state-of-the-art impossibilities for both the social welfare and gains-from-trade even with *independent* distributions:

Theorem:

- There exist independent distributions of the seller and buyer's values such that no deterministic Bayesian incentive-compatible mechanism provides an approximation ratio better than 2-approximation to the optimal gains from trade.
- There exist independent distributions of the seller and buyer's values such that no deterministic Bayesian incentive-compatible mechanism provides an approximation ratio better than 1.113-fraction to the optimal social welfare.

Open Questions and Future Directions. In this paper, we analyzed the power of incentive-compatible mechanisms for bilateral trade. We proved that the buyer-offering mechanism provides an approximation ratio of $\frac{e}{e-1} \approx 1.582$ even if the values are drawn from joint distributions. We proved that this ratio is optimal for one-sided mechanisms and that dominant strategy mechanisms cannot guarantee any fixed fraction of the welfare at all. However, there is a gap between this ratio and our impossibility result of $1 + \frac{\ln 2}{2} \approx 1.346$[6].

[1]There are some exceptions. For example, McAfee and Reny [16] provide conditions on the distribution in which the payment of the buyer equals his value if relaxing the individual rationality condition is allowed. Equilibria in some related information models are analyzed in [17]. Finally, Robust mechanisms for bilateral trade are studied in [13].

[2]The buyer-offering mechanism achieves this approximation ratio for every distribution for which it is defined. There are distributions for which the buyer-offering mechanism is undefined (i.e., there is a series of offers that approach the maximum profit but never attain it). For these distributions, we show that for every $\varepsilon > 0$, there is a slight variation of the buyer-offering mechanism that guarantees an $\frac{e}{e-1} + \varepsilon$ approximation.

[3]Consider a seller with a fixed value $s = 0$ and a buyer that its value is distributed by an equal revenue distribution in $[1, k]$. The optimal welfare is $\ln k$ whereas the welfare of the seller offering mechanism is only 1.

[4]A simple observation is that mechanisms in which both sides have a dominant strategy are fixed price mechanisms [3].

[5]Our impossibilities also apply to mechanisms which are a probability distribution over deterministic ex-post individually rational mechanisms, like the random offerer mechanism [9]. To see this, consider such a mechanism A. A simple averaging argument shows that in the support of A there must be a (deterministic and ex-post individually rational) mechanism A' that its approximation ratio is at least the approximation ratio of A. Our impossibilities directly apply to A', and the approximation ratio of A is no better than the approximation ratio of A'.

[6]This impossibility result may be found in the full version of this paper.

We leave closing this gap as an open question. It will also be interesting to understand the power of Bayesian incentive compatible mechanisms for independent distributions and determine whether they are more powerful than deterministic mechanisms [7, 12].

Another question that remains open is to determine whether Bayesian incentive-compatible mechanisms can give a constant approximation to the gains-from-trade even when the values are drawn from a joint distribution.

Lastly, all our impossibility results for Bayesian incentive-compatible mechanisms hold only for deterministic Bayesian incentive-compatible mechanisms. In the full version of this paper, we show that there exist distributions for which randomized Bayesian incentive-compatible mechanisms outperform deterministic Bayesian incentive-compatible mechanisms. An important future direction is understanding the power of randomized Bayesian incentive-compatible mechanisms in all models discussed in the paper.

Structure of the Paper. In Section 2 we give the necessary preliminaries. In Section 3, we prove that the buyer-offering mechanism provides an $\frac{e}{e-1}$-approximation, even for correlated distributions. Section 4 shows that $\frac{e}{e-1}$ is the best ratio achievable by one-sided dominant strategy mechanism. In Appendix A, we show that no two-sided dominant strategy incentive compatible mechanism provides a bounded approximation ratio. Our impossibilities for Bayesian incentive-compatible mechanisms may be found in the full version of this paper.

2 THE SETTING

In the bilateral trade problem, we have two agents: the seller and the buyer. The seller owns an indivisible item and his value for it is s. The buyer's value for the item is b. The values b and s are drawn from a joint distribution $\mathcal{F}$.

A (direct) deterministic *mechanism M* for the bilateral trade problem consists of two functions $M = (x, p)$. For every tuple of seller and buyer values (s, b) in the support of $\mathcal{F}$, $x(s, b) = 1$ if a trade occurs and $x(s, b) = 0$ otherwise. If there is a trade, $p(s, b)$ specifies the price that the buyer pays for the item and the payment that the seller gets for it. We require that $b \geq p(s, b) \geq s$. These restrictions on the payment are called *ex-post individual rationality*.

The *optimal welfare* of a joint distribution $\mathcal{F}$ is $OPT_{\mathcal{F}} = \mathbb{E}_{(s,b)\sim\mathcal{F}}[\max\{b, s\}]$. The *welfare* of a mechanism $M = (x, p)$ is $M_{\mathcal{F}} = \mathbb{E}_{(s,b)\sim\mathcal{F}}[(x(s, b) \cdot b + (1 - x(s, b)) \cdot s]$.

For a joint distribution $\mathcal{F}$, the *approximation ratio of a mechanism $M = (x, p)$ to the optimal welfare* is $\frac{OPT_{\mathcal{F}}}{M_F}$.

For a joint distribution $\mathcal{F}$, we denote by $\mathcal{F}_{s|b}$, the conditional cumulative distribution function of the seller, given that the buyer's value is b. Similarly, we denote by $\mathcal{F}_{b|s}$, the conditional cumulative distribution function of the buyer, given that the seller's value is s. We denote by $\mathbb{1}_A$ the indicator function for the event A. For example, we will use $\mathbb{1}_{[s>b]}$ to denote the indicator function for the event that the value of the seller s is larger than the value of the buyer b.

In this paper, we consider several notions of incentive compatibility of a mechanism $M = (x, p)$. We start by defining incentive compatibility for only one of the players:

- **Seller Dominant Strategy Incentive Compatibility:** for every s, s', b:

$$x(s,b)\cdot p(s,b)+(1-x(s,b))\cdot s \geq x(s',b)\cdot p(s',b)+(1-x(s',b))\cdot s.$$

- **Seller Bayesian Incentive Compatibility:** for every s, s', b:

$$\mathop{\mathbb{E}}_{b\sim\mathcal{F}_{b|s}}[x(s,b)\cdot p(s,b) + (1 - x(s,b))\cdot s] \geq \mathop{\mathbb{E}}_{b\sim\mathcal{F}_{b|s}}[x(s',b)\cdot p(s',b) + (1 - x(s',b))\cdot s].$$

- **Buyer Dominant Strategy Incentive Compatibility:** for every b, b', s:

$$x(s,b)\cdot(b - p(s,b)) \geq x(s,b')\cdot(b - p(s,b')).$$

- **Buyer Bayesian Incentive Compatibility:** for every b, b', s:

$$\mathop{\mathbb{E}}_{s\sim\mathcal{F}_{s|b}}[x(s,b)\cdot(b - p(s,b))] \geq \mathop{\mathbb{E}}_{s\sim\mathcal{F}_{s|b}}[x(s,b')\cdot(b - p(s,b'))].$$

We will say that a mechanism is *Dominant Strategy (Bayesian) Incentive Compatible* if it is dominant strategy (Bayesian) incentive compatible for both the buyer and the seller. A mechanism is *one-sided dominant strategy incentive compatible* if it is dominant strategy incentive compatible for at least one of the players and Bayesian incentive compatible for the other.

A distribution $\mathcal{F}_{b|s}$ is *equal revenue distribution* for a seller with value s, if, for every value of p in the support of the buyer's distribution $\mathcal{F}_{b|s}$, the expected payment to the seller from a take-it-or-leave-it offer of price p to the buyer is the same.

For example, consider a seller with a value 0, and a buyer with distribution $\mathcal{F}_{b|0}(b)$:

$$\mathcal{F}_{b|0}(b) = \begin{cases} 1 - \frac{1}{b} & b \geq 1; \\ 0 & b < 1. \end{cases}$$

The expected payment to the seller from any take-it-or-leave-it offer of price $1 \geq p$ to the buyer is 1.

The expected profit of a buyer with value b, from a take-it-or-leave-it offer of price p to the seller is $\mathcal{F}_{s|b}(p)) \cdot (b - p)$.

A distribution $\mathcal{F}_{s|b}$ is *equal profit distribution* for a buyer with value b, if, for every value of p in the support of the seller's distribution $\mathcal{F}_{s|b}$, the expected profit of the buyer from a take-it-or-leave-it offer of price p to the seller is the same. For example, consider a buyer with value 1, and a seller with distribution $\mathcal{F}_{s|1}(s)$:

$$\mathcal{F}_{s|1}(s) = \begin{cases} \frac{1}{e(1-s)} & 0 \leq s \leq 1 - \frac{1}{e}; \\ 1 & s > 1 - \frac{1}{e}. \end{cases}$$

The expected profit of the buyer from any take-it-or-leave-it offer of price $p \in [0, 1 - \frac{1}{e}]$ to the seller is $\frac{1}{e}$.

In the *buyer-offering mechanism*, a buyer with value b makes a take-it-or-leave-it offer of price p to the seller, where the price p maximizes the buyer's profit, i.e., maximizes $(b - p) \cdot \mathcal{F}_{s|b}(p)$.

3 AN $\frac{e}{e-1}$-APPROXIMATION FOR CORRELATED VALUES

In this section, we prove that the buyer-offering mechanism provides an $\frac{e}{e-1} \approx 1.58$ approximation to the optimal welfare, even if the values are drawn from a joint distribution. Recall that even when the values are drawn from independent distributions, the best currently known approximation mechanisms achieve a close approximation ratio of ≈ 1.38 [7, 12]. Our approximation guarantee holds for all distributions for which the buyer-offering mechanism is well defined (i.e., there always exists an offer that maximizes the profit), but note that there are distributions for which the buyer-offering mechanism is not well defined. For example, consider the following joint distribution $\mathcal{F}^{\varepsilon}$ for some small $0 < \varepsilon < 1$. In $\mathcal{F}^{\varepsilon}$, the buyer has only one value 1, and the seller's value is supported on the interval $(0, \frac{e-\varepsilon-1}{e-\varepsilon}]$, with marginal distribution function $\mathcal{F}_s^{\varepsilon}$:

$$\mathcal{F}_s^{\varepsilon}(s) = \begin{cases} \frac{1+\varepsilon(1-s)}{e(1-s)} & s \in (0, \frac{e-\varepsilon-1}{e-\varepsilon}]; \\ 1 & s \geq \frac{e-\varepsilon-1}{e-\varepsilon}. \end{cases}$$

Let $f(p) = (1-p) \cdot F_s^{\varepsilon}(p)$ be a function that denotes the expected profit of the buyer from a take-it-or-leave-it offer of p to the seller. Then, $f(0) = 0$, and for every $p \in (0, \frac{e-\varepsilon-1}{e-\varepsilon}]$, we get $f(p) = \frac{1+\varepsilon(1-p)}{e}$. Observe that the derivative of f for values $p \in (0, \frac{e-\varepsilon-1}{e-\varepsilon}]$ is $\frac{-\varepsilon}{e}$, which is negative. Hence, the buyer's expected profit from a take-it-or-leave-it offer of p to the seller is a strictly decreasing function in the interval $p \in (0, \frac{e-\varepsilon-1}{e-\varepsilon}]$. Since the interval is open at 0, the function does not have a maximum. Thus, the buyer-offering mechanismis not defined for this distribution.

We prove that the buyer-offering mechanism provides an $\frac{e}{e-1}$-approximation for all distributions for which it is defined. For the remaining distributions, we show that a slight variant of the buyer-offering mechanism provides a similar approximation ratio:

Theorem 3.1.

(1) For every joint distribution $\mathcal{F}$ of the values of the buyer and seller, the buyer-offering mechanism provides an $\frac{e}{e-1}$ approximation to the optimal welfare if the buyer-offering mechanism is well-defined.

(2) For every joint distribution $\mathcal{F}$ of the values of the buyer and seller and every $\varepsilon > 0$, there is a one-sided dominant strategy mechanism that provides an $\frac{e}{e-1} + O(\varepsilon)$ approximation to the optimal welfare.

We first prove the first part of the theorem. We then use the first part to prove the second part. After establishing the theorem, we discuss extending the result to double auctions (Subsection 3.4).

3.1 Proof of Theorem 3.1: Part I

Fix a value b of the buyer in the support of $\mathcal{F}$. Denote by $\mathcal{F}_{s|b}$ the distribution of the seller's value given that the value of the buyer is b. We will show that, for every b, the approximation ratio of the buyer-offering mechanism when the value of the buyer is always b and the value of the seller is always $\mathcal{F}_{s|b}$ is $\frac{e}{e-1}$. This immediately implies that the approximation ratio of the buyer-offering mechanism when the values are drawn from $\mathcal{F}$ is $\frac{e}{e-1}$.

Let p_b be the price that the buyer offers when his value is b. Since the value b is fixed, to simplify notation, we drop the subscripts from p_b and $\mathcal{F}_{s|b}$, and simply write p and $\mathcal{F}$ instead.

We now bound from above the approximation ratio of the buyer-offering for the distribution $\mathcal{F}$ (i.e., the expected optimal welfare divided by the expected welfare of the buyer-offering mechanism).

$$\begin{aligned} &\frac{b \cdot \mathcal{F}(b) + \mathbb{E}_{s\sim\mathcal{F}}[\mathbb{1}_{[s>b]} \cdot s]}{b \cdot \mathcal{F}(p) + \mathbb{E}_{s\sim\mathcal{F}}[\mathbb{1}_{[p<s\leq b]} \cdot s] + \mathbb{E}_{s\sim\mathcal{F}}[\mathbb{1}_{[s>b]} \cdot s]]} \\ &\leq \frac{b \cdot \mathcal{F}(b)}{b \cdot \mathcal{F}(p) + \mathbb{E}_{s\sim\mathcal{F}}[\mathbb{1}_{[p<s\leq b]} \cdot s]}. \end{aligned} \tag{1}$$

Let $q_1 = \mathcal{F}(p)$ and let $q_2 = \mathcal{F}(b)$. We have that $q_1 \leq q_2$ (since $p \leq b$). Observe that if $q_2 = q_1$, the approximation ratio is 1, as needed. Therefore, we assume that $q_1 < q_2$. Rewriting the RHS of (1) we have:

$$\frac{b \cdot \mathcal{F}(b)}{b \cdot \mathcal{F}(p) + \mathbb{E}_{s\sim\mathcal{F}}[\mathbb{1}_{[p<s\leq b]} \cdot s]} = \frac{b \cdot q_2}{b \cdot q_1 + \mathbb{E}_{s\sim\mathcal{F}}[s|p < s \leq b] \cdot (q_2 - q_1)}.$$

For fixed q_1, q_2 and b, this expression is maximized when $\mathbb{E}_{s\sim\mathcal{F}}[s|p < s \leq b] \cdot (q_2 - q_1)$ is minimized. Therefore, in Lemma 3.2, we bound from below the expression $\mathbb{E}_{s\sim\mathcal{F}}[s|p < s \leq b] \cdot (q_2 - q_1)$ to achieve an upper bound on the approximation ratio.

Lemma 3.2. $\mathbb{E}_{s\sim\mathcal{F}}[s|p < s \leq b] \cdot (q_2 - q_1) \geq q_1 \cdot (b-p) \cdot \ln \frac{q_1}{q_2} + b \cdot (q_2 - q_1)$.

Proof. Recall that the price p maximizes the buyer's profit for the distribution $\mathcal{F}$. We use this to bound from above $\mathcal{F}(p')$ for $p < p' < b$. Let $u = q_1 \cdot (b-p)$ be the expected profit of the buyer when the price is p. Then:

$$u \geq \mathcal{F}(p')(b-p') \iff \mathcal{F}(p') \leq \frac{u}{(b-p')}. \tag{2}$$

For $p' = b - \frac{u}{q_2}$, we get that the bound (2) is tight and equal to q_2[7]. Let f be the seller's probability density function given that the buyer's value is b. We want to bound from below the expression:

$$\mathbb{E}_{s\sim\mathcal{F}}[\mathbb{1}_{[p<s\leq b]} \cdot s] = \int_p^b s f(s)\, ds \underbrace{=}_{\text{integration by parts}} s\mathcal{F}(s)\Big|_p^b - \int_p^b \mathcal{F}(s)\, ds \geq b \cdot q_2 - p \cdot q_1 - \int_p^b \mathcal{F}(s)\, ds.$$

We use the following bounds on $\mathcal{F}(s)$: for $p \leq s \leq b - \frac{u}{q_2}$ we have $\mathcal{F}(s) \leq \frac{u}{b-s}$ (by Inequality (2)) and for $b - \frac{u}{q_2} \leq s \leq b$ we have

[7]Observe that $b - \frac{u}{q_2} \geq p$ since $b - \frac{q_1(b-p)}{q_2} \underbrace{\geq}_{q_1 \leq q_2} p$.

$\mathcal{F}(s) \leq q_2$. Now,

$$\mathbb{E}_{s\sim\mathcal{F}}[\mathbb{1}_{[p<s\leq b]} \cdot s] b \cdot q_2 - p \cdot q_1 - \int_{p}^{b-\frac{u}{q_2}} \frac{u}{b-s}\, ds - \int_{b-\frac{u}{q_2}}^{b} q_2\, ds$$
$$= b \cdot q_2 - p \cdot q_1 - u + u \cdot \ln(b-s)\Big|_{p}^{b-\frac{u}{q_2}}$$
$$= b \cdot q_2 - p \cdot q_1 - u + u \cdot \ln\left(\frac{\frac{u}{q_2}}{b-p}\right)$$
$$= q_1(b-p)\ln\left(\frac{q_1}{q_2}\right) + b \cdot (q_2 - q_1).$$

□

Overall, we get that the approximation ratio is bounded from above by

$$\frac{b \cdot q_2}{b \cdot q_1 + q_1 \cdot (b-p) \cdot \ln\frac{q_1}{q_2} + b \cdot (q_2 - q_1)} \underbrace{\leq}_{p\geq 0,\ \ln\frac{q_1}{q_2}\leq 0} \frac{b \cdot q_2}{b \cdot q_1 \cdot \ln\frac{q_1}{q_2} + b \cdot q_2} = \frac{1}{\frac{q_1}{q_2} \cdot \ln\frac{q_1}{q_2} + 1} \underbrace{\leq}_{8} \frac{e}{e-1}.$$

3.2 Proof of Theorem 3.1: Part II

We consider a slight variant of the buyer-offering mechanism. Given $\varepsilon > 0$, set $\delta = \varepsilon \cdot OPT$, where OPT is the expected value of the optimal social welfare. The ε-buyer offering mechanism makes a profit-maximizing take-it-or-leave-it offer p for the seller, where p belongs to the set of offers that are a multiple of δ. I.e., $p = k \cdot \delta$ for some $k \in \mathbb{N}$. This mechanism is Bayesian incentive compatible since the set of possible prices for a buyer with value b contains at most $\lceil \frac{b}{k} \rceil$ elements. Since this set is finite for every value b, it has a maximum-profit element. In addition, the seller obviously has a dominant strategy.

As in the first part, we prove the approximation guarantee for every value b of the buyer. Let $\mathcal{F}_s(s)$ be the marginal distribution of the seller given b. Obtain a "discretized" $\mathcal{F}_s'(s)$ by "pushing" the mass of $\mathcal{F}_s(s)$ in all points that are not a multiple of δ to the nearest (from above) multiple of δ. Note that the buyer offering mechanism is now defined since the support of $\mathcal{F}_s'(s)$ is finite. Thus, there is an offer with a maximum profit since for every offer that is not in the support there is an offer in the support with at least the same profit. Furthermore, observe that for every b, if the buyer offering mechanism makes an offer p when the marginal distribution is $\mathcal{F}_s'(s)$ then p is also the offer that ε-buyer offering mechanism makes when the marginal distribution is $\mathcal{F}_s(s)$. Observe that $\mathcal{F}_s(s)$ and $\mathcal{F}_s'(s)$ are very close to each other, and thus the expected welfare that the buyer offering mechanism provides for $\mathcal{F}_s'(s)$ and the expected welfare of the ε-buyer offering mechanism for $\mathcal{F}_s(s)$ differ only by δ. The second part of the theorem now follows since, by the first part, the buyer-offering mechanism provides an $\frac{e}{e-1}$ approximation for $\mathcal{F}_s'(s)$ and because the optimal welfare of $\mathcal{F}_s'(s)$ and the optimal welfare of $\mathcal{F}_s(s)$ differ only by δ.

3.3 Tightness of Analysis

We now present an instance of a distribution $\mathcal{F}$ where the buyer-offering mechanism yields an approximation ratio no better than $\frac{e}{e-1}$ to the optimal welfare. Subsequently, in Subsection 4, we establish a more robust result that states that no one-sided dominant strategy mechanism can offer an approximation ratio better than $\frac{e}{e-1}$. Since the buyer-offering mechanism is a one-sided dominant strategy mechanism, the result of Subsection 4 also implies that the analysis of the buyer-offering mechanism is precise. Nevertheless, we provide a direct analysis of the buyer-offering mechanism in this section for simplicity.

Let $\mathcal{F}$ be a joint distribution over the buyer and seller values in which the value of the buyer is always 1, and the value of the seller is in $[0, \frac{e-1}{e}]$. The seller's value is distributed as follows:

$$\mathcal{F}_s(s) = \begin{cases} \frac{\frac{1}{e}}{1-s} & s \in [0, \frac{e-1}{e}]; \\ 1 & s > \frac{e-1}{e}. \end{cases}$$

Note that the seller's distribution is an equal profit distribution for a buyer with value of 1. I.e., for the buyer, every price $p \in [0, \frac{e-1}{e}]$ yields the same expected profit. Thus, by tie-breaking, we may assume that the buyer-offering mechanism uses price 0 (alternatively, to eliminate the use of tie-breaking, one can change $\mathcal{F}_s$ and slightly increase the probability that the seller's value is 0 – the analysis remains almost identical).

We now analyze the approximation ratio of the buyer-offering mechanism for the distribution $\mathcal{F}$:

$$\frac{OPT_{\mathcal{F}}}{ALG_{\mathcal{F}}} = \frac{b \cdot \Pr_{\mathcal{F}_s}(s \leq p) + b \cdot \Pr_{\mathcal{F}_s}(p < s \leq b)}{b \cdot \Pr_{\mathcal{F}_s}(s \leq p) + \mathbb{E}_{s\sim\mathcal{F}_s}[s | p < s \leq b] \cdot \Pr_{\mathcal{F}_s}(p < s \leq b)}$$
$$= \frac{1}{\frac{1}{e} + \mathbb{E}_{s\sim\mathcal{F}_s}[s | 0 < s \leq 1] \cdot (1 - \frac{1}{e})}$$
$$= \frac{1}{\frac{1}{e} + \frac{1}{e} \cdot \ln\frac{1}{e} + 1 - \frac{1}{e}} = \frac{e}{e-1}.$$

3.4 Double Auctions

In double auctions there are multiple buyers and sellers, each seller i owns a single item and his value for it is s_i, all items are identical, and each buyer j wants one unit and his value for it is b_j. Similarly, to bilateral trade, we wish to approximate the optimal welfare, which in the double auction case with k sellers is equal to the sum of the k largest values among the sellers and buyers.

Our positive result for bilateral trade also implies a positive result for double auctions. Similarly to the work of Babaioff et al. [1], to obtain a solution for the double auction case, we combine McAfee's trade reduction mechanism [14] with the mechanism we use for the bilateral trade case (the buyer offering) in the following manner: compute the maximal number of efficient trades (where trade is efficient only if the value of the buyer is larger than the value of the seller). Run the trade reduction mechanism if there are at least two efficient trades. If there is only one efficient trade, run the buyer offering mechanism for the bilateral trade problem with the buyer

[8] Recall that $q_1 \leq q_2$, and so $0 \leq \frac{q_1}{q_2} \leq 1$. Thus, the function $f(x) = \frac{1}{1+x\ln x}$ is maximized when $x = \frac{1}{e}$, and its maximal value is $\frac{e}{e-1}$.

being the highest-value buyer and the seller being the lowest-value seller. The distribution over the seller's value is the conditional distribution over the value of this seller given all values except his own and that he is the lowest value seller and the additional requirement that the price must be at least as large as the value of the second highest buyer.

Similarly to [1], this mechanism is Bayesian incentive compatible and ex-post individually rational. Observe that the approximation ratio of this mechanism is at least 2. If there are at least two efficient trades, the approximation ratio of the trade reduction mechanism is $1-\frac{1}{k}$ where k is the number of efficient trades, which gives us a 2 approximation guarantee. If there is only one efficient trade, we get an approximation guarantee that is at least as good as the approximation guarantee of the buyer-offering mechanism for the bilateral trade case, which is $\frac{e}{e-1}$. Lastly, if the number of efficient trades is 0, we do nothing and get an optimal approximation.

4 THE LIMITS OF ONE-SIDED DOMINANT STRATEGY MECHANISMS

We now prove that no one-sided dominant strategy mechanism can provide an approximation ratio better than $\frac{e}{e-1}$. This shows, in particular, the optimality of the buyer-offering mechanism as its approximation ratio is $\frac{e}{e-1}$ (Theorem 3.1). It also shows, for example, that taking the better of the buyer-pricing mechanism and a fixed price mechanism does not improve the approximation ratio. In Subsection 4.1, we present a specific distribution for which every one-sided dominant strategy mechanism does not provide an approximation ratio better than $\frac{e}{e-1}$. We provide the formal analysis of the impossibility in Subsection 4.2.

4.1 The Distribution $\mathcal{F}^{k,\varepsilon}$

For $k\in\mathbb{N},\varepsilon>0$, let $\mathcal{F}^{k,\varepsilon}$ be a joint distribution over the buyer and seller values in which the buyer receives values in $[1,k]\cup\{k+1\}$ with probability density function of $\frac{1}{(b+\varepsilon)^2}$ for $1\le b\le k$ and of $\frac{1}{k+\varepsilon}+\frac{\varepsilon}{1+\varepsilon}$ for $b=k+1$. For every value b in $[1,k]\cup\{k+1\}$, the seller's value is distributed according to an (almost) equal profit distribution of the buyer. I.e., when the buyer's value is b, the cumulative distribution function of the value of the seller is $\mathcal{F}^{k,\varepsilon}_{s|b}$:

$$\mathcal{F}^{k,\varepsilon}_{s|b}(s)=\begin{cases}\frac{\frac{b}{e}}{b-s+\varepsilon} & 0\le s\le b\cdot\frac{e-1}{e}+\varepsilon;\\ 1 & s>b\cdot\frac{e-1}{e}+\varepsilon.\end{cases}$$

Note that the buyer's marginal distribution always sums up to 1:

$$\mathcal{F}^{k,\varepsilon}_{b}(k+1)=\Pr_{b\sim\mathcal{F}^{k,\varepsilon}_{b}}(b=k+1)+\int_1^k\frac{1}{(b+\varepsilon)^2}\,db$$

$$=\frac{1}{k+\varepsilon}+\frac{\varepsilon}{1+\varepsilon}+-\frac{1}{b+\varepsilon}\bigg|_1^k=\frac{1}{k+\varepsilon}+\frac{\varepsilon}{1+\varepsilon}-\frac{1}{k+\varepsilon}+\frac{1}{1+\varepsilon}=1.$$

Moreover, for every value in the buyer's support, the conditional distribution of the seller sums up to 1: $\mathcal{F}^{k,\varepsilon}_{s|b}(b\cdot\frac{e-1}{e}+\varepsilon)=\frac{\frac{b}{e}}{b-(b\cdot\frac{e-1}{e}+\varepsilon)+\varepsilon}=1$.

This distribution has two useful properties. The first is that when the seller's value is 0, the buyer's distribution is very close to an equal revenue distribution (it is implied by the marginal probability density function of the buyer that is close to $\frac{1}{b^2}$). The second property is that for every value b in the buyer's support, the seller's distribution is very close to an equal profit distribution.

Fix some mechanism M that is one-sided dominant strategy for one of the players. If M is dominant strategy for the buyer, our bound on the welfare is achieved by utilizing the first property. In this case, only the seller's value can affect the value of the offer. Given that the seller's value is 0, the buyer's distribution is very close to an equal revenue distribution. We show that when the seller's value is 0, the seller will strictly prefer higher prices. Since the mechanism is Bayesian incentive compatible for the seller, the highest take-it-or-leave it offer is when $s=0$. Very roughly speaking, this implies that the welfare of the mechanism is low: if the highest offer is low, then the contribution to the welfare of trades when $s=0$ is high, but no trade is done for larger values of the seller, which happens with significant probability. On the other hand, if the value of the highest offer p is large, trade is less likely in the $s=0$ case.

If M is dominant strategy for the seller, we utilize the second property. Now, only the buyer's value can affect the offer price. However, the buyer strictly prefers lower offers as lower prices will yield higher profit. Since the mechanism is incentive compatible for the buyer, the offer price is the same for every value b of the buyer (otherwise, the buyer will prefer the lower offer and deviate from his equilibrium strategy). Then, every mechanism for $\mathcal{F}^{k,\varepsilon}$ that offers a take-it-or-leave-it-offer p to the seller (one-sided dominant strategy for the buyer) is a fixed price mechanism. It only remains to show that a fixed price mechanism has low welfare.

4.2 Analysis of One-Sided Mechanisms for the Distribution $\mathcal{F}^{k,\varepsilon}$

Theorem 4.1. *Let $k\ge 2$ and $\varepsilon>0$. Every one-sided dominant strategy mechanism for $\mathcal{F}^{k,\varepsilon}$ provides an approximation ratio of at least $\frac{e}{e-1}$ as ε approaches 0 and k approaches ∞.*

To prove this theorem we use the family of allocation rules x_p (Definition 4.2). We show that the welfare of every one-sided dominant strategy mechanism for the seller is at most the welfare of x_p, for some p (Claim 4.4). Similarly, we show that the welfare of every one-sided dominant strategy mechanism for the buyer is no better than the welfare of x_p, for some p (Claim 4.3). We conclude the proof of the theorem by bounding the approximation ratio of every possible allocation rule x_p (Lemma 4.5).

Definition 4.2. *For every $p\ge 0$, let x_p be the following allocation rule for the distribution $\mathcal{F}^{k,\varepsilon}$:*

$$x_p(b,s)=\begin{cases}1 & b=k+1;\\ 1 & s=0 \text{ and } k\ge b\ge p;\\ 0 & s=0 \text{ and } b<p \text{ and } b\ne k+1;\\ 1 & s\ne 0 \text{ and } s\le p;\\ 0 & \text{otherwise.}\end{cases}$$

Observe that for some values of p, this allocation rule x_p, is not implementable. It is used simply to bound the welfare of every one-sided dominant strategy mechanism. The next three claims suffice to prove the theorem. Their proofs can be found in Subsections 4.2.1, 4.2.2, and 4.2.3.

Claim 4.3. *Let $k \geq 2$ and $\varepsilon > 0$. Let M be a one-sided dominant strategy mechanism for the buyer. There exists $p \geq 0$, such that the welfare of M is at most the welfare of x_p, both with respect to the distribution $\mathcal{F}^{k,\varepsilon}$.*

Claim 4.4. *Let $k \geq 2$ and $\varepsilon > 0$. Let M be a one-sided dominant strategy mechanism M for the seller. There exists $p \geq 0$, such that the welfare of M is at most the welfare of x_p, both with respect to the distribution $\mathcal{F}^{k,\varepsilon}$.*

Claim 4.5. *Fix $p \geq 0$. When ε approaches 0 and k approaches ∞, the allocation rule x_p provides an approximation ratio no better than $\frac{e}{e-1}$ to the optimal welfare for the distribution $\mathcal{F}^{k,\varepsilon}$.*

4.2.1 Mechanisms with Dominant Strategy for the Buyer (Proof of Claim 4.3).

Every one-sided dominant strategy mechanism for the buyer is a take-it-or-leave-it offer for the buyer, where the price may only depend on the seller's value. Fix such a mechanism $M = (x, p)$ for the distribution $\mathcal{F}^{k,\varepsilon}$. For every value s in the seller's support, denote the take-it-or-leave-it offer of the mechanism M by p_s. Now, by Lemma 4.6, the price p_0 offered to the buyer when the seller's value is 0, should be no lower than any price in $\{p_s | s \in [0, (k+1) \cdot \frac{e-1}{e} + \varepsilon]\}$. This is true since if there is a price $p'_s > p_0$, and the seller prefers higher prices for p_0 (Lemma 4.6), the seller will play the strategy that sets the price to p'_s, in contradiction to the incentive compatibility of the mechanism. Recall that the item can be traded only if the seller's value is at most the price. Thus, the mechanism M can only trade when $s < p_s \leq p_0$. Now, let $p = p_0$, and consider the allocation rule x_p (Definition 4.2). Observe that the allocation rule x_p trades the item in every instance that M trades the item and might trade the item when M does not. Therefore, the welfare of x_p is at least the welfare of M.

Lemma 4.6. *For $s = 0$ and $p \in [1, k]$, the expected profit of the seller from a take-it-or-leave-it offer with price p to the buyer is higher than a take-it-or-leave-it offer with price $p' < p$.*

Proof of Lemma 4.6. Recall that the conditional cumulative distribution function of the buyer, given that the seller's value is 0 is denoted by $\mathcal{F}^{k,\varepsilon}_{b|0}$. The profit of the seller with value 0 from a take-it-or-leave-it offer of $p \in [1, k] \cup \{k+1\}$ is $(1 - \mathcal{F}^{k,\varepsilon}_{b|0}(p)) \cdot p$. To analyze this expression, we first provide an explicit expression for $1 - \mathcal{F}^{k,\varepsilon}_{b|0}(p)$, where $p \in [1, k]$. By definition, we have:

$$
\begin{aligned}
1 - \mathcal{F}^{k,\varepsilon}_{b|0}(p) &= \left(\int_p^k \frac{f^{k,\varepsilon}(0,b)}{\Pr_{s \sim \mathcal{F}^{k,\varepsilon}_s}(s=0)}\, db \right) + \frac{\Pr_{\mathcal{F}^{k,\varepsilon}}(0, k+1)}{\Pr_{s \sim \mathcal{F}^{k,\varepsilon}_s}(s=0)} \\
&= \frac{1}{e \cdot \Pr_{s \sim \mathcal{F}^{k,\varepsilon}_s}(s=0)} \left(\frac{k+1}{k+1+\varepsilon} \left(\frac{1}{k+\varepsilon} + \frac{\varepsilon}{1+\varepsilon} \right) \right. \\
&\qquad \left. + \int_p^k \frac{b}{(b+\varepsilon)^3}\, db \right) \\
&= \frac{1}{e \cdot \Pr_{s \sim \mathcal{F}^{k,\varepsilon}_s}(s=0)} \left(\frac{k+1}{k+1+\varepsilon} \left(\frac{1}{k+\varepsilon} + \frac{\varepsilon}{1+\varepsilon} \right) + - \frac{\varepsilon + 2b}{2(b+\varepsilon)^2} \Bigg|_p^k \right) \\
&= \frac{1}{e \cdot \Pr_{s \sim \mathcal{F}^{k,\varepsilon}_s}(s=0)} \cdot \left(\frac{k+1}{k+1+\varepsilon} \left(\frac{1}{k+\varepsilon} + \frac{\varepsilon}{1+\varepsilon} \right) \right. \\
&\qquad \left. + \left(-\frac{\varepsilon+2k}{2(k+\varepsilon)^2} + \frac{\varepsilon+2p}{2(p+\varepsilon)^2} \right) \right).
\end{aligned}
$$

Let $g(p) = p \left(\frac{k+1}{k+1+\varepsilon} \left(\frac{1}{k+\varepsilon} + \frac{\varepsilon}{1+\varepsilon} \right) + \left(-\frac{\varepsilon+2k}{2(k+\varepsilon)^2} + \frac{\varepsilon+2p}{2(p+\varepsilon)^2} \right) \right)$. We now prove that g is strictly increasing by showing that its derivative is positive. We get that the profit of the seller with value 0 from a take-it-or-leave-it offer of $p \in [1, k]$ is strictly increasing in p, as $g(p) = p \cdot (1 - \mathcal{F}^{k,\varepsilon}_{b|s=0}(b)) \cdot e \cdot \Pr_{s \sim \mathcal{F}^{k,\varepsilon}_s}(s=0)$.

$$
\begin{aligned}
g'(p) &= \frac{k+1}{k+1+\varepsilon} \left(\frac{1}{k+\varepsilon} + \frac{\varepsilon}{1+\varepsilon} \right) + \left(-\frac{\varepsilon+2k}{2(k+\varepsilon)^2} + \frac{\varepsilon+2p}{2(p+\varepsilon)^2} \right) \\
&\qquad - p \frac{p}{(\varepsilon+p)^3} \\
&= \frac{-k\varepsilon + \varepsilon - \varepsilon^2}{2(k+1+\varepsilon)(k+\varepsilon)^2} + \frac{\varepsilon(k+1)}{(1+\varepsilon)(k+1+\varepsilon)} + \frac{\varepsilon+2p}{2(p+\varepsilon)^2} \\
&\qquad - \frac{p^2}{(\varepsilon+p)^3} \\
&= \frac{(-k\varepsilon + \varepsilon - \varepsilon^2)(1+\varepsilon) + 2\varepsilon(k+1)(k+\varepsilon)^2}{2(k+1+\varepsilon)(k+\varepsilon)^2(1+\varepsilon)} + \frac{\varepsilon+2p}{2(p+\varepsilon)^2} \\
&\qquad - \frac{p^2}{(\varepsilon+p)^3} \\
&> \frac{(1+\varepsilon)(\varepsilon - \varepsilon^2) + k\varepsilon(2-1-\varepsilon)}{2(k+1+\varepsilon)(k+\varepsilon)^2(1+\varepsilon)} + \frac{\varepsilon+2p}{2(p+\varepsilon)^2} - \frac{p^2}{(\varepsilon+p)^3} \\
&> \frac{\varepsilon+2p}{2(p+\varepsilon)^2} - \frac{p^2}{(\varepsilon+p)^3}.
\end{aligned}
$$

Where in the last inequality we assume that $\varepsilon - \varepsilon^2 > 0$, since we can choose $\varepsilon > 0$ to be as small as we want. We show that $\frac{\varepsilon+2p}{2(p+\varepsilon)^2} - \frac{p^2}{(\varepsilon+p)^3} > 0$, which implies that $g'(p) > 0$ for every $p \in [1, k]$:

$$\frac{\varepsilon + 2p}{2(p+\varepsilon)^2} > \frac{p^2}{(\varepsilon+p)^3} \iff 2p^2 + 3p\varepsilon + \varepsilon^2 > 2p^2$$
$$\iff 3p\varepsilon + \varepsilon^2 > 0.$$

This proves that the expected profit of the seller with value 0 from a take-it-or-leave-it offer of $p \in [1,k]$ is smaller than his expected profit from an offer of $p' \in [1,k]$ that is smaller than p. Observe that the expected profit of the seller with value 0 from a take-it-or-leave-it offer of $p < 1$ is even smaller than his expected profit from an offer of price 1 (as reducing the offer below 1, does not increase the probability that the buyer will buy the item). The profit of an offer $p = k+1$ is even larger profit than an offer of k, as the probability $1 - \mathcal{F}^{k,\varepsilon}_{b|s=0}(k+1)$ is $1 - \mathcal{F}^{k,\varepsilon}_{b|s=0}(k)$ but the price is larger $(k+1 > k)$. This concludes the proof of the lemma. □

4.2.2 *Mechanisms with Dominant Strategy for the Seller (Proof of Claim 4.4).*

Every one-sided dominant strategy mechanism for the seller is a take-it-or-leave-it offer to the seller, where the offer depends only on the value of the buyer. Fix a mechanism $M = (x,p)$ for the distribution $\mathcal{F}^{k,\varepsilon}$. For every $b \in [1,k] \cup \{k+1\}$, the mechanism offers a price p_b. By Lemma 4.7, the expected profit of the buyer with value $b \in [1,k] \cup \{k+1\}$ is higher when p_b is smaller. Since the mechanism is Bayesian incentive compatible for the buyer we claim that it must be a fixed price, i.e., p_b is equal for every $b \in [1,k] \cup \{k+1\}$. This is true since if there are two different values $p_b > p_{b'}$ and the buyer strictly prefers lower prices (Lemma 4.7), the buyer will play the strategy that sets the price to $p_{b'}$, in contradiction to the incentive compatibility of the mechanism. Let p be the fixed price and consider the allocation rule x_p (see Definition 4.2). Recall that the item can be traded only if the seller's value is at most the price. Thus, the mechanism M only trades the item when $s \le p$. Observe that the allocation rule x_p trades the item in every instance that M trades the item and might trade the item when M does not. Thus, the welfare of x_p is at least the welfare of M.

Lemma 4.7. *For $b \in [1,k] \cup \{k+1\}$ and $p \in [0, b \cdot \frac{e-1}{e} + \varepsilon]$, the expected profit of the buyer from a take-it-or-leave-it offer with price p to the seller is higher than a take-it-or-leave-it offer with price $p' > p$.*

Proof of Lemma 4.7. Let $b \in [1,k] \cup \{k+1\}$. The conditional cumulative distribution of the seller's value is $\mathcal{F}^{k,\varepsilon}_{s|b}$. Now, for every $p \in [0, b \cdot \frac{e-1}{e} + \varepsilon]$, the expected profit of the buyer from a take-it-or-leave-it offer of price p to the seller is $(b-p) \cdot \frac{b}{e(b-p+\varepsilon)}$. We define the function $g(p) = (b-p) \cdot \frac{b}{e(b-p+\varepsilon)}$ for every $p \in [0, b \cdot \frac{e-1}{e} + \varepsilon]$, and show that it is a strictly decreasing function. By definition, the function $g(p)$ is the expected profit of the buyer from a take-it-or-leave-it offer of price p to the seller. Intuitively, if $\mathcal{F}^{k,\varepsilon}_{s|b}$ was exactly the equal profit distribution, i.e., $\mathcal{F}^{k,\varepsilon}_{s|b}(p) = \frac{b}{e(b-p)}$ for every $p \in [0, b \cdot \frac{e-1}{e} + \varepsilon]$, then g was a constant function with value $\frac{b}{e}$. However, since $\mathcal{F}^{k,\varepsilon}_{s|b}$ is a slightly skewed equal profit distribution, lower prices yield strictly higher profits, and so g is a strictly decreasing function. Formally, $g(p)$ is a decreasing function since its derivative is negative for every value p in its range:

$$g'(p) = \frac{b}{e}\left(\frac{-\left(b-p+\varepsilon\right) - -\left(b-p\right)}{\left(b-p+\varepsilon\right)^2}\right) = \frac{-b\varepsilon}{e\left(b-p+\varepsilon\right)^2} \underbrace{<}_{b>0} 0.$$

□

4.2.3 *Bounding the Approximation Ratio (Proof of Claim 4.5).*

We start with computing the optimal welfare of the distribution $\mathcal{F}^{k,\varepsilon}$. Observe that in every instance in the distribution $\mathcal{F}^{k,\varepsilon}$, the buyer has a higher value than the seller, thus trade occurs in the optimal allocation rule. The optimal welfare is therefore:

$$(k+1)\left(\frac{1}{k+\varepsilon} + \frac{\varepsilon}{1+\varepsilon}\right) + \int_1^k \frac{1}{(b+\varepsilon)^2} \cdot b\, db$$
$$= (k+1)\left(\frac{1}{k+\varepsilon} + \frac{\varepsilon}{1+\varepsilon}\right) + \frac{\varepsilon}{\varepsilon+b} + \ln(b+\varepsilon)\Big|_1^k$$
$$= (k+1)\left(\frac{1}{k+\varepsilon} + \frac{\varepsilon}{1+\varepsilon}\right) + \varepsilon\left(\frac{1}{k+\varepsilon} - \frac{1}{1+\varepsilon}\right) + \ln\left(\frac{k+\varepsilon}{1+\varepsilon}\right)$$
$$= \left(\frac{k+1+\varepsilon}{k+\varepsilon} + \frac{k\varepsilon}{1+\varepsilon}\right) + \ln\left(\frac{k+\varepsilon}{1+\varepsilon}\right).$$

This expression approaches $\ln k + 1 + \frac{1}{k}$ as ε approaches 0.

We consider the possible values of the buyer's value and compute the contribution of each value to the expected welfare when the item is traded according to x_p. When $b = k+1$, x_p always sells the item (as does the optimal allocation rule). This contributes $(k+1)\left(\frac{1}{k+\varepsilon} + \frac{\varepsilon}{1+\varepsilon}\right)$ to the expected welfare.

For every $b \in [1,k]$, with probability $\frac{b}{e(b+\varepsilon)}$, the seller's value is 0. Then, according to x_p, the item is traded only when $b \ge p$. These instances contribute $\int_p^k f^{k,\varepsilon}_b(b) \cdot f^{k,\varepsilon}_s(0) \cdot b\, db$ to the expected welfare. When $b \in [1,k]$, and the seller's value is larger than 0, the item is traded only when $s \le p$ according to x_p. These instances contribute
$\int_1^k f^{k,\varepsilon}_b(b)[(b \cdot (\mathcal{F}^{k,\varepsilon}_{s|b}(p) - \mathcal{F}^{k,\varepsilon}_{s|b}(0)) + \int_p^{b\cdot\frac{e-1}{e}} s \cdot f^{k,\varepsilon}_{s|b}(s)\, ds)]\, db$
to the expected welfare. Thus, the expected welfare is the sum of these three expressions:

$$(k+1)\left(\frac{1}{k+\varepsilon} + \frac{\varepsilon}{1+\varepsilon}\right) \tag{3}$$

$$\int_p^k f^{k,\varepsilon}_b(b) \cdot f^{k,\varepsilon}_s(0) \cdot b\, db \tag{4}$$

$$\int_1^k f^{k,\varepsilon}_b(b)[(b \cdot (\mathcal{F}^{k,\varepsilon}_{s|b}(p) - \mathcal{F}^{k,\varepsilon}_{s|b}(0)) + \int_p^{b\cdot\frac{e-1}{e}} s \cdot f^{k,\varepsilon}_{s|b}(s)\, ds)]\, db \tag{5}$$

Next, we consider each expression separately, bound it from above, and take its limit as ε goes to 0. We start with the first expression (3):

$$\lim_{\varepsilon \to 0^+} (k+1)\left(\frac{1}{k+\varepsilon} + \frac{\varepsilon}{1+\varepsilon}\right) = 1 + \frac{1}{k}.$$

We continue with the second expression (4). By the dominated convergence theorem, we can swap the order of the integral and the limit operator[9].

$$\lim_{\varepsilon\to 0^+}\int_p^k f_b^{k,\varepsilon}(b)\cdot f_s^{k,\varepsilon}(0)\cdot b\,db$$
$$=\lim_{\varepsilon\to 0^+}\int_p^k \frac{1}{(b+\varepsilon)^2}\cdot\frac{b}{e\cdot(b+\varepsilon)}\cdot b\,db$$
$$=\lim_{\varepsilon\to 0^+}\int_p^k \frac{b^2}{e\cdot(b+\varepsilon)^3}\,db=\int_p^k\frac{b^2}{eb^3}\,db=\frac{1}{e}\ln\left(\frac{k}{p}\right).$$

Finally, for the third expression (5), we first break it into three expressions:

$$\int_1^k f_b^{k,\varepsilon}(b)\left(b\cdot\left(\mathcal{F}_{s|b}^{k,\varepsilon}(p)-\mathcal{F}_{s|b}^{k,\varepsilon}(0)\right)+\int_p^{b\cdot\frac{e-1}{e}} s\cdot f_{s|b}^{k,\varepsilon}(s)\,ds\right)db=$$
$$\int_1^{(p-\varepsilon)\cdot\frac{e}{e-1}}\left(1-\frac{b}{e(b+\varepsilon)}\right)\frac{b}{(b+\varepsilon)^2}\,db$$
$$+\int_{(p-\varepsilon)\cdot\frac{e}{e-1}}^k\frac{1}{(b+\varepsilon)^2}b\left(\frac{b}{e(b-p+\varepsilon)}-\frac{b}{e(b+\varepsilon)}\right)db$$
$$+\int_{(p-\varepsilon)\cdot\frac{e}{e-1}}^k\frac{1}{(b+\varepsilon)^2}\int_p^{b\cdot\frac{e-1}{e}+\varepsilon} s\cdot\frac{b}{e(b+\varepsilon-s)^2}\,ds\,db.$$

Then, we bound each expression and take its limit as ε approaches 0. Again, by the dominated convergence theorem, we can swap the order of the integral and the limit operator[9].

$$\lim_{\varepsilon\to 0^+}\int_1^{(p-\varepsilon)\cdot\frac{e}{e-1}}(1-\frac{b}{e(b+\varepsilon)})\frac{b}{(b+\varepsilon)^2}\,db=(1-\frac{1}{e})\ln\left(\frac{pe}{e-1}\right).$$
$$\lim_{\varepsilon\to 0^+}\int_{(p-\varepsilon)\cdot\frac{e}{e-1}}^k\frac{1}{(b+\varepsilon)^2}b(\frac{b}{e(b-p+\varepsilon)}-\frac{b}{e(b+\varepsilon)})\,db$$
$$=\int_{p\cdot\frac{e}{e-1}}^k\frac{1}{b^2}b(\frac{b}{e(b-p)}-\frac{b}{eb})\,db$$
$$=\frac{1}{e}\left(\ln\left(\frac{k-p}{\frac{p}{e-1}}\right)-\ln\left(\frac{k}{p\frac{e}{e-1}}\right)\right)$$
$$=\frac{1}{e}\left(\ln(k-p)+\ln(e-1)-\ln k+\ln\left(\frac{e}{e-1}\right)\right).$$
$$\lim_{\varepsilon\to 0^+}\int_{(p-\varepsilon)\cdot\frac{e}{e-1}}^k\frac{1}{(b+\varepsilon)^2}\int_p^{b\cdot\frac{e-1}{e}+\varepsilon}s\cdot\frac{b}{e(b+\varepsilon-s)^2}\,ds\,db$$
$$\leq\lim_{\varepsilon\to 0^+}\int_{(p-\varepsilon)\cdot\frac{e}{e-1}}^k\frac{1}{(b+\varepsilon)^2}\int_0^{b\cdot\frac{e-1}{e}+\varepsilon}s\cdot\frac{b}{e(b+\varepsilon-s)^2}\,ds\,db$$
$$=\int_{p\cdot\frac{e}{e-1}}^k\frac{1}{b^2}\int_0^{b\cdot\frac{e-1}{e}}s\cdot\frac{b}{e(b-s)^2}\,ds\,db$$
$$=\frac{1}{e}\int_{p\cdot\frac{e}{e-1}}^k\frac{1}{b}\left(\frac{b}{b-s}+\ln(b-s)\right)\Bigg|_0^{b\cdot\frac{e-1}{e}}$$
$$=\frac{1}{e}\int_{p\cdot\frac{e}{e-1}}^k\frac{1}{b}\left(e-1+\ln\left(\frac{1}{e}\right)\right)db=\frac{e-2}{e}\ln\left(\frac{k}{\frac{pe}{e-1}}\right).$$

Now, by summing all the expressions together, we get that as ε approaches 0, the welfare of x_p approaches:

$$1+\frac{1}{k}+\frac{1}{e}\ln\left(\frac{k}{p}\right)+(1-\frac{1}{e})\ln\left(\frac{ep}{e-1}\right)+\frac{1}{e}(\ln(k-p)-\ln k$$
$$+\ln(e-1)+\ln\left(\frac{e}{e-1}\right))+(1-\frac{2}{e})(\ln k-\ln p-\ln\left(\frac{e}{e-1}\right))$$
$$=1+\frac{1}{k}+\ln k(1-\frac{2}{e})+\frac{1}{e}\ln(k-p)+\ln p(-\frac{1}{e}+1-\frac{1}{e}-1+\frac{2}{e})$$
$$+\ln\left(\frac{e}{e-1}\right)(1-\frac{1}{e}+\frac{1}{e}-1+\frac{2}{e})+\frac{\ln(e-1)}{e}$$
$$\underbrace{\leq}_{p\geq 0}1+\frac{1}{k}+(1-\frac{1}{e})\ln k+\frac{2}{e}\ln\left(\frac{e}{e-1}\right)+\frac{\ln(e-1)}{e}$$
$$\leq 1+\frac{1}{k}+0.6+(1-\frac{1}{e})\ln k.$$

Thus, as ε approaches 0, the approximation ratio approaches $\frac{\ln k+1+\frac{1}{k}}{(1-\frac{1}{e})\ln k+0.6+1+\frac{1}{k}}$. As k approaches ∞, the upper bound approaches $\frac{e}{e-1}$.

A THE LACK OF POWER OF DOMINANT-STRATEGY MECHANISMS

Consider the discrete joint distribution $\mathcal{F}_k$. In $\mathcal{F}_k$, the set of possible values of the buyer is $\{0,k,k^2,\ldots,k^k\}$ and the set of possible values of the seller is $\{0,1,\ldots,k^{k-1}\}$. The probability that an instance (s,b) occurs is:

$$p_k(s,b)=\begin{cases}\frac{1}{b} & s=\frac{b}{k}\wedge b\in\{k,k^2,\ldots,k^k\};\\ 1-\sum_{i=1}^k\frac{1}{k^i}-\varepsilon & b=0\wedge s=0;\\ \frac{\varepsilon}{k(k+1)} & \begin{array}{l}(b\in\{0,k,k^2,\ldots,k^k\})\wedge\\(s\in\{0,1,k,k^2,\ldots,k^{k-1}\})\wedge\\(s\neq\frac{b}{k}\vee s\neq 0\vee b\neq 0);\end{array}\\ 0 & \text{otherwise.}\end{cases}$$

Where $0<\varepsilon<\frac{1}{k^k}$.

Theorem A.1. *Every dominant strategy incentive compatible mechanism for $\mathcal{F}_k$ provides an approximation ratio of $\Omega(k)$ in $\mathcal{F}_k$.*

[9]All the expressions we consider are bounded (for example, by expected welfare of the distribution, which is at most $\frac{2}{k+2}+\ln(k+1)$). In addition, the sequence of functions f_ε converges point-wise to the function f_0, for every f that we consider.

As discussed in the introduction, the impossibility immediately implies to "universally truthful" mechanisms, i.e., probability distributions over dominant-strategy incentive compatible mechanisms (that is, a mechanism that chooses a fixed price at random, which is that way that the state-of-the-art mechanisms for *independent* values [7, 12] are stated).

Proof of Theorem A.1. We prove that for every large enough $k \in \mathbb{N}$, every dominant strategy incentive compatible mechanism for $\mathcal{F}_k$ provides an approximation ratio that is no better than $\frac{k}{4}$. Recall that every dominant strategy incentive compatible mechanism is a fixed price mechanism [3]. Thus, we fix a mechanism M for the distribution $\mathcal{F}_k$ and denote its fixed price by p.

Lemma A.2. *There are at most two values of $b \in \{k, k^2, \ldots, k^k\}$ such that $b \geq p$ and $\frac{b}{k} \leq p$.*

Proof of Lemma A.2. Let $b_p \in \{k, k^2, \ldots, k^k\}$ be the smallest value that is at least p. If no such value exists, then every value $b \in \{k, k^2, \ldots, k^k\}$ is smaller than p, and thus no such value b satisfies the condition in the statement ($b \geq p$), so we are done.

Notice that for any value of b satisfying the condition, it must be that $b \geq b_p$. Therefore, we only consider such values of b from now on. If $\frac{b_p}{k} > p$, then $\frac{b}{k} > p$ for all $b \geq b_p$, and the lemma follows immediately. Hence, assume that $\frac{b_p}{k} \leq p$. For every value of $b \geq b_p \cdot k^2$, we have $\frac{b}{k} \geq b_p \cdot k > p$, so at most two values of b (b_p and $b_p \cdot k$) can satisfy $\frac{b}{k} \leq p$. □

Consider an instance $(\frac{b}{k}, b)$ in the support of the distribution $\mathcal{F}_k$, with $b \in \{k, k^2, \ldots, k^k\}$, its probability is $\frac{1}{b}$. The contribution of this instance to the expected welfare of a mechanism is $\frac{1}{b} \cdot b = 1$ if the item is traded in this instance, and $\frac{1}{b} \cdot \frac{b}{k} = \frac{1}{k}$ if the item is not traded. The expected welfare of instance where $s \neq \frac{b}{k}$ is at most $\frac{\varepsilon}{k(k+1)} \cdot k^k < \frac{1}{k(k+1)}$, and there are $k(k+1)$ such instances, so it can only increase the expected welfare by 1. Therefore, the optimal welfare of the distribution $\mathcal{F}_k$ is at least k. By Lemma A.2, in the mechanism M, there are at most two values of b in $\{k, k^2, \ldots, k^k\}$ for which the item can be traded. Therefore the expected welfare of M is at most: $2 + \frac{k-2}{k} + 1$. When $k \geq 3$ the approximation ratio of M is bounded from above by:

$$\frac{k}{3 + \frac{(k-2)}{k}} = \frac{k}{\frac{4k-2}{k}} = \frac{k^2}{4k-2} > \frac{k}{4}.$$

□

REFERENCES

[1] Moshe Babaioff, Yang Cai, Yannai A. Gonczarowski, and Mingfei Zhao. 2018. The Best of Both Worlds: Asymptotically Efficient Mechanisms with a Guarantee on the Expected Gains-From-Trade. In *Proceedings of the 2018 ACM Conference on Economics and Computation* (Ithaca, NY, USA) *(EC '18)*. Association for Computing Machinery, New York, NY, USA, 373. https://doi.org/10.1145/3219166.3219203

[2] Moshe Babaioff, Shahar Dobzinski, and Ron Kupfer. 2021. A Note on the Gains from Trade of the Random-Offerer Mechanism. arXiv:2111.07790 [cs.GT]

[3] Liad Blumrosen and Shahar Dobzinski. 2014. Reallocation Mechanisms. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation* (Palo Alto, California, USA) *(EC '14)*. Association for Computing Machinery, New York, NY, USA, 617. https://doi.org/10.1145/2600057.2602843

[4] Liad Blumrosen and Shahar Dobzinski. 2021. (Almost) efficient mechanisms for bilateral trading. *Games and Economic Behavior* 130 (2021), 369–383. https://doi.org/10.1016/j.geb.2021.08.011

[5] Liad Blumrosen and Yehonatan Mizrahi. 2016. Approximating Gains-from-Trade in Bilateral Trading. In *Web and Internet Economics*, Yang Cai and Adrian Vetta (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 400–413.

[6] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. 2017. Approximating Gains from Trade in Two-Sided Markets via Simple Mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation* (Cambridge, Massachusetts, USA) *(EC '17)*. Association for Computing Machinery, New York, NY, USA, 589–590. https://doi.org/10.1145/3033274.3085148

[7] Yang Cai and Jinzhao Wu. 2023. On the Optimal Fixed-Price Mechanism in Bilateral Trade. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (Orlando, FL, USA) *(STOC 2023)*. Association for Computing Machinery, New York, NY, USA, 737–750. https://doi.org/10.1145/3564246.3585171

[8] Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. 2016. Approximately efficient double auctions with strong budget balance. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms* (Arlington, Virginia) *(SODA '16)*. Society for Industrial and Applied Mathematics, USA, 1424–1443.

[9] Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. 2022. Approximately Efficient Bilateral Trade. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing* (Rome, Italy) *(STOC 2022)*. Association for Computing Machinery, New York, NY, USA, 718–721. https://doi.org/10.1145/3519935.3520054

[10] Yumou Fei. 2022. Improved Approximation to First-Best Gains-from-Trade. In *Web and Internet Economics*, Kristoffer Arnsfelt Hansen, Tracy Xiao Liu, and Azarakhsh Malekian (Eds.). Springer International Publishing, Cham, 204–218.

[11] Zi Yang Kang, Francisco Pernice, and Jan Vondrák. 2022. Fixed-Price Approximations in Bilateral Trade. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2964–2985. https://doi.org/10.1137/1.9781611977073.115

[12] Zhengyang Liu, Zeyu Ren, and Zihe Wang. 2023. Improved Approximation Ratios of Fixed-Price Mechanisms in Bilateral Trades. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (Orlando, FL, USA) *(STOC 2023)*. Association for Computing Machinery, New York, NY, USA, 751–760. https://doi.org/10.1145/3564246.3585160

[13] Komal Malik. 2022. Optimal Robust Mechanism in Bilateral Trading. arXiv:2212.14367 [econ.TH]

[14] R.Preston McAfee. 1992. A dominant strategy double auction. *Journal of Economic Theory* 56, 2 (1992), 434–450. https://doi.org/10.1016/0022-0531(92)90091-U

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

[16] R. Preston McAfee and Philip J. Reny. 1992. Correlated Information and Mecanism Design. *Econometrica* 60, 2 (1992), 395–421. http://www.jstor.org/stable/2951601

[17] Daniel Motoc. 2021. *The Bilateral Trade Problem: The Case of Correlated Values*. Undergraduate thesis. University of Michigan.

[18] Roger B Myerson and Mark A Satterthwaite. 1983. Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29, 2 (1983), 265–281. https://doi.org/10.1016/0022-0531(83)90048-0

Received 12-NOV-2023; accepted 2024-02-11