Fixed-Price Approximations to Optimal Efficiency in Bilateral Trade[*]

Zi Yang Kang[†] and Jan Vondrák[‡]

This version: September 2019

Abstract

This paper studies fixed-price mechanisms in bilateral trade with *ex ante* symmetric agents. We show that the optimal price is particularly simple: it is exactly equal to the mean of the agents' distribution. The optimal price guarantees a worst-case performance of at least 1/2 of the first-best gains from trade, regardless of the agents' distribution. We also show that the worst-case performance improves as the number of agents increases, and is robust to various extensions. Our results offer an explanation for the widespread use of fixed-price mechanisms for size discovery, such as in workup mechanisms and dark pools.

JEL classification: D47, D82

Keywords: fixed-price mechanisms, bilateral trade, robustness, Myerson–Sattherthwaite

1 Introduction

In this paper, we study the performance of fixed-price mechanisms in the canonical context of bilateral trade *à la* Myerson and Satterthwaite (1983), in which a seller and a buyer bargain over a single indivisible good. Agents are *ex ante* symmetric and have independent private values for the good. In their seminal work, Myerson and Satterthwaite (1983) show that the first-best outcome cannot generally be achieved. They characterize the second-best Bayesian mechanism, which

* We are very grateful to Alex Bloedel, Jeremy Bulow, Gabriel Carroll, Shahar Dobzinski, Darrell Duffie, Fuhito Kojima, Paul Milgrom, Mike Ostrovsky, Ilya Segal, Andy Skrzypacz, Takuo Sugaya, and Bob Wilson for helpful comments and suggestions.

† Stanford Graduate School of Business, 655 Knight Way, Stanford, CA 94305; zykang@stanford.edu.

‡ Stanford Department of Mathematics, 450 Serra Mall, Stanford, CA 94305; jvondrak@stanford.edu.

is often strategically complicated for agents. By contrast, fixed-price mechanisms are strategy-proof. Yet other questions arise: How close to the efficient outcome can fixed-price mechanisms achieve? What information about the agents is required to set the optimal price? As espoused by the Wilson (1987) doctrine, designing mechanisms that perform well in reality entails a tradeoff between efficiency, strategic simplicity and informational simplicity.

This paper derives two key features of fixed-price mechanisms that are important from the perspective of a market designer. First, the optimal fixed-price mechanism is particularly simple: it is exactly equal to the mean of the agents' distribution of values. Second, the optimal price guarantees a worst-case performance of at least $1/2$ of the first-best gains from trade, regardless of the agents' distribution. These results are robust to a number of extensions and variations of the bilateral trade model.

As an example, suppose that seller and buyer have values drawn i.i.d. from the uniform distribution on $[0, 1]$. The first-best gains from trade are $1/6$, since it is the expected difference between two i.i.d. draws, whenever the buyer is higher. By contrast, with a price of $1/2$, trade occurs precisely when the seller's value is below $1/2$ and the buyer's value is above $1/2$. Thus the gains from trade realized in expectation are $0.25 \times (3/4 - 1/4) = 1/8$, which constitutes $3/4$ ($= 75\%$) of the first-best gains from trade.

For the sake of comparison, the second-best Bayesian mechanism dictates trade whenever the buyer's value exceeds the seller's value by at least $1/4$. Such a mechanism is no longer strategically simple. Unlike fixed-price mechanisms, it is no longer a dominant strategy for agents to report their values truthfully; moreover, each agent is required to have the correct beliefs (and higher-order beliefs) about the other agent. The design of the second-best Bayesian mechanism also uses more distributional knowledge than the mean. The mechanism realizes in expectation gains from trade of $9/64$. Thus, despite the additional strategic and informational complexity of the second-best Bayesian mechanism, it achieves a welfare improvement of only $3/32$ ($= 9.375\%$) of the first-best gains from trade, over setting a price of $1/2$.

Our results show that this example is not mere coincidence. Fixed-price mechanisms continue to perform well even when agents' values are not uniformly distributed. We also show that the worst-case performance of fixed-price mechanisms improves as the market size grows. Moreover, the optimal price converges to the competitive price, which is equal to the median of the agents' distribution given the assumption that agents are *ex ante* symmetric.

Our focus on fixed-price mechanisms is motivated by the increasing use of such mechanisms in financial markets, namely size discovery mechanisms such as workup mechanisms and dark pools. Size discovery mechanisms derive prices from the lit exchange order flow, but differ in the way they do so. Fleming and Nguyen (2018) find that workup trades accounts for about 60% of trading volume of on-the-run notes in the U.S. Treasury market, while Duffie and Zhu (2017) report that dark pools account for about 15% of trading volume in the U.S. equity markets.[1] Given the size of these markets, designing fixed-price mechanisms that perform well is clearly important.

Our paper makes three different contributions. First, our results offer the explanation that the simplicity of designing the optimal price and robust worst-case efficiency guarantee could underlie the widespread use of fixed-price mechanisms in practice. Second, we give normative advice and intuition for how the optimal price changes with market size: roughly, the mean price is approximately optimal in thin markets, whereas the median (competitive) price is approximately in thick markets.

Finally, we extend the tools that can be used to analyze problems in robust mechanism design. Our methods are fairly general and can accommodate various extensions to the model. These include assuming more knowledge about the dispersion in agents' values, estimating the maximum welfare loss under a fixed-price mechanism, allowing for other mechanisms in addition to fixed-price mechanisms, and allowing agents to be asymmetric.

Our work is closest to the literature on robust mechanism design, beginning with the work of Bergemann and Morris (2005) and Chung and Ely (2007). Optimal mechanisms and worst-case performances have been characterized in different settings, such as moral hazard by Chassang (2013) and Carroll (2015), and price discrimination by Bergemann and Schlag (2011), Bergemann et al. (2015) and Carroll and Segal (2018). A recent strand of this literature has studied robustness in situations where agents' distributions are not known, including work by Wolitzky (2016), Suzdaltsev (2018) and Carrasco et al. (2018). More broadly, similar problems have been studied in the computer science literature, such as by Blumrosen and Dobzinski (2014), Blumrosen and Dobzinski (2016), Blumrosen and Mizrahi (2016), Colini-Baldeschi et al. (2016) and Brustle et al. (2017). We give a more detailed discussion about how our paper is related to these papers in Section 6.

Our focus on fixed-price mechanisms also connects this paper with the broader literatures on mechanism design, implementation and strategy-proofness. Since the seminal paper of Myerson and Satterthwaite (1983), Hagerty and Rogerson (1987) have shown that dominant-strategy mechanisms for bilateral trade that are budget-balanced *ex post* must essentially be fixed-price mechanisms. Drexl and Kleiner (2015) and Shao and Zhou (2016) show that, for certain classes of distributions, fixed-price mechanisms, together with "option mechanisms," are optimal even

[1] As Fleming and Nguyen (2018) note, the daily trading volume in the U.S. Treasury market exceeds $100 billion, while that of the U.S. equity markets is about $200 billion.

when the *ex post* budget-balanced condition is relaxed to a no-deficit condition. Our methods can be extended to show that these option mechanisms also perform well, which we consider in Section 5.

The rest of this paper is organized as follows. In Section 2, we describe our model. Section 3 presents our main results. In Section 4, we analyze a natural variation of our model in large markets, and show that the worst-case performance of the optimal price improves as the market size grows. Section 5 shows how our results persist under various extensions; Section 6 concludes with a discussion on how our results can be interpreted to inform market design.

2 Model

There are two components to our model: (i) a bilateral trade game between a seller and a buyer; and (ii) a mechanism design problem of choosing a fixed price for the bilateral trade game.

2.1 Bilateral trade

We adopt the canonical bilateral trade setting of Myerson and Satterthwaite (1983). A seller owns a single indivisible good which a buyer wants to buy. The seller and the buyer are privately informed of their respective values, S and B, for the good; S and B are independent, nonnegative random variables with cumulative distribution function F. For ease of exposition, we assume that F is supported on $[0, 1]$. We denote by $\Delta([0, 1])$ the set of all such distributions; in particular, we allow F to be a discrete probability distribution. Our results extend straightforwardly to all distributions with finite and nonzero mean.

Both agents are risk-neutral and have payoffs that are quasilinear in money. If trade occurs at a price of p, the seller receives a payoff of p while the buyer receives $B - p$. If no trade occurs, the seller receives S while the buyer receives zero.

Our assumption that seller and buyer have *ex ante* symmetric values differs from the setting of Myerson and Satterthwaite (1983) in which agents have different distributions. This assumption is motivated by our application to size discovery mechanisms in financial markets such as dark pools, where the same market participants may either buy or sell depending on their inventories. In this setting, we interpret the seller and buyer as liquidity traders, whose inventory positions are independent of their values of the traded asset.[2] The assumption that agents are *ex ante*

[2] As Zhu (2014) notes: "exchanges are more attractive to informed traders, and dark pools are more attractive to uninformed traders," due to the fact that execution is not guaranteed in dark pools.

symmetric is widely used in the literature on size discovery mechanisms, including by Degryse et al. (2009), Zhu (2014) and Duffie and Zhu (2017).

2.2 Mechanism design

Given the bilateral trade game between the seller and the buyer, we consider the problem faced by a benevolent mechanism designer ("Designer"). The Designer chooses a fixed price $p \in \mathbb{R}_+$, which is allowed to depend on the distribution F. We write p instead of $p(F)$ when confusion is unlikely to result, with the understanding that p is a function of F. Under the fixed-price mechanism p, trade occurs if and only if $B > p \geq S$.[3]

The Designer maximizes the expected gains from trade, $\Gamma(p; F)$, in the bilateral trade game, which can be expressed as

$$\Gamma(p; F) := \mathbb{E}[(B - S) \cdot \mathbb{1}_{B>p\geq S}].$$

3 Optimal Fixed-Price Mechanism

3.1 Main results

Our first result characterizes the optimal fixed-price mechanism:

Proposition 1. The optimal price p^* is the mean of the distribution F:

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

Proof. For any price $p > 0$, the expected gains from trade realized by the fixed-price mechanism p are

$$\Gamma(p; F) = (\mathbb{E}[S] - p) \cdot F(p) + \int_0^p F(x) \, \mathrm{d}x.$$

Letting $p^* := \mathbb{E}[S]$,

$$\Gamma(p^*; F) - \Gamma(p; F) = \int_p^{p^*} [F(x) - F(p)] \, \mathrm{d}x.$$

This expression is nonnegative because F is nondecreasing; hence p^* is optimal. □

[3] Our results on performance do not depend on how ties are resolved here. We adopt this formulation because it simplifies notation: $F(p) = \mathbb{P}[S \leq p] = 1 - \mathbb{P}[B > p]$, so we do not have to worry about limits as either S or B approach p. Changing how ties are resolved will change our characterization of the optimal fixed-price mechanism; but this becomes irrelevant under weak regularity assumptions on F, such as if F is atomless.

Even though the Designer's chosen fixed-price mechanism may depend arbitrarily on F, Proposition 1 shows that the mean is a sufficient statistic for determining the optimal price. Interestingly, dark pools often peg prices to the midpoint of the exchange bid and offer, which can be viewed as an estimator of the mean of the agents' values.[4]

Given the simplicity of both the statement and proof of Proposition 1, we view it as rather surprising that it has not appeared at this level of generality much earlier in the literature, at least to our knowledge. Shao and Zhou's (2016) Theorem 1 documents the optimality of the mean price in a closely related setting; however, their setting assumes that the agents' distribution satisfies various other regularity conditions,[5] and the derivation of their result involves taking a first-order condition. By contrast, we emphasize that Proposition 1 does not depend on regularity properties of F, and so holds for discrete probability distributions in particular.[6]

If regularity conditions on F are imposed, then we can show the stronger result of uniqueness:

Corollary 1. Suppose that F is differentiable in a neighborhood of its mean, so that its density 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.

Indeed, under the assumptions of Corollary 1, F is strictly increasing in a neighborhood of its mean; hence $\Gamma(p^*; F) - \Gamma(p; F) > 0$ for any $p \neq p^*$. Corollary 1 applies under the usual assumption that F is continuously differentiable with positive density.

We now analyze the worst-case performance of the optimal fixed-price mechanism. *A priori*, we expect that fixed-price mechanisms will incur welfare loss due to the impossibility result of Myerson and Satterthwaite (1983). Our main result of this section shows that, while fixed-price mechanisms may lead to inefficiency, the welfare loss cannot be too large:

[4] Gârleanu and Pedersen's (2004) Proposition 3, for instance, characterizes conditions under which the midpoint of the bid and offer prices is an unbiased estimator of the mean of the agents' values.

[5] Specifically, they assume that the distribution is continuously differentiable and satisfies an increasing hazard rate condition as well as a decreasing reverse hazard rate condition. In Section 5 we will consider their setting and show how our main result, Theorem 1, extends.

[6] This is relevant because, as we show in the proof outline of Theorem 1, the worst-case distributions for the performance of the optimal price are discrete distributions.

Theorem 1. For any distribution F with mean μ, denote the first-best gains from trade by $\Gamma^{\mathrm{FB}}(F)$, defined by $\Gamma^{\mathrm{FB}}(F) := \mathbb{E}[(B-S)\cdot \mathbb{1}_{B>S}$. Then the gains from trade that the optimal fixed-price mechanism achieves are bounded below:

$$\max_{p\in\mathbb{R}_+}\Gamma(p;F)=\Gamma(\mu;F)\geq\left[1-\sqrt{1-\frac{\Gamma^{\mathrm{FB}}(F)}{\mu\,(1-\mu)}}\right]\mu\,(1-\mu)\quad\text{for any } F\in\Delta([0,1]) \text{ with mean } \mu.$$

A simple interpretation of Theorem 1 can be obtained by performing a Taylor expansion on the lower bound:

$$\left[1-\sqrt{1-\frac{\Gamma^{\mathrm{FB}}(F)}{\mu\,(1-\mu)}}\right]\mu\,(1-\mu)=\frac{\Gamma^{\mathrm{FB}}(F)}{2}+\frac{\left[\Gamma^{\mathrm{FB}}(F)\right]^2}{8\mu\,(1-\mu)}+\mathcal{O}(\left[\Gamma^{\mathrm{FB}}(F)\right]^3).$$

This implies the following:[7]

Corollary 2. For any distribution F with mean μ, the optimal fixed-price mechanism achieves at least 1/2 of the first-best gains from trade $\Gamma^{\mathrm{FB}}(F)$:

$$\max_{p\in\mathbb{R}_+}\Gamma(p;F)=\Gamma(\mu;F)\geq\frac{1}{2}\Gamma^{\mathrm{FB}}(F)\quad\text{for any } F\in\Delta([0,1]) \text{ with mean } \mu.$$

3.2 Method

We now describe the proof outline of Theorem 1. Technical details are relegated to Appendix A.

Our approach begins with the observation that the gains from trade under the mean price μ can be expressed as

$$\Gamma(\mu;F)=\int_0^\mu F(x)\ \mathrm{d}x=\int_0^\mu (\mu-x)\ \mathrm{d}F(x).$$

The worst-case guarantee of Theorem 1 can be found by maximizing first-best gains from trade, holding fixed the mean and gains from trade under the mean price:

$$\sup_{F\in\Delta([0,1])}\Gamma^{\mathrm{FB}}(F)\quad\text{subject to}\quad\begin{cases}\int_0^1 x\ \mathrm{d}F(x) & =\mu,\\ \int_0^\mu (\mu-x)\ \mathrm{d}F(x) & =\eta.\end{cases}\tag{P}$$

The constraints of the maximization problem (P) are linear with piecewise linear weights.

[7] Notably, all higher-order terms must be positive in the Taylor expansion, which converges for all $0<\Gamma^{\mathrm{FB}}(F)<1$.

We now argue that the first-best gains from trade, subject to the constraints of (P), are highest when the agents' distribution F is supported on at most 3-points. This argument consists of two key lemmas.

The first lemma states that the first-best gains from trade are increasing in the convex order.[8]

Lemma 1. If $F \geq_{\text{cx}} G$, then $\Gamma^{\text{FB}}(F) \geq \Gamma^{\text{FB}}(G)$.

The second lemma states that in a maximization program with $n-1$ linear constraints with piecewise linear weights, any continuous objective function that is increasing in the convex order is maximized on a set of discrete distributions.

Lemma 2. Suppose $h : \Delta([0,1]) \to \mathbb{R}$ is increasing in the convex order and continuous with respect to the supremum norm. Let $g_1, \ldots, g_n : \mathbb{R}_+ \to \mathbb{R}$ be continuous piecewise linear functions, such that g_j consists of k_j pieces. Fix $\gamma_1, \ldots, \gamma_n \in \mathbb{R}$, and denote by $\Delta_k([0,1])$ the set of distributions in $\Delta([0,1])$ that are supported on at most k points. Then the value of the following two maximization problems are the same:

$$\text{(i)} \quad \sup_{F \in \Delta([0,1])} h(F) \quad \text{subject to} \quad \begin{cases} \displaystyle\int_0^1 x \, \mathrm{d}F(x) & = \mu, \\ \displaystyle\int_0^1 g_j(x) \, \mathrm{d}F(x) & = \gamma_j, \text{ for } j = 1, \ldots, n. \end{cases}$$

$$\text{(ii)} \quad \sup_{F \in \Delta_k([0,1]), \text{ where } k = \sum_{j=1}^n k_j} h(F) \quad \text{subject to} \quad \begin{cases} \displaystyle\int_0^1 x \, \mathrm{d}F(x) & = \mu, \\ \displaystyle\int_0^1 g_j(x) \, \mathrm{d}F(x) & = \gamma_j, \text{ for } j = 1, \ldots, n. \end{cases}$$

Moreover, when solving the maximization problem (ii), it suffices to consider $F \in \Delta_k([0,1])$ with masses only on the boundaries of the pieces of $g_1, \ldots, g_n$.

[8] For any two distributions F, G, the convex order $\geq_{\text{cx}}$ is defined as follows:

$$F \geq_{\text{cx}} G \iff \int v(x) \, \mathrm{d}F(x) \geq \int v(x) \, \mathrm{d}G(x) \text{ for any convex function } v(\cdot).$$

See, for example, Chapter 3 of Shaked and Shanthikumar (2007).

Lemma 2 reduces the infinite-dimensional optimization problem to a finite-dimensional optimization problem. Instead of (P), it suffices for us to solve:

$$\sup_{F\in\Delta_3([0,1])} \Gamma^{\mathrm{FB}}(F) \quad \text{subject to} \quad \begin{cases} \int_0^1 x \, \mathrm{d}F(x) & = \mu, \\ \int_0^\mu (\mu - x) \, \mathrm{d}F(x) & = \eta. \end{cases} \tag{P'}$$

By Lemma 2, it suffices to consider 3-point distributions with masses on $\{0, \mu, 1\}$. Any such distribution F can be written as follows:

$$F(x) = \frac{\eta}{\mu} \cdot \mathbb{1}_{x\geq 0} + \left[1 - \frac{\eta}{\mu\,(1-\mu)}\right] \cdot \mathbb{1}_{x\geq\mu} + \frac{\eta}{1-\mu} \cdot \mathbb{1}_{x\geq 1}.$$

Therefore (P') admits the solution:

$$\sup_{F\in\Delta_3([0,1])} \Gamma^{\mathrm{FB}}(F) = 2\eta - \frac{\eta^2}{\mu\,(1-\mu)}.$$

Finally, Theorem 1 from straightforward algebraic manipulation, which we show in Appendix A:

Lemma 3. Let $0 \leq \eta \leq \mu\,(1-\mu)$ such that

$$\Gamma^{\mathrm{FB}}(F) \leq 2\eta - \frac{\eta^2}{\mu\,(1-\mu)}.$$

Then

$$\eta \geq \left[1 - \sqrt{1 - \frac{\Gamma^{\mathrm{FB}}(F)}{\mu\,(1-\mu)}}\right] \mu\,(1-\mu)\,.$$

3.3 Discussion

Our proof of Theorem 1 extends the tools developed in the literature on Bayesian persuasion. In contrast to Bayesian persuasion problems, our objective function $\Gamma^{\mathrm{FB}}(F)$ is neither linear nor convex in F. We show in Lemma 2 that this can be overcome if the objective function $\Gamma^{\mathrm{FB}}(F)$ is increasing in the convex order and the weights in the constraints are piecewise linear.

An interesting related result of McAfee (2008) is that setting a price equal to the median of F achieves an identical worst-case performance of 1/2 of the first-best gains from trade. Because agents are *ex ante* symmetric, the median price p_{median} is the competitive price: $F(p_{\text{median}}) =$

$1 - F(p_{\text{median}})$. McAfee's (2008) result is based on the elegant observation that, for any F,

$$\begin{aligned}\Gamma(p_{\text{median}}; F) &= \frac{1}{2}\int_{p_{\text{median}}}^{\infty} [1 - F(x)] \ \mathrm{d}x + \frac{1}{2}\int_{0}^{p_{\text{median}}} F(x) \ \mathrm{d}x \\ &\geq \frac{1}{2}\int_{0}^{\infty} F(x) [1 - F(x)] \ \mathrm{d}x = \frac{1}{2}\Gamma^{\text{FB}}(F).\end{aligned}$$

While this analysis is simple, it does not generalize to prices other than the median of the distribution. Moreover, by our Proposition 1, the mean price weakly dominates the median price. By contrast, our analysis in Theorem 1 shows that this bound cannot be achieved except when $\Gamma^{\text{FB}}(F) = 0$.

In reality, the Designer may wish to exclude such pathological distributions from consideration when evaluating the worst-case performance. This consideration motivates our first extension in Section 5, where we only consider distributions with mean absolute deviation that are bounded from below, away from zero. The assumption that the mean absolute deviation is bounded away from zero ensures that first-best gains from trade are always positive and also bounded away from zero. We show in Section 5 that Theorem 1 can be strengthened in this setting: the worst-case performance improves as the mean absolute deviation of the distribution increases.

While Theorem 1 quantifies the worst-case performance of fixed-price mechanisms in terms of the first-best gains from trade, this approach requires knowing the first-best gains from trade to begin with. In our second extension in Section 5, we show how the same approach can be used to establish an upper bound on the welfare loss under the optimal fixed-price mechanism. This upper bound can be expressed as either a function of the mean or the variance of the distribution F, which provides a useful way of estimating the worst-case welfare loss when only simple statistics of F are known.

We consider two additional extensions in Section 5, which lie beyond our application to size discovery mechanisms but demonstrate how our methods apply to more general settings. While Theorem 1 gives the worst-case performance for fixed-price mechanisms, we show that under some regularity conditions on the agents' distribution, the same result applies to any dominant-strategy incentive-compatible mechanism that does not run a budget deficit. Finally, we show that Theorem 1 can be extended to settings in which there is some asymmetry between the agents.

Finally, in other papers in the robust mechanism design literature, it is often assumed that the Designer observes only some moments of the agents' distribution (*e.g.*, only the mean). By contrast, we assumed in our model that the Designer observes the entire distribution F. This strengthens the result of our Proposition 1: the mean price is optimal among *all* fixed-price mechanisms that could depend arbitrarily on F, and not only among those that could depend on F through only its mean.

4 Fixed-Price Mechanisms in Large Markets

While Proposition 1 and Theorem 1 respectively characterize the optimal price and its worst-case performance in bilateral trade, dark pools allow trades to happen between many different sellers and buyers. In this section, we show that as the market size grows, not only do welfare losses vanish as previous papers show, but the optimal price also converges to the competitive price. All proofs are in Appendix A.

4.1 Model

We consider a market with N_S sellers and N_B buyers. Each seller owns a single indivisible good, and each buyer has unit demand. As before, agents are risk-neutral and have quasilinear payoffs. Agents have independent private values drawn from an identical distribution F. In contrast to our previous analysis, we impose in this section the assumption that F is twice continuously differentiable with positive density f on its support $[0,1] \subset \mathbb{R}_+$. We denote this family of distributions by $\mathcal{F}$.

In this market, the Designer chooses a fixed price $p \in \mathbb{R}_+$, which is allowed to depend on the distribution F. Having observed p, agents decide whether or not to trade. If there are more buyers than sellers willing to trade, the Designer randomly selects buyers so that the market clears; the case when there are more sellers than buyers willing to trade is symmetric. The Designer maximizes the gains from trade, denoted by $\Gamma(p; F, N_S, N_B)$.

4.2 Results

Our first result characterizes the optimal fixed-price mechanism in large markets:

Proposition 2. Let μ denote the mean of the distribution F. The optimal price p^* satisfies:

$$p^* = \mu + \frac{\zeta'(F(p^*); N_S, N_B)}{\zeta(F(p^*); N_S, N_B)} \cdot \Gamma(p^*; F, 1, 1).$$

Here, the "large-market scaling function" $\zeta(x; N_S, N_B)$ is defined by

$$\zeta(x; N_S, N_B) := \frac{1}{x\,(1-x)} \sum_{m=1}^{N_B} \sum_{n=1}^{N_S} \min\{m, n\} \binom{N_B}{m} \binom{N_S}{n} \left[1 - F(p)\right]^{N_S+m-n} \left[F(p)\right]^{N_B-m+n}.$$

Proposition 2 generalizes Proposition 1 by deriving the dependence of the optimal price p^* on the market size. The deviation of p^* from the mean depends on the market size and the agents' distribution.

We now study the case when markets are balanced: $N_S = N_B$. The assumption of balanced markets is consistent with our assumption that agents are *ex ante* symmetric, with inventory position independent of their private valuation of the good. *Ex ante*, agents have equal probability of being a seller or a buyer, so as the number of agents grows, the seller pool and the buyer pool are equal in size. For ease of notation, we write $\Gamma_N(p^*; F) = \Gamma(p^*; F, N, N)$.

If the market is balanced and the agents' distribution is symmetric (*e.g.*, uniform distributions and symmetric triangular distributions), then the mean price remains optimal:

Corollary 3. Suppose that $N_S = N_B$, and that F is symmetric with mean μ. Then the mean price $p^* = \mu$ is optimal.

Intuitively, the assumptions of Corollary 3 ensure that, conditional on trading, agent values are still symmetrically distributed around the mean. Thus the optimal price is the mean.

Our next result shows that, in balanced markets, the worst-case performance of the optimal fixed-price mechanism improves as market size increases:

Proposition 3. Suppose that $N_S = N_B = N$, and let $\Gamma_N^{\text{FB}}(F)$ denote the first-best gains from trade. For any distribution F satisfying $\Gamma_1^{\text{FB}}(F) > 0$, the worst-case performance of the optimal fixed-price mechanism is bounded below by an increasing fraction of $\Gamma_N^{\text{FB}}(F)$:

$$\max_{p \in \mathbb{R}_+} \Gamma_N(p; F) \geq \kappa(N) \cdot \Gamma_N^{\text{FB}}(F).$$

The fraction $\kappa(N)$ is increasing in N and satisfies $\lim_{N \to \infty} \kappa(N) = 1$.

Proposition 3 indicates that thicker markets improve the average welfare per agent in the worst-case. As the market grows, the gains from trade realized per agent under the optimal fixed-price mechanism increases and converges to the first-best gains from trade.

4.3 Implications for market design

Proposition 1 and Proposition 2 characterize how the optimal price should depend on market size. In contrast to Proposition 1, however, Proposition 2 shows that the optimal design in large markets requires more distributional knowledge than only the mean. Two questions arise: First, how should the price be determined when the Designer has only simple statistics of the agents' distribution, such as the mean and the median? Second, how much can errors in design affect the performance of fixed-price mechanisms? These questions are important in practical market design, such as determining how size discovery mechanisms should derive prices from lit exchanges.

We answer both questions under the assumption that markets are balanced. However, our results can be extended to more general settings as well.[9]

The answer to the first question is the simple heuristic of choosing the mean price in thin markets, but the competitive price in thick markets:

Proposition 4. Suppose that $N_S = N_B = N$. For any distribution F, as $N \to \infty$, the optimal fixed-price mechanism converges to the competitive price at rate $\mathcal{O}(N^{-1/4})$.

Proposition 4 complements the intuition that the welfare loss from setting the competitive price vanishes as the market grows. It shows that, not only does efficiency converge, but prices do as well. For large markets, the rate of $\mathcal{O}(N^{-1/4})$ ensures that the optimal price is close to the competitive price.

Because this heuristic yields only an approximately optimal price, it becomes important to quantify the welfare loss that might result from setting an approximately optimal price rather than the optimal price. The answer to the second question shows that small errors lead to small welfare loss:

Proposition 5. Suppose that $N_S = N_B = N$, and fix $\varepsilon > 0$. Given a distribution F, let the optimal fixed-price mechanism be p^*. For any price p such that $|p - p^*| \leq \varepsilon$,

$$\Gamma_N(p^*; F) - \Gamma_N(p; F) \leq C\varepsilon \cdot \max_{x:|x-p^*|\leq\varepsilon} |f(x)| \cdot \Gamma_N^{\text{FB}}(F).$$

Here, $C > 0$ is a constant independent of F and N.

[9] In more general settings, a required assumption concerns how the rate at which market tightness, N_S/N_B, grows as the total number of agents increases. Typically, the rate at which market tightness grows can be determined by taking "replicas" of an exchange with a finite number of agents.

The key point of Proposition 5 is that the worst-case welfare loss from a small error in price is small, and depends on the distribution F only through its density in a neighborhood of the optimal price. Intuitively, this is because setting the price precisely equal to the optimum is important exactly when the agents' values are concentrated around the optimal price. Nevertheless, the expected gains from trade are continuous in price, and so the welfare loss cannot be too large.

5 Extensions

In this section, we analyze four variations of our model. In each extension, we show how analogs of our main results continue to hold using essentially the same methods presented in Section 3. All proofs are in Appendix A.

5.1 Distributions with bounded mean absolute deviation

While the Designer may have some uncertainty over the agents' distributions, the Designer may deem unrealistic distributions that assign essentially unit probability that any agent's value for the good is fixed. Such distributions underlie the worst-case bound of Corollary 2, which can be attained only when $\Gamma^{\mathrm{FB}}(F) = 0$.

To exclude these distributions, we adopt the same model as Section 2, and additionally require distributions to have mean absolute deviation bounded away from zero. That is, we consider distributions F that have positive mean μ such that

$$\int_0^1 |\mu - x| \, \mathrm{d}F(x) \geq \alpha \quad \text{for some } \alpha > 0.$$

In this constraint, the weight function $x \mapsto |\mu - x|$ is continuous and piecewise linear. Moreover, as this constraint binds, Lemma 2 applies. We obtain the following analog of Theorem 1:

Theorem 2. For any distribution F with mean absolute deviation bounded below by $\alpha > 0$, the optimal fixed-price mechanism achieves at least $1/(2 - 2\alpha)$ of the first-best gains from trade:

$$\max_{p \in \mathbb{R}_+} \Gamma(p; F) = \Gamma(\mu; F) \geq \frac{1}{2 - 2\alpha} \Gamma^{\mathrm{FB}}(F).$$

The mean absolute deviation is a measure of statistical dispersion of agent values. Theorem 2 shows that as dispersion increases, the performance of fixed-price mechanisms improve relative to first-best gains from trade. Increased dispersion allows for the agents' values to be further apart with higher probability, allowing any fixed price to better capture the potential gains from trade.

5.2 Estimating welfare loss

To estimate the worst-case performance of Theorem 1, the Designer is required to know the first-best gains from trade. However, the same approach as Theorem 1 in fact allows the Designer to estimate the welfare loss as a function of only the mean or the variance of the distribution F. This is relevant in practical market design, such as when the Designer may want to establish a bound on the worst-case welfare loss while only knowing simple statistics of F.

Formally, using the same model as Section 3, our proof outline of Theorem 1 implies that the maximum welfare loss is bounded above:

Theorem 3. For any distribution F with mean μ, the maximum welfare loss under the optimal fixed-price mechanism is:

$$\Gamma^{\mathrm{FB}}(F) - \Gamma(\mu; F) \leq \frac{\mu\,(1-\mu)}{4} \leq \frac{1}{16}.$$

Moreover, if F has variance σ^2 and is continuous, then the maximum welfare loss also satisfies:

$$\Gamma^{\mathrm{FB}}(F) - \Gamma(\mu; F) \leq \frac{1}{4}\sigma^2 + \frac{1}{8\sqrt{3}}.$$

Theorem 3 shows that the absolute welfare loss is small relative to the mean and variance of F. This extends our understanding of Theorem 1: even though the optimal fixed-price mechanism may achieve as small as $1/2$ of the first-best gains from trade, the absolute welfare loss is nonetheless still small, therefore supporting the practical use of fixed-price mechanisms.

5.3 Relaxing budget balance

Because of applications to size discovery mechanisms, our main results in Section 3 focused on fixed-price mechanisms, which are *ex post* budget-balanced. For more general applications, the budget balance assumption may not be important. We show in this subsection how our methods may be extended to more general classes of mechanisms when the budget balance condition is relaxed.

Specifically, we assume that the *ex post* budget-balanced condition is relaxed to an *ex post* no-deficit condition. We assume an additional regularity condition on the agents' value distribution:

Assumption 1. The distribution F is continuously differentiable with density f, such that the hazard rate is increasing but the reverse hazard rate is decreasing; that is:

$$\theta \mapsto \frac{f(\theta)}{1-F(\theta)} \quad \text{is increasing and} \quad \theta \mapsto \frac{f(\theta)}{F(\theta)} \quad \text{is decreasing.}$$

Shao and Zhou (2016) show that, under Assumption 1, the optimal dominant-strategy mechanism that satisfies the *ex post* no-deficit condition must be either: (i) a fixed-price mechanism with price equal to the mean μ of the distribution F; or (ii) an option mechanism, which dictates trade if and only if the seller's value exceeds μ. Their elegant characterization of the optimal mechanism replaces Proposition 1, which allows Theorem 1 to extend easily to this setting:

Theorem 4. For any distribution F satisfying Assumption 1, the gains from trade $\Gamma^{\text{DS, ND}}(F)$ that the optimal dominant-strategy mechanism that also satisfies the *ex post* no-deficit condition achieves are bounded below:

$$\Gamma^{\text{DS, ND}}(F) \geq \left[1-\sqrt{1-\frac{\Gamma^{\text{FB}}(F)}{\mu\,(1-\mu)}}\right]\mu\,(1-\mu) \quad \text{for any } F \in \Delta([0,1]) \text{ with mean } \mu.$$

Theorem 4 shows that the *ex post* budget balance condition does not play a substantial role in pinning down the worst-case performance of fixed-price mechanisms in Theorem 1. Instead, welfare loss in the worst case can be attributed to only the restriction to dominant-strategy mechanisms.

5.4 Asymmetric agents

While agents were assumed to be *ex ante* symmetric in our model due to applications to financial markets, we relax the assumption in this subsection, allowing the seller's distribution F_S to differ from the buyer's distribution F_B. We assume F_S and F_B to be supported on $[0,1]$. Except for this difference, we maintain the same model of bilateral trade as in Section 2.

Our main results extend naturally to this setting when the asymmetry between agents is small:

Proposition 6. Suppose that $||\, F_S - F_B \,||_\infty \leq \varepsilon$, where ε is sufficiently small.[10] Then the worst-case performance of either agent's mean, μ_S or μ_B, achieves $[1/2 + \mathcal{O}(\varepsilon)]$ of the first-best gains

[10] Here, $||\cdot||_\infty$ denotes the supremum norm: $||\, G \,||_\infty = \sup_{x\in\mathbb{R}_+} |G(x)|$ for a given function G.

from trade:

$$\begin{cases} \Gamma(\mu_S; F_S, F_B) \geq \left[\dfrac{1}{2} + \mathcal{O}(\varepsilon)\right] \Gamma^{\text{FB}}(F_S, F_B), \\ \Gamma(\mu_B; F_S, F_B) \geq \left[\dfrac{1}{2} + \mathcal{O}(\varepsilon)\right] \Gamma^{\text{FB}}(F_S, F_B). \end{cases}$$

Here, $\Gamma(p; F_S, F_B)$ denotes the gains from trade realized by the fixed-price mechanism p in expectation, and $\Gamma^{\text{FB}}(F_S, F_B)$ denotes the first-best gains from trade.

When the asymmetry between agents is small, either of μ_S or μ_B is close to the optimal price. On the other hand, because asymmetry between agents is small, the optimal price achieves close to 1/2 of the first-best gains from trade even in the worst case.

When the asymmetry between agents is large, it is known that the optimal price may no longer perform well relative to first-best gains from trade in the worst case. We give a more detailed discussion of this case in Appendix B.2.

6 Discussion

We have presented in this paper a simple model that suggests why fixed-price mechanisms are used commonly in practice. In addition to being dominant-strategy incentive-compatible, fixed-price mechanisms are simple to design and have robust worst-case efficiency guarantees. Strikingly, while the competitive price may be approximately efficient in thick markets, the mean price is approximately optimal instead in thin markets.

Our methods enable us to provide a tight characterization of the worst-case performance, which distinguishes our work from much of the related literature. Analyzing the worst-case performance of mechanisms is difficult because it entails solving a nontrivial infinite-dimensional minimax problem. Previous work includes that of Arnosti et al. (2016) (Theorem 2), which analyzes the worst-case performance of an auction mechanism over power-law distributions. By parametrizing power-law distributions, their analysis simplifies to a finite-dimensional optimization problem. Carrasco et al. (2018) consider the worst-case optimal mechanism that maximizes revenue in the classic monopoly problem over buyer distributions, assuming that the seller monopolist knows the first N moments of the buyer's distribution.

Beyond the economic literature, our work is directly related to numerous papers in the computer science literature. The closest paper to ours in this literature is that of Blumrosen and Dobzinski (2016), who also study fixed-price mechanisms in the bilateral trade problem.

However, they study what fraction of first-best *welfare* fixed-price mechanisms can achieve, as opposed to first-best gains from trade. While they allow asymmetry between agents to be arbitrarily large, they are unable to provide a tight characterization of the worst-case performance. Instead, they prove a lower bound on the worst-case performance by explicitly constructing a mechanism that achieves that performance.[11] Various papers in the computer science literature have taken a similar approach to this problem as well as others, including work by Blumrosen and Dobzinski (2014), Blumrosen and Mizrahi (2016), Colini-Baldeschi et al. (2016) and Brustle et al. (2017).

Returning to our application to financial markets, we now discuss other considerations that are important to size discovery mechanisms in practice, as well as the limitations of our results for this purpose.

An important practical consideration for the design of size discovery mechanisms is how these mechanisms interact with the lit exchanges in which price discovery occurs. We have modeled our agents as liquidity traders in these markets, who are unsure of the "right" price and so behave as price-takers. We have shown that designing the optimal price for this setting is simple and has robust worst-case performance guarantees. Of course, size discovery mechanisms are attractive to informed traders as well, who may prefer size discovery mechanisms to lit exchanges for strategic price impact avoidance. This substitution may affect the reliability of the price discovery mechanism in providing the "right" price, and hence can result in welfare loss. Pioneering work on this tradeoff—including work by Zhu (2014), Duffie and Zhu (2017) and Antill and Duffie (2018)—shows that the market designer has to incentivize informed traders to participate in the price discovery mechanism, and to mitigate the incentives of informed traders to exploit the size discovery mechanism for price impact avoidance.

Relatedly, dynamic considerations are also important for the efficiency of size discovery mechanisms. We took an extreme stance in our static setting, where unrealized trades are simply foregone. In practice, agents trade dynamically, and unrealized trades from a single iteration of a size discovery mechanism can be realized either in the next iteration or in the lit exchange, subject to delay costs. The possibility of future trade motivates important questions in the design of financial markets: What is the optimal frequency of trade? How does competition between lit and dark exchanges impact efficiency?

Ultimately, these questions aim to address the issue of whether the amalgam of price discovery and size discovery mechanisms in our financial markets is efficient. Clearly, in this

[11] While their setting and methods differ from ours, their methods are rather interesting. We thus provide a separate discussion and analysis of their setting and methods in Appendix B.2.

context, our Theorem 1 should not be taken literally. Although a constant-factor approximation is often celebrated in the computer science literature, a worst-case performance of 1/2 is often unsatisfactory for economic applications, especially if the economic efficiency at stake is as large as it is in the financial market. The extensions that we studied suggest that the worst-case performance binds not due to our benchmark of first-best gains from trade, but rather due to unrealistic assumptions on the set of distributions allowed. A more reasonable estimate would account for statistical dispersion in the agents' values. With Theorem 2, we showed how the worst-case performance can be improved; but the same analysis also applies straightforwardly to other measures of statistical dispersion such as the interquartile range, different quantiles, and the median absolute deviation.

With these issues in mind, then, our paper provides a simple framework for analyzing the efficiency of size discovery mechanisms. While abstractions from informational asymmetry and dynamics simplify the much broader problem of financial market design, they provide sharp analytical results that suggest that, at least on a heuristic level, size discovery mechanisms can perform reasonably well even when very little information is known about the underlying distribution of traders' values.

On a more general level, designing mechanisms that are robust to the distribution of agents' values is important beyond just financial market design. Our approach has been different from that of many others in the robust mechanism design literature: instead of modeling robustness directly through a worst-case objective, we have chosen to study the worst-case performance of a classical Bayesian objective. As the various extensions in Section 5 illustrate, this approach is flexible and can be adapted to study more complicated problems.

References

Agarwal, R., N. S. Barnett, P. Cerone, and S. S. Dragomir (2005). A survey on some inequalities for expectation and variance. *Computers & Mathematics with Applications 49*(2-3), 429–480.

Antill, S. and D. Duffie (2018). Augmenting markets with mechanisms. *Working paper*.

Arnosti, N., M. Beck, and P. Milgrom (2016). Adverse selection and auction design for internet display advertising. *American Economic Review 106*(10), 2852–2866.

Bergemann, D., B. Brooks, and S. Morris (2015). The limits of price discrimination. *American Economic Review 105*(3), 921–57.

Bergemann, D. and S. Morris (2005). Robust mechanism design. *Econometrica 73*(6), 1771–1813.

Bergemann, D. and K. Schlag (2011). Robust monopoly pricing. *Journal of Economic Theory 146*(6), 2527–2543.

Blumrosen, L. and S. Dobzinski (2014). Reallocation mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC'14)*, New York, NY, USA, pp. 617. ACM.

Blumrosen, L. and S. Dobzinski (2016). (Almost) efficient mechanisms for bilateral trading. Working Paper.

Blumrosen, L. and Y. Mizrahi (2016). Approximating gains-from-trade in bilateral trading. In *International Conference on Web and Internet Economics (WINE)*, pp. 400–413. Springer.

Brustle, J., Y. Cai, F. Wu, and M. Zhao (2017). Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 589–590. ACM.

Carrasco, V., V. F. Luz, N. Kos, M. Messner, P. Monteiro, and H. Moreira (2018). Optimal selling mechanisms under moment conditions. *Journal of Economic Theory 177*, 245–279.

Carroll, G. (2015). Robustness and linear contracts. *American Economic Review 105*(2), 536–563.

Carroll, G. and I. Segal (2018). Robustly optimal auctions with unknown resale opportunities. *Review of Economic Studies*, forthcoming.

Chassang, S. (2013). Calibrated incentive contracts. *Econometrica 81*(5), 1935–1971.

Chung, K.-S. and J. C. Ely (2007). Foundations of dominant-strategy mechanisms. *Review of Economic Studies 74*(2), 447–476.

Colini-Baldeschi, R., B. de Keijzer, S. Leonardi, and S. Turchetta (2016). 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, pp. 1424–1443. SIAM.

Degryse, H., M. Van Achter, and G. Wuyts (2009). Dynamic order submission strategies with competition between a dealer market and a crossing network. *Journal of Financial Economics 91*(3), 319–338.

Drexl, M. and A. Kleiner (2015). Optimal private good allocation: The case for a balanced budget. *Games and Economic Behavior 94*, 169–181.

Duffie, D. and H. Zhu (2017). Size discovery. *Review of Financial Studies 30*(4), 1095–1150.

Fleming, M. J. and G. Nguyen (2018). Price and size discovery in financial markets: Evidence from the u.s. treasury securities market. *Review of Asset Pricing Studies*, forthcoming.

Gârleanu, N. and L. H. Pedersen (2004). Adverse selection and the required return. *Review of Financial Studies 17*(3), 643–665.

Gresik, T. A. and M. A. Satterthwaite (1989). The rate at which a simple market converges to efficiency as the number of traders increases: An asymptotic result for optimal trading mechanisms. *Journal of Economic Theory 48*(2), 304–332.

Hagerty, K. M. and W. P. Rogerson (1987). Robust trading mechanisms. *Journal of Economic Theory 42*(1), 94–107.

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

Myerson, R. B. and M. A. Satterthwaite (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory 29*(2), 265–281.

Shaked, M. and J. G. Shanthikumar (2007). *Stochastic Orders.* Springer Science & Business Media.

Shao, R. and L. Zhou (2016). Optimal allocation of an indivisible good. *Games and Economic Behavior 100*, 95–112.

Suzdaltsev, A. (2018). Distributionally robust pricing in auctions. *Working paper*.

Wilson, R. (1987). Game-theoretic analyses of trading processes. In T. F. Bewley (Ed.), *Advances in Economic Theory*, pp. 33–70. Cambridge: Cambridge University Press.

Wolitzky, A. (2016). Mechanism design with maxmin agents: Theory and an application to bilateral trade. *Theoretical Economics 11*(3), 971–1004.

Zhu, H. (2014). Do dark pools harm price discovery? *Review of Financial Studies 27*(3), 747–789.

A Appendix: Omitted Proofs

A.1 Proof of Theorem 1

Theorem 1. For any distribution F with mean μ, the gains from trade that the optimal fixed-price mechanism achieves are bounded below:

$$\max_{p\in\mathbb{R}_+}\Gamma(p;F)=\Gamma(\mu;F)\geq\left[1-\sqrt{1-\frac{\Gamma^{\mathrm{FB}}(F)}{\mu\,(1-\mu)}}\right]\mu\,(1-\mu)\quad\text{for any } F\in\Delta([0,1]) \text{ with mean } \mu.$$

As argued in the main text, we begin by proving Lemmas 1 and 2. We then use the result of both lemmas to prove Theorem 1, with the help of algebraic manipulation in Lemma 3.

Lemma 1. If $F\geq_{\mathrm{cx}}G$, then $\Gamma^{\mathrm{FB}}(F)\geq\Gamma^{\mathrm{FB}}(G)$.

Proof of Lemma 1. Since $F\geq_{\mathrm{cx}}G$, for any z:

$$\int_0^1 (b-z)_+\,\mathrm{d}F(b)\geq\int_0^1 (b-z)_+\,\mathrm{d}G(b).$$

Consider the function:

$$\phi_F: z\mapsto\int_0^1 (b-z)_+\,\mathrm{d}F(b).$$

Since $z\mapsto(b-z)_+$ is convex, therefore ϕ_F as the expectation of convex functions is convex. Consequently:

$$\Gamma^{\mathrm{FB}}(F)=\int_0^1\phi_F(z)\,\mathrm{d}F(z)\geq\int_0^1\phi_F(z)\,\mathrm{d}G(z)\geq\int_0^1\phi_G(z)\,\mathrm{d}G(z)=\Gamma^{\mathrm{FB}}(G).$$

□

Lemma 2. Suppose $h:\Delta([0,1])\to\mathbb{R}$ is increasing in the convex order and continuous with respect to the supremum norm. Let $g_1,\ldots,g_n:\mathbb{R}_+\to\mathbb{R}$ be continuous piecewise linear functions, such that g_j consists of k_j pieces. Fix $\gamma_1,\ldots,\gamma_n\in\mathbb{R}$, and denote by $\Delta_k([0,1])$ the set of distributions in $\Delta([0,1])$ that are supported on at most k points. Then the value of the following two maximization problems are the same:

$$
\text{(i)} \quad \sup_{F \in \Delta([0,1])} h(F) \quad \text{subject to} \quad \begin{cases} \displaystyle\int_0^1 x \; \mathrm{d}F(x) & = \mu, \\ \displaystyle\int_0^1 g_j(x) \; \mathrm{d}F(x) & = \gamma_j, \; \text{for } j = 1, \dots, n. \end{cases}
$$

$$
\text{(ii)} \quad \sup_{F \in \Delta_k([0,1]), \text{ where } k = \sum_{j=1}^n k_j} h(F) \quad \text{subject to} \quad \begin{cases} \displaystyle\int_0^1 x \; \mathrm{d}F(x) & = \mu, \\ \displaystyle\int_0^1 g_j(x) \; \mathrm{d}F(x) & = \gamma_j, \; \text{for } j = 1, \dots, n. \end{cases}
$$

Moreover, when solving the maximization problem (ii), it suffices to consider $F \in \Delta_k([0,1])$ with masses only on the boundaries of the pieces of $g_1, \dots, g_n$.

Proof of Lemma 2. Let $F \in \Delta([0,1])$ satisfy

$$
\int_0^1 x \; \mathrm{d}F(x) = \mu \quad \text{and} \quad \int_0^1 g_i(x) \; \mathrm{d}F(x) = \gamma_i.
$$

Because h is continuous with respect to the L^1 norm, we may assume without loss of generality that F is a smooth (*e.g.*, $\mathcal{C}^2$) function supported on $[0,1]$; otherwise, we can approximate any distribution arbitrarily well using such smooth functions. We proceed by constructing $\widetilde{F} \in \Delta_k([0,1])$, where $k = \sum_{j=1}^n k_j$, such that $h(F) \le h(\widetilde{F})$ and

$$
\int_0^1 x \; \mathrm{d}\widetilde{F}(x) = \mu \quad \text{and} \quad \int_0^1 g_j(x) \; \mathrm{d}\widetilde{F}(x) = \gamma_j \quad \text{for } j = 1, \dots, n.
$$

Let $0 = s_1 < s_2 < \cdots < s_k = 1$ be the boundaries of the pieces of the piecewise linear functions $g_1(x), \dots, g_n(x)$. For ease of exposition, we have assumed here that the boundaries are distinct; a similar argument holds if the boundaries are not distinct. Define m_i $(i = 1, \dots, k)$ by

$$
\begin{cases} m_1 & = \dfrac{1}{s_2 - s_1} \displaystyle\int_{s_1}^{s_2} (s_2 - x) \; \mathrm{d}F(x), \\ m_k & = \dfrac{1}{s_k - s_{k-1}} \displaystyle\int_{s_{k-1}}^{s_k} (x - s_{k-1}) \; \mathrm{d}F(x), \\ m_i & = \dfrac{1}{s_i - s_{i-1}} \displaystyle\int_{s_{i-1}}^{s_i} (x - s_{i-1}) \; \mathrm{d}F(x) + \dfrac{1}{s_{i+1} - s_i} \displaystyle\int_{s_i}^{s_{i+1}} (s_{i+1} - x) \; \mathrm{d}F(x) \quad \text{for } i = 2, \dots, k-1. \end{cases}
$$

We construct $\widetilde{F}$ as follows:

$$\widetilde{F}(x) = \sum_{i=1}^{k} m_i \cdot \mathbb{1}_{x \le s_i}.$$

By definition, $m_1, \ldots, m_k$ satisfy $0 \le m_1, \ldots, m_k \le 1$ and $m_1 + \cdots + m_k = 1$; hence $\widetilde{F}$ is a valid k-point cumulative distribution function: $\widetilde{F} \in \Delta_k([0,1])$. We claim that $\widetilde{F} \ge_{\text{cx}} F$. By Theorem 3.A.1 of Shaked and Shanthikumar (2007), $\widetilde{F} \ge_{\text{cx}} F$ if and only if

$$\int_0^z \widetilde{F}(x)\ \mathrm{d}x \ge \int_0^z F(x)\ \mathrm{d}x \quad \text{for any } z \in [0,1].$$

Observe that by our definition of $m_1, \ldots, m_k$,

$$\begin{aligned}\int_{s_i}^{s_{i+1}} \widetilde{F}(x)\ \mathrm{d}x &= (m_1 + \cdots + m_i) \cdot (s_{i+1} - s_i) \\ &= F(s_i) \cdot (s_{i+1} - s_i) + \int_{s_i}^{s_{i+1}} (s_{i+1} - x)\ \mathrm{d}F(x) = \int_{s_i}^{s_{i+1}} F(x)\ \mathrm{d}x \quad \text{for } i = 1, \ldots, k-1.\end{aligned}$$

Therefore, our construction ensures that

$$\int_0^{s_i} \widetilde{F}(x)\ \mathrm{d}x = \int_0^{s_i} F(x)\ \mathrm{d}x \quad \text{for } i = 1, \ldots, k.$$

Moreover, $\widetilde{F}(x)$ is constant on each interval $[s_i, s_{i+1})$; hence

$$\frac{\partial}{\partial z} \int_{s_i}^{z} \left[\widetilde{F}(x) - F(x)\right]\ \mathrm{d}x = \widetilde{F}(s_i) - F(z) \quad \text{is decreasing for } z \in [s_i, s_{i+1}).$$

Thus, on $[s_i, s_{i+1})$,

$$z \mapsto \int_{s_i}^{z} \left[\widetilde{F}(x) - F(x)\right]\ \mathrm{d}x \quad \text{is concave.}$$

This function achieves its minimum at either s_i or s_{i+1}; also, its value at either endpoint is zero. This shows that

$$\int_{s_i}^{z} \widetilde{F}(x)\ \mathrm{d}x \ge \int_{s_i}^{z} F(x)\ \mathrm{d}x \quad \text{for } z \in [s_i, s_{i+1}).$$

We thus conclude that $\widetilde{F} \ge_{\text{cx}} F$, as claimed. Since h is increasing in the convex order, this implies that $h(F) \le h(\widetilde{F})$. So it remains to verify that $\widetilde{F}$ satisfies

$$\int_0^1 x\ \mathrm{d}\widetilde{F}(x) = \mu \quad \text{and} \quad \int_0^1 g_j(x)\ \mathrm{d}\widetilde{F}(x) = \gamma_j \quad \text{for } j = 1, \ldots, n.$$

We use the key observation that

$$\int_{s_i}^{s_{i+1}} x \ \mathrm{d}F(x) = \frac{s_i}{s_{i+1} - s_i} \int_{s_i}^{s_{i+1}} (s_{i+1} - x) \ \mathrm{d}F(x) + \frac{s_{i+1}}{s_{i+1} - s_i} \int_{s_i}^{s_{i+1}} (x - s_i) \ \mathrm{d}F(x).$$

Summing over both sides for $i = 1, \ldots, k-1$, we get

$$\mu = \int_0^1 x \ \mathrm{d}F(x) = \sum_{i=1}^{k-1} \int_{s_i}^{s_{i+1}} x \ \mathrm{d}F(x) = \sum_{i=1}^{k} s_i m_i = \int_0^1 x \ \mathrm{d}\widetilde{F}(x).$$

More generally, for any constants α_i, β_1:

$$\int_{s_i}^{s_{i+1}} (\alpha_i + \beta_i x) \ \mathrm{d}F(x) = \frac{\alpha_i + \beta_i s_i}{s_{i+1} - s_i} \int_{s_i}^{s_{i+1}} (s_{i+1} - x) \ \mathrm{d}F(x) + \frac{\alpha_i + \beta_i s_{i+1}}{s_{i+1} - s_i} \int_{s_i}^{s_{i+1}} (x - s_i) \ \mathrm{d}F(x).$$

Therefore, for any continuous piecewise linear function $g : [0,1] \to \mathbb{R}$ such that $g(x) = \alpha_i + \beta_i x$ on $[s_i, s_{i+1}]$, we sum over both sides of the above expression for $i = 1, \ldots, k-1$ to get

$$\int_0^1 g(x) \ \mathrm{d}F(x) = \sum_{i=1}^{k-1} \int_{s_i}^{s_{i+1}} (\alpha_i + \beta_i x) \ \mathrm{d}F(x) = \sum_{i=1}^{k} g(s_i) m_i = \int_0^1 g(x) \ \mathrm{d}\widetilde{F}(x).$$

□

Lemma 3. Let $0 \leq \eta \leq \mu(1-\mu)$ such that

$$\Gamma^{\mathrm{FB}}(F) \leq 2\eta - \frac{\eta^2}{\mu(1-\mu)}.$$

Then

$$\eta \geq \left[1 - \sqrt{1 - \frac{\Gamma^{\mathrm{FB}}(F)}{\mu(1-\mu)}}\right] \mu(1-\mu).$$

Proof of Lemma 3. We view the given hypothesis as a quadratic inequality in η. By the quadratic formula, the solution to this inequality must satisfy

$$\frac{1 - \sqrt{1 - \frac{\Gamma^{\mathrm{FB}}(F)}{\mu(1-\mu)}}}{\frac{1}{\mu(1-\mu)}} \leq \eta \leq \frac{1 + \sqrt{1 - \frac{\Gamma^{\mathrm{FB}}(F)}{\mu(1-\mu)}}}{\frac{1}{\mu(1-\mu)}} \implies 1 - \frac{\eta}{\mu(1-\mu)} \leq \sqrt{1 - \frac{\Gamma^{\mathrm{FB}}(F)}{\mu(1-\mu)}}.$$

The implication follows from the latter inequality. □

Proof of Theorem 1. We now combine the results of Lemmas 1, 2 and 3 to prove Theorem 1. By Lemma 1, Γ^{FB} is increasing in the convex order. Moreover, if there exists a sequence $F_m \to F$ in the supremum norm, then $\lim_{m\to\infty} F_m(x) = F(x)$ for all points $x \in [0,1]$ at which F is continuous; hence $\mathrm{d}F_m$ converges weakly to $\mathrm{d}F$. In particular, for any bounded continuous function $\phi : [0,1] \to \mathbb{R}$,

$$\int_0^1 \phi(x)\ \mathrm{d}F_m(x) \to \int_0^1 \phi(x)\ \mathrm{d}F(x).$$

Therefore $\Gamma^{\text{FB}}(F_m) \to \Gamma^{\text{FB}}(F)$; hence Γ^{FB} is continuous in the supremum norm. Consequently, Lemma 2 applies. It suffices for us to solve:

$$\sup_{F\in\Delta_3([0,1])} \Gamma^{\text{FB}}(F) \quad \text{subject to} \quad \begin{cases} \displaystyle\int_0^1 x\ \mathrm{d}F(x) & = \mu, \\ \displaystyle\int_0^\mu (\mu - x)\ \mathrm{d}F(x) & = \eta. \end{cases}$$

Now, for any $F \in \Delta_3([0,1])$ with mean μ with masses on $\{0, \mu, 1\}$, we can parametrize F by:

$$F(x) = \frac{\eta}{\mu} \cdot \mathbb{1}_{x\geq 0} + \left[1 - \frac{\eta}{\mu\,(1-\mu)}\right] \cdot \mathbb{1}_{x\geq\mu} + \frac{\eta}{1-\mu} \cdot \mathbb{1}_{x\geq 1}.$$

Therefore (P') admits the solution:

$$\sup_{F\in\Delta_3([0,1])} \Gamma^{\text{FB}}(F) = 2\eta - \frac{\eta^2}{\mu\,(1-\mu)}.$$

Thus, for any $F \in \Delta([0,1])$, we have

$$\Gamma^{\text{FB}}(F) \leq 2\,\Gamma(\mu; F) - \frac{[\Gamma(\mu;F)]^2}{\mu\,(1-\mu)}.$$

By Lemma 3, the latter implies that

$$\Gamma(\mu; F) \geq \left[1 - \sqrt{1 - \frac{\Gamma^{\text{FB}}(F)}{\mu\,(1-\mu)}}\right] \mu\,(1-\mu)\,.$$

This completes the proof of Theorem 1. □

A.2 Proof of Proposition 2

Proposition 2. Let μ denote the mean of the distribution F. The optimal price p^* satisfies:

$$p^* = \mu + \frac{\zeta'(F(p^*); N_S, N_B)}{\zeta(F(p^*); N_S, N_B)} \cdot \Gamma(p^*; F, 1, 1).$$

Here, the "large-market scaling function" $\zeta(x; N_S, N_B)$ is defined by

$$\zeta(x; N_S, N_B) := \frac{1}{x\,(1-x)} \sum_{m=1}^{N_B} \sum_{n=1}^{N_S} \min\{m, n\} \binom{N_B}{m} \binom{N_S}{n} \left[1 - F(p)\right]^{N_S+m-n} \left[F(p)\right]^{N_B-m+n}.$$

Proof. The gains from trade under the fixed price p can be written as

$$\begin{aligned}
&\Gamma(p; F, N_S, N_B) \\
&= \{\mathbb{E}\left[\theta \,|\, \theta > p\right] - \mathbb{E}\left[\theta \,|\, \theta \leq p\right]\} \cdot \sum_{m=1}^{N_B} \sum_{n=1}^{N_S} \min\{m, n\} \binom{N_B}{m} \binom{N_S}{n} \left[1 - F(p)\right]^{N_S+m-n} \left[F(p)\right]^{N_B-m+n} \\
&= \frac{\Gamma(p; F, 1, 1)}{F(p)\left[1 - F(p)\right]} \cdot \sum_{m=1}^{N} \sum_{n=1}^{N} \min\{m, n\} \binom{N}{m} \binom{N}{n} \left[1 - F(p)\right]^{N+m-n} \left[F(p)\right]^{N-m+n}.
\end{aligned}$$

This can be simplified using the large-market scaling function defined in the statement of the proposition:

$$\Gamma(p; F, N_S, N_B) = \zeta(F(p); N_S, N_B) \cdot \Gamma(p; F, 1, 1).$$

Since F is assumed to be twice continuously differentiable with density f, taking the first-order condition yields

$$0 = \frac{\partial\, \Gamma(p^*; F, N_S, N_B)}{\partial p} = \zeta'(F(p^*); N_S, N_B) f(p^*) \cdot \Gamma(p^*; F, 1, 1) + \zeta(F(p^*); N_S, N_B) \cdot \frac{\partial\, \Gamma(p^*; F, 1, 1)}{\partial p}.$$

Denoting by μ the mean of F,

$$\frac{\partial\, \Gamma(p; F, 1, 1)}{\partial p} = (\mu - p) \cdot f(p).$$

Since f is assumed to be positive, the optimal price p^* satisfies

$$\zeta'(F(p^*); N_S, N_B) \cdot \Gamma(p^*; F, 1, 1) + (\mu - p^*) \cdot \zeta(F(p^*); N_S, N_B) = 0.$$

Rearranging yields the desired expression. Finally, setting $p^* = 0, 1$ is dominated by setting any $p^* \in (0, 1)$; hence the first-order condition must be satisfied by the optimal price p^*. □

Corollary 2. Suppose that $N_S = N_B$, and that F is symmetric with mean μ. Then the mean price $p^* = \mu$ is optimal.

Proof. For any N, the large-market scaling function $\zeta(x; N, N)$ is symmetric about $x = 1/2$, where it achieves its maximum. Since F is symmetric, $F(\mu) = 1/2$; hence Proposition 2 implies that the mean price $p^* = \mu$ is optimal. □

A.3 Proof of Proposition 3

Proposition 3. Suppose that $N_S = N_B = N$, and let $\Gamma_N^{\mathrm{FB}}(F)$ denote the first-best gains from trade. For any distribution F satisfying $\Gamma_1^{\mathrm{FB}}(F) > 0$, the worst-case performance of the optimal fixed-price mechanism is bounded below by an increasing fraction of $\Gamma_N^{\mathrm{FB}}(F)$:

$$\max_{p \in \mathbb{R}_+} \Gamma_N(p; F) \geq \kappa(N) \cdot \Gamma_N^{\mathrm{FB}}(F).$$

The fraction $\kappa(N)$ is increasing in N and satisfies $\lim_{N \to \infty} \kappa(N) = 1$.

Proof. The proposition follows straightforwardly from Theorem 6 of Gresik and Satterthwaite (1989), who show that

$$\left| \frac{1}{N} \max_{p \in \mathbb{R}_+} \Gamma_N(p; F) - \frac{1}{N} \Gamma_N^{\mathrm{FB}}(F) \right| \leq \mathcal{O}(N^{-1/2}).$$

That is, the average gains from trade realized per agent in the optimal fixed-price mechanism converges at a rate $\mathcal{O}(N^{-1/2})$ to the first-best. Therefore, there exists some constant $C > 0$ such that

$$\left| \frac{1}{N} \max_{p \in \mathbb{R}_+} \Gamma_N(p; F) - \frac{1}{N} \Gamma_N^{\mathrm{FB}}(F) \right| \leq C \cdot N^{-1/2} \implies \frac{1}{N} \max_{p \in \mathbb{R}_+} \Gamma_N(p; F) \geq \frac{1}{N} \Gamma_N^{\mathrm{FB}}(F) - C \cdot N^{-1/2}.$$

Note that $\Gamma_N^{\mathrm{FB}}(F) \geq N \, \Gamma_1^{\mathrm{FB}}(F)$ for any N, since the first-best allocation performs weakly better than splitting agents into N buyer-seller pairs and realizing the first-best outcome within each pair. Therefore:

$$\max_{p \in \mathbb{R}_+} \Gamma_N(p; F) \geq \left[1 - \frac{C \cdot N^{-1/2}}{\Gamma_1^{\mathrm{FB}}(F)} \right] \Gamma_N^{\mathrm{FB}}(F).$$

Thus we may choose $\kappa(N) = 1 - CN^{-1/2}/\Gamma_1^{\mathrm{FB}}(F)$, which is increasing in N and satisfies $\lim_{N\to\infty}\kappa(N) = 1$. □

A.4 Proof of Proposition 4

Proposition 4. Suppose that $N_S = N_B = N$. For any distribution F, as $N \to \infty$, the optimal fixed-price mechanism converges to the competitive price at rate $\mathcal{O}(N^{-1/4})$.

Proof. To show Proposition 4, we first derive properties that the large-market scaling function ζ_N must satisfy. Denote

$$\zeta_N(x) := \frac{1}{N}\zeta(x; N, N) = \frac{1}{Nx\,(1-x)}\sum_{m=1}^{N}\sum_{n=1}^{N}\min\{m,n\}\binom{N}{m}\binom{N}{n}(1-x)^{N+m-n}\,x^{N-m+n}.$$

Consider random variables $Z_1 \sim \mathsf{Bin}(N, x)$ and $Z_2 \sim \mathsf{Bin}(N, 1-x)$. Then in fact:

$$\mathbb{E}\left[\min\{Z_1, Z_2\}\right] = \sum_{m=1}^{N}\sum_{n=1}^{N}\min\{m,n\}\binom{N}{m}\binom{N}{n}(1-x)^{N+m-n}\,x^{N-m+n}.$$

Therefore, convergence properties of $\zeta_N(x)$ are equivalent to the concentration of $\min\{Z_1, Z_2\}$ around its mean as $N \to \infty$. Since $\zeta_N(x)$ is symmetric about $x = 1/2$, we assume without loss of generality that $x \in [0, 1/2]$.

We first apply concentration bounds on Z_1 and Z_2 separately. By Hoeffding's inequality, observe that Z_1 must be concentrated about its mean of Nx; that is, for any t:

$$\mathbb{P}\left[|Z_1 - Nx| > t\right] < 2\exp\left(-\frac{2t^2}{N}\right).$$

Likewise, Z_2 must be concentrated about its mean of $N(1-x)$:

$$\mathbb{P}\left[|Z_2 - N(1-x)\,| > t\right] < 2\exp\left(-\frac{2t^2}{N}\right).$$

Now, given that $x \in [0, 1/2]$ by assumption, we have:

$$\begin{aligned}\mathbb{P}\left[\left|\min\{Z_1, Z_2\} - \min\{Nx, N(1-x)\}\right| < t\right] &\geq \mathbb{P}\left[|Z_1 - Nx| < t\right] \cdot \mathbb{P}\left[|Z_2 - N(1-x)| < t\right] \\ &\geq \left[1 - 2\exp\left(-\frac{2t^2}{N}\right)\right]^2 \geq 1 - 4\exp\left(-\frac{2t^2}{N}\right).\end{aligned}$$

Thus $\min\{Z_1, Z_2\}$ must be concentrated around $\min\{Nx, N(1-x)\}$ as $N \to \infty$. Applying Jensen's inequality for $z \mapsto |z|$ yields, for any $t > 0$:

$$\begin{aligned}&\left|Nx(1-x)\,\zeta_N(x) - \min\{Nx, N(1-x)\}\right| \\ &= \left|\mathbb{E}\left[\min\{Z_1, Z_2\}\right] - \min\{Nx, N(1-x)\}\right| \\ &\leq \mathbb{P}\left[\left|\min\{Z_1, Z_2\} - \min\{Nx, N(1-x)\}\right| < t\right] \cdot t \\ &\qquad + \mathbb{P}\left[\left|\min\{Z_1, Z_2\} - \min\{Nx, N(1-x)\}\right| \geq t\right] \cdot \left|Nx(1-x)\,\zeta_N(x) - \min\{Nx, N(1-x)\}\right|.\end{aligned}$$

Thus, for $t > 0$ satisfying $\exp\left(-2t^2/N\right) < 1/4$:

$$\left|Nx(1-x)\,\zeta_N(x) - \min\{Nx, N(1-x)\}\right| \leq \frac{t}{1 - 4\exp\left(-\frac{2t^2}{N}\right)}. \tag{†}$$

Now, we prove the claim of Proposition 4. Fix a distribution F, and let $F(\mu) \in (0, 1)$. Suppose that $F(\mu) \leq 1/2$. Denote by p_N^* the optimal price when there are N agents on each side of the market. If $F(p_N^*) < F(\mu)$, since $\zeta_N(x)$ is increasing on $[0, 1/2]$ and decreasing on $[1/2, 1]$, hence Proposition 2 implies that:

$$\frac{\zeta_N'(F(p_N^*))}{\zeta_N(F(p_N^*))} \cdot \Gamma(p_N^*; F, 1) > 0 \implies p_N^* > \mu,$$

whence $F(p_N^*) > F(\mu)$, a contradiction. If $F(p_N^*) > 1/2$, then Proposition 2 implies that:

$$\frac{\zeta_N'(F(p_N^*))}{\zeta_N(F(p_N^*))} \cdot \Gamma(p_N^*; F, 1) < 0 \implies p_N^* < \mu,$$

whence $F(p_N^*) < F(\mu) \leq 1/2$, a contradiction. Thus $F(\mu) \leq 1/2$ implies that $F(p_N^*) \in [F(\mu), 1/2]$ for any N. Symmetrically, $F(\mu) \geq 1/2$ implies that $F(p_N^*) \in [1/2, F(\mu)]$ for any N.

By symmetry, we assume $F(\mu) \leq 1/2$ without loss of generality, so that $F(p_N^*) \in [F(\mu), 1/2]$ for any N. Then we may take $t = Cx\sqrt{N}/F(\mu)$, for some sufficiently large constant $C > 0$, in (†)

to get:

$$|Nx\,(1-x)\,\zeta_N(x) - \min\{Nx, N\,(1-x)\}| \le \frac{Cx\sqrt{N}}{F(\mu)\left[1 - 4\exp\left(-\frac{2C^2x^2}{[F(\mu)]^2}\right)\right]}.$$

Therefore, for any $x \in [F(\mu), 1/2]$, we have:

$$\left|\zeta_N(x) - \min\left\{\frac{1}{x}, \frac{1}{1-x}\right\}\right| \le \mathcal{O}(N^{-1/2}) \quad \text{uniformly in } x \in [F(\mu), 1/2].$$

For ease of notation, we will write

$$\zeta_\infty(x) := \min\left\{\frac{1}{x}, \frac{1}{1-x}\right\} \quad \text{and} \quad \overline{\Gamma}_\infty(p;F) := \Gamma_1(p;F)\cdot\zeta_\infty(x).$$

For ease of notation, we will denote normalize the gains from trade by the number of agents on each side of the market: $\overline{\Gamma}_N(p;F) = \Gamma_N(p;F)/N$ and $\overline{\Gamma}_N^{\mathrm{FB}}(F) = \Gamma_N^{\mathrm{FB}}(F)/N$. Let p_∞^* denote the competitive price. We can verify that $\overline{\Gamma}_\infty(p_\infty^*;F) = \overline{\Gamma}_\infty^{\mathrm{FB}}(F) = \lim_{N\to\infty}\overline{\Gamma}_N^{\mathrm{FB}}(F)$; that is, the competitive price p_∞^* is optimal and achieves the first-best gains from trade in the limit as $N \to \infty$.

Next, since $F \in \mathcal{F}$ is assumed to be twice continuously differentiable, so is $\overline{\Gamma}_\infty(p;F)$; hence Taylor's theorem applies. Since $\partial\overline{\Gamma}_\infty(p_\infty^*;F)/\partial p = 0$, hence:

$$\overline{\Gamma}_\infty(p_N^*;F) = \overline{\Gamma}_\infty(p_\infty^*;F) + \frac{(p_N^* - p_\infty^*)^2}{2}\cdot\frac{\partial^2}{\partial p^2}\overline{\Gamma}_\infty(p_\infty^*;F) + R_2(p_N^*)\cdot(p_N^* - p_\infty^*)^2,$$

where R_2 is the Peano form of the remainder such that $\lim_{p\to p_\infty^*} R_2(p) = 0$. On the other hand:

$$\begin{aligned}\left|\overline{\Gamma}_\infty(p_N^*;F) - \overline{\Gamma}_\infty(p_\infty^*;F)\right| &\le \left|\overline{\Gamma}_\infty(p_N^*;F) - \overline{\Gamma}_N(p_N^*;F)\right| + \left|\overline{\Gamma}_N(p_N^*;F) - \overline{\Gamma}_\infty(p_\infty^*;F)\right| \\ &\le \left|\overline{\Gamma}_\infty(p_N^*;F) - \overline{\Gamma}_N(p_N^*;F)\right| + \left|\overline{\Gamma}_N(p_\infty^*;F) - \overline{\Gamma}_\infty^{\mathrm{FB}}(F)\right|.\end{aligned}$$

By Theorem 6 of Gresik and Satterthwaite (1989), the latter term is

$$\left|\overline{\Gamma}_N(p_\infty^*;F) - \overline{\Gamma}_\infty^{\mathrm{FB}}(F)\right| \le \mathcal{O}(N^{-1/2}).$$

Moreover, as we showed earlier, the former term is

$$\left|\overline{\Gamma}_\infty(p_N^*;F) - \overline{\Gamma}_N(p_N^*;F)\right| = \overline{\Gamma}_1(p_N^*;F)\cdot|\zeta_N(F(p_N^*)) - \zeta_\infty(F(p_N^*))| \le \mathcal{O}(N^{-1/2}).$$

These imply that

$$\mathcal{O}(N^{-1/2}) \geq \left|\overline{\Gamma}_\infty(p_N^*; F) - \overline{\Gamma}_\infty(p_\infty^*; F)\right| = (p_N^* - p_\infty^*)^2 \left|\frac{1}{2}\frac{\partial^2}{\partial p^2}\overline{\Gamma}_\infty(p_\infty^*; F) + R_2(p_N^*)\right|.$$

Finally, because $R_2(p_N^*) \to 0$ as $N \to \infty$ and $\partial^2\overline{\Gamma}_\infty(p_\infty^*; F)/\partial p^2 > 0$, hence we get the desired result:

$$|p_N^* - p_\infty^*| \leq \mathcal{O}(N^{-1/4}).$$

□

A.5 Proof of Proposition 5

Proposition 5. Suppose that $N_S = N_B = N$, and fix $\varepsilon > 0$. Given a distribution F, let the optimal fixed-price mechanism be p^*. For any price p such that $|p - p^*| \leq \varepsilon$,

$$\Gamma_N(p^*; F) - \Gamma_N(p; F) \leq C\varepsilon \cdot \max_{x:|x-p^*|\leq\varepsilon} |f(x)| \cdot \Gamma_N^{\mathrm{FB}}(F).$$

Here, $C > 0$ is a constant independent of F and N.

Proof. As in Proposition 2, we express the gains from trade using the large-market scaling function:

$$\frac{1}{N}{}_N\Gamma(p; F) = \frac{1}{N}\Gamma_1(p; F) \cdot \zeta(F(p); N, N) = \Gamma_1(p; F) \cdot \zeta_N(F(p)).$$

Here, as in Proposition 4, $\zeta_N(x) := \zeta(x; N, N)/N$. Additionally, we define:

$$M_1(N) := \sup_{x\in[0,1]} |\zeta_N(x)| \quad \text{and} \quad M_2(N) := \sup_{x\in[0,1]} |\zeta_N'(x)|.$$

We can prove that $\zeta_{N+1}(x) \geq \zeta_N(x) \geq 0$ and that $\zeta_N(x)$ is symmetric about $x = 1/2$; moreover, $\lim_{N\to\infty} \zeta_N(x) = \min\{1/x, 1/(1-x)\}$ by the proof of Proposition 4. Consequently, $M_1(N) \leq 2$ for all N; similarly, $M_2(N) \leq 4$ for all N. Thus, for any price p:

$$\begin{aligned}\frac{1}{N}|\,\Gamma_N(p^*; F) - \Gamma_N(p; F)| &= |\zeta_N(F(p^*))\,\Gamma_1(p^*; F) - \zeta_N(F(p))\,\Gamma_1(p; F)| \\ &\leq \zeta_N(F(p)) \cdot |\,\Gamma_1(p^*; F) - \Gamma_1(p; F)| + \Gamma_1(p^*; F) \cdot |\zeta_N(F(p^*)) - \zeta_N(F(p))|.\end{aligned}$$

Observe that

$$\begin{aligned}|\,\Gamma_1(p^*;F)-\Gamma_1(p;F)| &= \left|(\mu-p)\,[F(p^*)-F(p)]+\int_p^{p^*}\,[F(x)-F(p^*)]\;\mathrm{d}x\right| \\ &\leq |F(p^*)-F(p)|\cdot|\mu-p^*| \leq |F(p^*)-F(p)|\cdot\Gamma_1^{\mathrm{FB}}(F).\end{aligned}$$

On the other hand,

$$|\zeta_N(F(p^*))-\zeta_N(F(p))| \leq M_2(N)\cdot|F(p^*)-F(p)|.$$

Therefore, for any price p such that $|p-p^*|<\varepsilon$:

$$\frac{1}{N}|\,\Gamma_N(p^*;F)-\Gamma_N(p;F)| \leq 6\varepsilon\cdot\max_{x:|x-p^*|\leq\varepsilon}|f(x)|\cdot\Gamma_1^{\mathrm{FB}}(F) \leq 6\varepsilon\cdot\max_{x:|x-p^*|\leq\varepsilon}|f(x)|\cdot\frac{1}{N}\,\Gamma_N^{\mathrm{FB}}(F).$$

The latter inequality follows because the first-best gains from trade with N buyers and N sellers cannot be less than pairing up the buyers and sellers *a priori* and realizing N times of the first-best gains from trade within each pair. □

A.6 Proof of Theorem 2

Theorem 2. For any distribution F with mean absolute deviation bounded below by $\alpha>0$, the optimal fixed-price mechanism achieves at least $1/\left(2-2\alpha\right)$ of the first-best gains from trade:

$$\max_{p\in\mathbb{R}_+}\Gamma(p;F)=\Gamma(\mu;F)\geq\frac{1}{2-2\alpha}\,\Gamma^{\mathrm{FB}}(F).$$

Proof. We solve the following maximization problem:

$$\sup_{F\in\Delta([0,1])}\Gamma^{\mathrm{FB}}(F)\quad\text{subject to}\quad\begin{cases}\displaystyle\int_0^1 x\;\mathrm{d}F(x) & =\mu,\\ \displaystyle\int_0^{\mu}(\mu-x)\;\mathrm{d}F(x) & =\eta,\\ \displaystyle\int_0^1|\mu-x|\;\mathrm{d}F(x) & \geq\alpha.\end{cases}$$

By the proof of Theorem 1, Lemmas 1 and 2 apply: we solve the following equivalent maximization problem, where $F \in \Delta_3([0,1])$ has mass only on $\{0, \mu, 1\}$:

$$\sup_{F \in \Delta_3([0,1])} \Gamma^{\mathrm{FB}}(F) \quad \text{subject to} \quad \begin{cases} \int_0^1 x \ \mathrm{d}F(x) & = \mu, \\ \int_0^\mu (\mu - x) \ \mathrm{d}F(x) & = \eta, \\ \int_0^1 |\mu - x| \ \mathrm{d}F(x) & \geq \alpha. \end{cases}$$

Consider the relaxed problem in which only the first two constraints bind. Then, as in the proof of Theorem 1, we have the parametrization

$$F(x) = \frac{\eta}{\mu} \cdot \mathbb{1}_{x \geq 0} + \left[1 - \frac{\eta}{\mu\,(1-\mu)}\right] \cdot \mathbb{1}_{x \geq \mu} + \frac{\eta}{1-\mu} \cdot \mathbb{1}_{x \geq 1},$$

for which we compute that

$$\frac{\Gamma(\mu; F)}{\Gamma^{\mathrm{FB}}(F)} = \frac{1}{2 - \frac{\eta}{\mu(1-\mu)}}.$$

Now, observe that the inequality constraint is

$$\alpha \leq \int_0^1 |\mu - x| \ \mathrm{d}F(x) = \eta + \int_\mu^1 (x - \mu) \ \mathrm{d}F(x) = 2\eta.$$

Since $z\,(1-z) \leq 1/4$ for any $z \in \mathbb{R}$, we have

$$\frac{\eta}{\mu\,(1-\mu)} \geq 2\alpha \implies \frac{\Gamma(\mu; F)}{\Gamma^{\mathrm{FB}}(F)} = \frac{1}{2 - \frac{\eta}{\mu(1-\mu)}} \geq \frac{1}{2 - 2\alpha}.$$

□

A.7 Proof of Theorem 3

Theorem 3. For any distribution F with mean μ, the maximum welfare loss under the optimal fixed-price mechanism is:

$$\Gamma^{\mathrm{FB}}(F) - \Gamma(\mu; F) \leq \frac{\mu\,(1-\mu)}{4} \leq \frac{1}{16}.$$

Moreover, if F has variance σ^2 and is continuous, then the maximum welfare loss also satisfies:

$$\Gamma^{\mathrm{FB}}(F) - \Gamma(\mu; F) \leq \frac{1}{4}\sigma^2 + \frac{1}{8\sqrt{3}}.$$

Proof. In the proof of Theorem 1, we showed that

$$\Gamma^{\mathrm{FB}}(F) \leq 2\,\Gamma(\mu; F) - \frac{[\Gamma(\mu; F)]^2}{\mu\,(1-\mu)}.$$

This yields the following expression for welfare loss:

$$\Gamma^{\mathrm{FB}}(F) - \Gamma(\mu; F) \leq \Gamma(\mu; F) - \frac{[\Gamma(\mu; F)]^2}{\mu\,(1-\mu)}.$$

For any $\eta \in \mathbb{R}_+$, observe that

$$\left[\frac{\eta}{\sqrt{\mu\,(1-\mu)}} - \frac{\sqrt{\mu\,(1-\mu)}}{2}\right]^2 \geq 0 \implies \eta - \frac{\eta^2}{\mu\,(1-\mu)} \leq \frac{\mu\,(1-\mu)}{4}.$$

Moreover, because $z\,(1-z) \leq 1/4$ for any $z \in \mathbb{R}_+$, the maximum welfare loss under the optimal fixed-price mechanism is:

$$\Gamma^{\mathrm{FB}}(F) - \Gamma(\mu; F) \leq \frac{\mu\,(1-\mu)}{4} \leq \frac{1}{16}.$$

Finally, if F has variance σ^2 and is continuous, then Theorem 3 of Agarwal et al. (2005) shows that

$$\mu\,(1-\mu) \leq \sigma^2 + \frac{1}{2\sqrt{3}}.$$

□

A.8 Proof of Theorem 4

Theorem 4. For any distribution F satisfying Assumption 1, the gains from trade $\Gamma^{\mathrm{DS,ND}}(F)$ that the optimal dominant-strategy mechanism that also satisfies the *ex post* no-deficit condition achieves are bounded below:

$$\Gamma^{\mathrm{DS,ND}}(F) \geq \left[1 - \sqrt{1 - \frac{\Gamma^{\mathrm{FB}}(F)}{\mu\,(1-\mu)}}\right]\mu\,(1-\mu) \quad \text{for any } F \in \Delta([0,1]) \text{ with mean } \mu.$$

Proof. By Theorem 1 of Shao and Zhou (2016), the optimal mechanism must either be a fixed-price mechanism or an option mechanism with the same price p^* equal to the mean μ of the distribution F. Both mechanisms yield the same gains from trade in expectation. Hence we can restrict attention to only fixed-price mechanisms; Theorem 1 then applies. $\square$

A.9 Proof of Proposition 6

Proposition 6. Suppose that $\| F_S - F_B \|_\infty \leq \varepsilon$, where ε is sufficiently small.Then the worst-case performance of either agent's mean, μ_S or μ_B, achieves $[1/2 + \mathcal{O}(\varepsilon)]$ of the first-best gains from trade:

$$\begin{cases} \Gamma(\mu_S; F_S, F_B) \geq \left[\frac{1}{2} + \mathcal{O}(\varepsilon)\right] \Gamma^{\text{FB}}(F_S, F_B), \\ \Gamma(\mu_B; F_S, F_B) \geq \left[\frac{1}{2} + \mathcal{O}(\varepsilon)\right] \Gamma^{\text{FB}}(F_S, F_B). \end{cases}$$

Here, $\Gamma(p; F_S, F_B)$ denotes the gains from trade realized by the fixed-price mechanism p in expectation, and $\Gamma^{\text{FB}}(F_S, F_B)$ denotes the first-best gains from trade.

Proof. Fix F_S and F_B that satisfy the assumptions. We compute that

$$\begin{aligned} \Gamma(\mu_S; F_S, F_B) &= [1 - F_B(\mu_S)] \int_0^{\mu_S} F_S(x)\ \mathrm{d}x + F_S(\mu_S) \int_{\mu_S}^1 [1 - F_B(x)]\ \mathrm{d}x \\ &= [1 - F_S(\mu_S)] \int_0^{\mu_S} F_S(x)\ \mathrm{d}x + F_S(\mu_S) \int_{\mu_S}^1 [1 - F_S(x)]\ \mathrm{d}x \\ &\quad + [F_S(\mu_S) - F_B(\mu_S)] \int_0^{\mu_S} F_S(x)\ \mathrm{d}x + F_S(\mu_S) \int_{\mu_S}^1 [F_S(x) - F_B(x)]\ \mathrm{d}x \\ &= \Gamma(\mu_S; F_S, F_S) + \mathcal{O}(\varepsilon). \end{aligned}$$

On the other hand,

$$\begin{aligned} \Gamma^{\text{FB}}(F_S, F_B) &= \int_0^1 F_S(x) [1 - F_B(x)]\ \mathrm{d}x = \int_0^1 F_S(x) [1 - F_S(x)]\ \mathrm{d}x + \int_0^1 F_S(x) [F_S(x) - F_B(x)]\ \mathrm{d}x \\ &= \Gamma^{\text{FB}}(F_S, F_S) + \mathcal{O}(\varepsilon). \end{aligned}$$

Therefore:

$$\frac{\Gamma(\mu_S; F_S, F_B)}{\Gamma^{\text{FB}}(F_S, F_B)} = \frac{\Gamma(\mu_S; F_S, F_S) + \mathcal{O}(\varepsilon)}{\Gamma^{\text{FB}}(F_S, F_S) + \mathcal{O}(\varepsilon)} \geq \frac{1}{2} + \mathcal{O}(\varepsilon).$$

An identical argument holds for μ_B. $\square$

B Online Appendix: Discussion on General Asymmetry

In this Appendix, we consider the model of Section 2, but now relax the requirement that both agents' values are drawn from the same distribution. Our main argument is that, when agents are asymmetric, the Designer can no longer design informationally simple or even approximately efficient mechanisms under the restriction of fixed-price mechanisms. Therefore, despite the fact that fixed-price mechanisms may be strategically simple for agents, it may be justified from the Designer's perspective to use more complicated mechanisms.

We modify our model as follows. We denote the seller's distribution by F_S and the buyer's distribution by F_B. We assume that $F_S, F_B \in \Delta(\mathbb{R}_+)$, where F_S and F_B have finite and nonzero mean.

Accordingly, we modify the notation used previously. We denote the gains from trade realized by the fixed-price mechanism p as

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

where the expectation is taken over S and B, drawn independently from F_S and F_B respectively. Likewise, we denote the first-best gains from trade by

$$\Gamma^{\mathrm{FB}}(F_S, F_B) := \mathbb{E}[(B - S) \cdot \mathbb{1}_{B>S}].$$

The main difficulty in extending our main results to the case of generally asymmetric agents is the following impossibility result, which has been documented by other various other authors (see, *e.g.*, Blumrosen and Dobzinski, 2016):

Observation. The optimal fixed-price mechanism can achieve an arbitrarily small fraction of the first-best gains from trade:

$$\inf_{F_S, F_B \in \Delta(\mathbb{R}_+)} \sup_{p \in \mathbb{R}_+} \frac{\Gamma(p; F_S, F_B)}{\Gamma^{\mathrm{FB}}(F_S, F_B)} = 0.$$

To see why this is allowed in when agent distributions are not identical, we note that asymmetry can concentrate the joint probability density, conditional on $B > S$, close to the diagonal $B = S$, as shown in Figure 1. However, any fixed-price mechanism can capture only a small fraction of gains from trade near the diagonal. Thus, in the worst case, even the optimal fixed-price mechanism can only capture an arbitrary small fraction of the potential gains from trade.

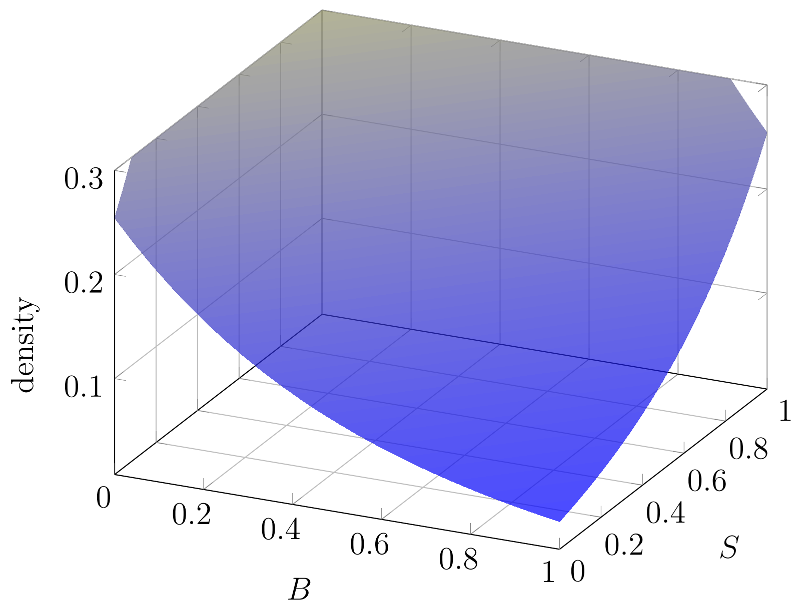

Figure 1: Joint density concentrated near the diagonal $B = S$ when conditioned on $B > S$.

There are two approaches to circumvent this impossibility result. First, we can consider the case where F_S and F_B are close to each other in the supremum norm. We refer to this case as one of "limited asymmetry." This was the approach presented in the extension in Section 5. In this Appendix, we consider the second approach, which sets a different benchmark than first-best gains from trade, namely first-best welfare. This benchmark has been used by various papers in the computer science literature, such as Colini-Baldeschi et al. (2016) and Blumrosen and Dobzinski (2016).

If we use the first-best welfare as a benchmark instead, then we are able to show strong results, especially in the case where the Designer knows the seller's distribution F_S but has complete uncertainty over the buyer's distribution F_B. However, compared to the first-best gains from trade, the first-best welfare is a less relevant benchmark for many economic purposes. The limitations of both approaches lead us to conclude that the Designer may well be justified in using more complicated mechanisms than fixed-price mechanisms when agents can be asymmetric.

We now consider the first-best welfare as our benchmark, instead of first-best gains from trade. Formally, we consider the agents' welfare

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

and the first-best welfare

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

Notably, observe that

$$\mathrm{W}(p; F_S, F_B) - \Gamma(p; F_S, F_B) = \mathrm{W}^{\mathrm{FB}}(F_S, F_B) - \Gamma^{\mathrm{FB}}(F_S, F_B) = \mathbb{E}[S].$$

Blumrosen and Dobzinski (2016) have shown that the optimal fixed-price mechanism p^* is a good approximation to first-best welfare. Precisely, they show that

$$\inf_{F_S, F_B \in \Delta(\mathbb{R}_+)} \frac{\mathrm{W}(p^*; F_S, F_B)}{\mathrm{W}^{\mathrm{FB}}(F_S, F_B)} \geq 1 - \frac{1}{e}, \quad \text{where } p^* \in \arg\max_{p \in \mathbb{R}_+} \Gamma(p; F_S, F_B).$$

Our first result shows that this bound is not tight:

Theorem 6. 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 first-best welfare. That is:

$$\inf_{F_S, F_B \in \Delta(\mathbb{R}_+)} \sup_{p \in \mathbb{R}_+} \frac{\mathrm{W}(p; F_S, F_B)}{\mathrm{W}^{\mathrm{FB}}(F_S, F_B)} \geq 1 - \frac{1}{e} + 0.0001.$$

Theorem 6 builds on the result by Blumrosen and Dobzinski (2016) as follows. Blumrosen and Dobzinski (2016) consider a mechanism that always achieves achieves the bound of $1 - 1/e$. Heuristically, our improvement can be achieved as follows:

- If $\mathbb{E}[(S - B)_+]$ is sufficiently small, then we construct a mechanism that exploits the severe asymmetry in the agents' distributions. This achieves a worst-case performance approximation to first-best welfare that is close to $3/4$.

- If $\mathbb{E}[(S - B)_+]$ is large, then the mechanism of Blumrosen and Dobzinski (2016) already achieves a worst-case performance approximation that is strictly higher than $1 - 1/e$.

In our view, an important shortcoming of Theorem 6 is that the mechanism is not informationally simple with respect to simple statistics of the distribution. In fact, it depends heavily on distributional knowledge. As a counterpoint to Theorem 6 and motivated by the mechanism proposed by Blumrosen and Dobzinski (2016), we consider mechanisms that use only distributional knowledge of the seller's distribution F_S. For this analysis, we assume that the agents' distributions are twice continuously differentiable with positive density: $F_S, F_B \in \mathcal{F}$.

Formally, we restrict the fixed-price mechanism that the Designer chooses to be independent of the buyer's distribution F_B, so that $p(F_S, F_B) = p(F_S)$. Under this restriction, $1 - 1/e$ is the best possible approximation to first-best welfare:

Theorem 7. Given distributions F_S and F_B, the Designer can achieve no better than $1 - 1/e$ of the first-best welfare if the price can only depend on F_S. That is:

$$\inf_{F_S, F_B \in \mathcal{F}} \sup_{p(F_S)} \frac{\mathrm{W}(G; F_S, F_B)}{\mathrm{W}^{\mathrm{FB}}(F_S, F_B)} = 1 - \frac{1}{e}.$$

B.1 Proof of Theorem 6

Theorem 8. 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 first-best welfare. That is:

$$\inf_{F_S, F_B \in \Delta(\mathbb{R}_+)} \sup_{p \in \mathbb{R}_+} \frac{\mathrm{W}(p; F_S, F_B)}{\mathrm{W}^{\mathrm{FB}}(F_S, F_B)} \geq 1 - \frac{1}{e} + 0.0001.$$

To prove Theorem 6, we consider first the simpler case where there is severe asymmetry between the agents, that is, when $\mathbb{E}[(S - B)_+] = 0$. This is dealt with in Section B.1.1. We show that, in this case, the Designer can achieve a good approximation to first-best welfare that is close to 0.75. We then proceed to show, in Section B.1.2, that this holds in general. We combine our results and give the complete proof of Theorem 6 in Section B.1.3.

B.1.1 Severe asymmetry: $\mathbb{E}[(S - B)_+] = 0$

Proposition 7. 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:

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

Proof. The result to Proposition 7 follows from Lemma 4, which we prove below. □

Lemma 4. Fix $\alpha > 0$. Given distributions F_S and F_B that satisfy $\mathbb{E}[(S - B)_+] \leq \alpha \cdot \mathrm{W}^{\mathrm{FB}}(F_S, F_B)$. 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

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

Remark. We note that Lemma 4 implies the result of Proposition 7 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

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

Observe that $\mathbb{P}[S > p^+] = \mathbb{P}[B < p^-] = 0 < \sqrt{10\alpha}$ and $p^+ - p^- = \sqrt{10\alpha}$, so Lemma 5 applies for any $\alpha > 0$. The result of Proposition 7 is obtained in the limit as $\alpha \to 0$.

Proof of Lemma 4. For ease of notation, given a fixed-price mechanism p, denote the welfare loss relative to first-best efficiency by $\mathrm{L}(p; F_S, F_B)$:

$$\mathrm{L}(p; F_S, F_B) := \mathrm{W}^{\mathrm{FB}}(F_S, F_B) - \mathrm{W}(p; F_S, F_B) = \mathbb{E}[(B - S) \cdot \mathbb{1}_{S < B \leq p}] + \mathbb{E}[(B - S) \cdot \mathbb{1}_{p < S < B}].$$

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

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

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

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

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

$$\mathbb{E}[B \cdot \mathbb{1}_{B\ge p^+} \mathbb{1}_{S\ge p^-}] \le (1-\beta)(1-\sigma) \cdot \mathbb{E}[B] \le \frac{1}{4} \cdot \mathbb{E}[B] \le \frac{1}{4} \cdot \mathrm{W}^{\mathrm{FB}}(F_S, F_B).$$

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

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

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

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

□

We now derive a sufficient condition for the hypotheses of Lemma 4 to hold:

Lemma 5. Fix $\alpha > 0$. Given distributions F_S and F_B that satisfy $\mathbb{E}[(S-B)_+] \le \alpha \cdot \mathrm{W}^{\mathrm{FB}}(F_S, F_B)$, suppose there exists p^* such that

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

Define

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

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

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

p^*, we have

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

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

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

□

Therefore, for small $\alpha > 0$, Lemma 5 guarantees a strict improvement over the $1 - 1/e$ bound (*i.e.*, that Lemma 4 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 6. 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

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

Proof. Define $q := \mathbb{P}[S \geq p^*] \leq 1/5$ and $r := \mathbb{E}[B \cdot \mathbb{1}_{B \leq p^*}]/\mathbb{E}[B]$. 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 $\mathrm{L}(p; F_S, F_B)$:

$$\mathrm{L}(p; F_S, F_B) := \mathrm{W}^{\mathrm{FB}}(F_S, F_B) - \mathrm{W}(p; F_S, F_B) = \mathbb{E}[(B-S) \cdot \mathbb{1}_{S<B\leq p}] + \mathbb{E}[(B-S) \cdot \mathbb{1}_{p<S<B}].$$

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

$$\begin{aligned}
\mathrm{L}(p^*; F_S, F_B) &\le \mathbb{E}[B \cdot \mathbb{1}_{B \le p^*} \mathbb{1}_{S < p^*}] + \mathbb{E}[B \cdot \mathbb{1}_{B > p^*} \mathbb{1}_{S \ge p^*}] \\
&= r\,(1-q) \cdot \mathbb{E}[B] + (1-r)\,q \cdot \mathbb{E}[B] \\
&= \left[\frac{1}{2} - 2\left(\frac{1}{2} - q\right)\left(\frac{1}{2} - r\right)\right] \mathbb{E}[B] \\
&\le \left[\frac{1}{2} - 2\left(\frac{1}{2} - \frac{1}{5}\right)^2\right] \mathbb{E}[B] = \frac{8}{25} \cdot \mathbb{E}[B] \le \frac{8}{25} \cdot \mathrm{W}^{\mathrm{FB}}(F_S, F_B).
\end{aligned}$$

Therefore

$$\mathrm{W}(p^*; F_S, F_B) = \mathrm{W}^{\mathrm{FB}}(F_S, F_B) - \mathrm{L}(p^*; F_S, F_B) \ge \frac{17}{25} \cdot \mathrm{W}^{\mathrm{FB}}(F_S, F_B).$$

□

B.1.2 General asymmetry

For the case where $\alpha = \mathbb{E}[(B - S)_+]$ is not small, we show for completeness:

Proposition 8 (Blumrosen and Dobzinski, 2016). For any given distributions F_S and F_B,

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

Proof. The proof is based almost entirely on the proof of Theorem 4.1 in Blumrosen and Dobzinski (2016). 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 $\widetilde{F}_S$ (replacing all values above b by b):

$$\widetilde{F}_S(x) = F_S(x) \cdot \mathbb{1}_{x < b} + \mathbb{1}_{x \ge b}.$$

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

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

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

By the same argument,

$$\mathrm{W}^{\mathrm{FB}}(\widetilde{F}_S, \Phi_b) = \mathrm{W}^{\mathrm{FB}}(F_S, \Phi_b) - \mathbb{P}[S > b] \cdot (\mathbb{E}_{S \sim F_S}[S \mid S > b] - b)\,.$$

Consider the distribution $G^*(x) = 1 + \log F_S(x)$, where $x \in [F_S^{-1}(1/e), F_S^{-1}(1)]$. Blumrosen and Dobzinski (2016) show that

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

Consequently, substitution reveals that

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

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

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

and

$$\mathrm{W}^{\mathrm{FB}}(F_S, F_B) = \mathbb{E}_{b \sim F_B}[\mathrm{W}^{\mathrm{FB}}(F_S, \Phi_b)].$$

Taking expectations in the above yields

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

□

B.1.3 Complete proof of Theorem 6

Proof. Let $\mathbb{E}[(S - B)_+] = \alpha \cdot \mathrm{W}^{\mathrm{FB}}$. If $\alpha \geq 0.0003$, then Proposition 8 yields

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

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 6 yields

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

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 5 shows that this is a sufficient condition to satisfy the hypotheses of Lemma 4, which implies:

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

Finally, Proposition 7 covers the case of $\alpha = 0$. □

B.2 Proof of Theorem 7

Theorem 9. Given distributions F_S and F_B, the Designer can achieve no better than $1 - 1/e$ of the first-best welfare if the price can only depend on F_S. That is:

$$\inf_{F_S, F_B \in \mathcal{F}} \sup_{p(F_S)} \frac{\mathrm{W}(G; F_S, F_B)}{\mathrm{W}^{\mathrm{FB}}(F_S, F_B)} = 1 - \frac{1}{e}.$$

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 (as a quantile of F_S) to maximize total welfare, 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.[12] This means that the first-best efficiency has value $\mathbb{E}[\max\{b, s\}] = 1$.

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

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

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, $x \cdot \mathbb{1}_{x<y}$ 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

$$\mathbb{E}[V(x, y)] = \mathbb{E}[(1 - y) + x \cdot \mathbb{1}_{x<y}] = \int_0^x (1 - y)\rho(y) \, \mathrm{d}y + \int_x^1 (1 - y + x)\rho(y) \, \mathrm{d}y. \qquad (\ddagger)$$

We posit that for an optimal Nature strategy $\rho(y)$, this quantity should be the same for every

[12] There is a lemma making this argument in Blumrosen and Dobzinski (2016) but since we are proving the opposite bound, this lemma is not formally needed here.

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 (‡) is constant for $x \in [1/e, 1]$. By differentiating (‡) with respect to x, we obtain (for $x \in [1/e, 1]$),

$$\int_x^1 \rho(y)\, \mathrm{d}y - x\, \rho(x) = 0.$$

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

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

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.[13] 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 strategy of Nature.

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 $1 - y$, $s = 1$.

[13] These arguments are not rigorous, but in any case we are just trying to guess the optimal form of Nature's strategy.

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

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

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

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

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

- $x \in [1/e, 1)$. In this range, we have $y \le x$ 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:

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

- $x = 1$. 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 $1 - 1/e$. □