

Fixed-Price Approximations in Bilateral Trade

Zi Yang Kang, Francisco Pernice, Jan Vondrák

August 2, 2021

Abstract

We consider the bilateral trade problem, in which two agents trade a single indivisible item. It is known that the only dominant-strategy truthful mechanism is the fixed-price mechanism: given commonly known distributions of the buyer's value B and the seller's value S , a price p is offered to both agents and trade occurs if $S \leq p < B$. The objective is to maximize either expected welfare $\mathbb{E}[S + (B - S)\mathbf{1}_{S \leq p < B}]$ or expected gains from trade $\mathbb{E}[(B - S)\mathbf{1}_{S \leq p < B}]$.

We determine the optimal approximation ratio for several variants of the problem. When the agents' distributions are identical, we show that the optimal approximation ratio for welfare is $\frac{2+\sqrt{2}}{4}$. The optimal approximation for gains from trade in this case was known to be $1/2$; we show that this can be achieved even with just 1 sample from the common distribution. We also show that a $3/4$ -approximation to welfare can be achieved with 1 sample from the common distribution. When agents' distributions are not required to be identical, we show that a previously best-known $(1 - 1/e)$ -approximation can be strictly improved, but $1 - 1/e$ is optimal if only the seller's distribution is known.

1 Introduction

In this paper, we study the performance of fixed-price mechanisms in the canonical context of bilateral trade, in which a buyer and seller bargain over a single indivisible good. As is standard in the literature, agents have independent private values for the good, B for the buyer and S for the seller, and the mechanism designer is given some information about the distributions from which these values are drawn. We consider two basic settings in this paper: one where the distributions of the agents' values are identical (the “symmetric” case), and one where the distributions are arbitrary (the “general” or “asymmetric” case). We measure the performance of a mechanism by considering either (i) the “gains from trade”: $B - S$ whenever trade occurs and 0 otherwise, or (ii) the “welfare”: B if trade occurs and S otherwise. Observe that the difference between welfare and gains from trade is simply the seller's value.

The “first-best” optimum is considered to be $\max\{B, S\}$ for welfare, and $\max\{B - S, 0\}$ for gains from trade. In their seminal work, Myerson and Satterthwaite [13] showed that no budget-balanced, individually rational Bayesian incentive-compatible (BIC) mechanism attains the first-best optimum in general. They also presented a “second-best” mechanism which is budget-balanced, individually rational, BIC, and achieves the best possible gains from trade attainable by any such mechanism. However, the second-best mechanism is complicated and difficult to implement in practical settings, so a recent growing literature seeks simpler mechanisms that maintain provably near-optimal approximation guarantees.

Motivated by the need for strategic simplicity, we focus on dominant-strategy incentive compatible (DSIC) mechanisms. A fundamental result of Hagerty and Rogerson [10] is that essentially every DSIC mechanism for bilateral trade has the following form: the mechanism designer chooses a (possibly random) price and offers trade at this price to the buyer and seller. Trade occurs if and only if the buyer's value exceeds the chosen price and the seller's value does not; this mechanism is DSIC. We refer to this as a *fixed-price* or *posted-price* mechanism. Given this explicit characterization of the allowed mechanisms under our incentive compatibility criterion, other questions arise. How close to the efficient outcome can fixed-price mechanisms get? What information about the agents is required to set the optimal price? In this paper we give a comprehensive answer to these questions in the special case of symmetric bilateral trade, and partial answers in the general case.

1.1 Prior Work

Since the seminal work of Myerson and Satterthwaite [13], a large literature has emerged that characterizes the optimal incentive-compatible mechanisms in various settings. When the strong budget balance condition is imposed, [10] shows that dominant-strategy mechanisms for bilateral trade must essentially be fixed-price mechanisms, which we study in this paper for their strategic simplicity. [7] and [16] show that fixed-price mechanisms can be optimal even when the budget balance condition is relaxed to a no-deficit condition.

The two papers most closely related to ours are [3] and [8]. The former shows that there is a randomized choice of a price based on the seller's distribution which provides a $(1 - 1/e)$ -approximation to the optimal welfare (i.e., the expected welfare provided by the mechanism is at least $1 - 1/e$ times the optimal welfare). This approximation improves the bounds shown in previous work by [4] and [6]. On the other hand, [8] shows that one can get a $1/2$ -approximation of welfare by just posting a sample from the seller distribution as a price (also in the asymmetric setting). Approximations to gains from trade have also been extensively studied in the literature [2, 5, 15, 12, 14, 11]. In particular, [12] shows that in the case of identical buyer's and seller's distributions, setting the price p to be the *median* of the distribution achieves a $1/2$ -approximation

to optimal gains from trade. More recently, [11] shows that choosing the *mean* also achieves a 1/2-approximation to optimal gains from trade, and this choice is in fact best possible among all fixed prices for a given distribution.

There have been various other results that show how first-best efficiency can be approximated when agents’ values are drawn from restricted families of distributions (i.e., distributions satisfying certain regularity conditions), such as those due to [1] and [5]. In this spirit, [16] demonstrates the optimality of the mean as a fixed-price mechanism in a bilateral trade setting for log-concave distributions that have an increasing hazard rate.¹ By contrast, our objective is to consider the most general class of distributions possible. As such, our results do not depend on distributional assumptions; all we require is that the distribution of agents’ values has finite and nonzero mean.

1.2 Our Results

We extend the state of the art for bilateral trade in the following ways:

1. We show that a 1/2-approximation of optimal gains from trade in the symmetric case (which was already known by using the mean or median of the distribution) can be also achieved with the knowledge of a single prior sample from the distribution; by posting a price equal to this sample. Interestingly, our mechanism always achieves exactly 1/2 of optimal gains from trade, for any distribution.
2. We prove that the same mechanism (posting a prior sample as a price) attains a 3/4-approximation to the optimal welfare in the symmetric case (compared to a 1/2-approximation by using a single sample in the asymmetric setting [8]). 3/4 is also tight for this mechanism.
3. We show that in the symmetric setting, a fixed-price mechanism using the mean as a price attains a $\frac{2+\sqrt{2}}{4}$ -approximation of the optimal welfare, which is the best possible among fixed-price mechanisms.
4. For general (asymmetric) bilateral trade, with arbitrary agents’ distributions, we prove that there is a $(1 - 1/e + \epsilon)$ -approximation to optimal welfare for some constant $\epsilon > 0$, improving the previously best-known $(1 - 1/e)$ -approximation [3].
5. We also show that the $1 - 1/e$ -approximation cannot be improved with the knowledge of the seller’s distribution only (which is what the mechanism of [3] uses).

Variant	Welfare approximation	Gains from trade approximation
Symmetric, full knowledge	$(2 + \sqrt{2})/4^*$	1/2* [11, 12]
Symmetric, 1 prior sample	$3/4^*$	$1/2^*$
Asymmetric, full knowledge	$1 - 1/e + \epsilon$	0* [3]
Asymmetric, 1 prior sample	1/2 [8]	0* [3]

Table 1: A survey of fixed-price mechanism approximations for bilateral trade. Our contributions are highlighted in magenta. Provably optimal results (within their variant) are marked with a *.

¹Another difference between the result in [16] and our Proposition 5 is that they do not restrict attention to posted-price mechanisms *ex ante* as we do, but rather consider dominant-strategy mechanisms with a no-deficit condition (as opposed to strong budget balance).

2 Model

Here we formally outline the model that was discussed above.

Bilateral Trade. The bilateral trade problem is simple: two agents, a buyer and a seller, value an item that is held by the seller at B and S respectively. The specific realizations of B and S are unknown to us, but we assume that we have $B \sim F_B$ and $S \sim F_S$, with B and S independent, for distribution functions F_S and F_B , about which we are given some information. For most of this paper, we assume that we have $F := F_B = F_S$.

Bilateral Trade Mechanisms. A bilateral trade mechanism consists of an allocation function $A : \mathbb{R} \times \mathbb{R} \rightarrow \{0, 1\}$, which takes in as input the reported valuations of the buyer and seller and outputs 1 if a transaction should occur or 0 otherwise, and a payment function $\Pi : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ which, if a trade occurs, determines the price at which the item should be transacted. By the revelation principle [9], we are interested in *incentive-compatible* mechanisms (A, Π) , where the agents are incentivized to report their true values. There are two standard notions of incentive compatibility that are considered in the literature: Bayesian incentive compatibility (BIC), and dominant-strategy incentive compatibility (DSIC). The former roughly says that reporting true values should be an optimal strategy for each agent, *in expectation* over the possible behavior of the other agent. The latter notion, which is what we will consider in this paper, means that reporting true values is *always* an optimal strategy for all agents, regardless of what the other agent does. Hence the notion that agents have incentive to act according to their true preferences is fully described by the DSIC property of a mechanism, which is formally defined as

$$\begin{cases} \Pi(s, b) - s \cdot A(s, b) \geq \Pi(s', b) - s \cdot A(s', b), \\ b \cdot A(s, b) - \Pi(s, b) \geq b \cdot A(s, b') - \Pi(s', b') \end{cases} \quad \text{for all } s, s', b, b' \in \mathbb{R}.$$

That is, both agents are always better off reporting their true preferences. It was shown by [10] that such mechanisms are essentially *fixed-price* mechanisms, where the payment function Π is taken to be a (possibly random) single value p , only depending on the distribution from which the valuations come from (and not on the valuations themselves). Then the allocation function is given by $A(b, s) = \mathbf{1}_{s \leq p \leq b}$: trade occurs if and only if the seller values the item less than p , and the buyer values it more.

Gains From Trade and Welfare. As discussed above, we consider two benchmarks in this paper: gains from trade and welfare. The optimal gains from trade, given a distribution function F from which the valuations are drawn, is defined as

$$\text{OPT-GFT}(F) = \mathbb{E}[\mathbf{1}_{B > S}(B - S)],$$

while the gains from trade achieved by a particular, non-random posted price p is given by

$$\text{GFT}(p, F) = \mathbb{E}[\mathbf{1}_{B \geq p > S}(B - S)].$$

Similarly, the corresponding welfare measures are defined as

$$\begin{aligned} \text{OPT-W}(F) &= \mathbb{E}[S] + \mathbb{E}[\mathbf{1}_{B > S}(B - S)] = \mathbb{E}[S] + \text{OPT-GFT}(F), \\ \text{W}(p, F) &= \mathbb{E}[S] + \mathbb{E}[\mathbf{1}_{B \geq p > S}(B - S)] = \mathbb{E}[S] + \text{GFT}(p, F). \end{aligned}$$

Some of our proofs in the next sections rely on explicit formulas for each of these quantities, which we now derive. The following derivation from [12] gives a formula for the expected optimal gains from trade: letting f denote the density of F , we have

$$\begin{aligned}
\text{OPT-GFT}(F) &= \mathbb{E}[\mathbf{1}_{B>S}(B - S)] \\
&= \int_0^\infty \int_0^b (b - s)f(s)f(b) \, ds \, db \\
&= \int_0^\infty \left[(b - s)F(s) \Big|_{s=0}^b + \int_0^b F(s) \, ds \right] f(b) \, db \\
&= \int_0^\infty \int_0^b F(s) \, ds \cdot f(b) \, db \\
&= \int_0^\infty F(s) \, ds \cdot \left[F(b) \Big|_{b=0}^\infty - \int_0^\infty F(b)^2 \, db \right] \\
&= \int_0^\infty F(s) \, ds - \int_0^\infty F(b)^2 \, db \\
&= \int_0^\infty F(x)(1 - F(x)) \, dx.
\end{aligned}$$

A similar calculation gives an analogous formula for the gains from trade attained by posting a price p . Indeed, we have

$$\begin{aligned}
\text{GFT}(p, F) &= \mathbb{E}[\mathbf{1}_{B \geq p > S}(B - S)] \\
&= \mathbb{E}[\mathbf{1}_{B \geq p > S}(B - p)] + \mathbb{E}[\mathbf{1}_{B \geq p > S}(p - S)] \\
&= \mathbb{E}[\mathbf{1}_{B \geq p}(B - p)]\mathbb{P}(S < p) + \mathbb{E}[\mathbf{1}_{p > S}(p - S)]\mathbb{P}(B \geq p) \\
&= F(p) \int_p^\infty f(b)(b - p) \, db + (1 - F(p)) \int_0^p f(s)(p - s) \, ds \\
&= F(p) \int_p^\infty (1 - F(x)) \, dx + (1 - F(p)) \int_0^p F(x) \, dx,
\end{aligned}$$

where the last equality is by integration by parts of each integral, as in the derivation of the formula for $\text{OPT-GFT}(F)$ above.

3 Main Results

We organize the results into three subsections. In the first subsection, we focus on the setting where only one sample from F is given, and show a $1/2$ -approximation of gains from trade and a $3/4$ -approximation of welfare. In the second subsection, we work in the setting of full access to the common distribution F , and show that posting the mean of F as a price gives a $(2 + \sqrt{2})/4$ -approximation of welfare, which is the best-possible in the symmetric model. Finally, in the third subsection we return to the asymmetric setting, where we now have different distributions F_S and F_B for the buyer and seller, respectively. In this setting we give a $(1 + 1/e + 0.0001)$ -approximation of welfare, showing that the mechanism of [4] is not optimal; however, we show that $1 - 1/e$ is indeed optimal if one uses only the seller distribution to set the price.

3.1 Single-Sample Approximations of Gains From Trade and Welfare

We now show that our posted-price mechanism with one sample achieves a $1/2$ -approximation of the optimal gains from trade. As mentioned above, [11] showed that this is in fact the best-possible worst-case approximation in the symmetric setting, even among DSIC mechanisms with full access to the distribution F . We emphasize that this is not a bound, but an exact computation of the approximation ratio. Formally, since our mechanism posts a sample from F as a price, we are interested in the ratio

$$\alpha := \frac{\mathbb{E}_{p \sim F}[\text{GFT}(p, F)]}{\text{OPT-GFT}(F)}.$$

As we show next, we can compute exactly that $\alpha = 1/2$, independently of F .

Theorem 1. *The symmetric bilateral trade mechanism which under a valuation distribution F posts a price $p \sim F$ achieves exactly $1/2$ of the optimal gains from trade.*

Proof. We can write $\mathbb{E}_{p \sim F}[\text{GFT}(p, F)]$ as

$$\int_0^\infty \left[F(p) \int_p^\infty (1 - F(x)) \, dx + (1 - F(p)) \int_0^p F(x) \, dx \right] f(p) \, dp.$$

We now define $\gamma_1 := \int_0^\infty f(p)F(p) \int_p^\infty (1 - F(x)) \, dx \, dp$ and $\gamma_2 := \int_0^\infty f(p)(1 - F(p)) \int_0^p F(x) \, dx \, dp$, which we will compute separately for simplicity, since we then have $\mathbb{E}_{p \sim F}[\text{GFT}(p, F)] = \gamma_1 + \gamma_2$. Let $\lambda := \int_0^\infty (1 - F(x)) \, dx$. Now for the first term, we have

$$\begin{aligned} \gamma_1 &= \int_0^\infty f(p)F(p) \int_p^\infty (1 - F(x)) \, dx \, dp \\ &= \int_0^\infty f(p)F(p) \left(\lambda - \int_0^p (1 - F(x)) \, dx \right) \, dp \\ &= \lambda \int_0^\infty f(p)F(p) \, dp - \int_0^\infty f(p)F(p) \int_0^p (1 - F(x)) \, dx \, dp \\ &= \frac{\lambda}{2} - \int_0^\infty f(p)F(p) \int_0^p (1 - F(x)) \, dx \, dp \\ &= \frac{\lambda}{2} - \frac{1}{2} F(p)^2 \int_0^p (1 - F(x)) \, dx \Big|_0^\infty + \frac{1}{2} \int_0^\infty F(p)^2 (1 - F(p)) \, dp \\ &= \frac{1}{2} \int_0^\infty F(p)^2 (1 - F(p)) \, dp, \end{aligned}$$

where we used that

$$\int_0^\infty f(x)F(x) \, dx = \frac{1}{2} F(x)^2 \Big|_0^\infty = 1/2.$$

Note $\int_0^a f(x)(1 - F(x)) \, dx = F(a)(1 - \frac{1}{2}F(a))$ by integrating by parts. Thus for the second term we can write

$$\begin{aligned} \gamma_2 &= \lim_{a \rightarrow \infty} \int_0^a f(p)(1 - F(p)) \int_0^p F(x) \, dx \, dp \\ &= \lim_{a \rightarrow \infty} \left[F(a)(1 - \frac{1}{2}F(a)) \int_0^a F(x) \, dx - \int_0^a F(p)^2 (1 - \frac{1}{2}F(p)) \, dp \right]. \end{aligned}$$

Now for $\theta := F(a)(1 - \frac{1}{2}F(a))$ we have $\theta \rightarrow 1/2$ as $a \rightarrow \infty$. Thus we have

$$\begin{aligned}\gamma_2 &= \lim_{a \rightarrow \infty} \left[\theta \int_0^a F(x) \, dx - \int_0^a F(p)^2 \left(1 - \frac{1}{2}F(p)\right) \, dp \right] \\ &= \int_0^\infty \left[\frac{1}{2}F(p) - F(p)^2 \left(1 - \frac{1}{2}F(p)\right) \right] \, dp \\ &= \frac{1}{2} \int_0^\infty F(p)(1 - F(p))^2 \, dp.\end{aligned}$$

Putting things together we get

$$\begin{aligned}\mathbb{E}_{p \sim F}[\text{GFT}(F; p)] &= \gamma_1 + \gamma_2 \\ &= \frac{1}{2} \int_0^\infty F(p)^2(1 - F(p)) \, dp + \frac{1}{2} \int_0^\infty F(p)(1 - F(p))^2 \, dp \\ &= \frac{1}{2} \int_0^\infty F(p)(1 - F(p))[F(p) + (1 - F(p))] \, dp \\ &= \frac{1}{2} \int_0^\infty F(p)(1 - F(p)) \, dp.\end{aligned}$$

Recalling that $\text{OPT-GFT}(F) = \int_0^\infty (1 - F(x))F(x) \, dx$, we get $\alpha = 1/2$, completing the proof. \square

Given Theorem 1, it is simple to extend our result to obtain a 3/4 approximation of optimal welfare. Recall the representations of the welfare measures that we derived in the previous section:

$$\begin{aligned}\text{OPT-W}(F) &= \mu + \text{OPT-GFT}(F), \\ \text{W}(F, p) &= \mu + \text{GFT}(F, p),\end{aligned}$$

where $\mu := \mathbb{E}[S] = \mathbb{E}[B]$. We can thus get the 3/4-approximation by applying Theorem 1:

Theorem 2. *The symmetric bilateral trade mechanism which under a valuation distribution F posts a price $p \sim F$ achieves a 3/4-approximation of the optimal welfare.*

Proof. Using that $\mathbb{E}_{p \sim F}[\text{GFT}(F, p)]/\text{OPT-GFT}(F) = 1/2$, we can write

$$\begin{aligned}\frac{\mathbb{E}_{p \sim F}[\text{W}(F, p)]}{\text{OPT-W}(F)} &= \frac{\mu + \mathbb{E}_{p \sim F}[\text{GFT}(F, p)]}{\mu + \text{OPT-GFT}(F)} \\ &= 1 - \frac{1}{2} \frac{\text{OPT-GFT}(F)}{\mu + \text{OPT-GFT}(F)} \\ &= 1 - \lim_{a \rightarrow \infty} \frac{1}{2} \frac{\int_0^a F(x)(1 - F(x)) \, dx}{\int_0^a x f(x) \, dx + \int_0^a F(x)(1 - F(x)) \, dx} \\ &= 1 - \lim_{a \rightarrow \infty} \frac{1}{2} \frac{\int_0^a F(x) \, dx - \int_0^a F(x)^2 \, dx}{aF(a) - \int_0^a F(x)^2 \, dx}.\end{aligned}$$

But note that we have $\int_0^a F(x) \, dx \leq a^{1/2}(\int_0^a F(x)^2 \, dx)^{1/2}$ by the Cauchy–Schwarz inequality. Thus letting $t := (\int_0^a F(x)^2 \, dx)^{1/2}$, and noting that $0 \leq t \leq a^{1/2}$ and $F(a) \leq 1$, we get

$$\frac{\int_0^a F(x) \, dx - \int_0^a F(x)^2 \, dx}{aF(a) - \int_0^a F(x)^2 \, dx} \leq \frac{a^{1/2}t - t^2}{aF(a) - t^2}$$

$$\begin{aligned}
&\leq \frac{1}{F(a)} \frac{t(a^{1/2} - t)}{(a^{1/2} + t)(a^{1/2} - t)} \\
&= \frac{1}{F(a)} \frac{t}{a^{1/2} + t} \\
&\leq \frac{1}{F(a)} \frac{1}{2} \\
&\xrightarrow{a \rightarrow \infty} 1/2.
\end{aligned}$$

Thus we have $\frac{\mathbb{E}_{p \sim F}[\mathbb{W}(F, p)]}{\text{OPT-W}(F)} \geq 1 - \frac{1}{2} \frac{1}{2} = 3/4$, as desired. \square

Remark 3. Unlike our previous result for gains from trade, this approximation ratio is a lower bound, and not an exact computation. Indeed, in the case of welfare, the approximation ratio of our mechanism varies with the distribution F . Moreover, in further contrast with Theorem 1, this 3/4-approximation is not the best-possible among the class of truthful mechanisms with full knowledge of F , since we show in the next subsection that the optimal ratio here is instead $(2 + \sqrt{2})/4 \approx 0.8536$.

Our remark notwithstanding, the approximation lower bound of 3/4 from Theorem 2 for our mechanism is tight, as we show next.

Lemma 4. The 3/4-approximation bound of Theorem 2 is tight.

Proof. Take the distribution functions $F_r(x) = \min\{x^r, 1\}$ parameterized by $r \in \mathbb{R}_+$. From the proof of Theorem 2, we get

$$\begin{aligned}
\frac{\mathbb{E}_{p \sim F}[\mathbb{W}(F_r, p)]}{\text{OPT-W}(F_r)} &= 1 - \frac{1}{2} \frac{\int_0^1 x^r (1 - x^r) dx}{r \int_0^1 x^r dx + \int_0^1 x^r (1 - x^r) dx} \\
&= 1 - \frac{1}{2} \frac{\frac{1}{r+1} - \frac{1}{2r+1}}{\frac{r}{r+1} + \frac{1}{r+1} - \frac{1}{2r+1}} \\
&\xrightarrow{r \rightarrow 0} 1 - \frac{1}{2} \frac{1}{2} = 3/4, \quad \text{as desired.} \quad \square
\end{aligned}$$

3.2 Best-Possible Approximation of Welfare

We begin by showing that the optimal posted price in the symmetric setting is given simply by the mean of the common distribution.

Proposition 5. The welfare-maximizing mechanism for symmetric bilateral trade is given by posting the mean of the common distribution F :

$$p^* = \mathbb{E}[S] = \mathbb{E}[B].$$

Proof. For any positive price $p > 0$, we may rewrite the welfare function as

$$\begin{aligned}
\mathbb{W}(p, F) &= \mathbb{E}[S] + \mathbb{E}[(B - S) \cdot \mathbf{1}_{B > p \geq p}] \\
&= \mathbb{E}[S] + \mathbb{E}[B \cdot \mathbf{1}_{B > p}] \cdot F(p) - \mathbb{E}[S \cdot \mathbf{1}_{S \leq p}] \cdot [1 - F(p)].
\end{aligned}$$

Because both agents' values are drawn from the same distribution,

$$\mathbb{E}[B \cdot \mathbf{1}_{B > p}] = \mathbb{E}[S \cdot \mathbf{1}_{S > p}] = \mathbb{E}[S] - \mathbb{E}[S \cdot \mathbf{1}_{S \leq p}].$$

Thus:

$$\begin{aligned}
\mathbb{W}(p, F) &= \mathbb{E}[S] + \mathbb{E}[B \cdot \mathbf{1}_{B > p}] \cdot F(p) - \mathbb{E}[S \cdot \mathbf{1}_{S \leq p}] \cdot [1 - F(p)] \\
&= \mathbb{E}[S] \cdot [1 + F(p)] - \mathbb{E}[S \cdot \mathbf{1}_{S \leq p}] \\
&= \mathbb{E}[S] \cdot [1 + F(p)] - p \cdot F(p) + \int_0^p F(s) \, ds.
\end{aligned}$$

Now, for any fixed $\delta \in (0, p)$,

$$\mathbb{W}(p \pm \delta; F) = \mathbb{E}[S] \cdot [1 + F(p \pm \delta)] - (p \pm \delta) \cdot F(p \pm \delta) + \int_0^{p \pm \delta} F(x) \, dx.$$

We thus compute that

$$\begin{cases}
\mathbb{W}(p^*; F) - \mathbb{W}(p^* - \delta; F) = [F(p^*) - F(p^* - \delta)] \cdot (\mathbb{E}[S] - p^*) - \delta \cdot F(p^* - \delta) + \int_{p^* - \delta}^{p^*} F(x) \, dx, \\
\mathbb{W}(p^*; F) - \mathbb{W}(p^* + \delta; F) = [F(p^*) - F(p^* + \delta)] \cdot (\mathbb{E}[S] - p^*) + \delta \cdot F(p^* + \delta) - \int_{p^*}^{p^* + \delta} F(x) \, dx.
\end{cases}$$

In the first expression, since F is non-decreasing,

$$\int_{p^* - \delta}^{p^*} F(x) \, dx \geq \int_{p^* - \delta}^{p^*} F(p^* - \delta) \, dx = \delta \cdot F(p^* - \delta).$$

Likewise, in the second expression, the same observation yields

$$-\int_{p^*}^{p^* + \delta} F(x) \, dx \geq -\int_{p^*}^{p^* + \delta} F(p^* + \delta) \, dx = -\delta \cdot F(p^* + \delta).$$

Therefore, with $p^* = \mathbb{E}[S] = \mathbb{E}[B]$, substitution yields

$$\begin{cases}
\mathbb{W}(p^*; F) - \mathbb{W}(p^* - \delta; F) \geq [F(p^*) - F(p^* - \delta)] \cdot (\mathbb{E}[S] - p^*) = 0, \\
\mathbb{W}(p^*; F) - \mathbb{W}(p^* + \delta; F) \geq [F(p^*) - F(p^* + \delta)] \cdot (\mathbb{E}[S] - p^*) = 0.
\end{cases}$$

Therefore the optimal fixed-price mechanism is given by $p^* = \mathbb{E}[S]$. \square

Proposition 5 shows that only the first moment of the agents' distribution is required by the mechanism designer to determine the optimal fixed price; all higher moments are inconsequential. Therefore, precise knowledge of the agents' distribution, other than the mean of the distribution, is irrelevant.

In general, the fixed-price mechanism $p^* = \mathbb{E}[S]$ is not uniquely optimal. Suppose, for example, that each agent's value is 0 with probability 1/2 and 1 with probability 1/2. Then any price $p \in (0, 1)$ is optimal. Nonetheless, under weak regularity conditions on F , the mean is the uniquely optimal price.

Corollary 6. *Suppose that F is differentiable in a neighborhood of its mean, so that its density $f = F'$ is positive in that neighborhood. Then the optimal fixed-price mechanism that sets the price as the mean of the distribution F is uniquely optimal.*

The proof of Corollary 6 can be directly obtained from the proof of Proposition 5 by noting that the inequalities must hold strictly under the stated assumptions; we thus omit the details. In particular, Corollary 6 applies under the usual assumption that F is continuously differentiable with positive density.

Despite Corollary 6, we proceed without additional regularity assumptions on F to obtain a general analysis. Given the mean $\mu := \mathbb{E}[S]$ of the agents' distribution, we can now determine the expected welfare that the optimal fixed-price mechanism achieves, $W(\mu; F)$:

$$\begin{aligned} W(\mu; F) &= \mu \cdot [1 + F(\mu)] - \mathbb{E}[S \cdot \mathbf{1}_{S \leq \mu}] \\ &= \mu + (\mu - \mathbb{E}[S | S \leq \mu]) \cdot F(\mu). \end{aligned}$$

An interesting observation here is that the maximum expected welfare achieved by *any* fixed-price mechanism depends on the agents' distribution F through only three quantities: (i) the mean of the distribution, μ ; (ii) the quantile of the mean, $F(\mu)$; and (iii) the mean of the distribution conditional on being no greater than μ , $\mathbb{E}[S | S \leq \mu]$. Consequently, any modification to F that preserves these three quantities will lead to the same maximum expected welfare.

This observation motivates the following approach to find the approximation ratio we achieve by posting the mean as a price which, as we showed, is the best-possible among DSIC mechanisms:

$$\inf_F \frac{W(\mu_F; F)}{\text{OPT-W}(F)} = \inf_F \sup_{p \in \mathbb{R}_+} \frac{W(p, F)}{\text{OPT-W}(F)}. \quad (\text{A})$$

Define the subspace of probability distributions that fixes the three quantities from above:

$$\Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma) := \{F \in \Delta_{L^1}(\mathbb{R}_+) : \mathbb{E}[S] = \mu, \mathbb{E}[S | S \leq \mu] = \mu_1, F(\mu) = \gamma\}.$$

To compute (A), we first compute

$$\inf_{F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)} \sup_{p \in \mathbb{R}_+} \frac{W(p, F)}{\text{OPT-W}(F)}. \quad (\text{A}')$$

The original problem (A) is thus equivalent to

$$\inf_{\substack{\mu, \mu_1, \gamma \in \mathbb{R}_+ \\ 0 \leq \mu_1 \leq \mu; \mu > 0; 0 < \gamma \leq 1}} \left\{ \inf_{F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)} \sup_{p \in \mathbb{R}_+} \frac{W(p, F)}{\text{OPT-W}(F)} \right\}.$$

The objective of this decoupling is to reduce the dimension of our optimization problem. Instead of minimizing over the infinite-dimensional space $\Delta_{L^1}(\mathbb{R}_+)$, our outer optimization problem is simplified to minimization over three variables. However, the inner optimization problem (A') still requires minimizing over the infinite-dimensional space $\Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$. To reduce the dimension of (A'), define the space of distributions *supported on at most 4 points* that fixes the three quantities:

$$\begin{aligned} \Delta_{L^1}^{(4)}(\mathbb{R}_+; \mu, \mu_1, \gamma) &:= \{F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma) : F(x) = q_0 \cdot \mathbf{1}_{x \geq 0} + q_1 \cdot \mathbf{1}_{x \geq x_1} \\ &\quad + q_2 \cdot \mathbf{1}_{x \geq x_2} + q_3 \cdot \mathbf{1}_{x \geq 1}\}. \end{aligned}$$

The probability masses are given by $0 \leq q_0, q_1, q_2, q_3 \leq 1$.

Lemma 7. *For any fixed $0 \leq \mu_1 \leq \mu \leq 1$, $\mu > 0$ and $0 < \gamma \leq 1$,*

$$\inf_{F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)} \sup_{p \in \mathbb{R}_+} \frac{W(p, F)}{\text{OPT-W}(F)} = \inf_{F \in \Delta_{L^1}^{(4)}(\mathbb{R}_+; \mu, \mu_1, \gamma)} \sup_{p \in \mathbb{R}_+} \frac{W(p, F)}{\text{OPT-W}(F)}.$$

We defer the proof of Lemma 7 to Appendix A and sketch the argument here. Fix $F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$. We may assume without loss of generality that F is a finite support discrete distribution because any $F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$ can be approximated arbitrarily well by a corresponding distribution in $\Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$ supported on finitely many points. Moreover, by rescaling if necessary, we may also assume that $F(1) = 1$.

To solve the inner problem (A'), our earlier observation shows that, given μ, μ_1 and γ , $W(p, F) = W(\mu; F)$ is constant for any $F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$. That is:

$$\inf_{F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)} \sup_{p \in \mathbb{R}_+} \frac{W(p, F)}{\text{OPT-W}(F)} = \inf_{F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)} \frac{\mu + (\mu - \mu_1) \cdot \gamma}{\text{OPT-W}(F)}.$$

Thus (A') can be solved by *maximizing* $\text{OPT-W}(F)$ over all $F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$. Note that, from the definition,

$$\text{OPT-W}(F) = \mathbb{E}[\max\{S, B\}] = \mathbb{E}[S] + \frac{1}{2} \mathbb{E}[|B - S|].$$

Therefore, $\text{OPT-W}(F)$ can be optimized by the following procedure:

- For any probability mass in $(0, \mu)$, split the mass into two equal masses. Move each mass in opposite directions, until one mass hits the boundary of the interval $(0, \mu)$.
- For any probability mass in $(\mu, 1)$, split the mass into two equal masses. Move each mass in opposite directions, until one mass hits the boundary of the interval $(\mu + \delta, 1)$, for some sufficiently small $\delta > 0$.²

Each of these operations is mean-preserving in the intervals $(0, \mu)$ and $(\mu, 1)$, so neither changes the quantities $\mu = \mathbb{E}[S]$, $\mu_1 = \mathbb{E}[S | S \leq \mu]$ and $\gamma = F(\mu)$. Therefore, the operations preserve the maximum expected welfare achieved by any fixed-price mechanism, $W(\mu; F) = \mu + (\mu - \mu_1) \cdot \gamma$. However, they increase $\mathbb{E}[|B - S|]$ in the respective intervals, and so the operations increase $\text{OPT-W}(F)$. Applying these operations recursively,³ we obtain a 4-point distribution $\tilde{F} \in \Delta_{L^1}^{(4)}(\mathbb{R}_+; \mu, \mu_1, \gamma)$ such that

$$\frac{W(p; \tilde{F})}{\text{OPT-W}(\tilde{F})} \leq \frac{W(p, F)}{\text{OPT-W}(F)}.$$

Since $\Delta_{L^1}^{(4)}(\mathbb{R}_+; \mu, \mu_1, \gamma) \subset \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$, this proves the result of Lemma 7.

As a consequence of Lemma 7, we can restrict attention to 4-point distributions when solving the minimax problem (A). Such distributions $F \in \Delta_{L^1}^{(4)}(\mathbb{R}_+; \mu, \mu_1, \gamma)$ can be written as⁴

$$F(x) = q_0 \cdot \mathbf{1}_{x \geq 0} + q_1 \cdot \mathbf{1}_{x \geq \mu} + q_2 \cdot \mathbf{1}_{x \geq \mu + \delta} + q_3 \cdot \mathbf{1}_{x \geq 1} \quad \text{for some } \delta \in (0, 1 - \mu).$$

Here, the probability masses q_0, q_1, q_2 and q_3 satisfy $q_0 + q_1 + q_2 + q_3 = 1$ and the three additional conditions:

$$\begin{cases} \mathbb{E}[S] = q_1 \mu + q_2 (\mu + \delta) + q_3 & = \mu, \\ \mathbb{E}[S | S \leq \mu] = \frac{q_1 \mu}{q_0 + q_1} & = \mu_1, \\ F(\mu) = q_0 + q_1 & = \gamma. \end{cases}$$

²Note that this $\delta > 0$ is required so that this operation does not change $F(\mu) = \gamma$.

³In the proof of Lemma 7 in Appendix A, we consider only the limiting distribution \tilde{F} obtained from recursive application of these operations; however, these operations motivate the construction of \tilde{F} in the proof.

⁴The definition of $\Delta_{L^1}^{(4)}(\mathbb{R}_+; \mu, \mu_1, \gamma)$ is slightly more general as it allows for any $F(x)$ in the form

$$F(x) = q_0 \cdot \mathbf{1}_{x \geq 0} + q_1 \cdot \mathbf{1}_{x \geq x_1} + q_2 \cdot \mathbf{1}_{x \geq x_2} + q_3 \cdot \mathbf{1}_{x \geq 1}.$$

However, from the proof sketch of Lemma 7 (and further justified in the proof given in Appendix A), we may narrow our attention to distributions F where $x_1 = \mu$ and $x_2 = \mu + \delta$ for sufficiently small $\delta > 0$.

Since μ , μ_1 and γ are given, we can rewrite F as follows:⁵

$$F(x) = \gamma \left(1 - \frac{\mu_1}{\mu}\right) \cdot \mathbf{1}_{x \geq 0} + \frac{\mu_1 \gamma}{\mu} \cdot \mathbf{1}_{x \geq \mu} + \frac{1 - \gamma - \mu + \mu_1 \gamma}{1 - \mu - \delta} \cdot \mathbf{1}_{x \geq \mu + \delta} + \frac{\mu \gamma - \mu_1 \gamma - \delta - \delta \gamma}{1 - \mu - \delta} \cdot \mathbf{1}_{x \geq 1}.$$

This parametrization of F allows us to easily compute $W(\mu; F)$ and $\text{OPT-W}(F)$. We have shown previously that $W(\mu; F)$ depends only on μ , μ_1 and γ :

$$W(\mu; F) = \mu + (\mu - \mu_1) \cdot \gamma.$$

On the other hand, $\text{OPT-W}(F)$ can be computed by observing that B, S must be either: (i) both no greater than μ ; (ii) both greater than μ ; or (iii) on different sides of μ . Thus:

$$\begin{aligned} \text{OPT-W}(F) &= \mathbb{E}[\max\{B, S\}] \\ &= \mathbb{E}[\max\{B, S\} \cdot \mathbf{1}_{B, S \leq \mu}] + \mathbb{E}[\max\{B, S\} \cdot \mathbf{1}_{B, S > \mu}] \\ &\quad + \mathbb{E}[\max\{B, S\} \cdot \mathbf{1}_{\max\{B, S\} > \mu} \mathbf{1}_{\min\{B, S\} \leq \mu}]. \end{aligned}$$

The first and third contributions can be computed separately:

$$\begin{cases} \mathbb{E}[\max\{B, S\} \cdot \mathbf{1}_{B, S \leq \mu}] = [F(\mu)]^2 \cdot \mu (1 - \mathbb{P}[B, S = 0 \mid B, S \leq \mu]) &= \mu_1 \gamma^2 \left(2 - \frac{\mu_1}{\mu}\right), \\ \mathbb{E}[\max\{B, S\} \cdot \mathbf{1}_{\max\{B, S\} > \mu} \mathbf{1}_{\min\{B, S\} \leq \mu}] = 2[F(\mu)] \cdot \mathbb{E}[S \cdot \mathbf{1}_{S > \mu}] &= 2\gamma(\mu - \gamma\mu_1). \end{cases}$$

The second contribution can be bounded above independently of δ :

$$\begin{aligned} \mathbb{E}[\max\{B, S\} \cdot \mathbf{1}_{B, S > \mu}] &= [1 - F(\mu)]^2 \cdot \{1 - (1 - \mu - \delta) \cdot \mathbb{P}[B = S = \mu + \delta \mid B, S > \mu]\} \\ &= (1 - \gamma)^2 - \frac{(1 - \gamma - \mu + \mu_1 \gamma)^2}{1 - \mu - \delta} \\ &\leq (1 - \gamma)^2 - \frac{(1 - \gamma - \mu + \mu_1 \gamma)^2}{1 - \mu}. \end{aligned}$$

Consequently, we obtain an upper bound for $\text{OPT-W}(F)$ that is independent of δ :

$$\text{OPT-W}(F) \leq \mu_1 \gamma^2 \left(2 - \frac{\mu_1}{\mu}\right) + 2\gamma(\mu - \gamma\mu_1) + (1 - \gamma)^2 - \frac{(1 - \gamma - \mu + \mu_1 \gamma)^2}{1 - \mu}.$$

Therefore, we have

$$\frac{W(\mu; F)}{\text{OPT-W}(F)} \geq \frac{\mu + (\mu - \mu_1) \gamma}{\mu_1 \gamma^2 \left(2 - \frac{\mu_1}{\mu}\right) + 2\gamma(\mu - \gamma\mu_1) + (1 - \gamma)^2 - \frac{(1 - \gamma - \mu + \mu_1 \gamma)^2}{1 - \mu}}.$$

This allows us to bound the value of the minimax problem (A) from below:

$$\begin{aligned} \inf_{F \in \Delta_{L^1}(\mathbb{R}_+)} \sup_{p \in \mathbb{R}_+} \frac{W(p, F)}{\text{OPT-W}(F)} &= \inf_{\substack{\mu, \mu_1 \in \mathbb{R}_+ \\ 0 \leq \mu_1 \leq \mu; \mu, \gamma \in (0, 1)}} \left\{ \inf_{F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)} \frac{W(\mu; F)}{\text{OPT-W}(F)} \right\} \\ &\geq \inf_{\substack{\mu, \mu_1 \in \mathbb{R}_+ \\ 0 \leq \mu_1 \leq \mu; \mu, \gamma \in (0, 1)}} \frac{\mu + (\mu - \mu_1) \gamma}{\mu_1 \gamma^2 \left(2 - \frac{\mu_1}{\mu}\right) + 2\gamma(\mu - \gamma\mu_1) + (1 - \gamma)^2 - \frac{(1 - \gamma - \mu + \mu_1 \gamma)^2}{1 - \mu}}. \end{aligned}$$

The lower bound in the above expression can be computed. Notably, optimization is carried out over a 3-dimensional space, and so elementary techniques apply. While we relegate the computational details to Appendix A, we state the result here:

⁵Details of this computation can be found in the proof of Lemma 7 given in Appendix A.

Lemma 8. *The minimax value of the game is bounded below:*

$$\inf_{F \in \Delta_{L^1}(\mathbb{R}_+)} \sup_{p \in \mathbb{R}_+} \frac{W(p, F)}{\text{OPT-W}(F)} \geq \frac{2 + \sqrt{2}}{4}.$$

This bound is, in fact, tight. To prove this, we demonstrate that the bound can be attained by explicitly constructing a minimizing sequence of distributions $\{F_n\}_{n=1}^\infty$. Formally, we consider

$$F_n(x) = (\sqrt{2} - 1) \left(1 - \frac{1}{n}\right) \cdot \mathbf{1}_{x \geq 0} + (2 - \sqrt{2}) \cdot \mathbf{1}_{x \geq 1/n} + \frac{1}{n} (\sqrt{2} - 1) \cdot \mathbf{1}_{x \geq 1}.$$

The mean of the distribution F_n is $\mu_n = 1/n$. Moreover, the quantile of the mean is $\gamma_n = F_n(\mu_n) = 1 - (\sqrt{2} - 1)/n$ and

$$\mu_{1,n} = \mathbb{E}[S | S \leq \mu_n] = \frac{2 - \sqrt{2}}{n - (\sqrt{2} - 1)}.$$

Thus

$$W(\mu_n; F_n) = \mu_n + (\mu_n - \mu_{1,n}) \gamma_n = \frac{\sqrt{2}}{n} - \frac{\sqrt{2} - 1}{n^2}.$$

On the other hand, we can compute that

$$\text{OPT-W}(F_n) = \frac{4(\sqrt{2} - 1)}{n} - \frac{4\sqrt{2} - 1}{n^2}.$$

Therefore, this sequence of distributions achieves the lower bound in the limit as $n \rightarrow \infty$:

$$\lim_{n \rightarrow \infty} \frac{W(\mu_n; F_n)}{\text{OPT-W}(F_n)} = \frac{\sqrt{2}}{4(\sqrt{2} - 1)} = \frac{2 + \sqrt{2}}{4}.$$

This proves our main result:

Theorem 9. *The minimax value of the game is*

$$\inf_{F \in \Delta_{L^1}(\mathbb{R}_+)} \sup_{p \in \mathbb{R}_+} \frac{W(p, F)}{\text{OPT-W}(F)} = \frac{2 + \sqrt{2}}{4}.$$

That is, for any distribution F , the designer can always select a price that achieves at least $(2 + \sqrt{2})/4 \approx 0.8536$ of the total expected welfare under the optimally efficient outcome.

3.3 Approximations of Welfare in the Asymmetric Model

We now relax the requirement that both agents' values are drawn from the same distribution. Let the seller's (resp. buyer's) value be drawn from F_S (resp. F_B). Accordingly, we modify the notation used previously. We denote the welfare function by

$$W(p; F_S, F_B) := \mathbb{E}[S + (B - S) \cdot \mathbf{1}_{B > p \geq S}],$$

where the expectation is taken over $S \sim F_S$ and $B \sim F_B$. Likewise, we denote the expected welfare under the first-best outcome by

$$\text{OPT-W}(F_S, F_B) := \mathbb{E}[S + (B - S) \cdot \mathbf{1}_{B > S}].$$

Therefore the best-possible posted-price approximation ratio is

$$\inf_{F_S, F_B \in \Delta_{L^1}(\mathbb{R}_+)} \sup_{p \in \mathbb{R}_+} \frac{W(p; F_S, F_B)}{\text{OPT-}W(F_S, F_B)}. \quad (\text{A-ASYM})$$

As mentioned previously, [3] shows that the value of the minimax problem (A-ASYM) is bounded below by $1 - 1/e$. This is done through explicit construction of a randomized mechanism that achieves this bound in expectation. While this mechanism requires the designer to use detailed distributional information, it has the interesting property that the designer requires knowledge of only the seller’s distribution and not the buyer’s. Towards leveraging this fact, we prove the following:

Proposition 10. *Given distributions F_S and F_B that satisfy $\mathbb{E}[(S - B)_+] = 0$, the designer can always select a price that achieves at least 3/4 of the optimal expected welfare. That is:*

$$\inf_{\substack{F_S, F_B \in \Delta_{L^1}(\mathbb{R}_+) \\ \mathbb{E}[(S - B)_+] = 0}} \sup_{p \in \mathbb{R}_+} \frac{W(p; F_S, F_B)}{\text{OPT-}W(F_S, F_B)} = \frac{3}{4}.$$

We defer the proof of this proposition and all other results in this subsection to the appendix. We interpret Proposition 10 as a statement of how asymmetry between the agents can lead to considerable, albeit bounded, welfare loss. In Proposition 10, the asymmetry is “worst possible” in the sense that the seller’s value is almost surely no greater than the buyer’s value. Yet, “worst possible” asymmetry is not the same as the worst-case scenario for the minimax problem (A-ASYM). Indeed, the designer can find a price that separates all the seller’s possible values from the buyer’s possible values; because trade does not occur when the buyer’s value is exactly equal to this price, inefficiency may result.⁶

When asymmetry is not as severe, we can improve the lower bound to the minimax problem (A-ASYM) through the mechanism proposed by [3]:

Proposition 11 (Theorem 4.1 in [3]). *For any given distributions F_S and F_B ,*

$$\sup_{p \in \mathbb{R}_+} \frac{W(p; F_S, F_B)}{\text{OPT-}W(F_S, F_B)} \geq 1 - \frac{1}{e} + \frac{1}{e} \cdot \mathbb{E}[(S - B)_+].$$

In view of Propositions 10 and 11, our strict improvement to the $1 - 1/e$ bound can be achieved heuristically as follows:

- If $\mathbb{E}[(S - B)_+]$ is sufficiently small, then we construct a mechanism that exploits the severe asymmetry in the agents’ distributions to get a performance ratio close to 3/4, as in Proposition 10.
- If $\mathbb{E}[(S - B)_+]$ is large, then the mechanism of [3] already ensures a performance ratio strictly higher than $1 - 1/e$.

Our improvement is documented by the following:⁷

⁶We emphasize that this result, along with our other results in this paper, is not an artifact of the tiebreak rule that we have adopted, which specifies that trade occurs if and only if $B > p \geq S$. For instance, if a different tiebreak rule were to be chosen such that trade occurs if and only if $B > p > S$, then a similar result obtains in the limit $\mathbb{E}[(S - B)_+] \rightarrow 0$ (e.g., let $F_B(x) = (1 - 1/n) \cdot \mathbf{1}_{x \geq 1/n - 1/n^2} + 1/n \cdot \mathbf{1}_{x \geq 1}$).

⁷We qualify our result in Theorem 12 by noting that we did not optimize for the strict constant-factor improvement of 0.0001, as our goal was to demonstrate that such an improvement exists. We do not believe that optimizing for this constant would yield a substantially higher improvement.

Theorem 12. *Given distributions F_S and F_B , the designer can always select a price that achieves at least $1 - 1/e + 0.0001$ of the optimal expected welfare. That is:*

$$\inf_{F_S, F_B \in \Delta_{L^1}(\mathbb{R}_+)} \sup_{p \in \mathbb{R}_+} \frac{W(p; F_S, F_B)}{\text{OPT-W}(F_S, F_B)} \geq 1 - \frac{1}{e} + 0.0001.$$

In contrast to Theorem 9, we interpret Theorem 12 as essentially a negative result in the design of informationally simple mechanisms that achieve close to optimal welfare in this setting. Despite the fact that our proposed mechanism uses distributional knowledge of F_S and F_B (and, in particular, substantially more knowledge of the distributions than just their means), we are unable to guarantee a substantially higher performance than we would have obtained by using only distributional knowledge of F_S .

As a counterpoint to Theorem 12 and motivated by the mechanism proposed by [3], we consider mechanisms that use only information about the seller’s distribution. The question that we ask here is: Is it possible to achieve a strictly better performance than $1 - 1/e$ of optimal welfare with any mechanism that uses *only* information about the seller’s distribution F_S ? The answer is no:

Theorem 13. *Given distributions F_S and F_B , for any (possibly randomized) mechanism that uses only quantile distributional knowledge of the seller’s distribution, the designer can achieve no better than $1 - 1/e$ of the optimal expected welfare.*

4 Concluding Remarks

Our results give a comprehensive description of the power of dominant-strategy truthful mechanisms for bilateral trade, in terms of two widely-studied objectives in mechanism design: welfare and gains from trade, together with the informational requirements needed to achieve those approximations.

We would like to distinguish our paper from previous work as follows. In our view, the main contribution of our paper is to show that fixed-price mechanisms can be effective *particularly in markets with symmetric agents*, for reasons other than strategic simplicity for agents. Two important properties of fixed-price mechanisms in such markets is that they are informationally simple from the designer’s perspective while being approximately efficient at the same time. We view robustness to different benchmarks (i.e., optimal welfare and optimal gains from trade) as an important criterion for “approximate efficiency.” In our view, previous approximability and inapproximability results in the asymmetric case, which crucially depend on benchmark adopted, constitute mainly a negative result for the asymmetric case.

A fundamental question remains in the asymmetric case, which is whether a constant-factor approximation to optimal gains from trade is possible with any Bayesian incentive-compatible mechanism. We already know that such a mechanism cannot be incentive-compatible in dominant strategies, i.e. it cannot be a fixed-price mechanism. There are other techniques to design Bayesian incentive-compatible mechanisms. In particular, the second-best mechanism is in some sense known, but its implicit description does not seem very useful for the derivation of approximation guarantees. It appears that a more specific Bayesian mechanism would be needed to resolve this question.

References

- [1] Nick Arnosti, Marissa Beck, and Paul Milgrom. “Adverse Selection and Auction Design for Internet Display Advertising”. In: *American Economic Review* 106.10 (2016), pp. 2852–2866.

- [2] Moshe Babaioff et al. “The Best of Both Worlds: Asymptotically Efficient Mechanisms with a Guarantee on the Expected Gains-From-Trade”. In: 2018, pp. 373–373.
- [3] Liad Blumrosen and Shahar Dobzinski. “(Almost) Efficient Mechanisms for Bilateral Trading”. 2016. URL: <https://arxiv.org/abs/1604.04876>.
- [4] Liad Blumrosen and Shahar Dobzinski. “Reallocation Mechanisms”. In: *Proceedings of the Fifteenth ACM Conference on Economics and Computation (EC’14)*. New York, NY, USA: ACM, 2014, p. 617.
- [5] Liad Blumrosen and Yehonatan Mizrahi. “Approximating Gains-from-Trade in Bilateral Trading”. In: *Proceedings of the 12th International Conference on Web and Internet Economics. WINE 2016*. 2016, pp. 400–413.
- [6] Riccardo Colini-Baldeschi et al. “Approximately Efficient Double Auctions with Strong Budget Balance”. In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA ’16)*. Philadelphia, PA, USA: SIAM, 2016, pp. 1424–1443.
- [7] Moritz Drexl and Andreas Kleiner. “Optimal Private Good Allocation: The Case for a Balanced Budget”. In: *Games and Economic Behavior* 94 (2015), pp. 169–181.
- [8] Paul Dütting et al. “Efficient Two-Sided Markets with Limited Information”. arXiv:2003.07503, to appear in *STOC 2021*.
- [9] Allan Gibbard. “Manipulation of Voting Schemes: A General Result”. In: *Econometrica* 41.4 (1973), pp. 587–601.
- [10] Kathleen M Hagerty and William P Rogerson. “Robust trading mechanisms”. In: *Journal of Economic Theory* 42.1 (1987), pp. 94–107. ISSN: 0022-0531.
- [11] Zi Yang Kang and J. Vondrák. “Fixed-Price Approximations to Optimal Efficiency in Bilateral Trade”. 2019. URL: <https://ssrn.com/abstract=3460336>.
- [12] R. McAfee. “The Gains from Trade Under Fixed Price Mechanisms”. In: *Applied Economics Research Bulletin* (2007).
- [13] Roger B Myerson and Mark A Satterthwaite. “Efficient mechanisms for bilateral trading”. In: *Journal of Economic Theory* 29.2 (1983), pp. 265–281.
- [14] Colini-Baldeschi R. et al. “Fixed Price Approximability of the Optimal Gain from Trade”. In: *Web and Internet Economics* (2017).
- [15] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. “SBBA: A Strongly-Budget-Balanced Double-Auction Mechanism”. In: 2016, pp. 260–272.
- [16] Ran Shao and Lin Zhou. “Optimal Allocation of an Indivisible Good”. In: *Games and Economic Behavior* 100 (2016), pp. 95–112.

A Omitted Proofs

Lemma 14. For any fixed $0 \leq \mu_1 \leq \mu \leq 1$, $\mu > 0$ and $0 < \gamma \leq 1$,

$$\inf_{F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)} \sup_{p \in \mathbb{R}_+} \frac{W(p; F)}{\text{OPT-W}(F)} = \inf_{F \in \Delta_{L^1}^{(4)}(\mathbb{R}_+; \mu, \mu_1, \gamma)} \sup_{p \in \mathbb{R}_+} \frac{W(p; F)}{\text{OPT-W}(F)}.$$

Proof. Fix $F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$. We may assume without loss of generality that F is a finite discrete distribution, otherwise we may approximate F arbitrarily well by such a distribution in

$\Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$.⁸ By rescaling if necessary, we may also assume that $F(1) = 1$. Let F put positive probability mass only on each of the points $x_0, x_1, \dots, x_N \in [0, 1]$.

Motivated by the procedure outlined in the proof sketch, we now construct $\tilde{F} \in \Delta_{L^1}^{(4)}(\mathbb{R}_+; \mu, \mu_1, \gamma)$ such that, for any $F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$,

$$\frac{W(\mu; \tilde{F})}{\text{OPT-W}(\tilde{F})} \leq \frac{W(\mu; F)}{\text{OPT-W}(F)}.$$

We write

$$\tilde{F}(x) = q_0 \cdot \mathbf{1}_{x \geq 0} + q_1 \cdot \mathbf{1}_{x \geq \mu} + q_2 \cdot \mathbf{1}_{x \geq \mu + \delta} + q_3 \cdot \mathbf{1}_{x \geq 1}.$$

To construct \tilde{F} , we choose q_0, q_1, q_2, q_3 and δ as follows:

- δ is defined by

$$\delta = \min\{x_i - \mu : x_i > \mu, i \in \{1, 2, \dots, N\}\}.$$

- q_0 and q_1 are defined by

$$\begin{cases} q_0 + q_1 & = \gamma, \\ (q_0 + q_1)\mu_1 & = q_1\mu. \end{cases}$$

That is, $q_0 = \gamma(1 - \mu_1/\mu)$ and $q_1 = \mu_1\gamma/\mu$.

- If $\mu + \delta < 1$, q_2 and q_3 are defined by

$$\begin{cases} q_2 + q_3 & = 1 - \gamma, \\ q_2(\mu + \delta) + q_3 & = \mu - (q_0 + q_1)\mu_1. \end{cases}$$

That is, $q_2 = (1 - \gamma - \mu + \mu_1\gamma)/(1 - \mu - \delta)$ and $q_3 = (\mu\gamma - \mu_1\gamma - \delta + \delta\gamma)/(1 - \mu - \delta)$.

If $\mu + \delta = 1$, q_2 and q_3 are defined by $q_2 = 0$ and $q_3 = 1 - \gamma$.

With \tilde{F} defined above, we first verify that $\tilde{F} \in \Delta_{L^1}^{(4)}(\mathbb{R}_+; \mu, \mu_1, \gamma)$. Indeed:

$$\begin{cases} \tilde{F}(\mu) = q_0 + q_1 & = \gamma, \\ \mathbb{E}_{\tilde{F}}[S] = q_1\mu + q_2(\mu + \delta) + q_3 & = \mu, \\ \mathbb{E}_{\tilde{F}}[S | S \leq \mu] = \frac{q_1\mu}{q_0 + q_1} & = \mu_1. \end{cases}$$

Next, we show that, for any $F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$,

$$\text{OPT-W}(\tilde{F}) \geq \text{OPT-W}(F).$$

To do so, we compute that

$$\begin{aligned} \text{OPT-W}(F) &= \mathbb{E}_F[\max\{B, S\}] \\ &= \mathbb{E}_F[\max\{B, S\} \cdot \mathbf{1}_{B, S \leq \mu}] + \mathbb{E}_F[\max\{B, S\} \cdot \mathbf{1}_{B, S > \mu}] \\ &\quad + \mathbb{E}_F[\max\{B, S\} \cdot \mathbf{1}_{\max\{B, S\} > \mu} \mathbf{1}_{\min\{B, S\} \leq \mu}]. \end{aligned}$$

⁸For instance, one may approximate F by the finite discrete distributions corresponding to Riemann-sum approximations of $\mathbb{E}[S | S \leq \mu]$ (when restricted to $S \leq \mu$) and $\mathbb{E}[S | S > \mu]$ (when restricted to $S > \mu$).

That is, B, S must be either: (i) both no greater than μ ; (ii) both greater than μ ; or (iii) on different sides of μ . For the last case, observe that:

$$\mathbb{E}_F[\max\{B, S\} \cdot \mathbf{1}_{\max\{B, S\} > \mu} \mathbf{1}_{\min\{B, S\} \leq \mu}] = \mathbb{E}_{\tilde{F}}[\max\{B, S\} \cdot \mathbf{1}_{\max\{B, S\} > \mu} \mathbf{1}_{\min\{B, S\} \leq \mu}].$$

Therefore, it suffices to show that $\mathbb{E}_F[\max\{B, S\} \cdot \mathbf{1}_{B, S \leq \mu}] \leq \mathbb{E}_{\tilde{F}}[\max\{B, S\} \cdot \mathbf{1}_{B, S \leq \mu}]$ and that $\mathbb{E}_F[\max\{B, S\} \cdot \mathbf{1}_{B, S > \mu}] \leq \mathbb{E}_{\tilde{F}}[\max\{B, S\} \cdot \mathbf{1}_{B, S > \mu}]$.

To show the former result, observe that

$$\begin{cases} \mathbb{E}_F[\max\{B, S\} \cdot \mathbf{1}_{B, S \leq \mu}] = 2 \int_0^\mu x F(x) \, dF(x) & = \mu\gamma^2 - \int_0^\mu [F(x)]^2 \, dx, \\ \mathbb{E}_{\tilde{F}}[\max\{B, S\} \cdot \mathbf{1}_{B, S \leq \mu}] = (q_0 + q_1)^2 \mu - q_0^2 \mu & = \mu\gamma^2 - q_0^2 \mu. \end{cases}$$

However, from the definition of q_0 :

$$\int_0^\mu F(x) \, dx = \mu\gamma - \int_0^\mu x \, dF(x) = (\mu - \mu_1)\gamma = q_0\mu.$$

Therefore

$$q_0^2 \mu = \frac{1}{\mu} \left[\int_0^\mu F(x) \, dx \right]^2 \leq \int_0^\mu [F(x)]^2 \, dx.$$

This establishes that $\mathbb{E}_F[\max\{B, S\} \cdot \mathbf{1}_{B, S \leq \mu}] \leq \mathbb{E}_{\tilde{F}}[\max\{B, S\} \cdot \mathbf{1}_{B, S \leq \mu}]$. The latter result follows by a similar argument.

We have thus proved that $\text{OPT-W}(\tilde{F}) \geq \text{OPT-W}(F)$. Moreover, since $W(\mu; \tilde{F}) = W(\mu; F)$, we have

$$\frac{W(p; \tilde{F})}{\text{OPT-W}(\tilde{F})} \leq \frac{W(p; F)}{\text{OPT-W}(F)}.$$

Since $\tilde{F} \in \Delta_{L^1}^{(4)}(\mathbb{R}_+; \mu, \mu_1, \gamma) \subset \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)$, hence this establishes:

$$\inf_{F \in \Delta_{L^1}(\mathbb{R}_+; \mu, \mu_1, \gamma)} \sup_{p \in \mathbb{R}_+} \frac{W(p; F)}{\text{OPT-W}(F)} = \inf_{F \in \Delta_{L^1}^{(4)}(\mathbb{R}_+; \mu, \mu_1, \gamma)} \sup_{p \in \mathbb{R}_+} \frac{W(p; F)}{\text{OPT-W}(F)}.$$

□

Lemma 15. *The value of the minimax problem (A) is bounded below:*

$$\inf_{F \in \Delta_{L^1}(\mathbb{R}_+)} \sup_{p \in \mathbb{R}_+} \frac{W(p; F)}{\text{OPT-W}(F)} \geq \frac{2 + \sqrt{2}}{4}.$$

Proof. From the argument given in the paper, it suffices to show that

$$\inf_{\substack{\mu, \mu_1, \gamma \in \mathbb{R}_+ \\ 0 \leq \mu_1 \leq \mu; \mu, \gamma \in (0, 1)}} \frac{\mu + (\mu - \mu_1)\gamma}{\mu_1\gamma^2(2 - \frac{\mu_1}{\mu}) + 2\gamma(\mu - \gamma\mu_1) + (1 - \gamma)^2 - \frac{(1 - \gamma - \mu + \mu_1\gamma)^2}{1 - \mu}} = \frac{2 + \sqrt{2}}{4}.$$

We begin with a change of variables. Define $x := 1 - \mu_1/\mu \in [0, 1]$; the minimization problem thus becomes

$$\inf_{\substack{x \in [0, 1] \\ \mu, \gamma \in (0, 1)}} \frac{1 + \gamma x}{1 + 2\gamma x - \frac{\gamma^2 x^2}{1 - \mu}}.$$

Observe that, for $x \in [0, 1]$ and $\mu, \gamma \in (0, 1]$,

$$\frac{\partial}{\partial \mu} \left[\frac{1 + \gamma x}{1 + 2\gamma x - \frac{\gamma^2 x^2}{1-\mu}} \right] = \frac{\gamma^2 x^2 (1 + \gamma x)}{[1 - \mu + \gamma x(2 - 2\mu - \gamma x)]^2} \geq 0.$$

That is, the objective function is non-decreasing in μ .

Thus the infimum of the objective function is achieved in the limit as $\mu \rightarrow 0$:

$$\inf_{\substack{x \in [0,1] \\ \mu, \gamma \in (0,1]}} \frac{1 + \gamma x}{1 + 2\gamma x - \frac{\gamma^2 x^2}{1-\mu}} = \inf_{\substack{x \in [0,1] \\ \gamma \in (0,1]}} \frac{1 + \gamma x}{1 + \gamma x(2 - \gamma x)}.$$

We now make another change of variables. Define $y := \gamma x \in [0, 1]$. The minimization problem thus becomes a single-dimensional minimization problem:

$$\inf_{\substack{x \in [0,1] \\ \gamma \in (0,1]}} \frac{1 + \gamma x}{1 + \gamma x(2 - \gamma x)} = \inf_{y \in [0,1]} \frac{1 + y}{1 + y(2 - y)} = \frac{2 + \sqrt{2}}{4}.$$

□

Proposition 16 (Proposition 10 in the main text). *Given distributions F_S and F_B that satisfy $\mathbb{E}[(S - B)_+] = 0$, the designer can always select a price that achieves at least 3/4 of the total expected welfare under the first-best efficient outcome. That is:*

$$\inf_{\substack{F_S, F_B \in \Delta_{L^1}(\mathbb{R}_+) \\ \mathbb{E}[(S-B)_+] = 0}} \sup_{p \in \mathbb{R}_+} \frac{W(p; F_S, F_B)}{\text{OPT-W}(F_S, F_B)} = \frac{3}{4}.$$

Proof. The result to Proposition 10 follows from Lemma 17, which we prove below. □

Lemma 17. *Fix $\alpha > 0$. Given distributions F_S and F_B that satisfy $\mathbb{E}[(S - B)_+] \leq \alpha \cdot \text{OPT-W}(F_S, F_B)$. Suppose there exist p^+ and p^- such that:*

- (i) $\mathbb{P}[S > p^+] \leq \sqrt{10\alpha}$;
- (ii) $\mathbb{P}[B < p^-] \leq \sqrt{10\alpha}$; and
- (iii) $0 < p^+ - p^- \leq \sqrt{10\alpha}$.

Then the designer can set the price to be either p^+ or p^- to achieve an expected welfare of at least

$$\left(\frac{3}{4} - 2\sqrt{10\alpha} \right) \cdot \text{OPT-W}(F_S, F_B).$$

Remark 18. *We note that Lemma 17 implies the result of Proposition 10 as follows. Given distributions F_S and F_B that satisfy $\mathbb{E}[(S - B)_+] = 0$, there exists a price p^* for which $\mathbb{P}[S \leq p^*] = \mathbb{P}[B \geq p^*] = 1$. Fix $\alpha > 0$. Define*

$$\begin{cases} p^+ := p^* + \frac{1}{2}\sqrt{10\alpha} \cdot \text{OPT-W}(F_S, F_B), \\ p^- := p^* - \frac{1}{2}\sqrt{10\alpha} \cdot \text{OPT-W}(F_S, F_B). \end{cases}$$

Observe that $\mathbb{P}[S > p^+] = \mathbb{P}[B < p^-] = 0 < \sqrt{10\alpha}$ and $p^+ - p^- = \sqrt{10\alpha}$, so Lemma 19 applies for any $\alpha > 0$. The result of Proposition 10 is obtained in the limit as $\alpha \rightarrow 0$.

Proof. For ease of notation, given a fixed-price mechanism p , denote the welfare loss relative to first-best efficiency by $\text{LOSS}(p; F_S, F_B)$:

$$\text{LOSS}(p; F_S, F_B) := \text{OPT-W}(F_S, F_B) - W(p; F_S, F_B) = \mathbb{E}[(B - S) \cdot \mathbf{1}_{S < B \leq p}] + \mathbb{E}[(B - S) \cdot \mathbf{1}_{p < S < B}].$$

We begin by analyzing the welfare loss for p^+ and p^- . We have:

$$\begin{cases} \text{LOSS}(p^+; F_S, F_B) &= \mathbb{E}[(B - S) \cdot \mathbf{1}_{S < B \leq p^+}] + \mathbb{E}[(B - S) \cdot \mathbf{1}_{p^+ < S < B}] \\ &\leq \mathbb{E}[B \cdot \mathbf{1}_{B < p^+} \mathbf{1}_{S < p^-}] + \mathbb{E}[(B - S) \cdot \mathbf{1}_{p^- \leq S < B \leq p^+}] + \mathbb{E}[B \cdot \mathbf{1}_{p^+ < B} \mathbf{1}_{p^+ < S}], \\ \text{LOSS}(p^-; F_S, F_B) &= \mathbb{E}[(B - S) \cdot \mathbf{1}_{S < B \leq p^-}] + \mathbb{E}[(B - S) \cdot \mathbf{1}_{p^- < S < B}] \\ &\leq \mathbb{E}[B \cdot \mathbf{1}_{B < p^-} \mathbf{1}_{S < p^-}] + \mathbb{E}[(B - S) \cdot \mathbf{1}_{p^- \leq S < B \leq p^+}] + \mathbb{E}[B \cdot \mathbf{1}_{B \geq p^+} \mathbf{1}_{S \geq p^-}]. \end{cases}$$

Consider the terms $\mathbb{E}[B \cdot \mathbf{1}_{B < p^+} \mathbf{1}_{S < p^-}]$ and $\mathbb{E}[B \cdot \mathbf{1}_{B \geq p^+} \mathbf{1}_{S \geq p^-}]$, which could potentially be large (i.e., close to the value of $\text{OPT-W}(F_S, F_B)$). However, they cannot be both large. Indeed, let $\beta := \mathbb{E}[B \cdot \mathbf{1}_{B < p^+}] / \mathbb{E}[B]$ and $\sigma := \mathbb{P}[S < p^-]$. Because of the independence between B and S , we can write $\mathbb{E}[B \cdot \mathbf{1}_{B < p^+} \mathbf{1}_{S < p^-}] = \beta \sigma \cdot \mathbb{E}[B]$ and $\mathbb{E}[B \cdot \mathbf{1}_{B \geq p^+} \mathbf{1}_{S \geq p^-}] = (1 - \beta)(1 - \sigma) \cdot \mathbb{E}[B]$. We distinguish between two cases: $\beta + \sigma \leq 1$ or $\beta + \sigma > 1$. In the first case, setting a price p^+ yields

$$\mathbb{E}[B \cdot \mathbf{1}_{B < p^+} \mathbf{1}_{S < p^-}] \leq \beta \sigma \cdot \mathbb{E}[B] \leq \frac{1}{4} \cdot \mathbb{E}[B] \leq \frac{1}{4} \cdot \text{OPT-W}(F_S, F_B).$$

In the second case, setting a price p^- yields

$$\mathbb{E}[B \cdot \mathbf{1}_{B \geq p^+} \mathbf{1}_{S \geq p^-}] \leq (1 - \beta)(1 - \sigma) \cdot \mathbb{E}[B] \leq \frac{1}{4} \cdot \mathbb{E}[B] \leq \frac{1}{4} \cdot \text{OPT-W}(F_S, F_B).$$

The remaining terms can be bounded by $\mathcal{O}(\sqrt{\alpha}) \cdot \text{OPT-W}(F_S, F_B)$ as follows:

$$\begin{cases} \mathbb{E}[(B - S) \cdot \mathbf{1}_{p^- \leq S < B \leq p^+}] \leq p^+ - p^- &\leq \sqrt{10\alpha} \cdot \text{OPT-W}(F_S, F_B), \\ \mathbb{E}[B \cdot \mathbf{1}_{p^+ < B} \mathbf{1}_{p^+ < S}] \leq \mathbb{E}[B] \cdot \mathbb{P}[p^+ < S] &\leq \sqrt{10\alpha} \cdot \text{OPT-W}(F_S, F_B), \\ \mathbb{E}[B \cdot \mathbf{1}_{B < p^-} \mathbf{1}_{S < p^-}] \leq \mathbb{E}[B] \cdot \mathbb{P}[B < p^-] &\leq \sqrt{10\alpha} \cdot \text{OPT-W}(F_S, F_B). \end{cases}$$

Therefore, in either case, the welfare loss is bounded above by $(\frac{1}{4} + 2\sqrt{10\alpha}) \cdot \text{OPT-W}(F_S, F_B)$. Thus

$$\max_{p \in \{p^+, p^-\}} W(p; F_S, F_B) \geq \left(\frac{3}{4} - 2\sqrt{10\alpha} \right) \cdot \text{OPT-W}(F_S, F_B).$$

□

We now derive a sufficient condition for the hypotheses of Lemma 17 to hold:

Lemma 19. *Fix $\alpha > 0$. Given distributions F_S and F_B that satisfy $\mathbb{E}[(S - B)_+] \leq \alpha \cdot \text{OPT-W}(F_S, F_B)$, suppose there exists p^* such that*

$$\mathbb{P}[S \geq p^*] > \frac{1}{5} \quad \text{and} \quad \mathbb{P}[B \leq p^*] > \frac{1}{5}.$$

Define

$$\begin{cases} p^+ := p^* + \frac{1}{2} \sqrt{10\alpha} \cdot \text{OPT-W}(F_S, F_B), \\ p^- := p^* - \frac{1}{2} \sqrt{10\alpha} \cdot \text{OPT-W}(F_S, F_B). \end{cases}$$

Then $\mathbb{P}[S \geq p^+] \leq \sqrt{10\alpha}$ and $\mathbb{P}[B \leq p^-] \leq \sqrt{10\alpha}$.

Proof. Suppose to the contrary that $\mathbb{P}[S \geq p^+] > \sqrt{10\alpha}$. Since $\mathbb{P}[B \leq p^*] > 1/5$ by definition of p^* , we have

$$\mathbb{E}[(S - B)_+] \geq (p^+ - p^*) \cdot \mathbb{P}[S \geq p^+] \cdot \mathbb{P}[B \leq p^*] > \alpha \cdot \text{OPT-W}(F_S, F_B), \quad \text{a contradiction.}$$

Similarly, if $\mathbb{P}[B \leq p^-] > \sqrt{10\alpha}$, then since $\mathbb{P}[S \geq p^*] > 1/5$ by definition of p^* , we have

$$\mathbb{E}[(S - B)_+] \geq (p^* - p^-) \cdot \mathbb{P}[S \geq p^*] \cdot \mathbb{P}[B \leq p^-] > \alpha \cdot \text{OPT-W}(F_S, F_B), \quad \text{a contradiction.}$$

□

Therefore, for small $\alpha > 0$, Lemma 19 guarantees a strict improvement over the $1 - 1/e$ bound (i.e., that Lemma 17 applies) if there exists p^* such that

$$\mathbb{P}[S \geq p^*] > \frac{1}{5} \quad \text{and} \quad \mathbb{P}[B \leq p^*] > \frac{1}{5}.$$

What if such a p^* does not exist? The following result ensures that we can still guarantee a strict improvement over the $1 - 1/e$ bound:

Lemma 20. *Given distributions F_S and F_B , let $p^* := \inf\{p : \mathbb{P}[S \geq p] \leq 1/5\}$. If $\mathbb{P}[B \leq p^*] \leq 1/5$, then*

$$\text{W}(p^*; F_S, F_B) \geq \frac{17}{25} \cdot \text{OPT-W}(F_S, F_B).$$

Proof. Define $q := \mathbb{P}[S \geq p^*] \leq 1/5$ and $r := \mathbb{E}[B \cdot \mathbf{1}_{B \leq p^*}] / \mathbb{E}[B]$. Observe that

$$r \leq \mathbb{P}[B \leq p^*] \leq \frac{1}{5}.$$

As above, denote the welfare loss relative to first-best efficiency by $\text{LOSS}(p; F_S, F_B)$:

$$\text{LOSS}(p; F_S, F_B) := \text{OPT-W}(F_S, F_B) - \text{W}(p; F_S, F_B) = \mathbb{E}[(B - S) \cdot \mathbf{1}_{S < B \leq p}] + \mathbb{E}[(B - S) \cdot \mathbf{1}_{p < S < B}].$$

We bound $\text{LOSS}(p^*; F_S, F_B)$ from above as follows:

$$\begin{aligned} \text{LOSS}(p^*; F_S, F_B) &\leq \mathbb{E}[B \cdot \mathbf{1}_{B \leq p^*} \mathbf{1}_{S < p^*}] + \mathbb{E}[B \cdot \mathbf{1}_{B > p^*} \mathbf{1}_{S \geq p^*}] \\ &= r(1 - q) \cdot \mathbb{E}[B] + (1 - r)q \cdot \mathbb{E}[B] \\ &= \left[\frac{1}{2} - 2 \left(\frac{1}{2} - q \right) \left(\frac{1}{2} - r \right) \right] \mathbb{E}[B] \\ &\leq \left[\frac{1}{2} - 2 \left(\frac{1}{2} - \frac{1}{5} \right)^2 \right] \mathbb{E}[B] = \frac{8}{25} \cdot \mathbb{E}[B] \leq \frac{8}{25} \cdot \text{OPT-W}(F_S, F_B). \end{aligned}$$

Therefore

$$\text{W}(p^*; F_S, F_B) = \text{OPT-W}(F_S, F_B) - \text{LOSS}(p^*; F_S, F_B) \geq \frac{17}{25} \cdot \text{OPT-W}(F_S, F_B).$$

□

Proposition 21 (Proposition 11 in the main text; Theorem 4.1 in [3]). *For any given distributions F_S and F_B ,*

$$\sup_{p \in \mathbb{R}_+} \frac{\text{W}(p; F_S, F_B)}{\text{OPT-W}(F_S, F_B)} \geq 1 - \frac{1}{e} + \frac{1}{e} \cdot \mathbb{E}[(S - B)_+].$$

Proof. The proof is based almost entirely on the proof of Theorem 4.1 in [3]. Since the mechanism depends only on the seller's distribution, we can work with a fixed buyer's value and then take an expectation over the buyer at the end. Given F_S , we fix $b \in \mathbb{R}_+$ and consider the truncated seller's distribution \tilde{F}_S (replacing all values above b by b):

$$\tilde{F}_S(x) = F_S(x) \cdot \mathbf{1}_{x < b} + \mathbf{1}_{x \geq b}.$$

We denote by Φ_b the step function $\Phi_b(x) = \mathbf{1}_{x \geq b}$ (corresponding to a deterministic value of b). For any distribution G that depends only on F_S , we note that

$$\mathbb{E}_{p \sim G}[\mathbf{W}(p; \tilde{F}_S, \Phi_b)] = \mathbb{E}_{p \sim G}[\mathbf{W}(p; F_S, \Phi_b)] - \mathbb{P}[S > b] \cdot (\mathbb{E}_{S \sim F_S}[S | S > b] - b),$$

because the last term is exactly the expected value that is lost by modifying the seller's distribution to \tilde{F}_S . (Note that the trade never happens when $S > b$, so the outcome in this case is always S .) By the same argument,

$$\text{OPT-W}(\tilde{F}_S, \Phi_b) = \text{OPT-W}(F_S, \Phi_b) - \mathbb{P}[S > b] \cdot (\mathbb{E}_{S \sim F_S}[S | S > b] - b).$$

Consider the distribution $G^*(x) = 1 + \log F_S(x)$, where $x \in [F_S^{-1}(1/e), F_S^{-1}(1)]$. [3] show that

$$\mathbb{E}_{p \sim G^*}[\mathbf{W}(p; \tilde{F}_S, \Phi_b)] \geq \left(1 - \frac{1}{e}\right) \cdot \text{OPT-W}(\tilde{F}_S, \Phi_b).$$

Consequently, substitution reveals that

$$\begin{aligned} \mathbb{E}_{p \sim G^*}[\mathbf{W}(p; F_S, \Phi_b)] &\geq \left(1 - \frac{1}{e}\right) \cdot \text{OPT-W}(F_S, \Phi_b) + \frac{1}{e} \cdot \mathbb{P}[S > b] \cdot (\mathbb{E}_{S \sim F_S}[S | S \geq b] - b) \\ &= \left(1 - \frac{1}{e}\right) \cdot \text{OPT-W}(F_S, \Phi_b) + \frac{1}{e} \cdot \mathbb{E}[(S - b)_+]. \end{aligned}$$

Because G^* depends only on F_S , it follows by linearity of expectation that

$$\mathbb{E}_{p \sim G^*}[\mathbf{W}(p; F_S, F_B)] = \mathbb{E}_{b \sim F_B} \mathbb{E}_{p \sim G^*}[\mathbf{W}(p; F_S, \Phi_b)]$$

and

$$\text{OPT-W}(F_S, F_B) = \mathbb{E}_{b \sim F_B}[\text{OPT-W}(F_S, \Phi_b)].$$

Taking expectations in the above yields

$$\begin{aligned} \mathbb{E}_{p \sim G^*}[\mathbf{W}(p; F_S, F_B)] &\geq \left(1 - \frac{1}{e}\right) \cdot \text{OPT-W}(F_S, F_B) + \frac{1}{e} \cdot \mathbb{E}[(S - B)_+] \\ &= \left(1 - \frac{1 - \alpha}{e}\right) \cdot \text{OPT-W}(F_S, F_B). \end{aligned}$$

□

Theorem 22 (Theorem 12 in the main text). *Given distributions F_S and F_B , the designer can always select a price that achieves at least $1 - 1/e + 0.0001$ of the optimal expected welfare. That is:*

$$\inf_{F_S, F_B \in \Delta_{L^1}(\mathbb{R}_+)} \sup_{p \in \mathbb{R}_+} \frac{\mathbf{W}(p; F_S, F_B)}{\text{OPT-W}(F_S, F_B)} \geq 1 - \frac{1}{e} + 0.0001.$$

Proof. Let $\mathbb{E}[(S - B)_+] = \alpha \cdot \text{OPT-W}$. If $\alpha \geq 0.0003$, then Proposition 11 yields

$$\sup_{p \in \mathbb{R}_+} \frac{W(p; F_S, F_B)}{\text{OPT-W}(F_S, F_B)} \geq 1 - \frac{1}{e} + \frac{\alpha}{e} > 1 - \frac{1}{e} + 0.0001.$$

If $0 < \alpha < 0.0003$, then consider $p^* := \inf\{p : \mathbb{P}[S \geq p] \leq 1/5\}$. If $\mathbb{P}[B \leq p^*] \leq 1/5$, then Lemma 20 yields

$$\sup_{p \in \mathbb{R}_+} \frac{W(p; F_S, F_B)}{\text{OPT-W}(F_S, F_B)} \geq \frac{W(p^*; F_S, F_B)}{\text{OPT-W}(F_S, F_B)} \geq \frac{17}{25}.$$

Otherwise, if $\mathbb{P}[B \leq p^*] > 1/5$, then there exists some sufficiently small $\varepsilon > 0$ such that $p_0 = p^* - \varepsilon$ satisfies $\mathbb{P}[S \geq p_0] > 1/5$ and $\mathbb{P}[B \leq p_0] > 1/5$. Lemma 19 shows that this is a sufficient condition to satisfy the hypotheses of Lemma 17, which implies:

$$\sup_{p \in \mathbb{R}_+} \frac{W(p; F_S, F_B)}{\text{OPT-W}(F_S, F_B)} \geq \frac{3}{4} - 2\sqrt{10\alpha} > \frac{16}{25}.$$

Finally, Proposition 10 covers the case of $\alpha = 0$. □

Theorem 23 (Theorem 13 in the main text). *Given distributions F_S and F_B , for any (possibly non-deterministic) mechanism G (where we let G be a distribution in $[0, 1]$ and we set the price $F_S^{-1}(G)$) that uses only quantile distributional knowledge of the seller's distribution, the designer can achieve no better than $1 - 1/e$ of the total expected welfare under the first-best efficient outcome. That is:*

$$\inf_{F_S, F_B \in \Delta_{L^1}(\mathbb{R}_+)} \sup_{G \in \Delta([0, 1])} \frac{W(G; F_S, F_B)}{\text{OPT-W}(F_S, F_B)} = 1 - \frac{1}{e},$$

where $W(G; F_S, F_B)$ denotes the expected welfare that the mechanism G achieves.

Before we present a formal proof, let us discuss a game-theoretic intuition behind this result. We can view the situation as a game between two players, the designer and nature. The designer tries to select a parameter x to maximize efficiency, and nature tries to select a distribution F_S to thwart the designer's goal. Given the choice of x and F_S , the buyer's distribution F_B is considered to be worst possible with respect to the designer's outcome. Our goal is to prove that there is a strategy of nature such that no strategy of the designer achieves an approximation factor better than $1 - 1/e$.

Due to von Neumann's theorem, there are optimal *mixed strategies* Ξ and Φ . It is important to keep in mind that these are randomized strategies: in the case of the designer, this means a random choice of x ; in the case of nature, this means a random choice of F_S , (i.e., a probability distribution over cumulative distribution functions F_S , which is a more complicated object).

In order to simplify the game, let us make a few observations. Given x , F_S and F_B , the expected outcome is given by taking an expectation over the buyer's value b sampled from F_B (because there is no dependency between b and the choice of the price p and the seller's value s). Therefore, we might as well assume that the buyer's value is deterministic, namely, the worst possible value b , given x and F_S . Furthermore, for each choice of F_S and b , the values can be rescaled so that $b = 1$, without affecting the approximation ratio (i.e., the ratio of welfare relative to first-best efficiency). So we can assume without loss of generality that $b = 1$.

Further, given that $b = 1$, the seller's distribution can be truncated at 1: any probability mass above 1 can be moved to 1.⁹ This means that the first-best efficiency has value $\mathbb{E}[\max\{b, s\}] = 1$.

⁹There is a lemma making this argument in [3] but since we are proving the opposite bound, this lemma is not formally needed here.

Next, let us consider the strategy of nature. For a probability distribution F_S , if there is some mass between $(0, 1)$, it only decreases the approximation ratio if we push this probability mass towards 0 (the outcome possibly decreases, and the optimum is still 1). It is important here that the probability mass is not concentrated on a single point – relative comparisons between different possible values should still be non-trivial. However, we can assume for example that the probability mass below 1 is uniform between $[0, \varepsilon]$, with density y/ε . Considering this, the only important parameter that governs the seller’s distribution is the probability of s between equal to 1, $y = \mathbb{P}[s < 1]$.

Hence, the game we are considering has pure strategies x for the designer and y for nature. Randomized strategies are distributions over x and y . Given x, y the payoff function for the mechanism (ignoring terms proportional to ε) is

$$V(x, y) = (1 - y) + x \cdot \mathbf{1}_{x < y}.$$

This reflects the fact that with probability $1 - y$, the seller’s value is 1, in which case the outcome is certainly 1 (since the buyer’s value is also 1). Otherwise, the seller’s value is close to 0; then the trade occurs exactly when $s < p$ and $p < 1$, and the outcome in that case is 1; otherwise close to 0. The event $x < y$ is equivalent to the fact that $p < 1$, because $y = F_S(1)$. Given that $x < y$, the probability that $s < p$ is exactly x , because $p = F_S^{-1}(x)$. Therefore, $x \cdot \mathbf{1}_{x < y}$ is the contribution to the expected outcome in case the seller’s value is below 1.

We now derive the optimal mixed strategies. Let us assume that nature’s strategy is given by a probability density function $\rho(y)$. Then for a given (pure) designer’s strategy x , nature’s expected payoff is

$$\mathbb{E}[V(x, y)] = \mathbb{E}[(1 - y) + x \cdot \mathbf{1}_{x < y}] = \int_0^x (1 - y)\rho(y) \, dy + \int_x^1 (1 - y + x)\rho(y) \, dy. \quad (\dagger)$$

We posit that for an optimal nature strategy $\rho(y)$, this quantity should be the same for every x in the support of the optimal mechanism strategy. If not, then the designer’s strategy could be modified to achieve a better outcome, by picking the x maximizing the quantity above. We are trying to prove that the designer’s strategy is defined by $g(x) = 1/x$ for $x \in [1/e, 1]$. Hence, let us assume that the quantity in (\dagger) is constant for $x \in [1/e, 1]$. By differentiating (\dagger) with respect to x , we obtain (for $x \in [1/e, 1]$),

$$\int_x^1 \rho(y) \, dy - x \rho(x) = 0.$$

Note that for $x = 1 - \varepsilon$, we obtain

$$\int_{1-\varepsilon}^1 \rho(y) \, dy = (1 - \varepsilon) \rho(1 - \varepsilon).$$

This is a somewhat paradoxical conclusion. What this actually means is that the probability distribution cannot be fully defined by a density function; there must be a discrete probability mass at $x = 1$, which is equal to the density just below 1.¹⁰ Differentiating one more time, we get

$$-2\rho(x) - x \rho'(x) = 0.$$

This differential equation is easy to solve: the solution is $\rho(y) = C/y^2$ for $y \in (1/e, 1)$. There should also be a discrete probability mass at $y = 1$ equal to C . The normalization condition implies that $C = 1/e$. To complete the proof, we just have to show that there is no strategy of the designer that achieves a factor better than $1 - 1/e$ against this nature’s strategy.

¹⁰These arguments are not rigorous, but in any case we are just trying to guess the optimal form of nature’s strategy.

Proof. Motivated by the discussion above, we consider the following strategy of nature:

- With probability $1/e$, set $y = 1$.
- With probability $1 - 1/e$, sample $y \in [1/e, 1]$ with density $\rho(y) = 1/(ey^2)$.

Given y , nature's value s is distributed as follows:

- With probability y , $s \in [0, \varepsilon]$ uniformly at random.
- With probability $1 - y$, $s = 1$.

We claim that for any strategy of the designer, the approximation ratio is at most $1 - 1/e$ against this nature's strategy. Since mixed strategies are convex combinations of pure strategies, it is enough to consider pure strategies $x \in [0, 1]$.

As we argued above, given x and y , the expected value of the game, up to $\mathcal{O}(\varepsilon)$ terms, is

$$V(x, y) = (1 - y) + x \cdot \mathbf{1}_{x < y}.$$

We have the following cases:

- $x \in [0, 1/e]$. In this range, we certainly have $x < y$ (because y is always at least $1/e$). Thus, $V(x, y) = 1 - y + x$. In expectation over y , this quantity is

$$\begin{aligned} \mathbb{E}[V(x, y)] &= \mathbb{E}[1 - y + x] = \int_{1/e}^1 (1 - y + x) \cdot \frac{1}{ey^2} dy + \frac{1}{e} \cdot x \\ &= \left[-\frac{1+x}{ey} - \frac{1}{e} \ln y \right]_{y=1/e}^1 + \frac{1}{e} \cdot x = 1 - \frac{2}{e} + x. \end{aligned}$$

Since $x \leq 1/e$, we have $\mathbb{E}[V(x, y)] \leq 1 - 1/e$.

- $x \in [1/e, 1)$. In this range, we have $y \leq x$ or $y > x$ depending of the value of y . In the first case, the value of the game is $1 - y$ and in the second case it is $1 - y + x$. Thus we compute the expected value as follows:

$$\begin{aligned} \mathbb{E}[V(x, y)] &= \int_{1/e}^1 (1 - y) \cdot \frac{1}{ey^2} dy + \int_x^1 x \cdot \frac{1}{ey^2} dy + \frac{1}{e} \cdot x \\ &= \left[-\frac{1}{ey} - \frac{1}{e} \ln y \right]_{y=1/e}^1 + \left[-\frac{x}{ey} \right]_{y=x}^1 + \frac{1}{e} \cdot x \\ &= \left(1 - \frac{2}{e} \right) + \left(\frac{1}{e} - \frac{x}{e} \right) + \frac{x}{e} = 1 - \frac{1}{e}. \end{aligned}$$

- $x = 1$. Then we get the same expressions as above, except for the term $\frac{1}{e} \cdot x$, since the case of $y = 1$ does not contribute anything. Therefore, again the value is at most $1 - 1/e$.

We conclude that there no strategy of the designer that achieves an expected welfare of more than $1 - 1/e$ (neglecting $\mathcal{O}(\varepsilon)$ terms). The first-best efficiency is 1 (since the buyer's value is always 1) and hence the approximation factor cannot be better than $1 - 1/e$. \square