

Non-excludable Bilateral Trade between Groups

Yixuan Even Xu^{1*}, Hanrui Zhang², Vincent Conitzer³

¹Tsinghua University

²Simons Laufer Mathematical Sciences Institute

³Carnegie Mellon University

xuyx20@mails.tsinghua.edu.cn, {hanruiz1,conitzer}@cs.cmu.edu

Abstract

Bilateral trade is one of the most natural and important forms of economic interaction: A seller has a single, indivisible item for sale, and a buyer is potentially interested. The two parties typically have different, privately known valuations for the item, and ideally, they would like to trade if the buyer values the item more than the seller. The celebrated impossibility result by Myerson and Satterthwaite shows that any mechanism for this setting must violate at least one important desideratum. In this paper, we investigate a richer paradigm of bilateral trade, with many self-interested buyers and sellers on both sides of a single trade who cannot be excluded from the trade. We show that this allows for more positive results. In fact, we establish a dichotomy in the possibility of trading efficiently. If in expectation, the buyers value the item more, we can achieve efficiency in the limit. If this is not the case, then efficiency cannot be achieved in general. En route, we characterize trading mechanisms that encourage truth-telling, which may be of independent interest. We also evaluate our trading mechanisms experimentally, and the experiments align with our theoretical results.

1 Introduction

Bilateral trade is one of the most natural and important forms of economic interaction in human society: A seller has a single, indivisible item for sale, and a buyer is potentially interested. The two parties typically have different valuations for the item, e.g., the seller has a production cost, and the buyer can derive a certain utility from consuming the item. If the buyer values the item more than the seller, then trading in an appropriate way would benefit both parties: The buyer could offer to pay the seller a price strictly between both parties' valuations, giving both parties positive gain. On the other hand, if the seller values the item more, then a trade would hurt the overall utility of them, and thus should not take place. We can ensure that the trade takes place if and only if it is efficient by setting an appropriate price, *as long as both parties' valuations are public information*.

The situation is much subtler in the more realistic setting where the two parties' valuations are private. There, maximizing the gain from trade becomes a *mechanism design*

problem: We want to design a mechanism that, based on both parties' private valuations, decides whether they should trade, and if so, the amounts of money that the buyer pays and the seller receives. Such a mechanism can be viewed as a trusted mediator that facilitates trade. The key difficulty here is to encourage the buyer and the seller to truthfully reveal their private valuations, so that the mechanism can make the most socially efficient decisions, i.e., to trade whenever the buyer values the item more than the seller.

This basic and seemingly simple problem turns out to be surprisingly challenging. In fact, the seminal impossibility result of Myerson and Satterthwaite (1983) shows that in the above setting with private valuations, no mechanism simultaneously satisfies the following four natural properties:

- *(Bayesian) incentive compatibility*: In expectation, both parties maximize their respective utilities by truthfully reporting their valuations to the mechanism.
- *(Bayesian) individual rationality*: The expected values of participating are non-negative for both parties.
- *Weak budget balance*: If a trade happens, the amount of money that the buyer pays should be at least the amount that the seller receives, so that the mediator of the mechanism does not need to subsidize the trade.
- *Efficiency*: The two parties trade if and only if the buyer values the item more than the seller.

Requiring incentive compatibility is WLOG given the revelation principle, so if we want to implement the socially efficient outcome, then we have to (1) force at least one party to reveal their true valuation, (2) force at least one party to participate, or (3) subsidize the trade — none of which is feasible in practice. This gives us one of the most famous impossibility results in economics: *Efficient bilateral trade cannot be implemented in a feasible way*.

Interestingly, the above impossibility result is only partially echoed by real-life phenomena. Admittedly, negotiation between two individuals is often highly unpredictable and not fully efficient. However, when we look at trade between two sizable “organizations” (e.g., countries) each consisting of many self-interested “members” (e.g., citizens), it appears that the aggregate decisions made by these members more often result in an outcome that overall benefits members of both organizations due to the large number of participants and is thus more efficient. In such scenarios, the

*This work was done when Xu was a visiting intern at CMU. Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

multiplicity of members in the two parties seems to provide a way around the classical impossibility result by Myerson and Satterthwaite (1983), opening up a new design space.

In this paper, we investigate this richer and seemingly more permitting paradigm of bilateral trade, with many self-interested members on both sides of the trade. Roughly speaking, we consider a model where there are n buyers and n sellers on the two sides of the trade respectively¹, and all agents on each side share the same allocation and payment. For concreteness, consider the following example: Amateur soccer club B wants to buy the right to use a certain soccer field from amateur soccer club S. Members of the buying club B generally have different valuations for the right to the field, depending on how much each individual member enjoys training and playing. These valuations are drawn independently from a common prior distribution. The same also applies to the selling club S (with a different prior). If the two clubs close a deal, all members of the buying club B get to enjoy the right to the field, and at the same time, they all make the same payment in the form of raised membership fees. As for members of the selling club S, they lose the right to the field, but each member receives the same payment in the form of lowered membership fees. In other words, members on each side of the trade share the same allocation and payment, and no one can be excluded from the trade.

In the above setting, ideally, the two clubs should trade if and only if on average, members of the buying club B value the right to the field more than members of the selling club S. There are also other high-stakes real-world trading scenarios that are similar to the soccer club example, such as negotiations between countries, residential communities, or academic units (e.g., universities, schools, or departments). In such settings, the preferences of all members should be accounted for, so naturally, a trading mechanism should take into consideration the reported valuations of each individual member of the two sides of the trade, and based on all this information, try to make the optimal decision for the society. Our goal is to investigate the power and limitations of trading mechanisms in such settings, subject to participation and incentive constraints.

1.1 Our Results

We begin our investigation by looking at the gains-from-trade of the first-best mechanism. We show that when the buyers value the item more than the sellers in expectation, the first-best gains-from-trade (defined in Section 3) is proportional to the number of agents, and when the sellers value the item (weakly) more, even the first-best gains-from-trade per agent goes to zero (Lemma 4.1). This naturally results in two different cases: when the buyers value the item more, trading in the right way can benefit each individual agent significantly, whereas when the sellers value the item more, there is not much value to be gained in the first place. We will approach these two cases separately.

¹In most of our results, we treat n as the only asymptotic variable. Therefore, we assume the two parties consist of the same number of agents n for simplicity of presentation. Our results can also be adapted to the case when they are different.

With this in mind, we proceed by first looking at deterministic trading mechanisms, which are both more practical and structurally simpler. Our main result is a dichotomy: We show that if, in expectation, buyers value the item strictly more than sellers (i.e., we are in the first case discussed above), then essentially all desiderata discussed above can be achieved using a simple deterministic trading mechanism (Theorem 4.3). Alternatively, if sellers value the item at least as much as buyers in expectation (i.e., we are in the second case discussed above), then feasible mechanisms, in general, can only achieve an approximation to the socially optimal outcome, by a factor strictly smaller than 1. This holds even if we only require incentive compatibility (Theorem 4.4). So in short, the paradigm with multiple agents on both sides is in fact more permitting, but this makes a fundamental difference only when trading is beneficial in expectation (i.e., the first case). Nonetheless, when trading is not beneficial in expectation (i.e., the second case), there is little to be gained in the first place, so the fact that efficiency cannot be achieved in a feasible way does not cost much.

Technically, the above results are built on a powerful and intriguing characterization of feasible deterministic mechanisms: We show that any deterministic trading mechanism that is incentive compatible (i.e., encourages truth-telling) must proceed in a voting-like way (Theorem 4.1). In particular, if we require strong budget balance, then there must be a unique, predetermined price independent of the agents' valuations (Theorem 4.2). The mechanism decides whether to trade by asking each individual agent whether they want to trade at the above price. Moreover, in order to encourage truth-telling, this voting procedure has to be monotone (i.e., if the trade happens when a certain set of agents approve, then it must also happen when a superset of those agents approve). More generally, this characterization might also be useful in other mechanism design problems that involve certain forms of non-excludability.

Given the dichotomy in the case of deterministic mechanisms, one natural question is whether randomization can help. To this end, we characterize all incentive compatible randomized mechanisms that are "well-behaved" (i.e., twice continuously differentiable): Such mechanisms must be separable across different agents, in the sense that fixing the asks (resp. bids) of all sellers (resp. buyers), for each buyer (resp. seller), there is a component of the allocation function that depends only on the bid of that buyer (resp. seller), and the overall allocation function is the sum of all these components (Theorem 5.1). Somewhat counterintuitively, this characterization implies an even stronger negative result: When sellers in expectation value the item more, no "well-behaved" randomized mechanisms that encourage truth-telling can achieve any constant approximation against the socially optimal outcome (Theorem 5.2). In other words, randomization, at least when applied in a smooth way, cannot help in the paradigm of bilateral trade that we study.

Finally, we evaluate our trading mechanism with different classes of valuation distributions empirically. The experimental results align with our theoretical results. In particular, our mechanism consistently achieves all desiderata almost perfectly whenever the number of agents is large enough.

2 Related Work

Bilateral trade. Bilateral trade has been extensively studied in the classical setting where there is a single agent on each side of the trade. The seminal impossibility result of (Myerson and Satterthwaite 1983) shows that no mechanisms for bilateral trade can be both efficient and feasible (i.e., Bayesian incentive compatible, Bayesian individual rational, and weak budget balanced). As for approximation of gains-from-trade, (McAfee 2008) gives a 2-approximation of the first best when the value distribution has a higher median than the cost distribution. (Blumrosen and Mizrahi 2016) gives an ϵ -approximation assuming a monotone hazard rate of the value distribution. (Brustle et al. 2017) shows an unconditional 2-approximation of the second best. (Deng et al. 2022) shows an unconditional 8.23-approximation of the first best, later improved into a 3.15-approximation by (Fei 2022). Our results in this paper extend and complement existing research on bilateral trade.

Double auctions. A related line of work to our setting is double auctions, which also generalizes bilateral trade to the case when there are multiple buyers and sellers. There is a significant body of work on double auctions (Babaiouff et al. 2018; Babaiouff, Goldner, and Gonczarowski 2020; Balseiro et al. 2019; Brustle et al. 2017; Cai et al. 2021; McAfee 1992; Colini-Baldeschi et al. 2016, 2020, 2017). In the setting that we study, allocation and payment are shared across all agents, which is not the case in double auctions.

Non-excludable mechanism design. Another related line of work is non-excludable mechanism design. Here, “non-excludable” means that if something is allocated, then it is allocated to all agents participating in the mechanism, i.e., no agent can be excluded from the allocation (Samuelson 1954; Mugaev 1959; Hardin 1968; Ostrom 1990). A common example of non-excludable mechanisms is public projects, where the auctioneer allocates the same amount of the public goods to all players simultaneously. There is a rich body of work on this topic (Papadimitriou, Schapira, and Singer 2008; Buchfuhrer, Schapira, and Singer 2010; Dobzinski 2011; Dughmi, Roughgarden, and Yan 2016; Balseiro et al. 2021). In our setting, the allocation, payment, and receipt are similarly shared among all participants in each group. To the best of our knowledge, our work is the first to consider two sides trading in the context of non-excludable mechanism design with shared payment.

3 Preliminaries

A set of n sellers has a certain product that it can sell to a set of n buyers. Each buyer $i \in [n]$ has a private value v_i , and each seller $j \in [n]$ has a private cost c_j , for if the trade happens. We denote the vectors of buyer values and seller costs as \mathbf{v} and \mathbf{c} , respectively. According to the revelation principle, we can restrict our attention to incentive compatible mechanisms \mathcal{M} , where each buyer i submits a sealed bid b_i and each seller j submits a sealed ask a_j . Let \mathbf{b} be the vector of all buyer bids and \mathbf{a} be the vector of all seller asks. The mechanism then decides an allocation $x(\mathbf{b}, \mathbf{a}) \in [0, 1]$ shared by all participants, indicating the probability of trad-

ing, a payment $p(\mathbf{b}, \mathbf{a}) \in \mathbb{R}$ shared by all buyers, and a receipt $r(\mathbf{b}, \mathbf{a}) \in \mathbb{R}$ shared by all sellers.

We study this multiplayer bilateral trade (MBT) problem in the Bayesian setting, where we assume that the buyer values \mathbf{v} are drawn i.i.d. from a known bounded prior distribution F and the seller costs \mathbf{c} are drawn i.i.d. from a known bounded prior distribution G . Without loss of generality, let F and G be bounded on $[0, 1]$. We also use F and G to denote the cumulative distribution functions, i.e., $F(\tau) = \Pr_{v \sim F}[v \leq \tau]$, $G(\tau) = \Pr_{c \sim G}[c \leq \tau]$ ($\tau \in [0, 1]$).

For buyer i and seller j , the quasilinear utilities are defined as $u_i(\mathbf{b}, \mathbf{a}) = x(\mathbf{b}, \mathbf{a}) \cdot v_i - p(\mathbf{b}, \mathbf{a})$ and $u_j(\mathbf{b}, \mathbf{a}) = r(\mathbf{b}, \mathbf{a}) - x(\mathbf{b}, \mathbf{a}) \cdot c_j$ respectively. Incentive compatibility means that by truthfully revealing the private value or cost, each buyer or seller obtains the outcome maximizing his or her utility regardless of the bids and asks from others. For simplicity, we assume ties in utility are broken in favor of a trade. That is, if multiple outcomes share the same utility, the outcome with larger x is preferred by the participants. Then, in an incentive compatible mechanism, it is a dominant strategy to bid $\mathbf{b} = \mathbf{v}$ and ask $\mathbf{a} = \mathbf{c}$. As we consider an environment that requires shared allocation, payment, and receipt, we generalize the standard notion of individual rationality to one that considers each group as a whole, i.e., the total utility of buyers and sellers should be positive respectively.

For the bilateral trade problem, the mechanism designer is commonly concerned with the gains-from-trade (GFT), which refers to the expected utility gain from trading. Formally, for an incentive compatible mechanism \mathcal{M} , $\text{GFT} = \mathbf{E}_{\mathbf{v} \sim F^n, \mathbf{c} \sim G^n}[(\sum_i v_i - \sum_j c_j) \cdot x(\mathbf{v}, \mathbf{c})]$. To maximize GFT, ideally, we want the trade to happen whenever the total value of the buyers is greater than the total cost of the sellers. We call this optimal GFT the *first best* (FB). That is, FB equals

$$\mathbf{E}_{\mathbf{v} \sim F^n, \mathbf{c} \sim G^n} \left[\left(\sum_i v_i - \sum_j c_j \right) \mathbf{1} \left[\sum_i v_i \geq \sum_j c_j \right] \right].$$

We are interested in mechanisms that are incentive compatible (IC), individually rational (IR) in the above sense, weakly budget balanced (WBB), and *efficient in the limit*, meaning that as $n \rightarrow \infty$, the mechanism’s GFT ALG approaches FB, i.e., $\frac{\text{ALG}}{\text{FB}} = 1 - o(1)$.

4 Deterministic Mechanisms

In this section, we consider deterministic mechanisms for multiplayer bilateral trade (MBT). Namely, the allocation function $x(\mathbf{b}, \mathbf{a}) \in \{0, 1\}$ is either 0 or 1. Before we proceed, we first present the following Lemma 4.1. The proof follows standard calculation of concentration inequalities and is deferred to Appendix A.1.

Lemma 4.1. *When $\mathbf{E}_{v \sim F}[v] \leq \mathbf{E}_{c \sim G}[c]$, $\text{FB} = O(\sqrt{n})$. When $\mathbf{E}_{v \sim F}[v] > \mathbf{E}_{c \sim G}[c]$, $\text{FB} = \Omega(n)$.*

Lemma 4.1 naturally divides the MBT problem into two cases: when the buyers value the item more in expectation, and when the sellers value the item (weakly) more in expectation. In the latter case, the gains-from-trade per agent goes to zero even under the first-best mechanism.

With these two different cases in mind, we will approach them separately. We will first characterize the set of allocation functions that are implementable by an IC mechanism in Section 4.1. Using this characterization, we establish a dichotomy in MBT: We will present a positive result when buyers value the item more in expectation (Section 4.2), and then show in the case that sellers value the item (weakly) more in expectation, no mechanism can, in general, achieve all the desiderata (Section 4.3).

4.1 Characterization of Incentive Compatible Mechanisms

The following Theorem 4.1 is our characterization of the set of allocation functions that are implementable by an IC mechanism. If we additionally require the mechanism to be strongly budget balanced (SBB), i.e., the total payment of the buyers always equals the total receipt of the sellers, then it can be further simplified as described in Theorem 4.2.

At a high level, our characterization shows that any deterministic trading mechanism that is incentive compatible must proceed in a voting-like way. In particular, if we require strong budget balance, then there must be a unique, predetermined price independent of the agents' valuations. The way the mechanism decides whether to trade is by asking each individual agent whether they want to trade at the above price, which can also be viewed as running a voting procedure. Moreover, in order to encourage truth-telling, this voting procedure has to be monotone (i.e., if the trade happens when a certain set of agents approve, then it must also happen when a superset of those agents approve).

Theorem 4.1. *A deterministic allocation function $x(\mathbf{b}, \mathbf{a})$ for MBT can be implemented by an IC mechanism if and only if both of the following statements hold.*

- (a) $\forall \mathbf{a} \in [0, 1]^n$, there exists $\tau_{\mathbf{a}} \in [0, 1]$ and a monotone Boolean function $f_{\mathbf{a}} : \{0, 1\}^n \rightarrow \{0, 1\}$, such that $x(\mathbf{b}, \mathbf{a}) = f_{\mathbf{a}}(\mathbf{1}[b_1 \geq \tau_{\mathbf{a}}], \dots, \mathbf{1}[b_n \geq \tau_{\mathbf{a}}])$,
- (b) $\forall \mathbf{b} \in [0, 1]^n$, there exists $\theta_{\mathbf{b}} \in [0, 1]$ and a monotone Boolean function $g_{\mathbf{b}} : \{0, 1\}^n \rightarrow \{0, 1\}$ such that $x(\mathbf{b}, \mathbf{a}) = g_{\mathbf{b}}(\mathbf{1}[a_1 \leq \theta_{\mathbf{b}}], \dots, \mathbf{1}[a_n \leq \theta_{\mathbf{b}}])$.

Theorem 4.2. *A deterministic allocation function $x(\mathbf{b}, \mathbf{a})$ for MBT can be implemented by an IC and SBB mechanism if and only if there exists $\tau \in [0, 1]$ and a monotone Boolean function $f : \{0, 1\}^{2n} \rightarrow \{0, 1\}$ such that $x(\mathbf{b}, \mathbf{a}) = f(\mathbf{1}[b_1 \geq \tau], \dots, \mathbf{1}[b_n \geq \tau], \mathbf{1}[a_1 \leq \tau], \dots, \mathbf{1}[a_n \leq \tau])$.*

To give proofs to both of the theorems, we first invoke Myerson's Lemma from (Myerson 1981).

Lemma 4.2 (Myerson's Lemma (Myerson 1981)). *Consider a single-parameter environment with n bidders where bidder i has a private value $v_i \sim F_i$. Let the utility of bidder i be $u_i(\mathbf{b}) = v_i \cdot x_i(\mathbf{b}) - p_i(\mathbf{b})$. Here \mathbf{b} are the bidders' bids and $x_i(\mathbf{b}), p_i(\mathbf{b})$ are defined by the mechanism. Then*

- (a) Allocation function \mathbf{x} is implementable by an IC mechanism if and only if for every $i \in [n]$, $x_i(\mathbf{b})$ is non-decreasing in b_i .

- (b) If $\forall i \in [n]$, $x_i(\mathbf{b})$ is non-decreasing in b_i , then mechanism (\mathbf{x}, \mathbf{p}) is IC if and only if

$$p_i(b_i, \mathbf{b}_{-i}) = b_i \cdot x_i(b_i, \mathbf{b}_{-i}) - \int_{z=-\infty}^{b_i} x_i(z, \mathbf{b}_{-i}) dz + h_i(\mathbf{b}_{-i}) \quad \forall i \in [n].$$

Here, h_i is a function unrelated to b_i for each $i \in [n]$.

Note that in the context of Lemma 4.2, the allocation x_i and payment p_i functions for each bidder i can be different, but in our MBT setting, the allocation function x is shared by all participants, and the payment p and receipt r functions are shared by all buyers and sellers respectively. This further restricts the range of feasible allocation functions. A translated version of Lemma 4.2 in our setting is presented as Lemma 4.3 below with its proof deferred to Appendix A.2.

Lemma 4.3. *If an allocation function $x(\mathbf{b}, \mathbf{a})$ for MBT can be implemented by an IC mechanism (x, p, r) , then all of the following claims hold.*

- (a) $x(\mathbf{b}, \mathbf{a})$ is non-decreasing in b_i for all $i \in [n]$ and non-increasing in a_j for all $j \in [n]$.
- (b) There exist functions h_i for each $i \in [n]$ such that

$$p(b_i, \mathbf{b}_{-i}, \mathbf{a}) = b_i \cdot x(b_i, \mathbf{b}_{-i}, \mathbf{a}) - \int_{z=-\infty}^{b_i} x(z, \mathbf{b}_{-i}, \mathbf{a}) dz + h_i(\mathbf{b}_{-i}, \mathbf{a}) \quad \forall i \in [n].$$

- (c) There exist functions h_j for each $j \in [n]$ such that

$$r(\mathbf{b}, a_j, \mathbf{a}_{-j}) = a_j \cdot x(\mathbf{b}, a_j, \mathbf{a}_{-j}) + \int_{z=a_j}^{+\infty} x(\mathbf{b}, z, \mathbf{a}_{-j}) dz + h_j(\mathbf{b}, \mathbf{a}_{-j}) \quad \forall j \in [n].$$

- (d) $\{(\mathbf{b}, \mathbf{a}) \mid x(\mathbf{b}, \mathbf{a}) = 1\}$ is a closed set.

Applying Lemma 4.3, we will be able to prove Lemma 4.4 which relates the allocation function to the payment and receipt functions. Its proof can be found in Appendix A.3.

Lemma 4.4. *If an allocation function $x(\mathbf{b}, \mathbf{a})$ for MBT can be implemented by an IC mechanism (x, p, r) , then for any $0 \leq \alpha \leq \beta \leq 1$ and $i, j \in [n]$*

- (a) $x(\alpha, \mathbf{b}_{-i}, \mathbf{a}) = x(\beta, \mathbf{b}_{-i}, \mathbf{a})$
 $\Rightarrow p(\alpha, \mathbf{b}_{-i}, \mathbf{a}) = p(\beta, \mathbf{b}_{-i}, \mathbf{a})$,
- (b) $x(\mathbf{b}, \alpha, \mathbf{a}_{-j}) = x(\mathbf{b}, \beta, \mathbf{a}_{-j})$
 $\Rightarrow r(\mathbf{b}, \alpha, \mathbf{a}_{-j}) = r(\mathbf{b}, \beta, \mathbf{a}_{-j})$.

Note that for SBB mechanisms, $p(\mathbf{b}, \mathbf{a}) = r(\mathbf{b}, \mathbf{a})$. By iteratively applying Lemma 4.4 to each applicable coordinate, we can see the following Corollary 4.1 holds.

Corollary 4.1. *If an allocation function $x(\mathbf{b}, \mathbf{a})$ for MBT can be implemented by an IC mechanism (x, p, r) , then for any $\mathbf{b}^{(0)} \preceq \mathbf{b}^{(1)}, \mathbf{a}^{(0)} \succeq \mathbf{a}^{(1)}$,*

- (a) $x(\mathbf{b}^{(0)}, \mathbf{a}) = x(\mathbf{b}^{(1)}, \mathbf{a}) \Rightarrow p(\mathbf{b}^{(0)}, \mathbf{a}) = p(\mathbf{b}^{(1)}, \mathbf{a})$,
- (b) $x(\mathbf{b}, \mathbf{a}^{(0)}) = x(\mathbf{b}, \mathbf{a}^{(1)}) \Rightarrow r(\mathbf{b}, \mathbf{a}^{(0)}) = r(\mathbf{b}, \mathbf{a}^{(1)})$.

Moreover, if $x(\mathbf{b}, \mathbf{a})$ for MBT can be implemented by an IC and SBB mechanism,

$$(c) \quad x(\mathbf{b}^{(0)}, \mathbf{a}^{(0)}) = x(\mathbf{b}^{(1)}, \mathbf{a}^{(1)}) \Rightarrow p(\mathbf{b}^{(0)}, \mathbf{a}^{(0)}) \\ = p(\mathbf{b}^{(1)}, \mathbf{a}^{(1)}) = r(\mathbf{b}^{(0)}, \mathbf{a}^{(0)}) = r(\mathbf{b}^{(1)}, \mathbf{a}^{(1)}).$$

Here \preceq and \succeq mean componentwise \leq and \geq respectively.

Let $\vec{0}, \vec{1}$ be vectors of n zeros and n ones respectively. Next, we will present two lemmas connecting back to Theorem 4.1 and Theorem 4.2 with their proofs deferred to Appendices A.4 and A.5.

Lemma 4.5. *If a deterministic allocation function $x(\mathbf{b}, \mathbf{a})$ for MBT can be implemented by an IC mechanism, then the following statements hold.*

- (a) For any $\mathbf{a} \in [0, 1]^n$, let $\tau_{\mathbf{a}} = p(\vec{1}, \mathbf{a}) - p(\vec{0}, \mathbf{a})$, then

$$\mathbf{1} \left[\alpha \geq \tau_{\mathbf{a}} \right] = \mathbf{1} \left[\beta \geq \tau_{\mathbf{a}} \right] \Rightarrow x(\alpha, \mathbf{b}_{\cdot}, \mathbf{a}) = x(\beta, \mathbf{b}_{\cdot}, \mathbf{a}).$$
- (b) For any $\mathbf{b} \in [0, 1]^n$, let $\theta_{\mathbf{b}} = r(\mathbf{b}, \vec{0}) - r(\mathbf{b}, \vec{1})$, then

$$\mathbf{1} \left[\alpha \leq \theta_{\mathbf{b}} \right] = \mathbf{1} \left[\beta \leq \theta_{\mathbf{b}} \right] \Rightarrow x(\mathbf{b}, \alpha, \mathbf{a}_{\cdot}) = x(\mathbf{b}, \beta, \mathbf{a}_{\cdot}).$$

There is also a version of Lemma 4.5 for IC and SBB mechanisms with an almost identical proof.

Lemma 4.6. *If a deterministic allocation function $x(\mathbf{b}, \mathbf{a})$ for MBT can be implemented by an IC and SBB mechanism (x, p, r) , then let $\tau = p(\vec{1}, \vec{0}) - p(\vec{0}, \vec{1})$, $\forall 0 \leq \alpha \leq \beta \leq 1$,*

- (a) $\mathbf{1} \left[\alpha \geq \tau \right] = \mathbf{1} \left[\beta \geq \tau \right] \Rightarrow x(\alpha, \mathbf{b}_{\cdot}, \mathbf{a}) = x(\beta, \mathbf{b}_{\cdot}, \mathbf{a})$,
- (b) $\mathbf{1} \left[\alpha \leq \tau \right] = \mathbf{1} \left[\beta \leq \tau \right] \Rightarrow x(\mathbf{b}, \alpha, \mathbf{a}_{\cdot}) = x(\mathbf{b}, \beta, \mathbf{a}_{\cdot})$.

With all of the above lemmas, we are able to prove Theorem 4.1 and Theorem 4.2. In fact, the necessity of Theorem 4.1 (resp. Theorem 4.2) is already implied by Lemma 4.5 (resp. Lemma 4.6). The proofs of the sufficiency of these theorems can be found in Appendices A.6 and A.7.

4.2 When Expected Value is Greater than Expected Cost: A Simple Efficient Mechanism

Theorems 4.1 and 4.2 provide a characterization of the design space. Under this design space, we will first show a simple mechanism that is IC, SBB, IR in the limit and efficient in the limit when the expected value is greater than the expected cost, i.e., $\mathbf{E}_{v \sim F}[v] > \mathbf{E}_{c \sim G}[c]$.

Algorithm 1: Forced-Trade Mechanism

Trade if and only if $\mathbf{E}_{v \sim F}[v] > \mathbf{E}_{c \sim G}[c]$, with both payment and receipt being $0.5(\mathbf{E}_{v \sim F}[v] + \mathbf{E}_{c \sim G}[c])$. I.e.,

$$x(\mathbf{b}, \mathbf{a}) = \mathbf{1} \left[\mathbf{E}_{v \sim F}[v] > \mathbf{E}_{c \sim G}[c] \right], \\ p(\mathbf{b}, \mathbf{a}) = r(\mathbf{b}, \mathbf{a}) = 0.5(\mathbf{E}_{v \sim F}[v] + \mathbf{E}_{c \sim G}[c]) \cdot x(\mathbf{b}, \mathbf{a}).$$

We show Algorithm 1 achieves all desiderata in the limit.

Theorem 4.3. *Algorithm 1 is IC and SBB. When $\mathbf{E}_{v \sim F}[v] > \mathbf{E}_{c \sim G}[c]$, as $n \rightarrow \infty$, Algorithm 1 is IR with probability $1 - e^{-\Omega(n)}$, and $\frac{\text{ALG}}{\text{FB}} = 1 - e^{-\Omega(n)}$ where ALG is its GFT.*

The proof of Theorem 4.3 can be found in Appendix A.8.

4.3 When Expected Value is No Greater than Expected Cost: No Efficient Mechanisms

Theorem 4.3 shows that when the expected value is greater than the expected cost, the extremely simple Algorithm 1 achieves all desiderata in the limit. In our next Theorem 4.4, we will additionally show that when this is not the case, no mechanisms can be both IC and efficient in the limit.

Theorem 4.4. *When $\mathbf{E}_{v \sim F}[v] \leq \mathbf{E}_{c \sim G}[c]$, no IC mechanisms for MBT can be efficient in the limit.*

The proof of Theorem 4.4 is built on our characterization (Theorem 4.1) and deferred to Appendix A.9. By Theorem 4.4, when the expected value is no greater than the expected cost, the multiplicative approximation ratio of any IC deterministic mechanism is $1 - \Omega(1)$, i.e., no deterministic IC mechanism can be efficient in the limit. However, recall that Lemma 4.1 shows that in this case, the first best is $o(n)$. This implies that *additively*, Algorithm 1 does not lose much compared to the first best even if it does not trade at all.

5 Randomized Mechanisms

In this section, we move on to consider twice continuously differentiable randomized mechanisms for multiplayer bilateral trade (MBT). Namely, the allocation function $x(\mathbf{b}, \mathbf{a}) \in [0, 1]$ and x is twice continuously differentiable. Note that Lemma 4.1 still holds in this context. Therefore, we will take a similar approach with Section 4. We will still first characterize the set of allocation functions that are implementable by an IC mechanism in Section 5.1. When the expected value is greater than the expected cost, Algorithm 1 in Section 4.2 is still applicable as the allocation function of Algorithm 1 is also twice continuously differentiable. And when the expected value is not greater than the expected cost, we show in Section 5.2 that no twice continuously differentiable randomized mechanism can achieve all the desiderata, similarly to Section 4. In other words, completely smooth randomness cannot help in MBT. Surprisingly, our negative result is stronger than Theorem 4.4.

5.1 Characterization of Incentive Compatible Mechanisms

We characterize the set of all twice continuously differentiable randomized allocation functions that are implementable by an IC mechanism in the following Theorem 5.1.

From a general perspective, our characterization shows that such mechanisms must be separable across different agents, which means that fixing the asks (resp. bids) of all sellers (resp. buyers), for each buyer (resp. seller), there is a component of the allocation function that depends only on the bid of that buyer (resp. seller), and the overall allocation function is the sum of all these components

Theorem 5.1. *A twice continuously differentiable randomized allocation function $x(\mathbf{b}, \mathbf{a})$ for MBT can be implemented by an IC mechanism if and only if both of the following statements hold.*

- (a) For all $\mathbf{a} \in [0, 1]^n$, there exist n non-decreasing differentiable functions $f_{\mathbf{a}, i} : [0, 1] \rightarrow \mathbb{R}, \forall i \in [n]$, such that $x(\mathbf{b}, \mathbf{a}) = f_{\mathbf{a}, 1}(b_1) + f_{\mathbf{a}, 2}(b_2) + \dots + f_{\mathbf{a}, n}(b_n)$.

(μ_F, μ_G)		(0.6, 0.4)		(0.55, 0.45)		(0.51, 0.49)	
Distribution	n	IR	Efficiency	IR	Efficiency	IR	Efficiency
Normal	5	0.804033	0.992350	0.539705	0.899567	0.302261	0.359552
	100	1.000000	1.000000	0.995614	0.999999	0.510591	0.870741
	10000	1.000000	1.000000	1.000000	1.000000	1.000000	1.000000
Uniform	5	0.684952	0.972358	0.464824	0.823521	0.288943	0.290061
	100	0.999992	1.000000	0.969823	0.999900	0.445284	0.788148
	10000	1.000000	1.000000	1.000000	1.000000	0.999989	1.000000
Bernoulli	5	0.465557	0.809241	0.351963	0.558944	0.268607	0.150281
	100	0.966587	0.999779	0.749410	0.975987	0.382188	0.512118
	10000	1.000000	1.000000	1.000000	1.000000	0.955986	0.999752
Mixed	5	0.566107	0.909725	0.401324	0.694859	0.277582	0.208867
	100	0.997254	0.999999	0.870417	0.997094	0.382296	0.655883
	10000	1.000000	1.000000	1.000000	1.000000	0.997250	1.000000

Table 1: Performance of Algorithm 1 under various conditions. IR is the empirical probability of Algorithm 1 being IR and efficiency is the empirical expected GFT/FB. All data are averaged over 10^6 runs.

(b) For all $\mathbf{b} \in [0, 1]^n$, there exist n non-increasing differentiable functions $g_{\mathbf{b},j} : [0, 1] \rightarrow \mathbb{R}, \forall j \in [n]$, such that $x(\mathbf{b}, \mathbf{a}) = g_{\mathbf{b},1}(a_1) + g_{\mathbf{b},2}(a_2) + \dots + g_{\mathbf{b},n}(a_n)$.

Below, we will first provide two lemmas. Theorem 5.1 can then be proved by combining these two lemmas.

Lemma 5.1. For a twice continuously differentiable function $f(x_1, x_2, \dots, x_n) : \mathbb{R}^n \rightarrow \mathbb{R}$, the following statements (i) and (ii) are equivalent.

(i) There exist n differentiable functions $f_i : \mathbb{R} \rightarrow \mathbb{R}$, that $f(x_1, x_2, \dots, x_n) = \sum_{i=1}^n f_i(x_i)$.

(ii) For any $i, j \in [n], i \neq j, \frac{\partial^2 f}{\partial x_i \partial x_j}(x_1, x_2, \dots, x_n) = 0$.

Lemma 5.2. If a twice continuously differentiable randomized allocation function $x(\mathbf{b}, \mathbf{a})$ for MBT can be implemented by an IC mechanism, then the followings hold.

(a) Fix $\mathbf{a} \in [0, 1]^n$,

$$\frac{\partial^2 x}{\partial b_{i_1} \partial b_{i_2}}(\mathbf{b}, \mathbf{a}) = 0, \forall i_1, i_2 \in [n], i_1 \neq i_2.$$

(b) Fix $\mathbf{b} \in [0, 1]^n$,

$$\frac{\partial^2 x}{\partial a_{j_1} \partial a_{j_2}}(\mathbf{b}, \mathbf{a}) = 0, \forall j_1, j_2 \in [n], j_1 \neq j_2.$$

With Lemmas 5.1 and 5.2, we are ready to show Theorem 5.1. The detailed proofs of Lemmas 5.1 and 5.2 and Theorem 5.1 are deferred to Appendices B.1 to B.3.

5.2 When Expected Value is No Greater than Expected Cost: No Efficient Mechanisms

At first glance, one might think twice continuously differentiable randomized mechanisms should be more powerful than deterministic mechanisms because the choice of allocation function seems much wider. However, we show in the following Theorem 5.2 that on the contrary, these mechanisms are not even capable of being a constant approximation of FB in the limit.

Theorem 5.2. When $\mathbf{E}_{v \sim F}[v] \leq \mathbf{E}_{c \sim G}[c]$, no IC twice continuously differentiable mechanisms for MBT can have a GFT that is a constant approximation of FB in the limit.

The proof of Theorem 5.2 can be found in Appendix B.4.

6 Experiments

In the previous sections, we have shown that the simple Algorithm 1 achieves all desiderata in the limit asymptotically when $\mathbf{E}_{v \sim F}[v] > \mathbf{E}_{c \sim G}[c]$ and no mechanisms can do so when $\mathbf{E}_{v \sim F}[v] \leq \mathbf{E}_{c \sim G}[c]$. In this section, we study the non-asymptotic performance of Algorithm 1 by running experiments on various generated datasets. Only simulated data without real individuals are involved in the experiments. As Algorithm 1 is always IC and SBB, we focus on studying the empirical probability that Algorithm 1 is IR and its empirical efficiency on datasets of practically relevant sizes.

Experiment setup. We implemented Algorithm 1 in C++. The source code can be found at <https://github.com/YixuanEvenXu/bilateral-trade>. The datasets we used were generated as follows. The distributions F and G are both supported on $[0, 1]$. In every dataset, F and G belong to one of the 4 families of distributions given below and were truncated to the specified support whenever necessary.

- (1) Normal: F, G are normal distributions $\mathcal{N}(\mu, \sigma^2)$ where $\mu = \mu_F$ or μ_G and $\sigma = 0.2$.
- (2) Uniform: F, G are uniform distributions $\mathcal{U}[\mu - r, \mu + r]$ where $\mu = \mu_F$ or μ_G and $r = 0.4$.
- (3) Bernoulli: F, G are Bernoulli distributions $\text{Ber}(\mu)$ where $\mu = \mu_F$ or μ_G .
- (4) Mixed: F, G are mixed distributions with $\frac{1}{3}$ probability of (1), (2), and (3) each.

For each type of distribution above, we performed various tests with $n = \{5, 100, 10000\}$ and $(\mu_F, \mu_G) = \{(0.6, 0.4), (0.55, 0.45), (0.51, 0.49)\}$. The complete experimental results are shown in Table 1. All of the data are averaged over 10^6 runs. We also highlight a few results in Figs. 1 and 2 to make it easier to interpret the results.

Comparing when $n = 5, 100$, and 10000. As shown in Table 1 and Fig. 1, Algorithm 1’s empirical probability of being IR and empirical efficiency generally increase as n scales up. This is consistent with the asymptotic guarantees of Algorithm 1 in Theorem 4.3. Moreover, as long as




Figure 1: Performance of Algorithm 1 when $(\mu_F, \mu_G) = (0.55, 0.45)$ and $n = 5, 100, 10000$. Error bars indicate standard error of mean. The empirical probability of being IR and the empirical efficiency converge to 1.0 as n increases. Algorithm 1 has a good performance for reasonably large n . Errors are negligible due to the large number of samples.

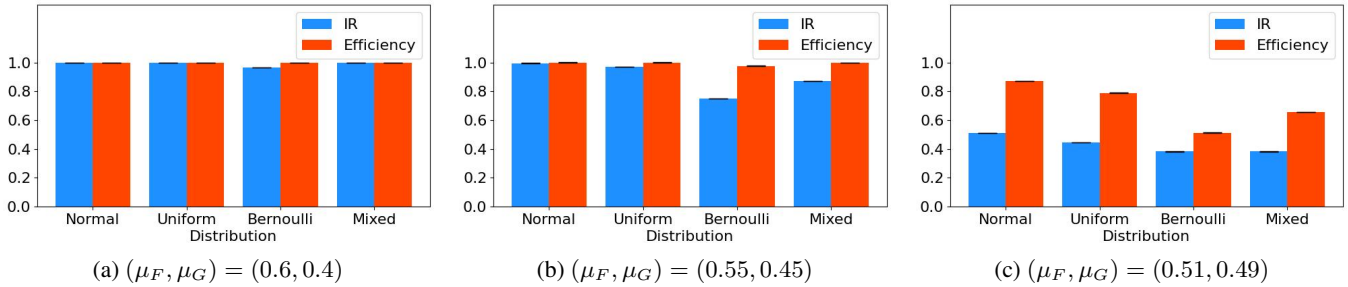


Figure 2: Performance of Algorithm 1 when $(\mu_F, \mu_G) = (0.6, 0.4), (0.55, 0.45), (0.51, 0.49)$ and $n = 100$. Error bars indicate standard error of mean. The empirical probability of being IR and the empirical efficiency decrease as $\mu_F - \mu_G$ gets smaller. When μ_F and μ_G are not too close, $n \geq 100$ is enough to guarantee good performance of Algorithm 1. Errors are negligible due to the large number of samples.

$\mu_F - \mu_G$ is not too small, n does not have to be very large to ensure a high probability of IR and efficiency.

Comparing when $(\mu_F, \mu_G) = (0.6, 0.4), (0.55, 0.45)$, and $(0.51, 0.49)$. As shown in Table 1 and Fig. 2, Algorithm 1’s empirical probability of being IR and empirical efficiency generally decrease as $\mu_F - \mu_G$ goes down. Note that Theorems 4.4 and 5.2 show that it is impossible to be efficient in the limit when $\mu_F \leq \mu_G$. The experiments imply that the closer an instance is to $\mu_F \leq \mu_G$, the harder it is to solve it well. Moreover, even if $\mu_F - \mu_G$ is extremely small, i.e., $(\mu_F, \mu_G) = (0.51, 0.49)$, with a sufficiently large n , say $n \geq 10000$, Algorithm 1’s empirical probability of IR and empirical efficiency can still be very close to 1.0.

Comparing different types of distributions. For each set of (n, μ_F, μ_G) , we can observe that the performance of Algorithm 1 is generally better on normal and uniform distributions, and worse on Bernoulli and mixed distributions. This is because the variances of normal and uniform distributions are lower than those of Bernoulli and mixed distributions, suggesting that the performance of Algorithm 1 is better if the distributions F, G have lower variances.

7 Discussions

In this paper, we conducted a thorough investigation of bilateral trade when there are multiple self-interested agents on both sides of the trade. Our results suggest that we can

achieve all desiderata in the limit, as long as buyers value the item more than sellers in expectation. Alternatively, when sellers value the item (weakly) more, efficiency is impossible with feasible deterministic mechanisms, and smooth randomness does not help, either. On the other hand, we showed that in these negative cases, even the first-best cannot obtain much value, so not much value is lost. These results extend and complement existing results in classical settings of bilateral trade. In addition, our characterizations of feasible mechanisms might be of independent interest in other mechanism design problems with shared allocation and payment.

While the setting we study is motivated by real-world applications, there are other settings that are also practically sensible. For example, in certain trading scenarios, it might be possible to exclude a subset of agents from the trade. Such settings are reminiscent of mechanism design for cost sharing (Moulin and Shenker 1992). One might also consider relaxed notions of incentive compatibility, or personalized payments, which would greatly expand the design space and allow for more efficient mechanisms. Another practical consideration is prior-independence: Can we design a mechanism that achieves efficiency independent of the value and cost distributions? Given our characterization, it is not hard to show that this cannot be done in the setting considered in this paper, but there might be possibilities if one, for example, considers a relaxed notion of incentive compatibility.

Acknowledgements

HZ was supported by the National Science Foundation under Grant No. DMS-1928930 and by the Alfred P. Sloan Foundation under grant G-2021-16778, while HZ was in residence at the Simons Laufer Mathematical Sciences Institute (formerly MSRI) in Berkeley, California, during the Fall 2023 semester. HZ and VC also thank the Cooperative AI Foundation, Polaris Ventures (formerly the Center for Emerging Risk Research), and Jaan Tallinn’s donor-advised fund at Founders Pledge for financial support.

References

- Babaioff, M.; Cai, Y.; Gonczarowski, Y. A.; and Zhao, M. 2018. The Best of Both Worlds: Asymptotically Efficient Mechanisms with a Guarantee on the Expected Gains-From-Trade. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 373–373.
- Babaioff, M.; Goldner, K.; and Gonczarowski, Y. A. 2020. Bulow-klemperer-style results for welfare maximization in two-sided markets. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2452–2471. SIAM.
- Balseiro, S. R.; Mirrokni, V.; Leme, R. P.; and Zuo, S. 2019. Dynamic Double Auctions: Towards First Best. In *Proceedings of the 2019 Annual ACM-SIAM Symposium on Discrete Algorithms*, 157–172.
- Balseiro, S. R.; Mirrokni, V.; Leme, R. P.; and Zuo, S. 2021. Non-excludable dynamic mechanism design. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1357–1373. SIAM.
- Blumrosen, L.; and Mizrahi, Y. 2016. Approximating gains-from-trade in bilateral trading. In *Web and Internet Economics: 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings 12*, 400–413. Springer.
- Brustle, J.; Cai, Y.; Wu, F.; and Zhao, M. 2017. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 589–590.
- Buchfuhrer, D.; Schapira, M.; and Singer, Y. 2010. Computation and incentives in combinatorial public projects. In *Proceedings of the 11th ACM conference on Electronic commerce*, 33–42.
- Cai, Y.; Goldner, K.; Ma, S.; and Zhao, M. 2021. On multi-dimensional gains from trade maximization. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1079–1098. SIAM.
- Colini-Baldeschi, R.; Goldberg, P.; de Keijzer, B.; Leonardi, S.; and Turchetta, S. 2017. Fixed price approximability of the optimal gain from trade. In *Web and Internet Economics: 13th International Conference, WINE 2017, Bangalore, India, December 17–20, 2017, Proceedings 13*, 146–160. Springer.
- Colini-Baldeschi, R.; Goldberg, P. W.; Keijzer, B. d.; Leonardi, S.; Roughgarden, T.; and Turchetta, S. 2020. Approximately efficient two-sided combinatorial auctions. *ACM Transactions on Economics and Computation (TEAC)*, 8(1): 1–29.
- Colini-Baldeschi, R.; Keijzer, B. d.; Leonardi, S.; and Turchetta, S. 2016. Approximately efficient double auctions with strong budget balance. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, 1424–1443. SIAM.
- Deng, Y.; Mao, J.; Sivan, B.; and Wang, K. 2022. Approximately efficient bilateral trade. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, 718–721.
- Dobzinski, S. 2011. An impossibility result for truthful combinatorial auctions with submodular valuations. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, 139–148.
- Dughmi, S.; Roughgarden, T.; and Yan, Q. 2016. Optimal mechanisms for combinatorial auctions and combinatorial public projects via convex rounding. *Journal of the ACM (JACM)*, 63(4): 1–33.
- Fei, Y. 2022. Improved approximation to first-best gains-from-trade. In *Web and Internet Economics: 18th International Conference, WINE 2022, Troy, NY, USA, December 12–15, 2022, Proceedings*, 204–218. Springer.
- Hardin, G. 1968. The Tragedy of the Commons. *Science*.
- McAfee, R. P. 1992. A dominant strategy double auction. *Journal of economic theory*, 56(2): 434–450.
- McAfee, R. P. 2008. The gains from trade under fixed price mechanisms. *Applied Economics Research Bulletin*, 1(1): 1–10.
- Moulin, H.; and Shenker, S. 1992. Serial cost sharing. *Econometrica: Journal of the Econometric Society*, 1009–1037.
- Mugrave, R. A. 1959. The theory of public finance: A study in public economy. *London and New York: McGraw-Hill*.
- Myerson, R. B. 1981. Optimal auction design. *Mathematics of operations research*, 6(1): 58–73.
- Myerson, R. B.; and Satterthwaite, M. A. 1983. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2): 265–281.
- Ostrom, E. 1990. *Governing the commons: The evolution of institutions for collective action*. Cambridge university press.
- Papadimitriou, C.; Schapira, M.; and Singer, Y. 2008. On the hardness of being truthful. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, 250–259. IEEE.
- Samuelson, P. A. 1954. The pure theory of public expenditure. *The review of economics and statistics*, 387–389.