Approximating Gains-from-Trade in Bilateral Trading

Liad Blumrosen* and Yehonatan Mizrahi†

Abstract

We consider the design of platforms that facilitate trade between a single seller and a single buyer. The most efficient mechanisms for such settings are complex and sometimes even intractable, and we therefore aim to design simple mechanisms that perform approximately well. We devise a mechanism that always guarantees at least $1/e$ of the optimal expected gain-from-trade for every set of distributions (assuming monotone hazard rate of the buyer's distribution). Our main mechanism is extremely simple, and achieves this approximation in Bayes-Nash equilibrium. Moreover, our mechanism approximates the optimal gain-from-trade, which is a strictly harder task than approximating efficiency. Our main impossibility result shows that no Bayes-Nash incentive compatible mechanism can achieve better approximation than $2/e$ to the optimal gain from trade. We also bound the power of Bayes-Nash incentive compatible mechanisms for approximating the expected efficiency.

1 Introduction

When we look at the global commerce landscape in the Internet era, we can see that most of the products and services are sold on platforms that involve users of different roles, usually sellers and buyers. In such environments, the "auctioneer" or the "social planner" is the platform designer and not any one of the sellers (as in classic auction settings). For example, online ads are sold via exchange markets where advertisers bid for ad slots and content providers seek to maximize profit. Another example is the recent Incentive Auctions run by the US FCC [1], where spectrum is traded between TV stations and wireless communication companies. Internet commerce giants like Amazon and eBay are essentially large-scale platforms that mitigate trade between sellers and buyers for a myriad of products, and Airbnb is a marketplace where travelers seek to purchase accommodation from various vendors. The design of such two-sided markets brings in major challenges for mechanism designers, and it has been the focus of a series of recent papers (e.g., [28, 25, 26, 11, 15, 14]).

In this paper we study the simplest two-sided market, known as the *Bilateral Trade* setting. In this setting, a single seller owns an item, and can consume it and gain a value s; a single buyer is interested in purchasing the item that can give him a value b. Since both values are private, agreeing on a price in an incentive-compatible mechanism may be hard. Indeed, the celebrated impossibility result by Myerson and Satterthwaite [22] claims that no Bayes-Nash incentive compatible mechanism can simultaneously achieve full efficiency (that is, perform a trade when $b > s$) and be *budget balanced* (BB) and *individually rational* (IR).[1] In situations where budget balance

*The Hebrew University, Jerusalem, Israel. blumrosen@gmail.com

†The Hebrew University, Jerusalem, Israel. yehonatan.mizrahi@mail.huji.ac.il

[1] A mechanism is *budget balanced* if the mechanism does not gain any profit nor requires any subsidies. A mechanism is *individually rational* if the utility of each player cannot decrease by participating in the mechanism. Formal definitions will be given later in the paper.

and individual rationality are hard constraints, one thus have to compromise and design mechanisms with approximate expected efficiency. In their original paper, Myerson and Satterthwaite [22] characterized the "second-best" mechanism, that is, the mechanism that maximizes efficiency subject to the BB and IR constraints. However, this second-best mechanism is often too complex to implement, as it involves solving a set of differential equations which is a challenging task in the bilateral-trade setting, and seems to be completely intractable when the setting is even slightly generalized. Moreover, even if one is able to implement it, determining how well this second-best mechanism performs, compared to the optimal ("first-best") efficiency, is not a trivial task.

There are two standard measures that quantify the efficiency of allocations in mechanisms. The first one is the expected *efficiency* (or social welfare), that is, the expected value of the player that obtains the item. The second measure is the expected *gain from trade* (GFT), which is the expected value of: $b - s$ when a trade happens, and 0 otherwise. While the maximal efficiency and the maximal gain-from-trade are achieved by the same allocation rule, it is clear that from an approximation perspective approximating the GFT is a harder task. Every c approximation to the gain-from-trade implies a c approximation to the expected efficiency, but the opposite does not hold (this easy observation will be discussed in the sequel of the paper). For example, think about an instance where both s and b are distributed over the support $[1, 2]$. Every mechanism clearly gains efficiency of at least 1 and of at most 2, and thus every mechanism guarantees $1/2$ approximation to the efficiency. However, designing a mechanism that attains $1/2$ of the expected GFT is completely non trivial. Approximating the GFT is a notoriously hard analytical problem, and in this paper we devise simple mechanisms that approximate this objective function.

1.1 Our results

A series of recent works compared the power of simple mechanisms and optimal (yet complex) mechanisms (e.g., [7, 19, 6, 17, 12, 8, 2, 24]). most of these results consider simple mechanisms that are *dominant-strategy* incentive compatible (DSIC). For the bilateral-trade problem, however, it was shown by Blumrosen and Dobzinski [5] that no DSIC mechanism can guarantee any constant approximation to the expected GFT. The weakness of DSIC mechanisms relates to the fact that they are essentially restricted to posting a single price to the agents, where this price cannot depend on the actual bids of the agents. In this paper, we devise a mechanism that achieves approximate efficiency in Bayes-Nash incentive compatibility (BNIC). This follows a recent line of research, mostly for combinatorial auction settings, that compared the power of simple BNIC mechanism to optimal outcomes (see, e.g., [9, 3, 27]). Our main result in this paper is a mechanism with extremely simple rules in which simple Bayes-Nash equilibrium strategies obtain a constant approximation ratio. This mechanism circumvents the DSIC limitations, and the final price may depend on the seller's value. More precisely, this mechanism admits a unique Bayes-Nash equilibrium with at least $1/e$ of the optimal ("first-best") GFT whenever the distribution of the buyer's value satisfies the monotone hazard rate (MHR) property (with no restrictions on the seller's distribution). We stress that, as we observe later in the paper, no DSIC mechanism can approximate the GFT even for distributions that satisfy the MHR condition.

Theorem 1: *When the distribution of the buyer's value satisfies the monotone hazard rate condition, there is a "simple" Bayes-Nash incentive-compatible, individually-rational and budget-balanced mechanism which always achieves at least a $\frac{1}{e}$-fraction of the optimal expected gain from trade.*

In this mechanism, the seller offers a take-it-or-leave-it price to the buyer, who then decides whether to accept it or not.[2] This mechanism is simple in several ways: first, the mechanism

[2]We note that our mechanism satisfies two stronger and desired versions of the above economic properties: it is

designer needs no distributional knowledge. The seller does need to know the distribution of the buyer in order to compute his optimal offer, but the buyer's strategy does not involve any distributional considerations. The computation required from the seller for computing her optimal offer is as complex as determining the monopoly price in the presence of a single buyer, which is known to have a simple closed-from formula and can be computed easily even in practical settings (e.g., [23]).

We note that this approximation result also implies the same approximation factor for the "second-best" mechanism.[3] That is, it follows that the expected gain-from-trade in the optimal BNIC mechanism cannot fall below a $1/e$ fraction of the optimal (first-best) gain-from-trade. Furthermore, the theorem demonstrates how this bound can be achieved even by simple, more practical, mechanisms.

We strengthen this approximation result in two respects. We first prove that the approximation ratio achieved by the mechanism is actually $\frac{1+c}{e}$, where $c \in [0, 1]$ is a constant that depends on the buyer's distribution (and more specifically, on the steepness of the virtual valuation function); for example, for the uniform distribution $c = 0.5$, so the approximation bound in this case is actually $\frac{1.5}{e} \cong 0.55$. We then prove that given a stronger condition on the buyer's distribution, namely, that the *hazard-rate is concave*, we can significantly improve the approximation bound for the GFT to $2/e \cong 0.74$.[4] We give an example for an MHR distribution with a non-concave hazard rate, for which the approximation achieved by our mechanism is strictly worse than $2/e$; therefore, the concavity assumption is necessary for the analysis of our mechanism.

Our main impossibility result in this paper shows that no Bayes-Nash incentive compatible mechanism can guarantee an approximation ratio better than $2/e$.

Theorem 2: *There is no Bayes-Nash incentive compatible, individually rational*[5] *and budget balanced mechanism that guarantees a* $\frac{2}{e}$*-fraction of the optimal expected gain from trade. Moreover, this holds even when both distributions admit the MHR property.*

Unlike the impossibility results for DSIC mechanisms ([4, 10]), there is no simple characterization for BNIC mechanisms; therefore, our proof relies on solving the complex "second-best" mechanism by [22] for carefully chosen distributions, and analyze its equilibrium properties. The buyer's distribution for which the bound is proven admits concave hazard rate, so this bound matches the above $2/e$ bound for this family of distributions.

Our final impossibility result bounds the power of BNIC mechanisms for approximating the expected efficiency (all the results described so far concerned approximating gains-from-trade). We show that no BNIC mechanism can guarantee better than a 0.93-approximation to the optimal efficiency. Although this bound appears to be weak compared to the other impossibility results, this is the strongest impossibility result for BNIC mechanisms that we are aware of. We know [5] that there are BNIC mechanisms (actually, even DSIC mechanisms) that achieve a $1 - 1/e$ approximation to the optimal efficiency. This leaves a considerable gap for BNIC mechanisms between 0.63 and 0.93.

---

strongly budget balanced, i.e., the sum of payments is always exactly zero; it is also *ex-post* individually rational, i.e., agents cannot lose in every instance and not only in expectation.

[3] Note that the characterization of the "second-best" mechanism by [22] requires that both agents have Myerson-regular distributions, while we require the stronger MHR assumption for the buyer and require nothing for the seller.

[4] Concavity of the hazard rate is satisfied by some standard distributions (e.g., exponential, Weibull(2,1), etc.), and does not hold for some other distributions (e.g, uniform on $[0, 1]$).

[5] This result considers the weaker version of interim IR, which makes the proof only harder.

1.2 More Related Work

McAfee [19] studied a similar problem to ours, i.e., how simple mechanisms can approximate the gain-from-trade in bilateral-trade settings. He proved that half of the expected gain from trade can be achieved via a DSIC mechanism for settings where the median of the buyer distribution is greater than the median of the seller's distribution. The mechanism simply posts any price between the medians as a take-it-or-leave-it offer to both agents. As mentioned, this bound cannot be generalized with DSIC mechanisms for general distributions [5], or even to MHR distributions. We overcome this impossibility by relaxing the incentive constraints from DSIC to BNIC. The Bilateral Trade problem for non quasi-linear settings was recently studied in [16].

Blumrosen and Dobzinski [4, 5] designed simple DSIC mechanisms that approximate the expected efficiency for Bilateral trade and more complex settings. [4, 5] were inspired by McAfee's work and used the medians of the distributions as a major design tool. [4, 5] showed how features that are used in mechanisms for Bilateral Trade can be used in more general exchange frameworks, and even constructed black-box reductions from other settings to Bilateral Trade. This highlights the importance of understanding the basic bilateral-trade problem for the design of more complex markets. Colini-Baldeschi et al. [10] further studied approximation mechanisms in exchange settings under strong budget balance, and proved, among other results, an impossibility result of 0.749 for the efficiency approximation obtained by DSIC mechanisms in the bilateral trade problem.

Two-sided markets have been extensively studied in the last three decades. McAfee [18] designed an elegant DSIC, BB and IR mechanism for two sided markets with homogenous goods, which is nearly efficient in large markets. Other work about asymptotic efficiency of two-sided markets include [25, 26, 11, 15]. Dutting et al. [14] developed a modular approach for the design of two-sided markets, based on the deferred-acceptance heuristics from [20].

We continue as follows: We present the model and a a brief survey of some relevant existing results in Section 2. Our main positive results are given in Section 3, and our negative results appear in Section 4.

2 Model

The bilateral trade problem involves two agents, a seller and a buyer. The seller owns one indivisible item from which he gains a value s. The buyer gains a value b from the same item after purchasing it. In fact, s and b are drawn from two independent distributions F_s and F_b which correspond to the two random variables S and B respectively. Each of the two agents does not know the realization of the other agent's value, but the distributions are public knowledge. In our analysis we shall assume the existence of the density functions f_s and f_b for the seller and the buyer respectively. Furthermore, we assume that both agents are risk neutral and that the prices and values are commensurable.

Based on their values, the seller and the buyer simultaneously report their bids, denoted by $\sigma(s)$ and $\beta(b)$ respectively, to the trading mechanism. The mechanism is defined by the two functions $t(\beta,\sigma)$ and $p(\beta,\sigma)$, both known to the agents, such that the item is transferred from the seller to the buyer at price $t(\beta,\sigma)$ with probability $p(\beta,\sigma)$. We will be focusing on deterministic mechanisms, such that the item is transferred iff $p(\beta,\sigma)=1$.

As previously mentioned, the two main measures that will be analyzed throughout this paper are the expected *gains from trade* and the expected *efficiency*. Given a mechanism $M=\langle t,p\rangle$ and two agents with distributions F_b and F_s, these two measures, denoted by $GFT_M^{F_b,F_s}$ and $EFF_M^{F_b,F_s}$

respectively, are defined as follows (when they are clear, the notations M,F_s or F_b are omitted):

$$GFT_M^{F_b,F_s} = E[(B-S)\cdot p(\beta(B),\sigma(S))]$$
$$EFF_M^{F_b,F_s} = E[B\cdot p(\beta(B),\sigma(S)) + S\cdot(1-p(\beta(B),\sigma(S)))]$$

From these definitions it becomes clear that $EFF_M^{F_b,F_s} = GFT_M^{F_b,F_s} + E[S]$. In the fully efficient case (i.e., when $p(\beta(b),\sigma(s)) = 1$ iff $b \geq s$), the measures are $GFT_{OPT}^{F_b,F_s} = E[\max\{B-S,0\}]$ and $EFF_{OPT}^{F_b,F_s} = E[\max\{B,S\}]$. We note that, by definition, maximizing GFT also implies maximizing efficiency. The fully efficient allocation is our benchmark for our approximation results; we say that for a pair of such distributions, a mechanism M achieves a k-approximation to the optimal GFT if $\frac{GFT_M^{F_b,F_s}}{GFT_{OPT}^{F_b,F_s}} \geq k$ and similarly for EFF, and we note that it always holds that $\frac{EFF_M^{F_b,F_s}}{EFF_{OPT}^{F_b,F_s}} \geq \frac{GFT_M^{F_b,F_s}}{GFT_{OPT}^{F_b,F_s}}$.[6]

We proceed with further preliminaries needed for our main results.

2.1 The Hazard Rate of a Distribution

We now present some definitions, properties and notations regarding the *Hazard Rate* of a general distribution F with density f that has a non-negative support. These are used in our main approximation results in the next section.

We begin by defining the *Hazard Rate* of such distribution by $h(x) = \frac{f(x)}{1-F(x)}$. The *Cumulative Hazard Function* of F is defined by $H(x) = -ln(1-F(x))$ for every $x \geq 0$ (which is not to the right of F's support). We note that $e^{-H(x)} = 1 - F(x)$, and that $H(0) = 0$. Differentiating yields $H'(x) = h(x)$, and we get that $H(x) = \int_0^x h(t)\,dt + k$ for some k. Placing $x = 0$ shows that $k = 0$.

We continue by defining the *Virtual Valuation Function* of an agent with such distribution by $\varphi(x) = x - \frac{1-F(x)}{f(x)}$.

Moreover, we also define the *Monotone Hazard Rate (MHR)* property of a distribution, which simply states that h is monotone non-decreasing. This property also implies that φ is monotone increasing, a state in which we often call F a *regular distribution*[7]. We note that in this case, since φ is strictly monotone, its inverse function exists.

In this paper, we only require such hazard rate assumptions for the buyer's distribution, and therefore when we use these notations they shall be associated with F_b.

2.2 Bayes-Nash IC: The Second-Best Mechanism

While it was proved in [22] that no IR and BB mechanism is fully efficient in BNIC, the same paper present a characterization of the mechanisms that maximize GFT subject to the IR and BB constraints. We will now describe this "second-best" mechanism for bilateral trade from [22], which is used later in our inapproximability results.

As stated [22], in order to derive the correct approximation results using this mechanism, we need to assume that the support of F_b is $[\underline{b},\overline{b}]$ or $[\underline{b},\infty)$ for some $\overline{b} \geq \underline{b} \geq 0$ and that the support of F_s is $[\underline{s},\overline{s}]$ or $[\underline{s},\infty)$ for some $\overline{s} \geq \underline{s} \geq 0$. As in [22], we assume regularity of the distributions, i.e., that the functions $b - \frac{1-F_b(b)}{f_b(b)}$ and $s + \frac{F_s(s)}{f_s(s)}$ are monotone increasing. Using the fact that this

[6]This follows from $EFF_M \cdot GFT_{OPT} \geq GFT_M \cdot EFF_{OPT}$ which is equivalent by definition to the inequality $(GFT_M + E[S]) \cdot GFT_{OPT} \geq GFT_M \cdot (GFT_{OPT} + E[S])$ that holds by $GFT_{OPT} \geq GFT_M$.

[7]Most of the literature assumes a weaker condition, that the φ is non-decreasing. In our paper we often use the inverse function of φ, and the notations become much simpler when φ is strictly increasing. Moreover, our main results consider MHR distributions that imply that φ is always strictly increasing.

mechanism is truthful, i.e., in a Bayes-Nash equilibrium $\beta(b) = b$ and $\sigma(s) = s$, the mechanism is defined by:

$$p^{\alpha}(\beta(b), \sigma(s)) = \begin{cases} 1 & if \quad s + \alpha \cdot \frac{F_s(s)}{f_s(s)} \leq b - \alpha \cdot \frac{1-F_b(b)}{f_b(b)} \\ 0 & \text{otherwise.} \end{cases}$$

The appropriate parameter is the unique (as proved in [22]) $\alpha \in (0,1]$ that solves the following equation, presented for the bounded supports case (and similar for the unbounded case):

$$\int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} \left(\left(b - \frac{1-F_b(b)}{f_b(b)} \right) - \left(s + \frac{F_s(s)}{f_s(s)} \right) \right) \cdot p^{\alpha}(b,s) f_b(b) f_s(s) \, ds db = 0$$

The appropriate payment function can be determined ad hoc, given the distributions. Nonetheless, we note that it is not necessary in order to analyze the GFT and EFF measures.

Finally, we denote this mechanism by the shorthand MS.

3 A Constant Approximation for the Gains from Trade

In this section we present a simple mechanism that approximates the optimal gains from trade for bilateral trade settings. The mechanism has no dominant-strategy equilibrium, and the results are achieved in Bayes-Nash equilibrium. We now define this mechanism, we call *Seller-Offering Mechanism* (abbreviated as SO).

The Seller-Offering (SO) mechanism:

- *The seller offers a take-it-or-leave-it price t to the buyer, who chooses whether to accept it or not.*
- *If the buyer accepts the price, a trade occurs at price t. Otherwise, no trade occurs and no payments are transferred.*

We note that at first glance, it seems as if this mechanism does not fall into formal model of bilateral trade mechanisms we defined earlier, since it is two-staged and not simultaneous. However, using $p(\beta, \sigma) = 1_{\{\beta \geq \sigma\}}(\beta, \sigma)$ and $t(\beta, \sigma) = \sigma$ in the original scheme yields the same results.

3.1 Some Technical Definitions and Observations

For our results in this section, it suffices to assume that the support of F_b is $[\underline{b}, \overline{b}]$ or $[\underline{b}, \infty)$ for some $\underline{b} \geq 0$, the support of F_s is contained in $[0, \infty)$, f_b is differentiable and F_b adheres to the MHR assumption.

The inverse virtual valuation $\varphi^{-1}(\cdot)$ turns out to be very useful in our analysis. This inverse function is not well defined for all possible values, therefore we frequently use its extension denoted by $\overline{\varphi^{-1}}(\cdot)$.

Definition 3.1. *Under the aforementioned assumptions, we define the Extended Inverse Virtual Valuation Function, $\overline{\varphi^{-1}}(x)$, to be the continuous extension of $\varphi^{-1}(x)$:*

Since $\varphi(x)$ is increasing, $\varphi^{-1}(x)$ is undefined for $x \leq \varphi(\underline{b})$, and in case F_b's support is $[\underline{b}, \overline{b}]$, it is also undefined for $x \geq \varphi(\overline{b}) = \overline{b}$. The left part is extended using $\overline{\varphi^{-1}}(x) = \underline{b}$ and the right part using $\overline{\varphi^{-1}}(x) = x$.[8]

We continue by showing some useful technical observations regarding these functions, used later in our proofs:

Observation 3.2. *For every x in their domain, it holds that $\varphi(x) \leq x$ and $\overline{\varphi^{-1}}(x) \geq x$.*

Proof. The first inequality follows from $\varphi(x) = x - \frac{1}{h(x)} \leq x$ since h is positive. The second follows from the fact that φ^{-1} is the reflection of φ with respect to the line y=x, and since the extension of it preserves the inequality. □

Observation 3.3. *For every $x \geq \varphi(\underline{b})$:*

1. *If $\overline{b} \geq x$ then $\overline{\varphi^{-1}}(x) - x = \frac{1-F_b\left(\varphi^{-1}(x)\right)}{f_b(\varphi^{-1}(x))} = \frac{1}{h(\varphi^{-1}(x))}$.*
2. *If $\overline{b} \leq x$ then $\overline{\varphi^{-1}}(x) - x = 0$.*

Proof. For the first case, it holds that $\overline{\varphi^{-1}}(x) - x = \varphi^{-1}(x) - \varphi\left(\varphi^{-1}(x)\right) = \varphi^{-1}(x) - \varphi^{-1}(x) + \frac{1-F_b\left(\varphi^{-1}(x)\right)}{f_b(\varphi^{-1}(x))} = \frac{1-F_b\left(\varphi^{-1}(x)\right)}{f_b(\varphi^{-1}(x))} = \frac{1}{h(\varphi^{-1}(x))}$ by the definitions of φ and h. For the second case, by the definition of the right extension of φ^{-1}, it holds that $\overline{\varphi^{-1}}(x) - x = x - x = 0$. □

Observation 3.4. *For every $\overline{b} \geq x \geq \varphi(\underline{b})$ it holds that $\frac{d\varphi^{-1}(x)}{dx} = \frac{1}{1+\frac{h'\left(\varphi^{-1}(x)\right)}{\left(h\left(\varphi^{-1}(x)\right)\right)^2}} \in [0,1]$ under the MHR assumption.*

Proof. We remind that $\varphi(x) = x - \frac{1}{h(x)}$. By the reciprocal rule, differentiating yields $\varphi'(x) = 1 - \frac{0-h'(x)}{(h(x))^2} = 1 + \frac{h'(x)}{(h(x))^2}$. Furthermore, $\frac{d\varphi^{-1}(x)}{dx} = \frac{1}{\varphi'(\varphi^{-1}(x))}$ by the derivative of an inverse function. Thus, the identity follows by plugging $\varphi^{-1}(x)$ in the derviative. We also note that by the MHR assumption, $\frac{h'\left(\varphi^{-1}(x)\right)}{(h(\varphi^{-1}(x)))^2} \geq 0$, hence $\frac{d\varphi^{-1}(x)}{dx} \in [0,1]$. □

3.2 Analysis of the Seller-Offering Mechanism

Although not admitting a dominant-strategy equilibrium, the above Seller-Offering mechanism induces quite straightforward Bayes-Nash equilibrium strategies for the agents. In equilibrium, the seller offers the monopoly price given his own value for the item, that is, $\overline{\varphi^{-1}}(s)$ (as in [21]), and the bidder will simply bid truthfully to accept the deal if its value exceeds the offered price.[9] This is an immediate application of Myerson's theory ([21]), but for completeness, a proof is given in Appendix A.1.

Proposition 3.5. *For every MHR distribution F_b for the buyer and every distribution F_s for the seller, the bids $\beta(b) = b$ and $\sigma(s) = \overline{\varphi^{-1}}(s)$ form a Bayes-Nash equilibrium in the Seller-Offering Mechanism.*

[8] We note that $\overline{\varphi^{-1}}(x)$ is defined for every x, even when F_b's support is $[\underline{b}, \infty)$, since $\varphi(x)$ is unbounded from above in that case. This can be seen by noticing that for every $y \in \mathbb{R}$, choosing $x > max\{\frac{1}{h(\underline{b})} + y + 1, \underline{b}\}$ yields $\varphi(x) = x - \frac{1}{h(x)} \geq \frac{1}{h(\underline{b})} + y + 1 - \frac{1}{h(\underline{b})} > y$ by the MHR assumption.

[9] Recall that φ denotes the virtual valuation of the buyer, and the seller use the details of this distribution to determine what price to post.

In the following lemma we present a convenient representation of GFT_{OPT} and GFT_{SO} which will be used in the main theorems. The representation of GFT_{OPT} is also shown in [19]. We prove this lemma in Appendix A.2.

Lemma 3.6. *For every MHR distribution F_b for the buyer and every distribution F_s for the seller, the following equalities hold:*

$$GFT_{OPT} = \int_0^\infty F_s(s) \cdot (1 - F_b(s))\, ds$$

$$GFT_{SO} = \int_0^\infty F_s(s) \cdot \left(1 + \frac{d\overline{\varphi^{-1}}(s)}{ds}\right) \cdot (1 - F_b(\overline{\varphi^{-1}}(s)))ds$$

We now turn to proving the main result of the paper, concerning the constant approximation guarantee obtained using the Seller-Offering mechanism. This approximation result is parameterized by a parameter c that describes the steepness of the buyer's virtual function.

Definition 3.7. *We define the Virtual Steepness Constant of an MHR distribution F with a differentiable density f by $c = min_s \frac{d\varphi^{-1}(s)}{ds}$. We note that c is in fact the reciprocal of the virtual valuation function's Lipschitz constant, since $min_s \frac{d\varphi^{-1}(s)}{ds} = min_s \frac{1}{\varphi'(\varphi^{-1}(s))} = \frac{1}{max_s \varphi'(s)}$.*

Our theorem shows that given the MHR condition on the buyer's valuation, our mechanism attains a $\frac{1+c}{e}$ fraction of the optimal gains-from-trade. Since by Observation 3.4 we have that $c \in [0, 1]$, this approximation is at least $\frac{1}{e}$ for all possible distributions.

Theorem 3.8. *For every MHR distribution F_b for the buyer and every distribution F_s for the seller, the Seller-Offering Mechanism obtains a $\frac{1+c}{e}$-approximation to the optimal gains from trade.*

Proof. We remind that in Lemma 3.6, we concluded that $GFT_{OPT} = \int_0^\infty F_s(s) \cdot (1 - F_b(s))\, ds$ and that $GFT_{SO} = \int_0^\infty F_s(s) \cdot (1 + \frac{d\overline{\varphi^{-1}}(s)}{ds}) \cdot (1 - F_b(\overline{\varphi^{-1}}(s)))ds$. We therefore analyze the relation between $(1 + \frac{d\overline{\varphi^{-1}}(s)}{ds}) \cdot (1 - F_b(\overline{\varphi^{-1}}(s)))$ and $(1 - F_b(s))$ for every $s \geq 0$.

If $s \geq \overline{b}$, both terms are 0 (we use Observation 3.2 for the first term). If $s \leq \varphi(\underline{b})$ then $(1 + \frac{d\overline{\varphi^{-1}}(s)}{ds}) \cdot (1 - F_b(\overline{\varphi^{-1}}(s))) = (1 + 0) \cdot (1 - F_b(\underline{b})) = 1 = (1 - F_b(s))$. The first equality follows from $\overline{\varphi^{-1}}(s) = \underline{b}$ for such s, and the last equality follows from $\underline{b} \geq \varphi(\underline{b})$ as noted in Observation 3.2. The ratio between these two terms is 1, which is greater than $\frac{1+c}{e}$.

We now focus on the case where $\overline{b} \geq s \geq max\{0, \varphi(\underline{b})\}$, such that $\overline{\varphi^{-1}}(s) = \varphi^{-1}(s)$, and we show that $e \cdot (1 - F_b(\overline{\varphi^{-1}}(s))) \geq 1 - F_b(s)$. The Cumulative Hazard Function H of the buyer is the integral of the monotone increasing function h, hence H is convex. Therefore, the line tangent to H at any point is below the function. In other words, fixing $x_0 \in [0, \overline{b}]$, for every $x \in [0, \overline{b}]$ it holds that $H(x) \geq H(x_0) + h(x_0) \cdot (x - x_0)$. By Observation 3.3, choosing $x_0 = \varphi^{-1}(s)$ we get that for every $x \in [0, \overline{b}]$, and specifically $x = s$, it holds that:

$$\begin{aligned} H(x) &\geq H(\varphi^{-1}(x)) + h(\varphi^{-1}(x)) \cdot (x - \varphi^{-1}(x)) = \\ &= H(\varphi^{-1}(x)) + h(\varphi^{-1}(x)) \cdot \left(-\frac{1}{h(\varphi^{-1}(x))}\right) = H(\varphi^{-1}(x)) - 1 \end{aligned}$$

Hence:

$$1 - F_b(s) = e^{-H(s)} \leq e^{-H(\varphi^{-1}(s))+1} = e \cdot e^{-H(\varphi^{-1}(s))} = e \cdot (1 - F_b(\overline{\varphi^{-1}}(s)))$$

The first and the last equalities in the last line follow from the definition of H as described in Section 2.1.

Concluding, we get that for every $s \geq 0$ it holds that

$$F_s(s) \cdot (1 + \frac{d\overline{\varphi^{-1}}(s)}{ds}) \cdot (1 - F_b(\overline{\varphi^{-1}}(s))) \geq F_s(s) \cdot \frac{1+c}{e} \cdot (1 - F_b(s))$$

Integrating both parts and by the monotonicity of the integral, we get that by Lemma 3.6

$$GFT_{SO} \geq \frac{1+c}{e} \cdot GFT_{OPT}$$

. ☐

We now proceed to proving an amplified version of this theorem. In the proof of Theorem 3.8 we relied on a linear approximation of H. The next theorem utilizes a quadratic approximation of H (via the Taylor expansion) to improve the bound, but requires an additional assumption, the concavity of h. With this additional assumption, the approximation can be improved to $2/e$.

The following technical lemma, which is proved in Appendix A.3, manifests the importance of the concavity assumption.

Lemma 3.9. *Let f be a twice differentiable function that has a concave derivative, and let $T(x)$ be a second degree Taylor polynomial at x_0, i.e., a quadratic approximation at this point. Then for every $x \leq x_0$ it holds that $T(x) \leq f(x)$.*

Theorem 3.10. *For every MHR distribution F_b with a concave hazard rate for the buyer and every distribution F_s for the seller, the Seller-Offering Mechanism obtains a $\frac{2}{e}$-approximation to the optimal gains from trade.*

Proof. We first note that since h is concave, it must hold that that $\underline{b} = 0$.[10]

Now, let $s \in [0, \overline{b}]$. By our assumptions, H is twice differentiable and the derivative of H is concave, thus by Lemma 3.9, if we take the second degree Taylor polynomial $T(x)$ (instead of the linear approximation as in Theorem 3.8) in $x_0 = \varphi^{-1}(s)$ (which is defined since $\underline{b} = 0$), for every $x \in [0, x_0]$, and specifically $x = s$, due to similar considerations used in Theorem 3.8:

$$\begin{aligned} H(s) &\geq T(s) \\ &= H\left(\varphi^{-1}(s)\right) + h\left(\varphi^{-1}(s)\right) \cdot \left(s - \varphi^{-1}(s)\right) + \frac{1}{2} \cdot h'\left(\varphi^{-1}(s)\right) \cdot \left(s - \varphi^{-1}(s)\right)^2 \\ &= H\left(\varphi^{-1}(s)\right) - 1 + \frac{1}{2} \cdot \frac{h'\left(\varphi^{-1}(s)\right)}{\left(h\left(\varphi^{-1}(s)\right)\right)^2} \end{aligned}$$

This time, as we deduced in Theorem 3.8, we get that:

$$1 - F_b(s) \leq e \cdot (1 - F_b(\overline{\varphi^{-1}}(s))) \cdot e^{-\frac{1}{2} \cdot \frac{h'(\varphi^{-1}(s))}{(h(\varphi^{-1}(s)))^2}}$$

[10]If $\underline{b} > 0$, the line connecting the points $\left(\frac{1}{2}\underline{b}, h\left(\frac{1}{2}\underline{b}\right)\right) = \left(\frac{1}{2}\underline{b}, 0\right)$ and $(\underline{b}, h(\underline{b})) = (\underline{b}, f(\underline{b}))$ on the graph of h is above the point $\left(\frac{3}{4}\underline{b}, h\left(\frac{3}{4}\underline{b}\right)\right) = \left(\frac{3}{4}\underline{b}, 0\right)$ contradicting h's concavity. For the sake of clearness we focus on this assumption, but inspecting the details of this proof shows that it can be generalized to the case where h is concave only in the union of F_b's and F_s's supports except the part which is greater than $\overline{b}$.

Therefore, it suffices to show that $\left(1+\frac{d\overline{\varphi^{-1}}(s)}{ds}\right)\cdot e^{\frac{1}{2}\cdot\frac{h'\left(\varphi^{-1}(s)\right)}{\left(h\left(\varphi^{-1}(s)\right)\right)^2}}\geq 2$. We remind that $\overline{\varphi^{-1}}\left(s\right)=\varphi^{-1}\left(s\right)$ for such s and that by Observation 3.4, it holds that $\frac{d\varphi^{-1}(s)}{ds}=\frac{1}{1+\frac{h'\left(\varphi^{-1}(s)\right)}{\left(h\left(\varphi^{-1}(s)\right)\right)^2}}$. Hence, denoting $x=h'\left(\varphi^{-1}\left(s\right)\right)$ and $y=h\left(\varphi^{-1}\left(s\right)\right)$:

$$\begin{aligned}\left(1+\frac{d\overline{\varphi^{-1}}\left(s\right)}{ds}\right)\cdot e^{\frac{1}{2}\cdot\frac{h'\left(\varphi^{-1}(s)\right)}{\left(h\left(\varphi^{-1}(s)\right)\right)^2}} &\geq \left(1+\frac{1}{1+\frac{x}{y^2}}\right)\cdot\left(1+\frac{1}{2}\cdot\frac{x}{y^2}\right)\\ &=\frac{2y^2+x}{y^2+x}\cdot\frac{2y^2+x}{2y^2}\\ &\geq\frac{\left(2y^2+x\right)^2}{\left(y^2+x\right)\cdot 2y^2+0.5x^2}=2\end{aligned}$$

In the first inequality we used the fact that $e^x\geq 1+x$ for every x.

It follows that for every $s\in[0,\bar{b}]$ it holds that

$$F_s(s)\cdot(1+\frac{d\overline{\varphi^{-1}}(s)}{ds})\cdot(1-F_b(\overline{\varphi^{-1}}(s)))\geq F_s(s)\cdot\frac{2}{e}\cdot(1-F_b(s))$$

, while for $s\geq\bar{b}$ both terms are 0. Integrating both parts, by the monotonicity of the integral, and by Lemma 3.6, we get that $GFT_{SO}\geq\frac{2}{e}\cdot GFT_{OPT}$. □

We can now use Theorem 3.10 to separate the power of DSIC and BNIC mechanisms in terms of approximating the gains from trade. The following proposition shows that there are instances where no DSIC mechanism can obtain a constant approximation to the gains from trade, but as the relevant distributions satisfy MHR and admit a concave hazard function, Theorem 3.10 implies the existence of BNIC mechanisms with $\frac{2}{e}$ approximation.

Proposition 3.11. *There exists a pair of distributions F_b for the buyer and F_s for the seller, for which no DSIC mechanism that is IR and BB can achieve a constant approximation to the optimal gains from trade, while there exists a BNIC mechanism that is IR and BB that does achieve a $2/e$-approximation to the optimal gains from trade for them.*

Proof. Consider the two distributions $F_b\sim \mathit{Exponential}\left(1\right)$ and $F_s\left(x\right)=\lambda\left(e^{x-t}-e^{-t}\right)$ with $\lambda=\frac{1}{1-e^{-t}}$ on the support $[0,t]$. In [5], Blumrosen and Dobzinski analyze the scenario in which $F_b\left(x\right)=\lambda\left(1-e^{-x}\right)$ on the support $[0,t]$ and F_s is the same as above. They show that every fixed price mechanism achieves at most $O\left(1/t\right)$-approximation to the optimal gains from trade in this case. By taking t that tends to infinity, this buyer's distribution converges to Exponential(1) while $1/t$ converges to 0. Alternatively, a direct calculation using the original distributions yields these results. Since it is well known that every DSIC mechanism that is IR and BB is a fixed price mechanism (see, e.g., [13, 10] and the references therein), the first part follows.

We note that in this case, $h\left(x\right)=1$ which is a constant function and therefore the MHR and concavity assumptions hold. Thus, by Theorem 3.10 the Seller-Offering Mechanism indeed obtains a $2/e$-approximation to the optimal gains from trade. □

Lastly, the following proposition signifies the necessity of h's concavity assumption for Theorem 3.10. We also show that the analysis of Theorem 3.10 is tight, and for some distributions (that satisfy MHR and concave hazard rate) our mechanism achieves exactly $2/e$ approximation. A proof can be found in Appendix A.4.

Proposition 3.12. *Using the Seller-Offering Mechanism:*

1. *There exists an MHR distribution F_b with a non-concave hazard rate, and a distribution F_s for the seller, such that the mechanism achieves an approximation to the optimal GFT which is strictly worse than $2/e$.*

2. *There exists an MHR distribution F_b with a concave hazard rate h and a distribution F_s for the seller, such that h is concave and $\frac{GFT_{SO}}{GFT_{OPT}} = \frac{2}{e}$.*

4 Inapproximability Results

In this section, we present impossibility results for approximating the gains from trade and efficiency using BNIC mechanism.

In the previous section, we presented an IR, BB and BNIC mechanism that guarantees a $1/e$-approximation to the optimal gains from trade for any pair of distributions under standard MHR assumptions. A question that naturally arises concerns the limitations of BNIC mechanisms in this our setting. The following theorem addresses that question and shows that no BNIC mechanism can maintain IR and BB and guarantee more than $2/e$ approximation. Moreover, this holds even when the distributions satisfy the MHR condition.[11] We also note that this result is proven for the case where the buyer's distribution has concave hazard rate, and thus it matches the positive result in Theorem 3.10 when this condition is satisfied.

Theorem 4.1. *No BNIC mechanism which is IR and BB can guarantee an approximation to the optimal gains from trade which is better than $2/e$. This holds even if both distributions satisfy the MHR condition.*

Proof. The proof relies on the Second-Best mechanism devised by Myerson and Satterthwaite in [22]. We show that for every $\epsilon > 0$ there exists a pair of distributions such that $\frac{GFT_{MS}}{GFT_{OPT}} < \frac{2}{e} + \epsilon$. Since by its definition, no BNIC mechanism which is IR and BB can achieve a better approximation than this mechanism for these distributions, the claim follows. In fact, the relevant distributions are exactly the ones used in Proposition 3.11, i.e., $F_b \sim \textit{Exponential}\,(1)$ and $F_s\,(x) = \lambda\left(e^{x-t} - e^{-t}\right)$ with $\lambda = \frac{1}{1-e^{-t}}$ on the support $[0,t]$. The distributions satisfy the MHR property. We remind that the second-best solution requires that $b - \frac{1-F_b(b)}{f_b(b)}$ and $s + \frac{F_s(s)}{f_s(s)}$ are monotone increasing, and indeed this property holds for $b - \frac{1-F_b(b)}{f_b(b)} = b - 1$ and $s + \frac{F_s(s)}{f_s(s)} = s + 1 - e^{-s}$.

As noted in Section 2, trade occurs whenever:

$$b - s \geq \alpha \cdot \left(\frac{F_s\,(s)}{f_s\,(s)} + \frac{1 - F_b\,(b)}{f_b\,(b)}\right) = \alpha \cdot \left(1 - e^{-s} + 1\right) \;\Leftrightarrow\; b \geq s + \alpha \cdot \left(2 - e^{-s}\right)$$

In that section, we also mentioned that the proper α is the one that solves:

$$\int_0^t \int_{s+\alpha\cdot(2-e^{-s})}^{\infty} \left((b-s) - \left(2 - e^{-s}\right)\right) \cdot e^{-b} \cdot \frac{e^{s-t}}{1-e^{-t}} dbds = 0$$

We now show that $\alpha^*\,(t) \underset{t\to\infty}{\longrightarrow} \frac{1}{2}$. One can verify that the last term tends to 0 as t tends to infinity. Thus, to find the proper α we instead seek for $\alpha = \alpha^*\,(t)$ for which it holds that:

$$\frac{\int_0^t \int_{s+\alpha\cdot(2-e^{-s})}^{\infty} (b-s) \cdot e^{-b} \cdot \frac{e^{s-t}}{1-e^{-t}} dbds}{\int_0^t \int_{s+\alpha\cdot(2-e^{-s})}^{\infty} (2-e^{-s}) \cdot e^{-b} \cdot \frac{e^{s-t}}{1-e^{-t}} dbds} \underset{t\to\infty}{\longrightarrow} 1$$

[11] This theorem holds for a weaker notion of interim individual rationality (as in [22]); This clearly strengthens the result.

Elementary integration shows that:

$$\int_{s+\alpha\cdot(2-e^{-s})}^{\infty}(b-s)\cdot e^{-b}\cdot\frac{e^{s-t}}{1-e^{-t}}db=\frac{e^{-\left(s+\alpha\left(2-e^{-s}\right)\right)}\left((1+2\alpha)\,e^{s}-\alpha\right)}{e^{t}-1}$$

and that:

$$\int_{s+\alpha\cdot(2-e^{-s})}^{\infty}\left(2-e^{-s}\right)\cdot e^{-b}\cdot\frac{e^{s-t}}{1-e^{-t}}db=\frac{e^{-\left(s+\alpha\left(2-e^{-s}\right)\right)}\left(2e^{s}-1\right)}{e^{t}-1}$$

Calculating the reduced integrals, for the numerator it holds that:

$$\begin{aligned}Numerator\,(\alpha,t)\cdot\left(e^{t}-1\right)&=\int_{0}^{t}e^{-\left(s+\alpha\left(2-e^{-s}\right)\right)}\left((1+2\alpha)\,e^{s}-\alpha\right)ds\underbrace{=}_{x=e^{s}}\\&=\int_{1}^{e^{t}}\frac{e^{-2\alpha+\frac{\alpha}{x}}\left((1+2\alpha)\,x-\alpha\right)}{x^{2}}dx=\\&=e^{-2\alpha}\cdot\left((1+2\alpha)\int_{1}^{e^{t}}\frac{e^{\frac{\alpha}{x}}}{x}dx-\alpha\int_{1}^{e^{t}}\frac{e^{\frac{\alpha}{x}}}{x^{2}}dx\right)\underbrace{=}_{t=\frac{\alpha}{x}}\\&=e^{-2\alpha}\cdot\left((1+2\alpha)\int_{1}^{e^{t}}\frac{e^{\frac{\alpha}{x}}}{x}dx+\alpha\cdot\frac{1}{\alpha}\int_{\alpha}^{\alpha/e^{t}}e^{t}dt\right)=\\&=e^{-2\alpha}\cdot\left((1+2\alpha)\int_{1}^{e^{t}}\frac{e^{\frac{\alpha}{x}}}{x}dx+\left(e^{\frac{\alpha}{e^{t}}}-e^{\alpha}\right)\right)\end{aligned}$$

While for the denominator it holds that:

$$\begin{aligned}Denominator\,(\alpha,t)\cdot\left(e^{t}-1\right)&=\int_{0}^{t}e^{-\left(s+\alpha\left(2-e^{-s}\right)\right)}\left(2e^{s}-1\right)ds\underbrace{=}_{x=e^{s}}\\&=\int_{1}^{e^{t}}\frac{e^{-2\alpha+\frac{\alpha}{x}}\left(2x-1\right)}{x^{2}}dx=\\&=e^{-2\alpha}\cdot\left(2\int_{1}^{e^{t}}\frac{e^{\frac{\alpha}{x}}}{x}dx-\int_{1}^{e^{t}}\frac{e^{\frac{\alpha}{x}}}{x^{2}}dx\right)=\\&=e^{-2\alpha}\cdot\left(2\int_{1}^{e^{t}}\frac{e^{\frac{\alpha}{x}}}{x}dx+\frac{1}{\alpha}\left(e^{\frac{\alpha}{e^{t}}}-e^{\alpha}\right)\right)\end{aligned}$$

If we denote $q\,(\alpha,t)=\int_{1}^{e^{t}}\frac{e^{\frac{\alpha}{x}}}{x}dx$, since $q\,(\alpha,t)\geq\int_{1}^{t}\frac{1}{x}dx$, by the comparison test it follows that $q\,(\alpha,t)\underset{t\to\infty}{\longrightarrow}\infty$.

Thus, assuming $\alpha^{*}\,(t)\underset{t\to\infty}{\longrightarrow}x$ (we see this assumption holds by the next consideration), we get that this ratio equals:

$$\begin{aligned}\frac{\int_{0}^{t}\frac{e^{-\left(s+\alpha^{*}(t)\left(2-e^{-s}\right)\right)}\left((1+2\alpha^{*}(t))e^{s}-\alpha^{*}(t)\right)}{e^{t}-1}ds}{\int_{0}^{t}\frac{e^{-\left(s+\alpha^{*}(t)(2-e^{-s})\right)}(2e^{s}-1)}{e^{t}-1}ds}&=\frac{(1+2\alpha^{*}\,(t))\cdot q\,(\alpha^{*}\,(t)\,,t)+\left(e^{\frac{\alpha^{*}(t)}{e^{t}}}-e^{\alpha^{*}(t)}\right)}{2\cdot q\,(\alpha^{*}\,(t)\,,t)+\frac{1}{\alpha^{*}(t)}\left(e^{\frac{\alpha^{*}(t)}{e^{t}}}-e^{\alpha^{*}(t)}\right)}=\\&=\alpha^{*}\,(t)+\frac{\alpha^{*}\,(t)\cdot q\,(\alpha^{*}\,(t)\,,t)}{2\alpha^{*}\,(t)\cdot q\,(\alpha^{*}\,(t)\,,t)+e^{\frac{\alpha^{*}(t)}{e^{t}}}-e^{\alpha^{*}(t)}}\underset{t\to\infty}{\longrightarrow}x+\frac{1}{2}\end{aligned}$$

By the definition of $\alpha \in (0,1]$, we conclude that $\alpha^*(t) \underset{t\to\infty}{\longrightarrow} x = \frac{1}{2}$. We also note that $GFT_{MS}(t) = Numerator(\alpha^*(t), t)$, and evaluating $GFT_{OPT}(t)$ by elementary integration:

$$GFT_{OPT}(t) = \int_0^t \int_s^\infty (b-s)\cdot 1 \cdot e^{-1\cdot b} \cdot \frac{e^{s-t}}{1-e^{-t}} dbds = \frac{t}{e^t - 1}$$

We therefore conclude that:

$$\frac{GFT_{MS}(t)}{GFT_{OPT}(t)} = e^{-2\alpha^*(t)}(1+2\alpha^*(t))\cdot \frac{q(\alpha^*(t),t)}{t} + \frac{e^{-2\alpha^*(t)}(e^{\frac{\alpha^*(t)}{e^t}} - e^{\alpha^*(t)})}{t} \underset{t\to\infty}{\longrightarrow} e^{-2\cdot 0.5}(1+2\cdot 0.5)\cdot 1 = \frac{2}{e}$$

For evaluating the first term, we used the fact that $q(0,t) \leq q(\alpha^*(t), t) \leq q(1,t)$. For every β, it holds that $\frac{\partial q(\beta,t)}{\partial t} = e^t \cdot \left(e^{\beta e^{-t}}/e^t\right) = e^{\beta e^{-t}}$, so by L'Hopital's rule we conclude that $\frac{q(\alpha^*(t),t)}{t}$ tends to 1 as t tends to infinity. □

We conclude by showing a similar result for the expected efficiency in the bilateral trade setting. As the previous proof illustrates, and as supported by simulations using various distributions, the Second-Best mechanism achieves a relatively low approximation to the optimal GFT when the buyer's values tend to be low and the seller's values tend to be high. Since this is normally associated with low expected gains from trade, and since $EFF = GFT + E[S]$, these scenarios often produce high approximation to the optimal efficiency. Thus, it seems that tackling the question of finding an approximation to that measure that cannot be guaranteed requires observing somewhat more balanced scenarios.

We remark that [10] studied this question for the DSIC case, and showed that no DSIC mechanism which is IR and BB can guarantee a 0.749-approximation to the optimal efficiency. The following theorem shows a similar result for the general case of BNIC mechanisms, and is proved in Appendix 4.2. While this bound appears to be weak compared to the bound on the GFT in Theorem 4.1, we are not aware of any stronger bound for this problem. The best positive result to date for this problem is by [5], which show a DSIC (and thus also BNIC) mechanism that guarantees about 0.63 fraction of the optimal efficiency.

Theorem 4.2. *No BNIC mechanism which is IR and BB can guarantee an approximation to the optimal efficiency which is better than* 0.934.

5 Conclusion

This paper considers the bilateral-trade problem, which is a fundamental problem in economics for more than three decades and it demonstrates the simplest form of two sided markets. We hope that developing understanding of this fundamental problem may also be helpful in the design of more general two sided markets.

Our main result is a mechanism that achieves at least $1/e$ fraction of the optimal gain from trade, assuming that the distribution of the buyer satisfies MHR. The mechanism is simple, Bayes-Nash incentive compatible, strongly budget balanced and ex-post individually rational. The bound also implies that the most efficient mechanism subject to the IR and BB, which was characterized in the seminal paper of [22], must also achieve at least the same fraction of the optimal gain-from-trade. Our main impossibility result shows that no BNIC mechanism can guarantee an approximation which is better than $2/e$.

The main open question that is raised in this paper is whether the MHR assumption (on the buyer's side) is really required for achieving a constant approximation to the gain from trade via BNIC mechanisms. In other words, is there a BNIC, IR and BB mechanism that guarantees a constant approximation to the gain from trade for all distributions? We note that Myerson and Satterthwaite's [22] characterization of the "second-best" mechanism was not general, and assumed that the distributions are regular (a slightly weaker assumption than MHR).

A second interesting open question concerns closing the relatively-wide gap between the lower and the upper bound for the efficiency-maximizing problem by DSIC mechanisms. The best currently known approximation for this problem is 0.63 ([5]), while our impossibility result gives a bound of 0.93. As these results are given for Bayes-Nash incentive compatible mechanisms, the analysis can be challenging.

References

[1] Lawrence Ausubel, Jon Levin, Paul Milgrom, and Ilya Segal. Incentive auction rules option and discussion, 2013. Appendix to the FCCs 28-Sep-2012 NPRM on Incentive Auctions.

[2] M. Babaioff, N. Immorlica, B. Lucier, and S. M. Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 21–30, 2014.

[3] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 700–709, 2011.

[4] Liad Blumrosen and Shahar Dobzinski. Reallocation mechanisms. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 617–617, 2014.

[5] Liad Blumrosen and Shahar Dobzinski. (almost) efficient mechanisms for bilateral trading, 2016. Working paper.

[6] Liad Blumrosen and Thomas Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *ACM Conference on Electronic Commerce*, 2008.

[7] Liad Blumrosen, Noam Nisan, and Ilya Segal. Auctions with severely bounded communications. *Journal of Artificial Intelligence Research*, 28:233–266, 2006.

[8] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM symposium on Theory of computing*, STOC '10, pages 311–320, 2010.

[9] George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. In *Automata, Languages and Programming: 35th International Colloquium*, ICALP 2008, pages 820–832, 2008.

[10] Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *SODA 2016*, 2016.

[11] Martin W. Cripps and Jeroen M. Swinkels. Efficiency of large double auctions. *Econometrica*, 74(1):47–92, 2006.

[12] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC '10, pages 129–138, 2010.

[13] Moritz Drexl and Andreas Kleiner. Optimal private good allocation: The case for a balanced budget. *Games and Economic Behavior*, 94:169 – 181, 2015.

[14] Paul Dutting, Tim Roughgarden, and Inbal Talgam-Cohen. Modularity and greed in double auctions. *EC'14*.

[15] Drew Fudenberg, Markus Mobius, and Adam Szeidl. Existence of equilibrium in large double auctions. *Journal of Economic Theory*, 133(1):550–567, 2007.

[16] Rod Garratt and Marek Pycia. Efficient bilateral trade, 2015. Working paper.

[17] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, EC '09, pages 225–234, 2009.

[18] Preston R. McAfee. Amicable divorce: Dissolving a partnership with simple mechanisms. *Journal of Economic Theory*, 56(2):266–293, 1992.

[19] Preston R. McAfee. The gains from trade under fixed price mechanisms. *Applied Economics Research Bulletin*, 1, 2008.

[20] Paul Milgrom and Ilya Segal. Deferred-acceptance heuristic auctions, 2013. Working paper, Stanford University.

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

[22] R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265 – 281, 1983.

[23] Michael Ostrovsky and Michael Schwarz. Reserve prices in internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, pages 59–60, 2011.

[24] Aviad Rubinstein and S. Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 377–394, 2015.

[25] Aldo Rustichini, Mark A Satterthwaite, and Steven R Williams. Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 62(5):1041–63, 1994.

[26] Mark A. Satterthwaite and Steven R Williams. The optimality of a simple market mechanism. *Econometrica*, 70(5):1841–1863, 2002.

[27] Vasilis Syrgkanis and Eva Tardos. Composable and efficient mechanisms. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 211–220, 2013.

[28] Robert Wilson. Incentive efficiency of double auctions. *Econometrica*, 53(5):1101–1115, 1985.

A Missing proofs from Section 3

A.1 Equilibrium in the Seller-Offering Mechanism

Proof of Proposition 3.5:

Proof. Clearly, truthful bidding is the dominant strategy of the buyer, since by bidding higher than his value, he does not affect the eventual payment, but might end up having a negative surplus. Similarly, by bidding less than his value, he does not affect the eventual payment, but might miss a profitable trade.
As for the seller, assuming the buyer bids truthfully, the expected surplus of the seller is:

$$R(\sigma) = (\sigma - s) \cdot P(\sigma \geq \beta(B)) = (\sigma - s) \cdot P(\sigma \geq B) = (\sigma - s) \cdot (1 - F_b(\sigma))$$

To the left of F_b's support, this is a linearly increasing function; in its support this is a differentiable function, and to the right of it, the value of the function is 0. If $s > \overline{b}$, choosing $\overline{\varphi^{-1}}(s) = s$ ensures a zero surplus which is the maximum in this case. Otherwise, choosing σ between s and the right end of the support ensures a positive surplus, so the maximizer of R is in F_b's support.
Differentiating yields $R'(\sigma) = (1 - F_b(\sigma)) - f_b(\sigma) \cdot (\sigma - s)$ which is negative when $\varphi(\sigma) - s > 0$. Since φ is monotone, if $s < \varphi(\underline{b})$, the function is decreasing in the support, and attains maximum at $\sigma = \underline{b} = \overline{\varphi^{-1}}(s)$. Otherwise, the function attains maximum at $\overline{\varphi^{-1}}(s) = \varphi^{-1}(s)$ by first order conditions.

□

A.2 Characterization of the Expected GFT

Proof of Lemma 3.6:

Proof. We start by proving the first identity. Under full information, trading takes place whenever $b \geq s \geq 0$, thus:

$$GFT_{OPT} = \int_0^\infty \int_s^\infty (b-s) f_b(b) f_s(s) db ds = \int_0^\infty \left(-(1-F_b(b)) \cdot (b-s)\Big|_s^\infty + \int_s^\infty (1-F_b(b)) db\right) \cdot f_s(s) ds =$$

$$= \int_0^\infty \left(\int_s^\infty (1-F_b(b)) db\right) \cdot f_s(s) ds = F_s(s) \cdot \int_s^\infty (1-F_b(b)) db \Bigg|_0^\infty + \int_0^\infty F_s(s) \cdot (1-F_b(s)) ds =$$

$$= \int_0^\infty F_s(s) \cdot (1-F_b(s)) ds$$

The third and the fifth steps follow from the fact that $E[B]$ is finite, which is due to the MHR property. To see this, note that for every x in F_b's support, it holds that $h(x) \geq h(\underline{b})$, i.e., $f_b(x) \geq f_b(\underline{b}) \cdot (1 - F_b(x))$, and for $x \geq \overline{b}$ (if exists) both sides equal zero, hence:

$$E[B] = \int_0^\infty (1-F_b(b)) db = \underline{b} + \int_{\underline{b}}^\infty (1-F_b(b)) db \leq \underline{b} + \frac{1}{f_b(\underline{b})} \cdot \int_{\underline{b}}^\infty f_b(b) db = \underline{b} + \frac{1}{f_b(\underline{b})} < \infty$$

The third step follows from the following consideration:

$E[B]=\int_0^\infty b\cdot f_b(b)db = -(1-F_b(b))\cdot b|_0^\infty + \int_0^\infty (1-F_b(b))db = -(1-F_b(b))\cdot b|_0^\infty + E[B]$ and since both sides are finite, the term $-(1-F_b(b))\cdot b|_0^\infty$ equals zero.

The fifth step follows from the following consideration:
$\int_s^\infty (1-F_b(b))db = E[B] - \int_0^s (1-F_b(b))db \underset{s\to\infty}{\longrightarrow} E[B]-E[B]=0.$

We advance to proving the second identity. Assuming both agents play by the strategies proved in Proposition 3.5, trading takes place whenever $b \geq \overline{\varphi^{-1}}(s) \geq 0$, thus:

$$
\begin{aligned}
GFT_{SO} &= \int_0^\infty \int_{\overline{\varphi^{-1}}(s)}^\infty (b-s) f_b(b) f_s(s) db ds = \\
&= \int_0^\infty (-(1-F_b(b))\cdot(b-s)|_{\overline{\varphi^{-1}}(s)}^\infty + \int_{\overline{\varphi^{-1}}(s)}^\infty (1-F_b(b))db)\cdot f_s(s)ds = \\
&= \int_0^\infty (\overline{\varphi^{-1}}(s)-s)\cdot(1-F_b(\overline{\varphi^{-1}}(s)))\cdot f_s(s)ds + F_s(s)\cdot \int_{\overline{\varphi^{-1}}(s)}^\infty (1-F_b(b))db \Bigg|_0^\infty + \\
&+ \int_0^\infty F_s(s)\cdot \frac{d\overline{\varphi^{-1}}(s)}{ds}\cdot(1-F_b(\overline{\varphi^{-1}}(s)))ds = F_s(s)\cdot(\overline{\varphi^{-1}}(s)-s)\cdot(1-F_b(\overline{\varphi^{-1}}(s)))\Big|_0^\infty - \\
&- \int_0^\infty F_s(s)\cdot((\frac{d\overline{\varphi^{-1}}(s)}{ds}-1)(1-F_b(\overline{\varphi^{-1}}(s))) - \frac{d\overline{\varphi^{-1}}(s)}{ds} f_b(\overline{\varphi^{-1}}(s))(\overline{\varphi^{-1}}(s)-s))ds + \\
&+ \int_0^\infty F_s(s)\cdot \frac{d\overline{\varphi^{-1}}(s)}{ds}\cdot(1-F_b(\overline{\varphi^{-1}}(s)))ds = \int_0^\infty F_s(s)\cdot(1+\frac{d\overline{\varphi^{-1}}(s)}{ds})\cdot(1-F_b(\overline{\varphi^{-1}}(s)))ds
\end{aligned}
$$

The fourth step follows from $0 \leq \int_{\overline{\varphi^{-1}}(s)}^\infty (1-F_b(b))db \leq \int_s^\infty (1-F_b(b))db \underset{s\to\infty}{\longrightarrow} 0$ since $\overline{\varphi^{-1}}(s) \geq s$ for every s by Observation 3.2.

The fifth step follows from the two following considerations: Firstly, by Observation 3.3, for every $\overline{b} \geq x \geq \varphi(\underline{b})$ it holds that $\overline{\varphi^{-1}}(x) - x = \frac{1-F_b(\varphi^{-1}(x))}{f_b(\varphi^{-1}(x))} = \frac{1}{h(\varphi^{-1}(x))}$ and for every $\overline{b} \leq x$ we get $\overline{\varphi^{-1}}(x) - x = x - x = 0$, so in both cases it is bounded above by some constant k by the MHR asumption. Hence, for a sufficiently large s it holds that:
$0 \leq F_s(s)\cdot(\overline{\varphi^{-1}}(s)-s)\cdot(1-F_b(\overline{\varphi^{-1}}(s))) \leq 1\cdot k\cdot(1-F_b(\varphi^{-1}(s))) \underset{s\to\infty}{\longrightarrow} 0.$
Secondly, due to the same reason, for $s \geq \varphi(\underline{b})$ it holds that $\frac{d\overline{\varphi^{-1}}(s)}{ds} f_b(\overline{\varphi^{-1}}(s))(\overline{\varphi^{-1}}(s)-s) = \frac{d\overline{\varphi^{-1}}(s)}{ds}(1-F_b(\overline{\varphi^{-1}}(s)))$ and for $s < \varphi(\underline{b})$ this also holds since $\overline{\varphi^{-1}}(s)$ is constant and $\frac{d\overline{\varphi^{-1}}(s)}{ds} = 0$. □

A.3 Quadratic Approximation

Proof of Lemma 3.9.

Proof. First, we remind that by definition, $T(x) = f(x_0) + f'(x_0)\cdot(x-x_0) + \frac{1}{2} f''(x_0)\cdot(x-x_0)^2$. Consider $g = (f-T)|_{x\leq x_0}$, that is $f-T$ restricted to $x \leq x_0$. We note that $g''(x) = f''(x) -$

$T''(x) = f''(x) - f''(x_0) \geq 0$, where the inequality follows from the assumption that f' is concave, and so f'' is decreasing. We conclude that g is convex, and it attains minimum at x_0 since $g'(x_0) = f'(x_0) - T'(x_0) = f'(x_0) - f'(x_0) = 0$.

Thus, for every $x \leq x_0$ it holds that $f(x) - T(x) \geq f(x_0) - T(x_0) = f(x_0) - f(x_0) = 0$, as desired. $\square$

A.4 Tightness of Analysis for the Seller-Offering mechanism.

Proof of Proposition 3.12.

Proof. Using this mechanism:

1. Consider $F_b(x) = 1 - \sqrt{\frac{1-x}{1+x}}$ in the support $[0,1]$. We remind that $1 - F_b(x) = e^{-H(x)}$, so it follows that $h(x) = \frac{d\left(\frac{1}{2}\cdot(ln(1+x)-ln(1-x))\right)}{dx} = \frac{1}{2}\cdot\left(\frac{1}{1+x} + \frac{1}{1-x}\right) = \frac{1}{1-x^2}$. Hence, $\varphi(x) = x + x^2 - 1$ and $\varphi^{-1}(x) = \frac{1}{2}\cdot\left(\sqrt{4x+5} - 1\right)$ with the derivative $\frac{1}{\sqrt{4x+5}}$.

   Computation shows that $\int_0^\infty (1 - F_b(s))ds = \int_0^1 \sqrt{\frac{1-s}{1+s}}ds = \frac{1}{2}\cdot(\pi - 2)$ while $\int_0^\infty (1 + \frac{d\overline{\varphi^{-1}}(s)}{ds}) \cdot (1 - F_b(\overline{\varphi^{-1}}(s)))ds = \int_0^1 (1 + \frac{1}{\sqrt{4s+5}}) \cdot \sqrt{\frac{1-\frac{1}{2}\cdot(\sqrt{4s+5}-1)}{1+\frac{1}{2}\cdot(\sqrt{4s+5}-1)}}ds = -\sqrt{\sqrt{5}-2} + \pi - cos^{-1}(\frac{1}{2}(1-\sqrt{5}))$. The ratio between these two terms is 0.7335 which is less than $2/e$.

   We note that by Lemma 3.6, the original integrals also include the $F_s(s)$ parameter, but choosing $F_s \sim Uniform[0,\epsilon]$ for ϵ that tends to 0 yields that the values of the GFT tend to the aforementioned values. Concretely, choosing $\epsilon = 0.001$ yields the desired results.

2. Consider $F_b \sim Exponential(1)$, such that $h(x) = 1$ is a constant function, hence the MHR and concavity assumptions hold. Moreover, $\varphi(x) = x - 1$ and thus $\varphi^{-1}(x) = x + 1$ with a derivative that equals to 1 for every x. Let $s \geq 0$. For such s, it holds that $1 - F_b(s) = e^{-s}$ while $(1 + \frac{d\overline{\varphi^{-1}}(s)}{ds}) \cdot (1 - F_b(\overline{\varphi^{-1}}(s))) = (1+1)\cdot e^{-(s+1)} = \frac{2}{e}\cdot(1 - F_b(s))$. Taking any proper distribution F_s, multiplying both terms by $F_s(s)$ and integrating yields the desired result.

$\square$

B Missing proofs from Section 4

B.1 Hardness of Efficiency Approximation with BNIC Mechanisms

Proof of Theorem 4.2.

Proof. This proof also relies on the Second-Best mechanism in a similar manner to Theorem 4.1. We show a pair of distributions for which this approximation to the optimal efficiency holds.

Consider $F_b \sim Uniform[0,1]$ and $F_s(x) = x^{\frac{1}{4}}$ on the support $[0,1]$. The requirements from the Seond-Best mechanism apply for these distributions since $b - \frac{1-F_b(b)}{f_b(b)} = 2b - 1$ and $s + \frac{F_s(s)}{f_s(s)} = s + \frac{s^{0.25}}{0.25\cdot s^{-0.75}} = 5\cdot s$ are monotone increasing.

By the definition of the mechanism, trading take place whenever

$$b - s \geq \alpha\cdot\left(\frac{F_s(s)}{f_s(s)} + \frac{1-F_b(b)}{f_b(b)}\right) = \alpha\cdot(4s + 1 - b) \;\Leftrightarrow\; b + \alpha\cdot b \geq s + 4\alpha\cdot s + \alpha \;\Leftrightarrow\; b \geq \frac{s + 4\alpha\cdot s + \alpha}{1+\alpha}$$

Where α solves the equation:

$$\int_0^{\frac{1}{1+4\alpha}} \int_{\frac{s+4\alpha\cdot s+\alpha}{1+\alpha}}^1 (b-s)\cdot 1\cdot \frac{1}{4}s^{-\frac{3}{4}}dbds = \int_0^{\frac{1}{1+4\alpha}} \int_{\frac{s+4\alpha\cdot s+\alpha}{1+\alpha}}^1 (4s+1-b)\cdot 1\cdot \frac{1}{4}s^{-\frac{3}{4}}dbds$$

The upper limit in the integral is $\frac{1}{1+4\alpha}$ since $\frac{s+4\alpha\cdot s+\alpha}{1+\alpha} < 1$ iff $s < \frac{1}{1+4\alpha}$ for $\alpha \in (0,1]$.
We now calculate these two integrals.
Calculating shows:

$$\begin{aligned}
LHS(\alpha) &= \int_0^{\frac{1}{1+4\alpha}} \int_{\frac{s+4\alpha\cdot s+\alpha}{1+\alpha}}^1 (b-s)\cdot 1\cdot \frac{1}{4}s^{-\frac{3}{4}}dbds = \\
&= \int_0^{\frac{1}{1+4\alpha}} \left(\frac{1}{4}s^{-\frac{3}{4}}\cdot\left(\frac{b^2}{2}-s\cdot b\right)\Big|_{\frac{s+4\alpha\cdot s+\alpha}{1+\alpha}}^1\right)ds = \\
&= \int_0^{\frac{1}{1+4\alpha}} \frac{1}{4}s^{-\frac{3}{4}}\cdot\left(\frac{1}{2}-s-\frac{1}{2}\cdot\left(\frac{s+4\alpha\cdot s+\alpha}{1+\alpha}\right)^2+s\cdot\frac{s+4\alpha\cdot s+\alpha}{1+\alpha}\right)ds = \\
&= -\frac{1}{8(1+\alpha)^2}\cdot\int_0^{\frac{1}{1+4\alpha}} \frac{(s(2\alpha-1)+2\alpha+1)\cdot(s(4\alpha+1)-1)}{s^{3/4}}ds = \\
&= -\frac{1}{8(1+\alpha)^2}\cdot\int_0^{\frac{1}{1+4\alpha}} \left(s^{5/4}\left(8\alpha^2-2\alpha-1\right)+\frac{-2\alpha-1}{s^{3/4}}+s^{1/4}\left(8\alpha^2+4\alpha+2\right)\right)ds = \\
&= \left(\frac{s^{9/4}\left(-8\alpha^2+2\alpha+1\right)}{18(1+\alpha)^2}+\frac{s^{1/4}(2\alpha+1)}{2(1+\alpha)^2}-\frac{s^{5/4}\left(4\alpha^2+2\alpha+1\right)}{5(1+\alpha)^2}\right)\Bigg|_0^{\frac{1}{1+4\alpha}} = \\
&= \frac{16\left(\frac{1}{4\alpha+1}\right)^{5/4}(\alpha(9\alpha+7)+1)}{45(1+\alpha)^2}
\end{aligned}$$

And:

$$\begin{aligned}
RHS(\alpha) &= \int_0^{\frac{1}{1+4\alpha}} \int_{\frac{s+4\alpha\cdot s+\alpha}{1+\alpha}}^1 (4s+1-b)\cdot 1\cdot \frac{1}{4}s^{-\frac{3}{4}}dbds = \\
&= \int_0^{\frac{1}{1+4\alpha}} \left(\frac{1}{4}s^{-\frac{3}{4}}\cdot(\left(4s\cdot b+b-\frac{b^2}{2}\right)\Big|_{\frac{s+4\alpha\cdot s+\alpha}{1+\alpha}}^1\right)ds = \\
&= \int_0^{\frac{1}{1+4\alpha}} \frac{1}{4}s^{-\frac{3}{4}}\cdot\left(4s+1-\frac{1}{2}-(4s+1)\cdot\frac{s+4\alpha\cdot s+\alpha}{1+\alpha}+\frac{1}{2}\cdot\left(\frac{s+4\alpha\cdot s+\alpha}{1+\alpha}\right)^2\right)ds = \\
&= -\frac{1}{8(1+\alpha)^2}\cdot\int_0^{\frac{1}{1+4\alpha}} \frac{(s(4\alpha+1)-1)\cdot(s(4\alpha+7)+1)}{s^{3/4}}ds = \\
&= -\frac{1}{8(1+\alpha)^2}\cdot\int_0^{\frac{1}{1+4\alpha}} \left(s^{5/4}\left(16\alpha^2+32\alpha+7\right)-\frac{1}{s^{3/4}}-6s^{1/4}\right)ds = \\
&= \left(-\frac{s^{9/4}\left(16\alpha^2+32\alpha+7\right)}{18(1+\alpha)^2}+\frac{s^{1/4}}{2(1+\alpha)^2}+\frac{3s^{5/4}}{5(1+\alpha)^2}\right)\Bigg|_0^{\frac{1}{1+4\alpha}} = \\
&= \frac{16\left(\frac{1}{4\alpha+1}\right)^{5/4}(5\alpha+2)}{45(1+\alpha)^2}
\end{aligned}$$

Equating both terms shows that the equivalent equation is $\alpha(9\alpha+7)+1=5\alpha+2$ with the solutions $\alpha_1=\frac{-1+\sqrt{10}}{9}$and $\alpha_2=\frac{-1-\sqrt{10}}{9}$. The former of these solution is the appropriate one since it is positive.

We now observe that $LHS(\alpha_1)=GFT_{MS}$, and placing α_1 yields that $GFT_{MS}=0.31887$. Moreover, we note that:

$$\begin{aligned}
GFT_{OPT} &= \int_0^1\int_s^1 (b-s)\cdot 1\cdot \frac{1}{4}s^{-\frac{3}{4}}dbds = \\
&= \frac{1}{4}\int_0^1 s^{-3/4}\left(\frac{1}{2}-\frac{s^2}{2}-s+s^2\right)ds = \\
&= \frac{1}{4}\int_0^1 \left(\frac{1}{2s^{3/4}}+\frac{s^{5/4}}{2}-s^{1/4}\right)ds = \\
&= \left(\frac{1}{4}\cdot\left(\frac{1}{2}\cdot\frac{s^{1/4}}{1/4}+\frac{1}{2}\cdot\frac{s^{9/4}}{9/4}-\frac{s^{5/4}}{5/4}\right)\right)\Bigg|_0^1 = \\
&= \frac{1}{4}\cdot\left(2+\frac{4}{18}-\frac{4}{5}\right)=\frac{16}{45}
\end{aligned}$$

and that: $E[s]=\int_0^1(1-F_s(s))\,ds=1-\frac{s^{5/4}}{5/4}\Big|_0^1=\frac{1}{5}$.

Thus, we conclude that: $\frac{EFF_{MS}}{EFF_{OPT}}\approx\frac{0.31887+0.2}{16/45+0.2}=0.934$ □