(Almost) Efficient Mechanisms for Bilateral Trading

Liad Blumrosen*   Shahar Dobzinski

August 20, 2018

Abstract

We study the bilateral trade problem: one seller, one buyer and a single, indivisible item for sale. It is well known that there is no fully-efficient and incentive compatible mechanism for this problem that maintains a balanced budget. We design simple and robust mechanisms that obtain approximate efficiency with these properties. We show that even minimal use of statistical data can yield good approximation results. Finally, we demonstrate how a mechanism for this simple bilateral-trade problem can be used as a "black-box" for constructing mechanisms in more general environments.

1   Introduction

In this paper we revisit the simple and well studied model of bilateral trade [22]. In the bilateral trade problem, a single seller is the owner of one indivisible item which can be traded with a single buyer. Both the seller and the buyer have privately known values for consuming the item; These values are unknown to the market designer who only has access to their probability distributions.

The bilateral trade problem is probably the simplest form of two sided markets and supply chains. Two-sided markets, where multiple strategic sellers and buyers compete on each side, have always been an important form of trade. The Internet era emphasized the importance of understanding such markets with the emergence of large-scale online platforms like eBay and Amazon, and, more recently, with the popularity of the sharing economy and two sided platforms like Uber and Airbnb. Our paper raises new insights regarding the simple bilateral-trade settings, and we will argue that this understanding could help in the design of more general markets either directly or indirectly.

We will have three requirements from a mechanism for the bilateral trade problem. The first one is (ex-post) *individual rationality*: the participation of the agents is voluntary and at any point they may leave the market and consume their initial endowments. The second requirement is (ex-post) *budget balance*: the mechanism is not allowed to subsidize the agents or to make any profits. The third requirement is *incentive compatibility* (IC). While most of the literature focuses on Bayes-Nash incentive compatibility, our mechanisms will possess the stronger (and thus more desired) property of dominant-strategy incentive compatibility (DSIC).

The seminal paper by Myerson and Satterthwaite [22] analyzes the bilateral trade problem. Their celebrated impossibility result states that even in this simple setting, there is no Bayes-Nash incentive compatible mechanism which is fully efficient, individually rational and budget balanced. Myerson and Satterswaite go on and characterize the "second best" mechanism, i.e., the mechanism that maximizes the expected total gains from trade subject to the individual rationality, budget balance, and Bayes-Nash incentive constraints. Notice that since maximizing the expected total gains from trade is equivalent to maximizing the expected

---

improvement in efficiency, Myerson and Satterthwaite characterize in fact the most efficient mechanism subject to individual rationality, budget balance and incentive constraints.

Unfortunately, this "second best" mechanism is not given as a closed-form formula, and it might be a non-trivial task to explicitly describe it and compute the expected efficiency even for simple distributions. Our main goal in this paper is to develop simple and practical mechanisms for the bilateral trade problem that will approximately maximize the expected efficiency. En route, we will obtain good bounds on the efficiency of Myerson and Satterthwaite's "second best" mechanisms.

1.1 The Main Results

Our work is inspired by a recent line of research that studies the power of simple mechanisms in comparison to the performance of the optimal, but complex, mechanisms (see, e.g., [5, 17, 13, 19, 20] and references within). We restrict our attention to mechanisms that are easier to understand and implement. Specifically, we focus on mechanisms that simply post a take-it-or-leave-it price; A trade occurs only if both agents accept this price. In such mechanisms, the agents have obvious dominant strategies, and in the equilibrium analysis we do not need to speculate whether the agents compute the equilibrium correctly or converge to a particular Bayes-Nash equilibrium.

The simplicity of our mechanisms comes at the obvious price of sub-optimality, as the efficiency of our mechanisms might be inferior to that of the second-best mechanism. However, we are able to quantify this loss and show that our simple mechanisms perform quite well even with respect to a stronger benchmark: the optimal ex-post efficiency (i.e., "first-best" efficiency or simply the expected value of the maximum between the seller's and the buyer's value for the item). We measure the quality of our mechanisms by the fraction of the optimal efficiency that they guarantee[1]. Our main result is as follows:

Theorem 1: For every pair of distributions, there is a take-it-or-or-leave-it price mechanism that achieves at least $1 - \frac{1}{e}$ of the optimal efficiency in dominant strategies. The mechanism is obviously individually rational and budget balanced.

We stress that this is a worst-case bound over all possible distributions; typically there is a price that achieves a much higher fraction of the optimal welfare. A recent impossibility result by [18] shows that no dominant-strategy mechanism which is also individually rational and budget balanced can guarantee more than 0.7385 of the optimal welfare (improving over a previous hardness result of 0.7485 that was provided by [9]). Together with our $1 - \frac{1}{e} \simeq 0.632$ bound this leaves a relatively small gap.[2] One immediate corollary of the above theorem is a bound on the performance of the most efficient (second-best) mechanism with Bayes-Nash incentive compatibility. That is, it shows that this mechanism (that was characterized by Myerson and Satterthwaite in [22]) achieves at least $1 - \frac{1}{e}$ of the optimal efficiency even for the worst pair of distributions.

Given two distributions, one can find the price that maximizes the expected revenue simply by considering all possible values in the support and calculating the expected revenue when the price equals this point. However, while one cannot ignore the wealth of distributional information that retailers have at their

---

[1]Another possible objective may be approximating the expected gains from trade (GFT). As we have noted, a mechanism that maximizes gain-from-trade is fully efficient as well. However, from an approximation point of view, the two objectives behave differently: indeed, any c approximation to the GFT is also at least a c approximation to the efficiency. However, the converse is false and in particular in Appendix B.1 we show that no DSIC, budget balance and individual rational mechanism can guarantee a constant approximation to the optimal GFT. On the brighter side, McAfee [20] provides a mechanism that achieves a $1/2$ approximation to the optimal GFT when the median of the seller's distribution is no greater than the median of the buyer, and more recent approximations are in [4, 10, 6].

[2]In an earlier version of this paper [3] we proved the existence of a price that guarantees 0.51 of the optimal welfare. The follow-up [9] improved this bound to 0.52, and here we improve it further to $1 - \frac{1}{e}$. Following our improvements, Kang and Vondrák [18] have further improved the approximation ratio to $1 - \frac{1}{e} + 0.0001$ by a much more complicated mechanism. As far as we know, we were the first to approximate efficiency for the general bilateral-trade problem.

disposal, this knowledge is often incomplete. In many scenarios, statistical knowledge is not available to the designers (for example, in new markets or markets that involve new participants) or it may be costly. In other cases, the designer can only accurately estimate some statistics of the distributions (expectation, median, etc.) rather than the full details of the distribution.

Our next result shows that the $1 - 1/e$ approximation can be achieved by a randomized mechanism that chooses the prices according to a carefully chosen distribution. Importantly, this improved approximation only needs access to the distribution of the seller.

Theorem 2: For every distribution of the seller, there is a distribution G of prices with the following property: if we post a take-it-or-leave-it price sampled from G then for *every* distribution of the buyer we obtain a fraction of at least $(1 - 1/e) \approx 0.63$ of the optimal efficiency.

Our main result, stated as Theorem 1, is actually a direct corollary of Theorem 2. Theorem 2 says that if prices are chosen at random according to some distribution, then the expected efficiency is at least $1 - 1/e$ of the optimal efficiency. It follows that for every distribution of the buyer, there must be one price that attains at least $1 - 1/e$ of the optimal efficiency as well. To explicitly find this price one needs to know the details of the distribution of the buyer as well.

We also present two simple mechanisms that achieve at least $1/2$ of the full efficiency but use very restricted distributional knowledge.[3] The first mechanism posts a price that is equal to the median value of the seller's distribution, and the other mechanism computes some sort of a weighted median based only on the distribution of the buyer and offers this price to the agents. We show that the $1/2$ bound is tight in the sense that no single take-it-or-leave-it price can achieve a better approximation knowing only the distribution of the seller or only the distribution of the buyer.

1.2 Using Solutions for Bilateral Trade as "Black-Boxes"

The above results study the fundamental problem of Bilateral Trade, and insights from the study of this basic problem should alleviate the design of mechanisms in more complex trade environments.[4] In the second part of the paper, we show how results for the bilateral trade problem can be directly applied to construct approximately-optimal mechanisms in other settings in a black box manner. More specifically, we show reductions[5] of the following form: given a mechanism that guarantees some α approximation to Bilateral Trade, we design a mechanism for the other problem that obtains $f(\alpha)$ approximation for some function f. We assume that we are given a DSIC, ex-post budget balanced and individually-rational mechanism as input, and the mechanism we output maintains these desired properties. Given an α-approximation mechanism for Bilateral Trade, we give the following reductions:

- *Partnership Dissolving.* Consider n players, each player i initially owns a fraction r_i of a divisible good. Players have linear valuations, such that the value for player i for a fraction c of the item is $c \cdot v_i$. This is the classic model by Cramton, Gibbons and Klemperer [12]. We design a mechanism that achieves an α fraction of the optimal efficiency. This result proves that the extreme-ownership

---

[3]At first glance, achieving 1/2 of the optimal efficiency looks trivial: we can allocate the item to the player with the higher expected efficiency and that should give us at least 1/2 of the optimal efficiency. However, this overly-optimistic argument fails since in bilateral trade the two agents are not symmetric: while it is easy to leave the item with the seller, we need to convince the seller to relinquish the item (via a sufficiently high payment) for the two agents to trade.

[4]In a companion paper [3], we show how ideas that stem from our understanding of the simple bilateral-trade setting can be used in the design of mechanisms in more general, multi dimensional exchange settings. In particular, we show in [3] how a careful use of distribution quantiles as prices can help approximate the optimal efficiency; In prior-free settings, quantile-like prices are artificially created by random sampling methods.

[5]A *reduction* is a fundamental and widely used technique in theoretical computer science. This is essentially an algorithm that transforms one problem into another problem such that any solution to the latter problem implies a solution to the original problem as well.

scenario of the partnership dissolving problem is actually the hardest to solve, and any approximation for the first implies the same approximation factor for the latter problem.

- *Divisible good with general monotone valuations.* Consider a seller and a buyer, where the seller initially owns a fully divisible good. The players have monotonically non-decreasing valuations. We design a mechanism that guarantees $\frac{\alpha}{\alpha+1}$ of the optimal efficiency in this setting.

- *Divisible good with convex valuations.* Consider a 2-player environment, where each player i initially holds some fraction r_i of a divisible good. The players have convex valuations, i.e., preferences with decreasing marginal valuations. We design a mechanism that guarantees $\frac{\alpha}{\alpha+1}$ of the optimal efficiency in this setting.

Applying our $(1 - 1/e)$-approximation mechanism mentioned above, we obtain a 0.39-approximation mechanisms for the two 2-player trade problems with a divisible good, and a $(1 - 1/e)$-approximation mechanism to the n-player partnership dissolving problem. Any future improvements to the $1 - 1/e$ bound would immediately imply an improvement to the above bounds as well.

More related research. The paper of McAfee [20] is closest in spirit to ours. McAfee shows that for distributions such that the median of the buyer is greater than the median of the seller, posting any price between these two medians achieves at least half of the optimal expected gains from trade. In a sense we show that this median technique can be developed to approximate the social welfare for every pair of distributions.

A recent paper by Garratt and Pycia [15] shows that with non quasi-linear preferences, e.g., with risk aversion and wealth effects, efficient trade is possible under some general conditions on the information structure. The partnership dissolving model was first studied by Cramton, Gibbons and Klemperer [12] who show that if the shares are close enough to equal holdings, there exists a fully efficient Bayes-Nash incentive compatible mechanism (see also [19], and a survey [21]).

Our paper is part of two notable research directions from the last decade. The first is a series of paper that show how simple, feasible mechanisms can approximate the results achieved by the optimal, yet often too complex, mechanisms (e.g., [5, 17, 13, 8, 1]). The second line of research studies "robust" mechanism design, which is an effort to design mechanisms that are less sensitive to the fine details of the specific environment that may not be available to the planner in practice (see the survey [2] and the many references within).

2 Model

The bilateral trade problem involves two agents, a seller and a buyer. The seller owns an indivisible item, and his valuation for consuming this item is s. If a trade occurs and the item is allocated to the buyer, then the buyer enjoys a value of b.

s and b are independently distributed according to distributions F_s and F_b, respectively. All our results hold for any pair of distributions, but for the simplicity of presentation we assume the existence of density functions f_s and f_b, for the seller and the buyer respectively, which are always positive on a support $[\underline{a}, \overline{a}]$ ($\underline{a} \geq 0$) and atomless.

For every realization of s and b, the deterministic mechanism takes s, b as input and determines the allocation $X(s, b) \in \{0, 1\}$, the payment p_s that the seller receives, and the payment p_b of the buyer. We set $X(s, b) = 0$ in the case of no trade (the seller keeps the item) and $X(s, b) = 1$ if the buyer receives the good.

In this paper we restrict our attention to ex-post budget balanced mechanisms in which for all s, b we have that $p_s = p_b$.[6] We can therefore denote the trade price by p. We also consider only ex-post individually rational mechanisms, in which $p = 0$ if there is no trade and $s \le p \le b$ otherwise.

The private information of the buyer is b and the private information of the seller is s. When trade occurs at price p, the seller's utility is p and the buyer's utility is $b - p$. With no trade, the seller's utility is s and the buyer's utility is 0.

Our mechanisms are dominant strategy incentive compatible (DSIC). In particular, we consider posted price mechanisms, in which the center posts a trade price p that is independent on the realizations of the valuations and trade occurs only if both players agree to that price. Posted price mechanisms are clearly DSIC, budget balanced and individually rational. In fact, since[7] $X(s, b) \in \{0, 1\}$ it is not hard to see that every DSIC, budget balanced and individually rational mechanism is a posted price mechanism, simply because p cannot depend on b nor on s.

A fully efficient mechanism will initiate a trade whenever $s > b$. Define the optimal efficiency as

$$E_{s \sim F_s, b \sim F_b}[\max\{s, b\}]$$

We measure the performance of our mechanisms by the fraction of the optimal efficiency that they obtain. Specifically, we denote the expected efficiency achieved by a mechanism with an allocation function X as

$$E_{s \sim F_s, b \sim F_b}[(1 - X(s, b)) \cdot s + X(s, b) \cdot b]$$

3 Warm Up: Simple Mechanisms for Bilateral Trade

We begin our investigations by introducing two simple mechanisms that guarantee at least half of the optimal efficiency. In the next sections we provide an improved approximation ratio that relies on ideas that are presented here in their most crystallized form.

The mechanisms that are developed in this section – as well as all of our mechanisms for the bilateral trade problem – simply set a price p. The item is traded if and only if $b \ge p$ *and* $s \le p$. In case of trade, the buyer pays p and the seller receives the same amount. Since the choice of p will depend only on the the distributions of the players and not on the actual realizations, this family obviously yields dominant strategy mechanisms that are individually rational and budget balanced.

Our first mechanism is the *median mechanism*[8] sets the price p to be the median of seller's distribution, i.e., the point M_s such that $F_b(M_s) = \frac{1}{2}$.

Theorem 3.1 *The median mechanism is DSIC, individually rational, budget balanced and always achieves at least $1/2$ of the optimal social welfare.*

Proof: By the discussion above, the mechanism is obviously DSIC, individually rational, and budget balanced. We now analyze its approximation ratio. We start with some notation. Since the trade price depends only on the distribution of the seller, to analyze the approximation ratio we may fix the value b of the buyer and show that for every fixed b the expected approximation ratio is $1/2$ in expectation over the seller's distribution. We divide our analysis to two disjoint cases. In both cases we use the observation that the expected optimal welfare is at most $\Pr[b \ge s] \cdot b + \Pr[s > b] \cdot E[s | s > b]$.

---

[6] This requirement is sometimes known as *strong* budget balance. Weak budget balance allows the mechanism to accumulate profit. We are able to prove our positive results even with the strong budget balance property.

[7] One can also relax our requirements and consider mechanisms with $X(s, b) \in [0, 1]$ with all definitions extended accordingly. However, Hagerty and Rogerson [16] essentially show that every optimal mechanism is a posted price mechanism. Other papers (e.g., [23, 14]) show the optimality of posted price mechanisms in related settings.

[8] As discussed in the introduction, this mechanism was has shown to approximate the GFT in some settings by [20].

1. $b \geq M_s$. In this case the item is sold with probability $\frac{1}{2}$, when $s \leq M_s$. We get that the approximation ratio of the mechanism (i.e., the expected optimal efficiency divided by the expected efficiency of the median mechanism) is at most:

$$\frac{\Pr[b > s] \cdot b + \Pr[s \geq b] \cdot E[s|s \geq b]}{\Pr[s < M_s] \cdot b + \Pr[s \geq M_s] \cdot E[s|s \geq M_s]}$$

$$= \frac{\Pr[b > s] \cdot b + \Pr[s \geq b] \cdot E[s|s \geq b]}{\Pr[s < M_s] \cdot b + \Pr[s \geq b] \cdot E[s|s \geq b] + \Pr[M_s < s < b] \cdot E[s|M_s < s < b]}$$

$$\leq \frac{\Pr[b \geq s] \cdot b + \Pr[s \geq b] \cdot E[s|s \geq b]}{\frac{1}{2} \cdot b + \Pr[s \geq b] \cdot E[s|s \geq b]} \leq 2$$

2. $b < M_s$. Here the item is never sold and the expected welfare is $E[s]$. We bound the expected approximation ratio:

$$\frac{\Pr[b \geq s] \cdot b + \Pr[s > b] \cdot E[s|s > b]}{E[s]} \leq \frac{\frac{1}{2}M_s + \Pr[s > b] \cdot E[s|s > b]}{\Pr[s \leq b] \cdot E[s|s \leq b] + \Pr[s > b] \cdot E[s|s > b]}$$

$$\leq \frac{\frac{1}{2}M_s + \Pr[s > b] \cdot E[s|s > b]}{\Pr[s > b] \cdot E[s|s > b]}$$

$$\leq \frac{\frac{1}{2}M_s + \frac{1}{2}M_s}{\frac{1}{2}M_s} = 2$$

where in the last transition we use the fact that $\Pr[s > b] \cdot E[s|s > b] = \Pr[M_s \geq s > b] \cdot E[s|M_s \geq s > b] + \Pr[s > M_s] \cdot E[s|s > M_s] \geq \frac{1}{2}M_s$.

□

One could hope that the "symmetric" mechanism that sets the trade price to be M_b (the median of the distribution of the buyer) will provide a good approximation ratio as well. Unfortunately, this is not the case. To see this, consider the distribution where the buyer's value is ϵ with probability $\frac{1}{2} + \epsilon$ and $t > 1$ with probability $\frac{1}{2} - \epsilon$. Let the seller's value be 1 with probability 1. When we set the trade price to be $M_b = 1$ the item is never sold since the price is always too low for the seller. The welfare of the mechanism is 1 while ideally we would like to sell the the item whenever the buyer's value is t and get an expected approximation ratio of $(\frac{1}{2} - \epsilon)t$. Letting t approach infinity and ϵ approach 0, we get that setting the trade price to be M_b does not guarantee any constant approximation ratio.

However, we do show that a more careful choice of a trade price p (using only the distribution of the buyer) results in an approximation ratio of $1/2$. The *weighted median mechanism* sets the trade price p to be the point W_b for which $F_b(W_b) \cdot E[b|b < W_b] = (1 - F_b(W_b)) \cdot E[b|b \geq W_b]$. In particular, since $E[b|b \geq W_b] \geq E[b|b < W_b]$ we get that $\Pr[b > W_b] \leq \frac{1}{2}$. In addition:

$$(1 - F_b(W_b)) \cdot E[b|b \geq W_b] = F_b(W_b) \cdot E[b|b < W_b] = \frac{E[b]}{2} \qquad (1)$$

Theorem 3.2 *The weighted median mechanism is DSIC, individually rational, budget balanced and always achieves at least $1/2$ of the optimal social welfare.*

Proof: Again, we only analyze the approximation ratio of the mechanism. Similarly to the analysis of the median mechanism, since the trade price p depends only on the buyer's distribution we may analyze the approximation ratio assuming that s is fixed. We divide our analysis into two cases:

6

1. $s < W_b$. In this case the item is sold whenever the buyer's value is at least W_b. The approximation ratio can be bounded by:

$$\frac{\Pr[b \geq s] \cdot E[b|b \geq s] + \Pr[s > b] \cdot s}{\Pr[b \geq W_b] \cdot E[b|b \geq W_b] + \Pr[b < W_b] \cdot s} \leq \frac{E[b] + s}{\frac{E[b]}{2} + \frac{s}{2}} \leq 2$$

where we use Equation (1) and $\Pr[b < W_b] \geq \frac{1}{2}$.

2. $s \geq W_b$. In this case the item is never sold. We will make use of $E[b] = \Pr[b > W_b] \cdot E[b|b > W_b] + \Pr[b \leq W_b] \cdot E[b|b \leq W_b] \leq \frac{1}{2} \cdot E[b] + W_b$ which implies that $\frac{E[b]}{2} \leq W_b \leq s$. Therefore:

$$\frac{\Pr[b \geq s] \cdot E[b|b \geq s] + \Pr[s > b] \cdot s}{s} \leq \frac{\Pr[b \geq W_b] \cdot E[b|b \geq W_b] + \Pr[s > b] \cdot s}{s}$$

$$\leq \frac{\frac{E[b]}{2} + s}{s} \leq \frac{s + s}{s} = 2$$

where we again use Equation (1).

$\square$

The two simple mechanisms that we presented in this section prove that one can achieve good approximation to the optimal efficiency while preserving budget balance and individual rationality. Moreover, this can be done having statistical knowledge only on one side of the market. It turns out that this is the best ratio that can be achieved via deterministic DSIC mechanisms (proof is in Appendix A).

Proposition 3.3 *No deterministic DSIC, individually rational and budget balanced mechanism can obtain an approximation ratio better than $1/2$ if the mechanism uses only the distribution of the seller or only the distribution of the buyer.*

4   A (1-1/e)-Approximation Mechanism

In Section 3 we showed that one can compute a price that depends only on the distribution of the seller and this price will guarantee at least half of the optimal efficiency no matter what the distribution of the buyer is. We showed that this bound of $\frac{1}{2}$ is tight for all deterministic dominant-strategy incentive-compatible mechanisms. However, this bound could be improved by randomly choosing a price. We will describe a mechanism that posts a price at random, and we will only require this mechanism to have access to the distribution of the seller. We will then conclude the existence of a deterministic posted-price mechanism that has access to both distributions.

We let $q(\cdot)$ denote the quantile function of the seller, that is, $F_s(q(x)) = x$ for every $x \in [\underline{a}, \overline{a}]$. Consider the following mechanism that posts a price to both players, but chooses the price randomly. We call this mechanism the *Random-Quantile mechanism*:

- *Choose a number $x \in [1/e, 1]$ according to the distribution with CDF $G(x) = \ln(e \cdot x)$.*

- *Set the price to be $q(x)$.*

Theorem 4.1 *The Random-Quantile mechanism is DSIC, individually rational, budget balanced, and achieves in expectation at least $1 - \frac{1}{e}$ fraction of the optimal efficiency for every pair of distributions F_s and F_b.*

Before providing the proof of the theorem, we note that our Random-Quantile mechanism was shown by [18] to provide the best approximation ratio among all "quantile" mechanisms: mechanisms that choose x according to some distribution, and then set the price to be $q(x)$.

Proof: We prove the theorem for every fixed value b of the buyer and it will clearly hold for every distribution of the buyer as well.

We will first prove that if we truncate the seller's distribution at b, the approximation ratio can only get worse.

Claim 4.2 *Let the buyer have a fixed value b, and let F_s^* be a distribution for the seller such that $F_s^*(x) = F_s(x)$ for $x < b$, $F_s^*(x) = 1$ for $x \geq b$. We denote by $RAND(F_s, b)$ the expected efficiency of the Random-Quantile mechanism for the distribution F_s for the seller and a fixed value b for the buyer. We denote by $OPT(F_s, b)$ the optimal efficiency for these distributions. Then,*

$$\frac{RAND(F_s, b)}{OPT(F_s, b)} \geq \frac{RAND(F_s^*, b)}{OPT(F_s^*, b)}$$

Proof: We first note that

$$RAND(F_s, b) - RAND(F_s^*, b) = (1 - F_s(b))\Big(E[s|s \geq b] - b\Big) \tag{2}$$

$$= OPT(F_s, b) - OPT(F_s^*, b) \tag{3}$$

To see this, note that that for all values of s that are below b, the expected efficiency is the same for F_s^* and F_s. For values s that are above b, both the optimal allocation and the Random-Quantile mechanism gain efficiency of s for F_s but of only b for F_s^*. Thus, the difference between the expected efficiency in the Random-Quantile mechanism and in the optimal efficiency is exactly $(1 - F_s(b))\Big(E[s|s \geq b] - b\Big)$.

Therefore, moving from $\frac{RAND(F_s, b)}{OPT(F_s, b)}$ to $\frac{RAND(F_s^*, b)}{OPT(F_s^*, b)}$ the enumerator and the denominator decrease by the same additive term. The claim follows as the two fractions are at most 1. $\square$

We now analyze the approximation ratio assuming a fixed buyer valuation b and the seller distribution F_s^*, for which $OPT = b$. This will also imply the same approximation bound for the original distributions by the above claim. Let $y = F_s(b)$ denote the probability that the seller's value is below b for F_s^*.

In our randomized mechanism, either the seller keeps the item or the buyer gets it. We will separately bound the expected contribution of the seller's value and the buyer's value in the mechanism and by the linearity of expectation get a bound on the overall efficiency. Recall that the seller accepts a price $q(x)$ with probability x by definition and that the density of the price distribution $G(x)$ is $\frac{1}{x}$. When the mechanism chooses a quantile x between $\frac{1}{e}$ and y, a trade occurs with probability x, and the realized efficiency is b. Therefore, the contribution of the buyer to the overall efficiency is at least $\int_{\frac{1}{e}}^{y} x \cdot b \cdot \frac{1}{x} dx$. On the other hand, note that unless the seller's value is strictly below b, the efficiency is b with probability 1. The seller has value b with probability $1 - y$, and therefore the contribution of the seller's value to the overall efficiency is at least $b \cdot (1 - y)$. We have that

$$RAND(F_s^*, b) \geq \int_{\frac{1}{e}}^{y} x \cdot b \cdot \frac{1}{x} dx + b \cdot (1 - y) \tag{4}$$

$$= \left(y - \frac{1}{e}\right) \cdot b + b \cdot (1 - y) \tag{5}$$

$$= (1 - \frac{1}{e}) \cdot b \tag{6}$$

8

$\square$

This shows that by knowing only the seller's distribution, one can build a price distribution that guarantees a $1 - 1/e$ fraction of the optimal efficiency regardless of the distribution of the buyer. An immediate corollary of this result is that for every distribution of the buyer there exists at least one price that achieves at least the same fraction of the optimal efficiency (as there always is a price in the support that the expected revenue it guarantees is higher than the expected revenue with respect to the distribution).

Corollary 4.3 *For every pair of distributions F_s, F_b, there exists a deterministic take-it-or-leave-it price that achieves at least $1 - 1/e$ of the optimal efficiency. This mechanism is dominant-strategy incentive compatible, budget balanced and (ex-post) individually rational.*

The above proposition is non-constructive: to compute such a price one needs to have access to the distribution of the buyer as well. This result was presented informally in the introduction as Theorem 1.

5 Black-Box Reductions

The Bilateral Trade problem is one of the most fundamental and well-studied problems in mechanism design. So far, we designed mechanisms that approximate the optimal social welfare for this problem. In this section, we show that the simple Bilateral Trade environment can be used as a building block for the construction of mechanisms in more general environments. We design the mechanisms using black-box reductions: we show that given any α approximation mechanism for Bilateral Trade we can construct a mechanism with an approximation ratio which is some function of this α that maintains all the desired economic properties.

In the rest of this section we will consider a more general model that reallocates a fully divisible good. In Section 5.1 we study the n-player partnership dissolving setting, where players initially own fractions of the good and have linear utilities. In Section 5.2 we study 2-player settings with general monotone valuations, and in Section 5.3 we discuss a similar 2-player setting with the restriction to convex valuations (but with an arbitrary initial allocation).

5.1 Dissolving Partnerships

In the *partnership dissolving* problem, there are n agents, each agent i owns a share r_i of an asset, and $\sum_{i=1}^{n} r_i = 1$. Each agent i has a value v_i for holding the entire asset, or a value $c \cdot v_i$ for holding a fraction $c \geq 0$ of the asset. Let $r_{max} = \max\{r_1, ..., r_n\}$ be the largest share held by an agent.

We show that any approximation mechanism for bilateral trade can be used for constructing a mechanism for partnership dissolving with exactly the same guarantee on the approximation ratio. This reduction from partnership dissolving to bilateral trade holds for all distributions of the buyers, and irrespectively of the size of the initial shares. The idea is that each agent sells his share to the other agents via a second-price auction, where the price taken from the bilateral trade mechanisms for this setting serves as a reserve price. We show that for each player, and thus for the whole economy, this mechanism can only improve the efficiency of the bilateral trade mechanism.

Theorem 5.1 *Let M be some DSIC, individually rational, and budget balanced mechanism for bilateral trade that achieves an α-approximation to the welfare. There is a DSIC, individually-rational and budget-balanced mechanism for partnership dissolving which also achieves an α-approximation to the optimal efficiency.*

Proof: In the proof we use the already-mentioned fact that any truthful mechanism for bilateral trade simply sets a trade price p that does not depend on the values of the bidders. We develop our mechanism for partnership dissolving in two stages.

9

First Stage: A Mechanism where only Bidder i may sell. We will "run" M with bidder i and a hypothetical buyer whose value is distributed according to the distribution of $\max_{k \neq i} v_k$. Let p be the price that M posts. Let M' be the following mechanism:

- Let $j \in \arg\max_{k \neq i}\{v_k\}$ and let m_2 be the second highest value of $v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n$, that is, $m_2 = \max_{k \in N \setminus \{j,i\}} v_k$.

- Let $p^* = \max\{p, m_2\}$. If $v_i \leq p^*$ and $v_j \geq p^*$ then bidder j pays to bidder i the amount of p^* and receives the item. Otherwise bidder i keeps his item and does not get paid.

To see that the mechanism is truthful, observe that if a sale is made neither the winning buyer nor bidder i cannot affect the price by changing their bid. A losing buyer k can only turn into a winner by overbidding the winning buyer and paying $p^* \geq v_j \geq v_k$, and therefore cannot gain a positive payoff. If the item is not sold, it is either because the seller's value exceeds p^* (and winning by underbidding induces a payment below v_i) or all of the buyers' values are below p^* (and again, overbidding results in a payment higher than v_j).

Now for the approximation ratio. Notice that whenever there is a trade in M there is a trade in M' (but the opposite is not true; for example, a trade occurs when $p < v_i \leq m_2$). M achieves an α-approximation to the optimal solution (that may only allocate bidder i's share) which equals $\max\{v_i, \max\{v_{-i}\}\} = \max_k\{v_k\}$. M' achieves at least the same expected welfare as M, thus it is an α-approximation to the optimal welfare (that may only allocate bidder i's share) as well.

The mechanism is individually rational since we always have that $v_i \leq p^* \leq v_j$. In addition, payments are transferred from one player to another, hence the mechanism is budget balanced.

Second Stage: The Final Mechanism. At an arbitrary order, use M' to sell to the other bidders the endowment r_i of each bidder i as a single indivisible item.

We now analyze the approximation ratio. Let $v_{max} = \max_k v_k$ be the highest value. By selling the endowment of bidder i the expected social welfare is at least $\frac{r_i v_{max}}{\alpha}$. Since the valuations of the bidders are linear, after selling all endowments the expected social welfare of at least $\frac{v_{max}}{\alpha}$.

The truthfulness of the mechanism also follows from the linearity of the valuations of the bidders: at every stage they will maximize their payoff from the item independently of the other sales. Therefore, truthfulness follows from the truthfulness of M'. Similarly, the mechanism is individually rational and budget balanced. $\qquad\square$

As our best approximation for the bilateral-trade problem is $1 - \frac{1}{e}$, the reduction guarantees the same approximation ratio for Partnership Dissolving.

Corollary 5.2 *There is a DSIC, individually rational, and budget balanced mechanism the Partnership Dissolving problem which is a $1 - \frac{1}{e}$ approximation to the optimal effciency.*

5.2 2-player Markets with Monotone Valuations

We now consider a 2-player market where a seller initially owns the whole quantity of a divisible good. The seller and the buyer may have arbitrary monotone valuations, that is, the only restriction on the valuations is that for fractions $x > y$ we have $v(x) \geq v(y)$. In this setting, we show that given an α-approximation mechanism to the bilateral trade problem, we can design a mechanism with an approximation guarantee of $\frac{\alpha}{1+\alpha}$ that preserves the economic properties of the original mechanism.

Theorem 5.3 *Let A be a budget balanced and DSIC mechanism for bilateral trade with one indivisible good that achieves an approximation ratio of $\alpha < 1$. There is a DSIC and budget balanced mechanism A' for bilateral trade with one divisible good that guarantees an approximation ratio of $\frac{\alpha}{\alpha+1}$, as long as the players' valuations are monotonically increasing.*

Proof: Let S be the expected contribution of the seller to the welfare-maximizing solution, and let B be the expected contribution of the buyer B. Thus, if we let OPT denote the expected optimal welfare, we have that $OPT = B + S$.

Consider some instance and let s be the value of the seller for getting all the item and similarly let b be the value of the buyer for getting all the item. Let D_s be the distribution of s given the seller's distribution and D_b the distribution of b. Our mechanism A' will simply run A to obtain a trade price p. If $b \geq p \geq s$ the item is re-allocated in full to the buyer and the buyer transfers a payment of p to the seller. Otherwise the seller keeps the full item. Notice that A' inherits from A its truthfulness, budget balance, and individual rationality.

We divide the analysis of the approximation ratio into two cases. In the first case we assume that $S \geq \frac{\alpha}{\alpha+1} OPT$. Observe that in this case the approximation ratio is $\frac{\alpha}{\alpha+1}$: the expected welfare of any mechanism is at at least $E[s]$ since every reallocation cannot decrease the welfare. To conclude this case, observe that clearly we have that by the monotonicity of the valuations, $E[s] \geq S$.

Thus assume that $S < \frac{\alpha}{\alpha+1} OPT$. In other words, $B \geq \frac{OPT}{\alpha+1}$. Notice that if we restrict our attention to instances where the item can only be reallocated in full, then the optimal welfare is at least $E[b] \geq B$ (again, using the monotonicity of the valuations). Thus, using the approximation guarantee of A, the expected welfare of A' is at least $\alpha B \geq \frac{\alpha}{\alpha+1} OPT$, as needed. $\qquad\square$

5.3 Trading a Divisible Good with Convex Valuations

We consider a 2-player market, where each player owns some arbitrary share of a fully divisible good. The initial shares of players $1, 2$ are denoted by r_1, r_2, where $r_1 + r_2 = 1$. Each player i has a valuation function $v_i : [0, 1] \to \mathbb{R}$, and for every x, y we define the marginal valuation $v_i(x|y) = v_i(x + y) - v_i(y)$. We assume that the valuation functions are normalized ($v_i(0) = 0$), non decreasing, and have decreasing marginal valuations (i.e., $v_i(\epsilon|x) \geq v_i(\epsilon|y)$ for every $\epsilon > 0, y > x$).[9]

In this model, we show that given any DSIC and budget balanced α-approximation mechanism, we can use it to create an $\frac{\alpha}{\alpha+1}$-approximation mechanism for the 2-player market with convex valuations.

Theorem 5.4 *Let A be a DSIC and budget balanced mechanism for bilateral trade that achieves an $\alpha < 1$ approximation to the optimal efficiency. There is a DSIC and budget balanced mechanism for the 2-player reallocation problem with arbitrary endowments when the players' valuations are monotone and have decreasing marginals that guarantees an approximation ratio of $\frac{\alpha}{\alpha+1}$.*

Proof: Consider the following mechanism: give the seller for free a fraction x of the item ($x \leq 1$ will be specified later, one of the players that has a share of at least x will be viewed as the seller and we will claim that there is such a player). Denote by D_s the marginal distribution that specifies the value of the seller for receiving an additional $y \in [0, 1-x]$ fraction of the item and by D_b the distribution that specifies the value of the buyer for receiving $y \in [0, 1-x]$ of the item. Let i be the player that received the item when running the mechanism A on the distributions D_s and D_b, and let $p_s \geq 0$ be the payment that the seller receives and p_b the amount that the buyer is charged. Our mechanism give player i a fraction of $1 - x$ of the item, additionally allocates the seller x, charges the buyer p_b and gives the seller a payment p_s. The mechanism is clearly truthful and budget balanced. All that is left is to analyze its approximation ratio.

Let S be the expected value of the seller for receiving all the item, and similarly let B denote the expected value of the buyer for receiving all the item. Clearly, $S + B$ is an upper bound on the expected value of the optimal solution. We will show that for $x = \alpha/(\alpha+1)$ the expected welfare of our mechanism is $x \cdot (S + B)$ and the claim regarding the approximation ratio will follow.

---

[9] When $v_i(\cdot)$ is twice differentiable, we simply assume that $v_i''(x) \leq 0$ for every x.

We first bound the contribution of the seller for receiving a fraction x of the item in the first step: Since the valuation of the seller exhibits decreasing marginals the expected contribution is at least $x \cdot S$. We now bound the contribution of the second step. Ideally, we could have calculated the expected optimal value OPT_2 of the secondary problem (with the distributions D_s and D_b) and claim that the when running the mechanism A the expected welfare is at least $\alpha \cdot OPT_2$. However, since we do not know how to compute OPT_2 exactly we will only give a rough bound to it. Specifically, we claim that $OPT_2 \geq (1 - x) \cdot B$, which is the welfare we when always allocating the buyer a fraction of $(1 - x)$ of the item, since the buyer's valuation function exhibits decreasing marginals.

We get that the overall expected welfare of our mechanism is $x \cdot S + \alpha(1 - x) \cdot B$. Thus, for $x = \frac{\alpha}{\alpha+1}$ the expected welfare is $\frac{\alpha S}{1+\alpha} + \frac{\alpha B}{1+\alpha}$, which gives us an $\frac{\alpha}{\alpha+1}$ approximation ratio. As there is always a player with at least $\alpha/(\alpha + 1)$ fraction of the good (this value is at most 1/2), this player will be the seller. $\qquad\square$

6 Discussion

In this paper, we constructed mechanisms that approximate efficiency in bilateral-trade settings, and showed how such mechanisms can be used as building blocks of mechanisms in more complex scenarios. Our main result states that there always exists a simple, fixed-price, dominant-strategy incentive compatible, strongly budget balanced and ex-post individually rational mechanism that achieves at least $1 - 1/e$ of the optimal efficiency. This mechanism computes a price as a function of the distribution functions of the seller.

We then show how this implies a 0.39 approximation for a similar setting with a divisible good and arbitrary monotone valuations, or a 2-player exchange setting for general convex valuations. We also prove the same $1 - 1/e$ approximation factor for the n player partnership-dissolving problem. These three approximation results are given as "black-box" reductions, such that if the $1 - 1/e$ bound is improved in the future (we know that it cannot be larger than 0.749 [9]), these bounds will be immediately improved as well.

A recent working paper [7] discussed a monopoly pricing problem, where the seller have a limited knowledge on the distribution of the buyer's value and the seller aims to maximize his worst-case profit. While no approximation results are given in [7], the structure of the optimal selling mechanism has similarities to our main mechanism (random prices from a distribution with a logarithmic shape). Future work should study whether this new approach can help proving the optimality of such mechanisms in efficiency maximizing scenarios like ours. Another very recent paper is [11] that studies equilibria in a game where a buyer chooses a distribution and the seller posts a price after observing this distribution. In equilibrium the seller posts a single deterministic price but the buyer chooses a distribution over the range $[\frac{1}{e}, 1]$, which again, may indicate that our main mechanism is optimal in some sense.

Acknowledgments

References

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

[2] Dirk Bergemann and Stephen Morris. An introduction to robust mechanism design. *Foundations and Trends in Microeconomics*, 8(3):169–230, 2013.

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

[4] Liad Blumrosen and Yehonatan Mizrahi. Approximating gains-from-trade in bilateral trading. In *Web and Internet Economics WINE*, pages 400–413, 2016.

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

[6] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 589–590, 2017.

[7] Vinicius Carrasco, Vitor Farinha Luz, Paulo Monteiro, and Humberto Moreira. Robust selling mechanisms, April 2015. Working paper.

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

[10] Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Fixed price approximability of the optimal gain from trade. In *Web and Internet Economics - WINE*, pages 146–160, 2017.

[11] Daniele Condorelli and Balazs Szentes. Buyer-optimal demand and monopoly pricing, February 2016. Working paper.

[12] Peter Cramton, Robert Gibbons, and Paul Klemperer. Dissolving a partnership efficiently. *Econometrica*, 55(3):615–32, 1987.

[13] 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.

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

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

[16] Kathleen M Hagerty and William P Rogerson. Robust trading mechanisms. *Journal of Economic Theory*, 42(1):94–107, 1987.

[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] Zi Yang Kang and Jan Vondrák. Strategy proof approximations of optimal efficiency in bilateral trade, 2018. Working paper.

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

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

[21] Benny Moldovanu. How to dissolve a partnership. *Journal of Institutional and Theoretical Economics (JITE)*, 158(1), 2002.

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

[23] Ran Shao and Lin Zhou. Optimal allocation of an indivisible good. *Games and Economic Behavior*, 100:95 − 112, 2016.

A   Deterministic 1/2 Approximation: Tightness

Proof of Proposition 3.3

We start with the first part. Consider the following distribution D_s of the seller: with probability $\frac{1}{2}$, v_s gets a value (uniformly at random) in $(0, \epsilon)$. With probability $\frac{1}{2}$, v_s gets a value (uniformly at random) between $(1, 1 + \epsilon)$. Observe that the median of D_s is 1.

Recall that every DSIC, individually rational and budget balanced mechanism in our setting is a posted price mechanism. Now there are two possible cases, depending on the trade price r:

1. $r \leq 1$: let $v_b = \infty$ with probability 1. The optimal solution always sells the item to the buyer, but the mechanism will sell the item with probability $\frac{1}{2}$. We get an approximation of 2 since $v_b >> v_s$.

2. $r > 1$: let $v_b = 1 - \epsilon$ with probability 1. The value of the optimal solution is at least $1 - \epsilon$ (always sell the item to the buyer). However, the mechanism sells the item only when $v_s \in (0, \epsilon)$, which happens with probability $\frac{1}{2}$. The approximation is 2 also in this case.

We now prove the second part. Consider the following distribution D_b of the buyer: with probability 0.99, v_b gets a value (uniformly at random) in $(1, 1 + \epsilon)$. With probability 0.01, v_b gets a value (uniformly at random) between $(100, 100 + \epsilon)$.

There are several possible cases, depending on the trade price r.

1. $r \leq 1 + \epsilon$: let $v_s = 1 + \epsilon$ with probability 1. The optimal solution sells the item to the buyer with probability $\frac{1}{1}00$ and the expected welfare is about 2. The mechanism that post the price r however will never sell the item and will generate an expected welfare of $1 + \epsilon$.

2. $1 + \epsilon < r \leq 100 + \epsilon$: let $v_s = 0$ with probability 1. The expected value of the optimal solution is about 2 (always sell the item to the buyer). However, the mechanism sells the item only when $v_b \in (100, 100 + \epsilon)$, so the expected welfare is about 1. The approximation is 2 also in this case.

3. $r > 100 + \epsilon$: let $v_s = 0$. The expected value of the optimal solution is 2, but the mechanism achieves welfare of 0.

B   Approximating the GFT

We first show that the GFT cannot be approximated by a budget balanced, individually rational and incentive compatible mechanism to within any constant factor. Notice that this result is proved via examples in which the optimal GFT is a negligible fraction of the social welfare, demonstrating the fact that when the expected GFT is a small fraction of the overall efficiency then approximating it may not be very informative. We

also make a straightforward observation that if the fraction of the optimal GFT out of the optimal overall efficiency is large, then an approximation to the optimal efficiency implies a good approximation ratio to the GFT as well (Appendix B.2).

B.1 An Impossibility Result for Approximating the GFT

Consider a buyer and a seller with values on the support $[0, ..., t]$, and let $\lambda = \frac{1}{1-e^{-t}}$. Let $F_b(x) = \lambda(1-e^{-x})$ with $f_b(x) = \lambda e^{-x}$ and $F_s(x) = \lambda(e^{x-t} - e^{-t})$ with $f_s(x) = \lambda e^{x-t}$.

Proposition B.1 *For the above distributions, every fixed price mechanism achieves at most $O(1/t)$ approximation to the optimal gain from trade.*

The optimal gain from trade (divided by the normalization factor of the distributions):

$$\frac{1}{\lambda^2}GFT^{\cdot} = \int_0^t \int_0^v (v-c)e^{c-t}dc\ e^{-v}dv$$
$$= \int_0^t \left[e^{c-t}(v-c+1)\right]_0^v\ e^{-v}dv$$
$$= \int_0^t \left(e^{v-t} - e^{-t}(v+1)\right)\ e^{-v}dv$$
$$= e^{-t}t - \int_0^t e^{-v-t}(v+1)dv$$
$$= e^{-t}t + \left[e^{-t-v}(v+2)\right]_0^t dv$$
$$= e^{-t}t + e^{-2t}(t+2) - 2e^{-t}$$
$$= \frac{t-2}{e^t} + \frac{t+2}{e^{2t}}$$

The gain from trade from posting a price p:

$$\frac{1}{\lambda^2}GFT(p) = \int_0^p \int_p^t (v-c)e^{-v}dv\ e^{c-t}dc$$
$$= \int_0^p \left[(c-v-1)e^{-v}\right]_p^t\ e^{c-t}dc$$
$$= \int_0^p \left((c-t-1)e^{-t} - (c-p-1)e^{-p}\right)\ e^{c-t}dc$$
$$= \int_0^p (c-t-1)e^{c-2t} - (c-p-1)e^{c-t-p}dc$$
$$= \left[e^{c-2t}(c-t-2)\right]_0^p - \left[e^{c-p-t}(c-p-2)\right]_0^p$$
$$= e^{p-2t}(p-t-2) - e^{-2t}(-t-2) + 2e^{-t} + e^{-p-t}(-p-2)$$
$$= \frac{t+2}{e^{2t}} + \frac{2}{e^t} - \frac{p+2}{e^{p+t}} - \frac{e^p(t+2-p)}{e^{2t}}$$
$$< \frac{t+2}{e^{2t}} + \frac{2}{e^t}$$

Overall, the optimal gain from trade is about $\frac{t-2}{e^t}$ which is $O(t)$ more than the gain from trade from any price p which is at most $\frac{2}{e^t}$. (Note that the terms that are $O(\frac{t}{e^{2t}})$ are negligible for large t's.)

B.2 A Simple Reduction from Efficiency to Gains From Trade

Observation B.2 *Let c be the fraction of the optimal expected gain-from-trade out of the optimal expected efficiency. If a mechanism achieves a x fraction of the optimal efficiency, then it achieves at least a $\frac{x+c-1}{c}$ fraction of the optimal gain-from-trade.*

Proof: We know that $\frac{OPTGFT}{OPT} = c$, where $OPTGFT$ is the optimal expected gain from trade (in the first-best outcome) and OPT is the optimal efficiency.

$OPT = OPTGFT + E[s]$, where $E[s]$ is the expected value of the seller. Thus, $\frac{E[s]}{OPT} = 1 - c$.

Now, the approximation we get for the GFT is ($MECH$ is the expected efficiency of the x-approximation mechanism):

$$\frac{MECH - E[s]}{OPT - E[s]} = \frac{\frac{MECH}{OPT} - \frac{E[s]}{OPT}}{1 - \frac{E[s]}{OPT}} \geq \frac{x - (1-c)}{c}$$

$\square$

For example, for the uniform distribution we have $OPTGFT = 1/6, OPT = 2/3$ and then $c = 1/4$. If we have a mechanism that gains us 90 percent of the welfare, we know it gains at least $(0.9 + 0.25 - 1)/0.25 = 0.6$ of the GFT.