



The Power of Two-Sided Recruitment in Two-Sided Markets

Yang Cai*

Yale University
New Haven, USA
yang.cai@yale.edu

Aranyak Mehta

Google Research
Mountain View, USA
aranyak@google.com

ABSTRACT

We consider the problem of maximizing the gains from trade (GFT) in two-sided markets. The seminal impossibility result by Myerson and Satterthwaite (1983) shows that even for bilateral trade, there is no individually rational (IR), Bayesian incentive compatible (BIC) and budget balanced (BB) mechanism that can achieve the full GFT. Moreover, the optimal BIC, IR and BB mechanism that maximizes the GFT is known to be complex and heavily depends on the prior.

In this paper, we pursue a Bulow-Klemperer-style question, i.e., does augmentation allow for prior-independent mechanisms to compete against the optimal mechanism? Our first main result shows that in the double auction setting with m i.i.d. buyers and n i.i.d. sellers, by augmenting $O(1)$ buyers and sellers to the market, the GFT of a simple, dominant strategy incentive compatible (DSIC), and prior-independent mechanism in the augmented market is at least the optimal in the original market, when the buyers' distribution first-order stochastically dominates the sellers' distribution. The mechanism we consider is a slight variant of the standard Trade Reduction mechanism due to McAfee (1992). For comparison, Babaioff, Goldner, and Gonczarowski (2020) showed that if one is restricted to augmenting only one side of the market, then $n(m + 4\sqrt{m})$ additional agents are sufficient for their mechanism to beat the original optimal and $\lfloor \log_2 m \rfloor$ additional agents are necessary for any prior-independent mechanism.

Next, we go beyond the i.i.d. setting and study the power of two-sided recruitment in more general markets. Our second main result is that for any $\varepsilon > 0$ and any set of $O(1/\varepsilon)$ buyers and sellers where the buyers' value exceeds the sellers' value with constant probability, if we add these additional agents into any market with arbitrary correlations, the Trade Reduction mechanism obtains a $(1-\varepsilon)$ -approximation of the GFT of the augmented market. Importantly, the newly recruited agents are agnostic to the original market.

*Supported by a Sloan Foundation Research Fellowship and the National Science Foundation Award CCF-1942583 (CAREER). This work was done in part while this author was a Visiting Faculty Researcher at Google Research, Mountain View.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0383-6/24/06

<https://doi.org/10.1145/3618260.3649669>

Christopher Liaw

Google Research
Mountain View, USA
cvliaw@google.com

Mingfei Zhao

Google Research
Mountain View, USA
mingfei@google.com

CCS CONCEPTS

- Theory of computation → Algorithmic mechanism design.

KEYWORDS

Two-sided Markets, Recruitment, Bulow-Klemperer

ACM Reference Format:

Yang Cai, Christopher Liaw, Aranyak Mehta, and Mingfei Zhao. 2024. The Power of Two-Sided Recruitment in Two-Sided Markets. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24), June 24–28, 2024, Vancouver, BC, Canada*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3618260.3649669>

1 INTRODUCTION

In this paper, we study the problem of maximizing the gains from trade (GFT) in two-sided markets. Two-sided markets are ubiquitous and have many practical applications; some major examples include the FCC spectrum auction and online marketplaces such as Uber, Lyft, and Airbnb. For example, in ride-sharing platforms, passengers (as the role of buyers) have some private value for obtaining transportation services and drivers (as the role of sellers) have some private cost for providing the necessary services. Mechanism design for two-sided markets poses additional challenges over its one-sided counterpart. In a one-sided market, the mechanism designer aims to maximize some objective (e.g. welfare or revenue) subject to a one-sided incentive-compatibility constraint. The seminal papers of Vickrey [32] and Myerson [26] described how to design mechanisms that achieve the optimal welfare and revenue for one-sided markets, respectively. However, in a two-sided market, one needs to ensure incentive compatibility for *both* sides of the market as well as to ensure that the mechanism itself does not run a deficit (called *budget balance*). The seminal impossibility result of Myerson and Satterthwaite [27] show that these additional constraints make the mechanism design problem much more challenging. In particular, even in the simplest setting with a single seller selling a single item to a single buyer (known as bilateral trade), no mechanism can achieve full efficiency while being Bayesian incentive-compatible (BIC), individually rational (IR), and budget balanced (BB). Myerson and Satterthwaite [27] also described the best BIC, IR and BB mechanism that maximizes efficiency in bilateral trade. However, the optimal mechanism is complex and heavily depends on the prior.

Motivated by the aforementioned challenges, there has been extensive research efforts and substantial progress in recent years investigating the two-sided market in the “simple versus optimal”

perspective, i.e. to show that a simple mechanism can approximate the performance of the optimal mechanism. A non-exhaustive list includes [1, 4–6, 9–14, 16, 18, 22, 24]. However, in many of these results, the mechanism designer requires *a priori* knowledge of both the buyers' and the sellers' distribution. Alternatively, the designer can increase the competition and thus her objective by recruiting more agents to the market. In single-item one-sided markets, the seminal work by Bulow and Klemperer [7] showed that the revenue of a second price auction with only a *single* additional participant from the same population is at least that of the optimal mechanism with the original set of participants. More recently, this result has been generalized to other one-sided market settings [3, 15, 17, 19–21, 23, 31]. Such results showcase how additional competition, coupled with a simple mechanism can be used to overcome the requirement of having precise knowledge of the underlying distributions and using the optimal-yet-complex mechanism. A natural question is whether such Bulow-Klemperer-type results also hold in the two-sided market settings.

In a recent paper, Babaioff, Goldner, and Gonczarowski [2] initiated this line of work to develop Bulow-Klemperer-type results for a fundamental single-parameter two-sided market setting called a *double auction*. In this problem there are n sellers that each hold an identical item and the value of each seller is drawn i.i.d. from some distribution F_S . There are m buyers that each wish to obtain one of these items and their value is drawn i.i.d. from another distribution F_B . They consider a variant of the natural, prior-independent Trade Reduction mechanism [25], which they call Buyer Trade Reduction (BTR). They show that when the two distributions are identical,¹ BTR with one additional buyer can achieve welfare at least the best welfare achievable in the original market even without the BIC, IR, and BB constraints (such a benchmark is called the *first best*). Note that the first best is exactly the welfare of the celebrated VCG mechanism, which is BIC and IR, but may violate the BB constraint. While this resolves the most basic case, it is rarely assumed that both distributions are identical. For example, a passenger in ride-sharing applications likely has higher value for obtaining transportation than the driver's cost for providing such transportation. When there are no assumptions on the distributions, [2] prove that no finite bound is possible. When the buyers' distribution first-order stochastically dominates² the sellers' distribution, they prove that $n(m+4\sqrt{m})$ additional buyers are sufficient for BTR to have welfare at least the first-best welfare in the original market when $m \geq n$.³

An immediate question from the result in [2] is whether the number of additional agents can be improved. Does a constant number of agents suffice for any number of buyers and sellers? The main difficulty turns out to come from the mechanism recruiting only one side of the market. In fact, their paper shows that $\lceil \log_2 m \rceil$ buyers are necessary if only extra buyers are recruited, even when there is a single seller. However in many situations, the mechanism designer is able to recruit both buyers and sellers. For example, in ride-sharing applications, recruiting both sides is very much feasible – more riders will use the platform with better marketing,

¹Note that no further assumptions are placed on this distribution, while similar results in one-sided markets make certain regularity assumption about the distribution.

²A distribution D first-order stochastically dominates D' if $\Pr_{x \sim D}[x \leq c] \leq \Pr_{x \sim D'}[x \leq c]$ for every c .

³Their result applies to the $m \leq n$ case using *Seller Trade Reduction*.

advertisement, or deals, and more drivers will adopt the platform with better incentives and marketing towards them. In this paper, we allow recruiting from *both* sides of the market. We show that with the Seller Trade Reduction (STR), a mechanism analogous to BTR, only $O(1)$ additional agents suffice. We give a formal definition of the mechanism in Section 1.1.

In the above result, we assume that all the agents are independent, that all the buyers are drawn from a common buyer distribution, that all the sellers are drawn from a common seller distribution, and that the buyer distribution first-order stochastically dominates the seller distribution. These are the same assumptions that were made by [2]. We next turn to the setting where we make minimal assumptions on the market and ask about the power of two-sided augmentation in such a general setting. Our second main result is that for any $\epsilon > 0$ and any set of $O(1/\epsilon)$ buyers and $O(1/\epsilon)$ sellers where the buyers' value exceeds the sellers' value with constant probability, the following holds. If we augment these buyers and sellers into any market then Trade Reduction achieves a $(1 - \epsilon)$ -approximation of the optimal efficiency of the augmented market. We stress that the augmentation requires zero knowledge of the original market. We also note that one-sided augmentation cannot be done in an agnostic manner. For example, suppose we augment the market with buyers that happens to have values less than all the sellers. It is not hard to see that for any prior-independent mechanism that is incentive-compatible, IR, and BB, its GFT remains unchanged after this augmentation.

To formally state our results, it is crucial to first discuss the measure of efficiency we adopt in this paper. There are two main measures of efficiency in two-sided markets. The first is the standard notion of *welfare* in the literature, which is equal to the sum of the value of all buyers and sellers that hold the items in the final allocation. The second is the *gain from trade (GFT)* which is the welfare of the final allocation minus the total value of sellers. At a high-level the GFT of a mechanism is a direct measure of the *additional* value of a mechanism. Note that when the set of sellers is fixed, maximizing welfare in a market where only the buyers are augmented (as in [2]) is identical to maximizing the GFT. In other words, there is no need to make a distinction between welfare and GFT. However, since we are interested in the problem where both sides of the market can be augmented, we focus on GFT as it is the more meaningful measure. As an extreme example, simply augmenting the market with additional sellers, and leaving them untraded, would increase the welfare while the GFT remains unchanged.

1.1 Our Results

We summarize prior results and our results in Table 1. Our first main result is Theorem 1.1 which states that if we use a simple and prior-independent mechanism, namely Seller Trade Reduction (STR)⁴, then augmenting both sides of the market by a constant number of participants has at least as much GFT as the optimal allocation before augmentation, assuming that the buyers' distribution first-order stochastically dominates the sellers' distribution. A formal definition of the mechanisms can be found in Section 2.

⁴The STR mechanism was introduced by [2] and is a variant of the Trade Reduction mechanism [25].

Table 1: Summary of our main results. The upper bounds state the number of additional agents suffices for a prior-independent mechanism (BTR or STR) to achieve GFT at least the first best. The lower bounds state the number of additional agents necessary for any anonymous and deterministic mechanism. Results with “buyer” listed indicate that only buyers can be augmented to the market. Note that the last row is for an approximation result instead of beating the GFT of the first best.

Distribution Assumptions	Upper Bound		Lower Bound
	Previous Work	This Work	
$F_B = F_S$	1 buyer [2, Theorem 1.1]	–	1 agent (trivial)
F_B FSD F_S	$n(m + 4\sqrt{m})$ buyers [2, Theorem 1.10]	$O(1)$ agents (Theorem 1.1)	$\lfloor \log_2 m \rfloor$ buyers [2, Theorem 5.1]
No assumption.	–	–	any finite number [2, Proposition 3.4]
$F_{B,i}^{-1}(1 - \gamma) \geq F_{S,j}^{-1}(\gamma)$ for new buyer i , new seller j . No assumptions on original market.	–	$O(1/\epsilon\gamma^2)$ agents for ($1 - \epsilon$)-approximation (Theorem 1.4)	$\Omega(1/\epsilon\gamma)$ agents for Trade Reduction (Proposition 1.6)

THEOREM 1.1. Consider the double auction with m i.i.d. buyers and n i.i.d. sellers. Suppose $m \geq n$ and the buyers’ distribution F_B first-order stochastically dominates the sellers’ distribution F_S . Then there is a global constant integer $c > 0$ such that the GFT of STR with $m + c$ buyers and $n + c$ sellers is at least the first-best GFT with m buyers and n sellers.

Remark 1.2. While Theorem 1.1 assumes that $m \geq n$, the result applies analogously to the case where $m \leq n$ using Buyer Trade Reduction, by negating the values/costs and swapping the role of buyers and sellers. See [2, Proposition A.1].

Remark 1.3. Another natural benchmark is to consider the per-unit GFT defined as the GFT divided by the number of items in the market. In the setting where we recruit only buyers, as in [2], the per-unit GFT objective is equivalent to the GFT objective. However, the per-unit GFT objective is a strictly stronger benchmark when one is also allowed to recruit sellers. Thus, it is natural to ask if recruiting $O(1)$ agents suffices for this stronger benchmark. In Appendix A, we build on the lower bound example in [2] to prove that if there are m buyers in the original market, $\Omega(\log m)$ agents are necessary for the per-unit GFT in the augmented market to weakly exceed the per-unit GFT in the original market.

Since the welfare of any mechanism is the GFT plus the sum of the seller values, our results immediately apply to the welfare objective, as the sum of seller values in the augmented market is at least the one in the original market.

For our second main result, we consider a setting where we make no assumptions on the original market and only fairly mild assumptions on the recruited agents’ distributions.

THEOREM 1.4. Fix $\gamma \in (0, 1/2]$ along with a set of c buyers with value distributions $F_{B,1}, \dots, F_{B,c}$ and a set of c sellers with value distributions $F_{S,1}, \dots, F_{S,c}$ such that all agents’ values are mutually independent and for all $i, j \in [c]$, we have $F_{B,i}^{-1}(1 - \gamma) \geq F_{S,j}^{-1}(\gamma)$. Fix any market M with arbitrary correlation between buyers and sellers. Suppose that we augment M by including the c buyers and c sellers described above. Let M' be the augmented market. Then the GFT of Trade Reduction is at least a $(1 - O(1/\gamma^2 c))$ -approximation to the GFT of the optimal matching in M' .

We stress that Theorem 1.4 makes no assumptions on M and that the value distributions of the agents that we augment into the market is completely agnostic of M .

There are several ways that one can interpret this result. The most obvious is that simply recruiting agents into the market suddenly makes a simple mechanism efficient. For example, a ride-sharing platform can simply recruit more drivers and more riders into the platform without any further market analysis. Next, in any large market, it is reasonable to assume that there must be a small subset of buyers with high value and a small subset of sellers that can produce goods at relatively low values. Our result implies that a simple mechanism is already efficient.

Remark 1.5. Earlier, we stated that a sufficient condition for Trade Reduction in an augmented market to obtain a $(1 - \epsilon)$ -approximation is that the buyers’ value exceeds the sellers’ value with constant probability. We note that this condition implies the c.d.f. condition in Theorem 1.4, up to a constant. Indeed, if buyer i ’s value exceeds seller j ’s value with probability at least γ then $F_{B,i}^{-1}(1 - \gamma/2) \geq F_{S,j}^{-1}(\gamma/2)$ (see [9, Lemma 3.1]).

Finally, we show that Theorem 1.4 is nearly tight but proving the following lower bound. The proof is omitted from the current version of the paper.

Proposition 1.6. For any $\gamma \in (0, 1/2)$, there exists a distributions F_B and F_S such that $F_B^{-1}(1 - \gamma) \geq F_S^{-1}(\gamma)$ and the following statement holds. If a market has c buyers whose value distributions are drawn from F_B and c sellers whose value distributions are drawn from F_S than running a Trade Reduction mechanism obtains a $(1 - \Omega(1/\gamma c))$ -approximation to the optimal GFT.

Due to space limitations, we do not provide complete proofs for all our results in this version of the paper. A complete version of the paper can be found on the arXiv <https://arxiv.org/abs/2307.03844>.

1.2 Additional Related Work

The paper that is mostly related to our work is [2]. They study Bulow-Klemperer-style results in two-sided market where one side of the market is augmented. When the buyer’s distribution is the same as the seller’s distribution, they prove that one additional

buyer is sufficient for BTR to achieve welfare at least the first-best welfare in the original market. They then study the problem with the stochastic dominance assumption, proving an upper bound of $4\sqrt{m}$ for a single seller and $n(m + 4\sqrt{m})$ for n sellers. They also provide lower bounds on the number of additional buyers required. Their lower bounds apply not only to BTR and STR, but also to any deterministic and prior-independent mechanisms. In this paper we study the same problem but allow both sides of the market to be augmented.

Approximations in two-sided markets. Despite the impossibility result by Myerson and Satterthwaite [27], many recent papers have successfully shown a multiplicative approximation to the first-best and second-best objective in various settings of two-sided markets. One line of work, which focuses on bilateral trade, aims to approximate the optimal welfare or GFT and to study the difference between the first-best and second-best [4, 6, 10, 11, 14, 22]. Another line of work studies the approximation problem in more general two-sided markets such as double auctions and multi-dimensional two-sided markets [1, 9, 12, 13, 16]. In sharp contrast to our paper, the mechanisms in all these works are not prior-independent: either the mechanism designer or the agents need to know the others' prior distributions. Another line of work provides *asymptotic* approximation guarantees in the number of items optimally traded for settings as general as multi-unit buyers and sellers and k types of items [1, 25, 29, 30]. Moreover, [24] consider a model of interactive communication in bilateral trade and prove that the efficient allocation is achievable with a smaller number of rounds of communication.

Bulow-Klemperer-style results in one-sided markets. There have been many Bulow-Klemperer-style results that aim to beat or approximate the optimal revenue in auctions with the recruitment of additional buyers. Results in single-dimensional settings include [15, 20, 21] for regular distributions, [31] for irregular distributions, and [23] for a dynamic single-item auction. Another line of work extend the results to multi-dimensional auctions, when buyers are unit-demand [28] and additive [3, 8, 17, 19]. Results in this paper (and [2]) show that Bulow-Klemperer-style results can also be derived in two-sided markets. We note that in the revenue-maximizing auction setting, it is clearly impossible to perform augmentation while being completely agnostic to the agents' distributions. On the other hand, one of our main result is that it is possible to perform augmentation in the efficiency-maximizing two-sided market setting while being completely agnostic to the market.

2 PRELIMINARIES

Double Auction and Gains From Trade. This paper focuses on the double auction setting, a two-sided market with m unit-demand buyers and n unit-supply sellers. Without loss of generality, we assume that $m \geq n$ (see Remark 1.2). All items are interchangeable and thus the value for each agent can be described as a scalar.

An allocation in a double auction is a (possibly random) set of n agents who hold the items. A buyer *trades* in the allocation if she holds the item and a seller *trades* if she does not hold the item. The *gains from trade* (GFT) of an allocation is defined as the difference

between the sum of all traded buyers' values and the sum of all traded sellers' values.

Mechanisms. We denote the buyer values by b_1, \dots, b_m and the seller values by s_1, \dots, s_n . We also write $\mathbf{b} = (b_1, \dots, b_m)$ and $\mathbf{s} = (s_1, \dots, s_n)$. A mechanism can be specified by, for each agents' profile (\mathbf{b}, \mathbf{s}) an allocation and a payment for each agent. We assume that all agents have quasi-linear utilities. Specifically, if a buyer trades in the mechanism, her utility is her value minus the payment for her. Similarly if a seller trades, her utility is the payment she receives minus her value. A mechanism is *Bayesian Incentive Compatible* (BIC) if every agent maximizes her *expected* utility (over all the other agents' randomness and the randomness of the mechanism) when she bids truthfully her value. In addition, it is *Dominant Strategy Incentive Compatible* (DSIC) if every agent maximizes her utility when she bids truthfully, no matter what the other agents report. We say that a mechanism is *individually rational* (IR) if every agent has non-negative utility when she bids truthfully, no matter what the other agents report. A mechanism is said to be *weakly budget-balanced* (WBB) if the sum of payment from the buyers is at least the sum of payment to the sellers for any agents' profile, i.e. the mechanism does not run a deficit.

First Best and Trade Reduction. Given any buyers' profile, the *first-best allocation* (also denoted by OPT) is the welfare-maximizing allocation under this profile (the allocation for the VCG mechanism). Formally, let $b^{(1)} \geq \dots \geq b^{(m)}$ be the buyer's bids ordered in the non-increasing order and $s^{(1)} \leq \dots \leq s^{(n)}$ be the seller's bids ordered in non-decreasing order. We abuse the notation and use $b^{(i)}$ and $s^{(i)}$ to represent the corresponding buyer and seller. The first-best allocation trades buyers $b^{(1)}, \dots, b^{(r)}$ with $s^{(1)}, \dots, s^{(r)}$, where $r = \max\{i \leq \min\{m, n\} : b^{(i)} \geq s^{(i)}\}$. We refer to r as the *optimal trade size*. Next, we define the trade reduction mechanism that we consider in this paper.

Definition 2.1 (Trade Reduction Mechanism [25]). *Let $u \in [0, 1]$ be a parameter. If $r < \min\{m, n\}$ and $b^{(r)} \geq u \cdot b^{(r+1)} + (1-u) \cdot s^{(r+1)} \geq s^{(r)}$ then TR trades buyers $b^{(1)}, \dots, b^{(r)}$ with $s^{(1)}, \dots, s^{(r)}$ at price $u \cdot b^{(r+1)} + (1-u) \cdot s^{(r+1)}$. Otherwise, the mechanism trades buyers $b^{(1)}, \dots, b^{(r-1)}$ with $s^{(1)}, \dots, s^{(r-1)}$ (if $r \leq 1$ then there is no trade). Each traded buyer pays $b^{(r)}$ and each traded seller receives $s^{(r)}$.*

Our first main result (Theorem 1.1) holds for a particular version of trade reduction (TR) where $u = 0$ which we refer to as seller's trade reduction (STR). We note that [2] also consider an asymmetric version of TR where they set $u = 1$; they refer to this version as buyer's trade reduction (BTR). Our second main result (Theorem 1.4) holds for all variants of TR in addition to the variant where we only utilize the "otherwise" part of the above mechanism. Specifically, we never trade buyer $b^{(r)}$ and seller $s^{(r)}$. Buyers $b^{(1)}, \dots, b^{(r-1)}$ are offered a price of $b^{(r)}$ and sellers $s^{(1)}, \dots, s^{(r-1)}$ are offered a price of $s^{(r)}$. The following lemma shows that Trade Reduction is an incentive-compatible mechanism.⁵

Lemma 2.2 ([2, Proposition C.1]). *TR is a deterministic, prior-independent mechanism and satisfies DSIC, IR, and WBB.*

⁵[2] prove this for STR but it is not difficult to adapt their proof of TR.

3 CONSTANT AGENTS SUFFICE TO BEAT FIRST-BEST WHEN F_B FSD F_S

For the rest of the paper, we focus on the i.i.d. setting and study the problem of beating the first-best GFT through augmentation. We prove that STR with $O(1)$ additional agents extracts at least as much GFT as the first-best allocation with the original set of agents (Theorem 1.1). Throughout this section we assume that buyer (resp. seller) values are drawn i.i.d. according to a common cumulative density function F_B (resp. F_S). For any quantile $q \in (0, 1)$, define the value $b(q)$ corresponding to quantile q as $b(q) = \inf\{x \mid \Pr_{b \sim F_B}[b \leq x] \geq q\}$. Similarly, define $s(q) = \inf\{x \mid \Pr_{s \sim F_S}[s \leq x] \geq q\}$. Clearly both $b(q)$ and $s(q)$ are non-decreasing in q . We say that F_B *first-order stochastically dominates* (FSD) F_S if for every $q \in (0, 1)$, $b(q) \geq s(q)$.

3.1 Proof Techniques

First, we present a high-level discussion about the proof techniques in this section. Notice that STR loses no more than a single trade from the first-best allocation in the augmented market. Thus a natural (but erroneous) starting point to prove Theorem 1.1 may be to (i) show that with only a constant number of new buyers and new sellers, at least one of the new buyers is eligible to trade with a new seller and (ii) show that if there is a trade between a new buyer and a new seller then the trade size must increase by 1 and thus STR performs at least as well as OPT. If the second statement were true then the proof should be relatively straightforward since the first statement happens with fairly high probability due to the stochastic dominance assumption. Unfortunately, the second statement is false and thus the first statement is not a sufficient condition for STR to outperform OPT. For an example where this happens, see Appendix B.5.

The message in the previous paragraph is that having additional trades among the new agents is not sufficient to guarantee that the optimal trade size increases. We would like to find an event such that the optimal trade size increases, which is sufficient for STR to outperforms OPT. Naively, we could simply consider the event where the optimal trade size does increase. However, the difficulty is in being able to lower bound the gain of the expected GFT restricted to this event and compare that with the loss of the expected GFT when this does not happen. In order to make the analysis more feasible, we consider more structured events that (i) make it possible to analyze the gain or loss in GFT and (ii) we can compare the probabilities of these events.

To make this formal, we use a coupling argument that was also used by Babaioff, Goldner, and Gonczarowski [2]. We first fix a set of quantiles and then assign these quantiles uniformly at random to the new and original buyers and sellers. However, the techniques in our paper and Babaioff, Goldner, and Gonczarowski [2] are otherwise very different. Babaioff, Goldner, and Gonczarowski [2] first consider the single seller and m buyers setting and proceed by showing that by adding a sufficient number of buyers it must be that (i) the GFT difference between the new and original optimal allocations is large and (ii) the GFT difference between the new optimal allocation and BTR is small. The only way for this to be possible is that the GFT of BTR must be large compared to the original optimal allocation. To handle the case with an arbitrary

number of sellers, they show that they can reduce the problem to the single seller case but this reduction incurs a linear overhead (in the number of sellers). In contrast, our argument directly compares the GFT difference between STR and OPT and show that this difference is net positive.

We now proceed with additional details on our argument. In the augmented market, $m + c$ buyers (including m original buyers and c augmented buyers) draw their values i.i.d. from F_B and $n + c$ sellers (including n original sellers and c augmented sellers) draw their values i.i.d. from F_S . Denote $N = m + n + 2c$ the total number of agents in the augmented market. We notice that the distribution of $b(q)$ (resp. $s(q)$) where q is drawn uniformly at random from $(0, 1)$ is exactly the distribution F_B (resp. F_S). We thus couple the random augmented market with the following random process: We draw N uniform quantiles from $(0, 1)$ and then assign these quantiles to all agents in the augmented market uniformly at random.

More specifically, denote q_1, \dots, q_N the N uniform quantiles in non-increasing order so that $q_1 \geq \dots \geq q_N$. Let $\mathbf{q} = (q_1, \dots, q_N)$. To avoid too many subscripts, we sometimes abuse notations and use $q(i)$ to denote q_i . These quantiles are assigned to all agents in the augmented market, including all original (called “old”) and augmented (called “new”) buyers and sellers. We notice that any two old buyers (or old sellers, new buyers, new sellers) are interchangeable, i.e. swapping their values will not change the GFT of the first-best allocation and STR in both the original and augmented market. Thus it suffices to consider any assignment from quantiles to those four labels. Formally, let $\pi: [N] \rightarrow \{\text{BO}, \text{BN}, \text{SO}, \text{SN}\}$ be a function that maps (quantile) indices to old buyers, new buyers, old sellers, and new sellers, respectively. Let $\Pi_{n,m,c} = \{\pi: |\pi^{-1}(\text{BO})| = m, |\pi^{-1}(\text{SO})| = n, |\pi^{-1}(\text{BN})| = |\pi^{-1}(\text{SN})| = c\}$ be the set of valid assignments. The assignment we choose is thus uniformly drawn from $\Pi_{n,m,c}$.

For any fixed quantiles \mathbf{q} and valid assignment π , $\text{STR}(\mathbf{q}, \pi)$ denotes the GFT of Seller Trade Reduction in the augmented market and $\text{OPT}(\mathbf{q}, \pi)$ denotes the GFT of the first-best allocation in the original market. Both values are well-defined since they are fully determined by the quantiles \mathbf{q} and the assignment π . Thus $\text{STR} = \text{STR}(m + c, n + c) = \mathbb{E}_{\mathbf{q}, \pi}[\text{STR}(\mathbf{q}, \pi)]$ and $\text{OPT} = \text{OPT}(m, n) = \mathbb{E}_{\mathbf{q}, \pi}[\text{OPT}(\mathbf{q}, \pi)]$.

To prove that STR is at least OPT, we would like to find an event such that the gain of the expected GFT (from first best to STR) restricted to this event can be lower bounded and compared with the loss of the expected GFT when the first-best allocation has more GFT than STR. To formalize the idea, we would like to construct two events \mathcal{E}_1 and \mathcal{E}_2 over the randomness of the assignment π such that:

- (1) For any \mathbf{q} , \mathcal{E}_1 is sufficient for $\text{STR}(\mathbf{q}, \pi) \geq \text{OPT}(\mathbf{q}, \pi)$. Moreover, $\mathbb{E}_\pi[\text{STR}(\mathbf{q}, \pi) - \text{OPT}(\mathbf{q}, \pi) | \mathcal{E}_1] \geq C(\mathbf{q})$ for some $C(\mathbf{q}) > 0$ (Lemma 3.3).
- (2) For any \mathbf{q} , \mathcal{E}_2 is necessary for $\text{OPT}(\mathbf{q}, \pi) > \text{STR}(\mathbf{q}, \pi)$. Moreover, $\mathbb{E}_\pi[\text{OPT}(\mathbf{q}, \pi) - \text{STR}(\mathbf{q}, \pi) | \mathcal{E}_2] \leq C(\mathbf{q})$ (Lemma 3.4 and Lemma 3.5).
- (3) $\Pr_\pi[\mathcal{E}_1] \geq \Pr_\pi[\mathcal{E}_2]$ (Lemma 3.6).

We notice that these conditions immediately proves Theorem 1.1 since

$$\text{STR} - \text{OPT} = \mathbb{E}_{\mathbf{q}, \pi}[\text{STR}(\mathbf{q}, \pi) - \text{OPT}(\mathbf{q}, \pi)]$$

$$\begin{aligned} &\geq \mathbb{E}_{\mathbf{q}}[\mathbb{E}_{\pi}[\text{STR}(\mathbf{q}, \pi) - \text{OPT}(\mathbf{q}, \pi)|\mathcal{E}_1] \cdot \Pr[\mathcal{E}_1] \\ &\quad + \mathbb{E}_{\pi}[\text{STR}(\mathbf{q}, \pi) - \text{OPT}(\mathbf{q}, \pi)|\mathcal{E}_2] \cdot \Pr[\mathcal{E}_2]] \geq 0. \end{aligned}$$

Here the first inequality follows from Property 2, which assert that for any \mathbf{q} , $\text{STR}(\mathbf{q}, \pi) \geq \text{OPT}(\mathbf{q}, \pi)$ when \mathcal{E}_2 does not happen. Thus, $\text{STR}(\mathbf{q}, \pi) \geq \text{OPT}(\mathbf{q}, \pi)$ on the event $\neg\mathcal{E}_1 \cap \neg\mathcal{E}_2$.

To construct the above events, we first break the set of quantiles into “buckets”. For some p , let I_1 correspond to the indices of the top p quantiles (i.e. high value agents) and J_1 correspond to the indices of the bottom p quantiles (i.e. low value agents).

As we will see below, the event \mathcal{E}_1 that we define ensures that the matching obtained by STR contains (i) at least one new buyer from I_1 and one new seller from J_1 and (ii) the other agents in the matching have GFT at least that of OPT. For the time-being, suppose that there were only one new buyer from I_1 and one new seller from J_1 . Then the new buyer would be a uniform random buyer from I_1 and the new seller would be a uniform random seller from J_1 . In particular, their contribution the GFT would be roughly $\mathbb{E}_{i,j} [b(q_i) - s(q_j)]$; this is state formally in Lemma 3.3. If there are multiple buyers and sellers in I_1 and J_1 , respectively, then one would expect that their contribution to the GFT would only increase. This suggests taking $C(\mathbf{q}) = \mathbb{E}_{i,j} [b(q_i) - s(q_j)]$. However, we note that p must be $\Theta(n)$ in order for the above argument to work. If $p \gg n$ then it becomes unlikely that new buyers in I_1 would be included in the first-best matching, let alone STR. On the other hand, if $p \ll n$ then it becomes too unlikely for new agents to actually be in I_1 or J_1 .

Analogously, it turns out that we can always upper bound the expected loss of GFT by the above choice of $C(\mathbf{q})$ provided $p \leq n$. For the event \mathcal{E}_2 , an obvious choice is to set $\mathcal{E}_2 = \neg\mathcal{E}_1$. However, when $n \ll m$, the event \mathcal{E}_2 becomes a very high probability event. For example, if $n = O(1)$ the probability that any new agent lands in $I_1 \cup J_1$ is $O(1/m)$ and so $\Pr[\mathcal{E}_2]$ would be $1 - O(1/m)$. To make this event smaller, we show that another necessary condition for OPT to perform better than STR is to have all the new sellers to be assigned the top $O(n)$ quantiles. If $n \ll m$ then this is a very unlikely event and we show that it is much smaller than $\Pr[\mathcal{E}_1]$.

Remark 3.1. Note that some of the proofs below require that m, n , and $m - n$ are larger than a constant. This is without loss of generality, since we can add a constant number of buyers and sellers and use the first-best GFT of the augmented market as the new benchmark.

3.2 Construction of the Events

In this section, we construct events \mathcal{E}_1 and \mathcal{E}_2 that satisfy the desired properties. For any valid assignment π , we denote $B_{\text{OLD}}^{\pi} = \pi^{-1}(\text{BO})$ the set of indices i such that the quantile q_i is assigned to an old buyer. Similarly, define $B_{\text{NEW}}^{\pi}, S_{\text{OLD}}^{\pi}, S_{\text{NEW}}^{\pi}$ as the sets for new buyers, old sellers and new sellers respectively. We omit the superscript π when the assignment is fixed and clear from context.

By adding a constant number of buyers and sellers, we assume that $m \geq n \geq 20$. Let $p = \lceil \frac{n}{10} \rceil \geq 2$. Define the sets

$$\begin{aligned} I_1 &= \{1, \dots, p\}, \quad I_2 = \{p+1, \dots, 2p\}, \\ J_1 &= \{N-p+1, \dots, N\}, \quad J_2 = \{N-2p+1, N-p\}. \end{aligned}$$

In other words, I_1 denotes the first p indices, I_2 denote the p indices after I_1 , J_1 denote the last p indices, and J_2 denote the p indices

before J_1 . It is straightforward to check that when $n \geq 20$, I_1, I_2, J_1, J_2 are all disjoint.

Claim 3.2. I_1, I_2, J_1, J_2 are all disjoint.

The good event \mathcal{E}_1 . Define the event \mathcal{E}_1 as the set of valid assignments π such that all of the properties below are satisfied:

- $|I_1 \cap B_{\text{NEW}}^{\pi}| \geq 2$, i.e. there are at least 2 new buyers in I_1 ;
- $|I_2 \cap B_{\text{OLD}}^{\pi}| \geq 1$, i.e. there are at least 1 old buyer in I_2 ;
- $|J_1 \cap S_{\text{NEW}}^{\pi}| \geq 2$, i.e. there are at least 2 new sellers in J_1 ;
- $|J_2 \cap S_{\text{OLD}}^{\pi}| \geq 1$, i.e. there are at least 1 old sellers in J_2 .

Here is the intuition for this event. We first show that every original buyer in $I_1 \cup I_2$ and every original seller in $J_1 \cup J_2$ trades in the original first-best allocation (Claim B.1). $|I_2 \cap B_{\text{OLD}}^{\pi}| \geq 1$ and $|J_2 \cap S_{\text{OLD}}^{\pi}| \geq 1$ ensure that the original first-best allocation contains at least one traded buyer from I_2 and one traded seller from J_2 . On top of it, the extra conditions $|I_1 \cap B_{\text{NEW}}^{\pi}| \geq 2$ and $|J_1 \cap S_{\text{NEW}}^{\pi}| \geq 2$ guarantee that the optimal trade size in the augmented market is increased by at least 2, with new buyers in I_1 and new sellers in J_1 joining in the trade. This suffices to not only show that STR has GFT more than the original first-best allocation, but also prove a lower bound on the gain using values of those new traded buyers/sellers. Formally, we prove the following lemma, whose proof is deferred to Subsection B.1.

Lemma 3.3. Fix any \mathbf{q} . We have that $\text{STR}(\mathbf{q}, \pi) \geq \text{OPT}(\mathbf{q}, \pi)$ for all $\pi \in \mathcal{E}_1$. Moreover,

$$\mathbb{E}_{\pi}[\text{STR}(\mathbf{q}, \pi) - \text{OPT}(\mathbf{q}, \pi)|\mathcal{E}_1] \geq \mathbb{E}_{i,j} [b(q_i) - s(q_j)]$$

where $i \sim I_1, j \sim J_1$ uniformly at random.

The bad event \mathcal{E}_2 . Next, we define the bad event \mathcal{E}_2 as $\neg\mathcal{E}_1 \cap \{\pi \in \Pi_{n,m,c} \mid S_{\text{NEW}}^{\pi} \subseteq [2n+2c]\}$. In other words, event \mathcal{E}_2 requires that (i) \mathcal{E}_1 does not happen and (ii) all new sellers are in the top $2n+2c$ quantiles. Lemma 3.4 shows that \mathcal{E}_2 is a necessary condition for OPT to obtain (strictly) more GFT than STR. We point out that \mathcal{E}_2 is not a necessary condition for OPT to outperform the classic Trade Reduction; an example can be found in Appendix B.4. Thus having STR is necessary for our proof.

Lemma 3.4. Fix any \mathbf{q} , we have $\text{STR}(\mathbf{q}, \pi) \geq \text{OPT}(\mathbf{q}, \pi)$ for all $\pi \notin \mathcal{E}_2$.

Next, we bound in Lemma 3.5 the loss in GFT conditioned on \mathcal{E}_2 , to match the lower bound proved in Lemma 3.3. To prove the lemma we use the following simple observation. The GFT loss between the original first best and STR is at most the loss between the augmented first best and STR, which is the value difference between the smallest traded buyer and the largest traded seller in the augmented market.

Lemma 3.5. For any \mathbf{q} , we have $\mathbb{E}_{\pi}[\text{OPT}(\mathbf{q}, \pi) - \text{STR}(\mathbf{q}, \pi)|\mathcal{E}_2] \leq \mathbb{E}_{i,j} [b(q_i) - s(q_j)]$ where $i \sim I_1, j \sim J_1$ uniformly at random.

The proofs of Lemma 3.4 and Lemma 3.5 can be found in Appendix B.2 and Appendix B.3, respectively.

Comparing probabilities of \mathcal{E}_1 and \mathcal{E}_2 . To complete the proof, it remains to show that $\Pr[\mathcal{E}_1] \geq \Pr[\mathcal{E}_2]$. For intuition, we consider two extremes. First, suppose that $n = m$, i.e. there are an equal number of buyers and sellers. Recall that $|I_1| = |I_2| = |J_1| = |J_2| =$

$p = \lceil \frac{n}{10} \rceil$. Assuming that $m \gg c$, we would have $|I_1|/N \approx 1/20$. In other words, if we take a random new buyer and assign it a uniformly random index from $[N]$, then with probability roughly $1/20$ it would land in I_1 . Since there are c new buyers, we have that $\mathbb{E}[|I_1 \cap B_{\text{NEW}}^\pi|] \approx c/20 \geq 2$ provided that $c \geq 40$. Thus by concentration, if c is a sufficiently large constant, then we expect that $|I_1 \cap B_{\text{NEW}}^\pi| \geq 2$ with probability at least $1 - \varepsilon$ for some small constant $\varepsilon > 0$. Similarly, we would have $|I_2 \cap B_{\text{OLD}}^\pi| \geq 1$, $|J_2 \cap S_{\text{OLD}}^\pi| \geq 1$, and $|J_1 \cap S_{\text{NEW}}^\pi| \geq 2$ each with probability at least $1 - \varepsilon$. By union bound the good event \mathcal{E}_1 happens with probability at least $1 - 4\varepsilon$ while the bad event $\mathcal{E}_2 \subseteq \neg \mathcal{E}_1$ happens with probability at most 4ε . This proves $\Pr[\mathcal{E}_1] \geq \Pr[\mathcal{E}_2]$ when $n = \Theta(m)$.

Now, let us consider the other extreme where $n \ll m$. In this case $|I_1|/N \approx \Omega(n/m)$ (and similarly for I_2, J_1, J_2). For any fixed agent, a random assignment would land the agent in I_1 with probability $\Omega(n/m)$. Thus, the probability of $|I_1 \cap B_{\text{NEW}}^\pi| \geq 2$ is $\Omega((n/m)^2)$. Similarly, the probability of $|J_1 \cap S_{\text{NEW}}^\pi| \geq 2$ is $\Omega(n/m)^2$. Moreover, the probability of the events $|J_2 \cap S_{\text{OLD}}^\pi| \geq 1$ and $|J_2 \cap B_{\text{OLD}}^\pi| \geq 1$ are both $\Omega(n/m)$. Note that this is a very conservative estimate obtained by considering the event that these quantities are equal to 1. We show that Thus, roughly speaking, the probability of \mathcal{E}_1 is at least the product of the probabilities of the four events, which indicates that $\Pr[\mathcal{E}_1] = \Omega(n/m)^6$.

On the other hand, the bad event \mathcal{E}_2 is a subset of the event that all the new sellers are in the top $2n + 2c$ quantiles. The probability that a new seller receives a uniform index and lands in $[2n + 2c]$ is $(2n + 2c)/(m + n + 2c) = \Theta(n/m)$. Thus, the probability that all the new sellers land in $[2n + 2c]$ is $\Theta((n/m)^c)$. Thus for a sufficiently large constant c , we have $\Pr[\mathcal{E}_2] \leq \Pr[\mathcal{E}_1]$.

Lemma 3.6. *Fix $c \geq 20000$ and suppose that $m \geq n + 2c$ and $n \geq c$. Then $\Pr[\mathcal{E}_1] \geq \Pr[\mathcal{E}_2]$.*

The proof of Lemma 3.6 requires a significant amount of work to make formal and is omitted from the present version of the paper.

PROOF OF THEOREM 1.1. The proof now follows by combining the previous four lemmas: Lemma 3.3, Lemma 3.4, Lemma 3.5, and Lemma 3.6. \square

4 MARKET AGNOSTIC RECRUITMENT

In this section, we prove that to obtain any constant approximation to the original market, it suffices to augment the market by a constant number of buyers and sellers, satisfying some mild conditions, and run the Trade Reduction mechanism.

A well-known observation is that the Trade Reduction mechanism loses at most one trade compared to the optimal allocation. Moreover, the trade that is lost is the least valuable trade. Thus, if the optimal allocation had at least k trades then the Trade Reduction mechanism is a $(1 - 1/k)$ -approximation to the optimal GFT. However, this is a conditional result and does not necessarily imply that the Trade Reduction mechanism is a good approximation to the optimal GFT.

In order to turn this conditional observation into a true approximation result, it would be sufficient to prove that the optimal GFT comes mainly from instances where there are a lot of trades. We do this using a coupling argument. Namely, for every instance I that may have a small number of trades, we map this instance into

many instances $f(I)_1, \dots, f(I)_T$ that certainly have a large number of trades and where for each $t \in [T]$, the optimal GFT in $f(I)_t$ exceeds the optimal GFT in I . A technical step here is that it is not sufficient to simply have T to be large; we require that the probability that we obtain the instance I to be much smaller than the probability of obtaining *at least one* of the instances $f(I)_1, \dots, f(I)_T$. We prove in Lemma 4.1 that such a mapping does exist. To summarize, we essentially show that (i) with high probability, we receive an instance where Trade Reduction is a good approximation to the optimal GFT and (ii) receiving an instance where Trade Reduction may not be a good approximation is a low probability event.

We now formalize the above argument. First, we require the following combinatorial lemma whose proof is omitted due to space constraints.

Lemma 4.1. *There are functions $\alpha(\gamma) = \Theta(\gamma^2)$ and $C(\gamma) = \Theta(1/\gamma^2)$ such that the following holds. For any $\gamma \in (0, 1/2]$, if $c \geq C(\gamma)$ then there exists T and a function $f: \binom{[c]}{\leq \alpha(\gamma) \cdot c} \rightarrow (2^{[c]})^T$ satisfying the following properties.*

- (1) *For every $t \in [T]$ and $S \in \binom{[c]}{\leq \alpha(\gamma) \cdot c}$ we have $|f_t(S)| \geq \gamma c/2$.*
- (2) *For every $t_1, t_2 \in [T]$ and $S_1, S_2 \in \binom{[c]}{\leq \alpha(\gamma) \cdot c}$, we have $f_{t_1}(S_1) \neq f_{t_2}(S_2)$ whenever $(t_1, S_1) \neq (t_2, S_2)$.*
- (3) *For every $S \in \binom{[c]}{\leq \alpha(\gamma) \cdot c}$, we have $c \cdot \gamma^{|S|} (1 - \gamma)^{c - |S|} \leq \sum_{t \in [T]} \gamma^{|f_t(S)|} (1 - \gamma)^{c - |f_t(S)|}$.*

For the proof, we need to define a bit of notation. We fix $\gamma \in (0, 1/2]$ and let α, C, T, f be as given by Lemma 4.1. Note that these parameters depend on γ but since γ is fixed for the proof, we omit the dependence on γ . Let m be the number of buyers in the original market and n be the number of sellers in the original market. We index the agents such that buyers $1, \dots, c$ and sellers $1, \dots, c$ are the new agents. Let $F_{B,1}, \dots, F_{B,c}$ be the value distributions for the new buyers and $F_{S,1}, \dots, F_{S,c}$ be the distributions for the new sellers. We note that they are mutually independent and independent of the distribution of the original market.

For a set of quantiles $\mathbf{q}_B = (q_B(1), \dots, q_B(m+c))$ and $\mathbf{q}_S = (q_S(1), \dots, q_S(n+c))$, define the random sets $B_+ = \{i \in [c] : q_B(i) \geq 1 - \gamma\}$ and $S_+ = \{j \in [c] : q_S(j) \leq \gamma\}$. We also define the following four events:

$$\begin{aligned} \mathcal{E}(1,1) &= \{|B_+| \geq \alpha c, |S_+| \geq \alpha c\}, \\ \mathcal{E}(1,0) &= \{|B_+| \geq \alpha c, |S_+| < \alpha c\}, \\ \mathcal{E}(0,1) &= \{|B_+| < \alpha c, |S_+| \geq \alpha c\}, \\ \mathcal{E}(0,0) &= \{|B_+| < \alpha c, |S_+| < \alpha c\}. \end{aligned}$$

Finally, for sets $B' \subseteq [c], S' \subseteq [c]$, we write

$$\text{OPT}(B', S') = \mathbb{E} \left[\text{OPT}(\mathbf{q}_B, \mathbf{q}_S) \middle| B_+ = B', S_+ = S' \right].$$

We define $\text{TR}(B', S')$ in a similar fashion. In addition, we write $\text{OPT}(\mathcal{E}(i,j)) = \mathbb{E} [\text{OPT}(B_+, S_+) \cdot \mathbb{1}[\mathcal{E}(i,j)]]$ and use similar notation for $\text{TR}(\mathcal{E}(1,1))$.

First, we have the straightforward observation that the optimal GFT is monotone in the set of buyers whose quantiles are above $1 - \gamma$ and the set of sellers whose quantiles are below γ .

Lemma 4.2. *If $B'' \supseteq B'$ and $S'' \supseteq S'$ then we have $\text{OPT}(B'', S'') \geq \text{OPT}(B', S')$.*

The following lemma is a well-known and follows from a simple observation that Trade Reduction loses the least valuable matching.

Lemma 4.3. *Let $k = \min\{|B'|, |S'|\}$. Then*

$$\text{TR}(B', S') \geq \left(1 - \frac{1}{k}\right) \text{OPT}(B', S').$$

Lemma 4.4. *If $c \geq C$ then $\text{OPT}(\mathcal{E}(1, 1)) \geq (1 - 3/c) \cdot \text{OPT}$.*

PROOF. Fix any $B' \subseteq B$ and $S' \subseteq S$ and let $\mathcal{E}(B', S') = \{B_+ = B' \setminus S_+, S_+ = S' \}$. Note that $\Pr[\mathcal{E}(B', S')] = \gamma^{|B'|} (1 - \gamma)^{c - |B'|} \gamma^{|S'|} (1 - \gamma)^{c - |S'|}$. We now consider three cases.

Case 1: $|B'| < \alpha c$ and $|S'| < \alpha c$. By Lemma 4.1, we have

$$\begin{aligned} & \sum_{t_1, t_2 \in [T]} \text{OPT}(f_{t_1}(B'), f_{t_2}(S')) \cdot \Pr[\mathcal{E}(f_{t_1}(B'), f_{t_2}(S'))] \\ &= \sum_{t_1, t_2 \in [T]} \left(\text{OPT}(f_{t_1}(B'), f_{t_2}(S')) \cdot \right. \\ & \quad \left. \gamma^{|f_{t_1}(B')|} (1 - \gamma)^{c - |f_{t_1}(B')|} \gamma^{|f_{t_2}(S')|} (1 - \gamma)^{c - |f_{t_2}(S')|} \right) \\ &\geq c \cdot \sum_{t_1 \in [T]} \left(\text{OPT}(f_{t_1}(B'), S') \cdot \right. \\ & \quad \left. \gamma^{|f_{t_1}(B')|} (1 - \gamma)^{c - |f_{t_1}(B')|} \gamma^{|S'|} (1 - \gamma)^{c - |S'|} \right) \\ &\geq c^2 \cdot \text{OPT}(B', S') \cdot \gamma^{|B'|} (1 - \gamma)^{c - |B'|} \gamma^{|S'|} (1 - \gamma)^{c - |S'|} \\ &= c^2 \cdot \text{OPT}(B', S') \cdot \Pr[\mathcal{E}(B', S')]. \end{aligned}$$

The first inequality uses Lemma 4.1 with S replaced by S' and Lemma 4.2 to show $\text{OPT}(f_{t_1}(B'), f_{t_2}(S')) \geq \text{OPT}(f_{t_1}(B'), S')$. The second inequality is similar which uses Lemma 4.1 with S replaced by B' and Lemma 4.2 to show $\text{OPT}(f_{t_1}(B'), S') \geq \text{OPT}(B', S')$. Observe that the first line is a lower bound on $\text{OPT}(\mathcal{E}(1, 1))$ (this uses the second assertion of Lemma 4.1). Thus, we can conclude that $\text{OPT}(\mathcal{E}(1, 1)) \geq c^2 \cdot \text{OPT}(\mathcal{E}(0, 0))$.

Case 2: $|B'| < \alpha c$ and $|S'| \geq \alpha c$. The calculation is similar to the first case. By Lemma 4.1 and Lemma 4.2, we have

$$\begin{aligned} & \sum_{t \in [T]} \text{OPT}(f_t(B'), S') \cdot \Pr[\mathcal{E}(f_t(B'), S')] \\ &\geq c \cdot \text{OPT}(B', S') \cdot \Pr[\mathcal{E}(B', S')]. \end{aligned}$$

We conclude that $\text{OPT}(\mathcal{E}(1, 1)) \geq c \cdot \text{OPT}(\mathcal{E}(0, 1))$.

Case 3: $|B'| \geq \alpha c$ and $|S'| < \alpha c$. This is analogous to the previous case and we get that $\text{OPT}(\mathcal{E}(1, 1)) \geq c \cdot \text{OPT}(\mathcal{E}(1, 0))$. \square

PROOF OF THEOREM 1.4. Note that on the event $\mathcal{E}(1, 1)$, the optimal trade size is at least αc and thus by Lemma 4.3, we have $\text{TR} \geq \text{TR}(\mathcal{E}(1, 1)) \geq (1 - 1/\alpha c) \text{OPT}(\mathcal{E}(1, 1))$. Next, by Lemma 4.4, we have $\text{OPT}(\mathcal{E}(1, 1)) \geq (1 - 3/c) \cdot \text{OPT}$. We conclude that

$$\text{TR} \geq (1 - 1/\alpha c) \cdot (1 - 1/3c) \cdot \text{OPT} \geq (1 - (3 + 1/\alpha)/c) \cdot \text{OPT}.$$

Recalling that $\alpha = \Theta(\gamma^2)$ completes the proof. \square

5 SUMMARY

In this paper we prove Bulow-Klemperer-style results in two-sided markets. When the buyer distribution FSD the seller distribution, we show that a deterministic, DSIC, IR, BB and prior-independent mechanism with constant additional agents achieved GFT at least

the first-best GFT in the original market. Here a constant number of buyers and sellers are both added to the market. While Babaioff, Goldner, and Gonczarowski [2] study the problem where only the larger side of the market is augmented (buyers are augmented with the assumption of $m \geq n$), it is an interesting direction to study the problem where only the smaller side of the market is allowed to augment. Intuitively, augmenting to the smaller side of the market is more efficient in increasing the trade size and GFT. Results in this direction yet remain open. Finally, we prove that adding independent agents agnostic to the (arbitrarily correlated) original market such that $F_{B,i}^{-1}(1 - \gamma) \geq F_{S,j}^{-1}(\gamma)$ help the prior-independent trade reduction mechanism obtain a $(1 - \varepsilon)$ -approximation to the optimal GFT. While we prove that $O(1/\varepsilon\gamma^2)$ agents suffices, the lower bound we construct requires only $\Omega(1/\varepsilon\gamma)$ agents. Closing this gap is also an interesting direction.

A COMPARISON WITH PER-UNIT GFT

In this section, we consider the per-unit GFT which is defined as the total GFT divided by the number of sellers. When there are m buyers and 1 seller, Babaioff, Goldner, and Gonczarowski [2, Theorem 5.1] give an example where if one is restricted to recruit only buyers than $\Omega(\log m)$ buyers are necessary for the per-unit GFT in the augmented market using a prior-independent mechanism to exceed the first-best GFT in the original market. A natural question is whether or not it is possible to recruit $O(1)$ sellers and $o(\log m)$ buyers so that the total number of additional agents is $o(\log m)$.

In this section, we provide an example which shows that $\Omega(\log m)$ additional agents are necessary for any prior independent mechanism to achieve at least the same per-unit GFT as the optimal allocation without augmentation. More specifically, we describe an instance with m buyers and 1 seller where if we add s sellers (for $1 \leq s < \sqrt{m}$) then we require $b \geq \Omega(s \log m)$ buyers just for the optimal per-unit GFT in the augmented market to exceed the optimal per-unit GFT in the original market.

The instance we consider is the following which is identical to the instance that appears in [2]. The buyer distribution, F_B , is as follows. With probability 0.5, the buyer value is 2 and otherwise, the buyer value is 0. For the seller distribution, F_S , we assume that the seller value is 1 with probability 0.5 and otherwise, it is equal to 0. Clearly, the buyer distribution FSD the seller distribution.

We assume that m is the number of original buyers and $n = 1$ is the number of original sellers. Let X_B be the number of original buyers with value 2. If the seller has value 0 then the optimal per-unit GFT is given $2 \min(X_B, 1)$ and if the seller has value 1 then the optimal (per-unit) GFT is given by $\min(X_B, 1)$. Since the buyer and seller values are independent, the expected per-unit GFT is

$$\begin{aligned} \text{OPT}(m, 1) &= 1.5 \cdot \mathbb{E}[\min(X_B, 1)] \\ &= 1.5 \cdot \Pr[X_B \geq 1] = 1.5 \cdot (1 - 2^{-m}). \end{aligned} \tag{A.1}$$

Now, let us assume we have $m + b$ buyers and $1 + s$ sellers. Note that for $k \leq s$, we have

$$\begin{aligned} \mathbb{E}[\text{OPT}(m + b, 1 + s) | X_B = k] &= k + \mathbb{E}[\min(X_S, k)] \\ &\leq s + \mathbb{E}[\min(X_S, b)] \\ &= 1.5s + 0.5 - 2^{-(s+1)}, \end{aligned}$$

where the last equality uses Claim A.1. We also have that

$$\mathbb{E} [\text{OPT}(m+b, 1+s) | X_B \geq 1+s] = 1.5(1+s).$$

Thus, we have

$$\begin{aligned} & \mathbb{E} [\text{OPT}(m+b, 1+s)] \\ &= \mathbb{E} [\text{OPT}(m+b, 1+s) | X_B \leq s] \Pr [X_B \leq s] \\ &+ \mathbb{E} [\text{OPT}(m+b, 1+s) | X_B \geq 1+s] \Pr [X_B \geq 1+s] \\ &\leq \left(1.5(s+1) - 1 - 2^{-(s+1)}\right) \cdot \Pr [X_B \leq s] \\ &+ 1.5(s+1) \cdot (1 - \Pr [X_B \leq s]) \\ &= 1.5(s+1) - \left(1 + 2^{-(s+1)}\right) \cdot \Pr [X_B \leq s] \\ &\leq 1.5(s+1) - \Pr [X_B \leq s]. \end{aligned}$$

Thus, the per-unit GFT satisfies

$$\frac{\mathbb{E} [\text{OPT}(m+b, 1+s)]}{1+s} \leq 1.5 - \frac{\Pr [X_B \leq s]}{1+s}. \quad (\text{A.2})$$

Comparing Eq. (A.1) and comparing Eq. (A.2), we have that a sufficient condition for the per-unit GFT with m buyers and 1 seller to be strictly larger than the per-unit GFT with $m+b$ buyers and $1+s$ sellers is if $2^{-m} < \frac{\Pr [X_B \leq s]}{1+s}$. Note that we have

$$\begin{aligned} \Pr [X_B \leq s] &\geq \Pr [X_B = s] = \binom{m+b}{s} 2^{-(m+b)} \\ &\geq \left(\frac{m+b}{s}\right)^s 2^{-(m+b)} > \left(\frac{m}{s}\right)^s 2^{-(m+b)}. \end{aligned}$$

Thus, a sufficient condition for the above inequality to hold is $2^{-m} < \frac{1}{s+1} \left(\frac{m}{s}\right)^s 2^{-(m+b)}$, which, after rearranging, is equivalent to $2^b < \frac{1}{s+1} \left(\frac{m}{s}\right)^s$. We conclude that $b \geq s \log_2(m/s) - \log_2(s+1)$ is necessary for the per-unit GFT in the augmented market to be at least the per-unit GFT in the original market. Note that this last bound implies that $\Omega(\log m)$ additional agents are required for the per-unit GFT in the augmented market to be at least the per-unit GFT in the original market. If $s = 1$ then we require $b \geq \Omega(\log m)$, and if $2 \leq s \leq \sqrt{m} - 1$ then the inequality requires $b \geq s \log_2(\sqrt{m}) - \log_2(\sqrt{m}) = \frac{s-1}{2} \log_2 m$.

Claim A.1. *Let $s \geq 1$ be an integer. Suppose that the random variable $X \sim \text{Binom}(0.5, s+1)$. Then $\mathbb{E} [\min(X, s)] = 0.5(s+1) - 2^{-(s+1)}$.*

PROOF. Note that we can write

$$\begin{aligned} \mathbb{E} [\min(X, s)] &= \mathbb{E} [\min(X, s+1)] - \Pr [X = s+1] \\ &= 0.5(s+1) - 2^{-(s+1)}, \end{aligned}$$

where in the second equality, we used that $\min(X, s+1) = X$ and $\mathbb{E} [X] = 0.5(s+1)$. \square

B MISSING PROOFS FROM SECTION 3

We prove Theorem 1.1 in this section. We first provide some notations used throughout this section. Let $i_1^\pi \leq \dots \leq i_{m+c}^\pi$ be the indices of all buyers, in an decreasing order of their quantiles and thus their values. Let $j_1^\pi \geq \dots \geq j_{m+c}^\pi$ be the indices of all sellers (in an increasing order of their quantiles and thus their values). Similarly, let $i_1^O \leq \dots \leq i_m^O$ be the indices of all old buyers, $i_1^N \leq \dots \leq i_c^N$

be the indices of all new buyers. Let let $j_1^N \geq \dots \geq j_n^O$ be the indices of all old sellers, $j_1^N \geq \dots \geq j_c^N$ be the indices of all new sellers.

B.1 Proof of Lemma 3.3

Claim B.1. *In the original first-best matching, every original buyer in $I_1 \cup I_2$ is matched and every original seller in $J_1 \cup J_2$ is matched. Similarly, in the augmented first-best matching, every buyer in $I_1 \cup I_2$ is matched and every seller in $J_1 \cup J_2$ is matched.*

PROOF. For the first statement, we prove only that every original buyer in $I_1 \cup I_2$ is matched. The proof for the sellers is analogous.

Let $k = |(I_1 \cup I_2) \cap B_{\text{OLD}}^\pi|$ be the number of original buyers in $I_1 \cup I_2$. Note that $k \leq 2 \cdot \lceil n/10 \rceil$. We prove that there are at least k original sellers outside of $I_1 \cup I_2$, i.e. $|S_{\text{OLD}}^\pi \cap ([N] \setminus (I_1 \cup I_2))| \geq k$. Indeed, we have

$$\begin{aligned} |S_{\text{OLD}}^\pi \cap ([N] \setminus (I_1 \cup I_2))| &= \underbrace{|S_{\text{OLD}}^\pi \cap [N]|}_{=n} - |S_{\text{OLD}}^\pi \cap (I_1 \cup I_2)| \\ &\geq n - 2 \left\lceil \frac{n}{10} \right\rceil \geq \frac{8n}{10} - 2 \geq \frac{2n}{10} + 2 \geq 2 \left\lceil \frac{n}{10} \right\rceil \geq k, \end{aligned}$$

where in the third inequality we use $n \geq 20$. Recall that i_k^O is the index of the k -th highest original buyer and j_k^O is the index of the k -th lowest original seller. The above argument immediately implies that $q(i_k^O) \geq q(j_k^O)$ and thus $b(q(i_k^O)) \geq s(q(j_k^O))$ since F_B FSD F_S . Thus there are at least k trades in the original first-best matching, which implies that every original buyer in $I_1 \cup I_2$ is matched.

The second statement follows from a similar argument. Let $k = |(I_1 \cup I_2) \cap (B_{\text{OLD}}^\pi \cup B_{\text{NEW}}^\pi)|$ be the number of buyers in $I_1 \cup I_2$. Note that $k \leq 2 \cdot \lceil n/10 \rceil$. We prove that there are at least k sellers outside of $I_1 \cup I_2$, i.e. $|(S_{\text{OLD}}^\pi \cup S_{\text{NEW}}^\pi) \cap ([N] \setminus (I_1 \cup I_2))| \geq k$. Indeed, we have

$$\begin{aligned} & |(S_{\text{OLD}}^\pi \cup S_{\text{NEW}}^\pi) \cap ([N] \setminus (I_1 \cup I_2))| \\ &= \underbrace{|(S_{\text{OLD}}^\pi \cup S_{\text{NEW}}^\pi) \cap [N]|}_{=n+c} - |(S_{\text{OLD}}^\pi \cup S_{\text{NEW}}^\pi) \cap (I_1 \cup I_2)| \\ &\geq n+c - 2 \left\lceil \frac{n}{10} \right\rceil \geq \frac{8n}{10} + c - 2 \geq \frac{2n}{10} + 2 \geq 2 \left\lceil \frac{n}{10} \right\rceil \geq k, \end{aligned}$$

where in the third inequality we use $n \geq 20$. Since F_B FSD F_S , every buyer in $I_1 \cup I_2$ has value no less than the cost of every seller outside of $I_1 \cup I_2$. Thus there are at least k trades in the augmented first-best matching, which implies that every buyer in $I_1 \cup I_2$ is matched. \square

PROOF OF LEMMA 3.3. Let π be any assignment in event \mathcal{E}_1 . Let T be the number of trades in the original first-best matching. We prove in the following claim that there are at least $T+2$ trades in the augmented first-best matching. An immediate consequence of this is that STR must have at least $T+1$ trades.

Claim B.2. *Recall that i_2^N is the index of the second-highest new buyer and j_2^N is the index of the second-lowest new seller. Then $b(q(i_2^N)) \geq b(q(i_{T+2})) \geq s(q(j_{T+2})) \geq s(q(j_2^N))$. Thus there are at least $T+2$ trades in the augmented first-best matching.*

PROOF. Let $i' \in I_2 \cap B_{\text{OLD}}$ be the index of any original buyer in I_2 (by definition of \mathcal{E}_1 there is at least one). By Claim B.1, i' is matched in the original first-best matching and thus $q(i') \geq q(i_T^O)$. By the property of \mathcal{E}_1 that $|I_1 \cap B_{\text{NEW}}| \geq 2$, we have $q(i_2^N) \geq q(i') \geq q(i_T^O)$ as $i_2^N \in I_1$ while $i' \in I_2$. Therefore, $q(i_2^N) \geq q(i_{T+2}) \geq q(i_T^O)$ as

both the highest and second-highest new buyer have quantile no less than $q(i_T^O)$. A similar argument shows that $q(j_2^N) \leq q(j_{T+2}) \leq q(j_T^O)$. We conclude that $b(q(i_2^N)) \geq b(q(i_{T+2})) \geq b(q(i_T^O)) \geq s(q(j_T^O)) \geq s(q(j_{T+2})) \geq s(q(j_2^N))$. The third inequality is because there are T trades in the original first-best matching. \square

Claim B.2 shows that STR trades the $T + 1$ highest buyers and the $T + 1$ lowest sellers (it may trade more). Therefore, $\text{STR}(\mathbf{q}, \pi) \geq \sum_{t=1}^{T+1} [b(q(i_t)) - s(q(j_t))]$. On the other hand, OPT trades the T highest original buyers with the T lowest original sellers and thus, $\text{OPT}(\mathbf{q}, \pi) = \sum_{t=1}^T [b(q(i_t^O)) - s(q(j_t^O))]$.

We claim that $\sum_{t=1}^{T+1} b(q(i_t)) - \sum_{t=1}^T b(q(i_t^O)) \geq b(q(i_1^N))$. This is because by Claim B.2, buyers i_1^N and i_2^N are among the top $T + 2$ highest-value buyers. So i_1^N must be in the top $T + 1$ highest-value buyers, which is contributed in the first term. Note that $\{i_1, \dots, i_{T+1}\} \setminus \{i_1^N\}$ correspond to the T highest value buyers excluding i_1^N and $\{i_1^O, \dots, i_T^O\}$ correspond to the top T highest value original buyers. Thus, we conclude that $\sum_{t=1}^{T+1} b(q(i_t)) - b(q(i_1^N)) \geq \sum_{t=1}^T b(q(i_t^O))$. By a similar argument, we have $\sum_{t=1}^{T+1} s(q(j_t)) - \sum_{t=1}^T s(q(j_t^O)) \leq s(q(j_1^N))$. Thus

$$\text{STR}(\mathbf{q}, \pi) - \text{OPT}(\mathbf{q}, \pi) \geq b(q(i_1^N)) - s(q(j_1^N)) \geq 0$$

It remains to lower bound the expected difference between $\text{STR}(\mathbf{q}, \pi)$ and $\text{OPT}(\mathbf{q}, \pi)$ conditioned on the event \mathcal{E}_1 . From the above inequality, $\text{STR}(\mathbf{q}, \pi) - \text{OPT}(\mathbf{q}, \pi)$ is lower bounded by the value of the highest new buyer subtracting the cost of the lowest new seller. We need the following definition.

Definition B.3. For any event \mathcal{E} over an assignment π , \mathcal{E} is swappable in a set S if: For every $\pi \in \mathcal{E}$, the assignment π' obtained by swapping the label for any two indices in S is also in \mathcal{E} . In other words, for every $\pi \in \mathcal{E}$ and every $i', i'' \in S$, if $\pi'(i') = \pi(i'')$, $\pi'(i'') = \pi(i')$, and $\pi'(i) = \pi(i)$ for $i \notin \{i', i''\}$ (clearly π' is also a valid assignment), then $\pi' \in \mathcal{E}$.

Lemma B.4. \mathcal{E}_1 is swappable in I_1 and it is swappable in J_1 .

PROOF. The lemma directly follows from the fact that swapping the label for any two indices in I_1 (or J_1) will not change the value of $|I_1 \cap B_{\text{NEW}}^\pi|$, $|I_2 \cap B_{\text{OLD}}^\pi|$, $|J_1 \cap S_{\text{NEW}}^\pi|$, $|J_2 \cap S_{\text{OLD}}^\pi|$. \square

Consider the following process that generates a random assignment π from \mathcal{E}_1 :

- (1) Choose an index i uniformly at random from I_1 and assign it to the “New Buyer” label. Choose an index j uniformly at random from J_1 and assign it to the “New Seller” label.
- (2) Denote $\Pi_{i,j}$ the set of valid assignments in set \mathcal{E}_1 such that i is assigned to the “New Buyer” label and j is assigned to the “New Seller” label. Draw an assignment π uniformly at random from $\Pi_{i,j}$ and assign the indices accordingly.

By Lemma B.4, we have that $|\Pi_{i',j'}| = |\Pi_{i'',j''}|$ for any indices $i', i'' \in I_1, j', j'' \in J_1$: For any assignment in $\Pi_{i',j'}$, we can swap the label between indices i', i'' and swap between j, j'' . This generates an assignment in $\Pi_{i'',j'}$ and vice versa. Moreover, for any valid assignment π , the number of “New Buyer” (or “New Seller”) labels is c . Hence, for every valid assignment π , $|\{i \in I_1, j \in J_1 : \pi \in \Pi_{i,j}\}| = c^2$. Thus the above random process chooses the assignment π uniformly at random from \mathcal{E}_1 .

For any realization of the above process, the value of the highest new buyer is at least $b(q_i)$ and the cost of the lowest new seller is at most $s(q_j)$. Thus the difference is at least $b(q_i) - s(q_j)$. Taking expectation over the random process, we have

$$\mathbb{E}_\pi [\text{STR}(\mathbf{q}, \pi) - \text{OPT}(\mathbf{q}, \pi) | \mathcal{E}_1] \geq \mathbb{E}_{i,j} [b(q_i) - s(q_j)]$$

where $i \sim I_1, j \sim J_1$ uniformly at random according to Step 1 of the process. \square

B.2 Proof of Lemma 3.4

PROOF OF LEMMA 3.4. We know from Lemma 3.3 that on the event \mathcal{E}_1 , we have $\text{STR}(\mathbf{q}, \pi) \geq \text{OPT}(\mathbf{q}, \pi)$. Hence it suffices to show the inequality on the event $\mathcal{E}' = \{\pi \in \Pi_{n,m,c} \mid S_{\text{NEW}}^\pi \cap [2n + 2c] \neq S_{\text{NEW}}^\pi\}$. Let OPT' be the first-best matching in the augmented market. For any $\pi \in \mathcal{E}'$, we consider two cases based on the number of trades in OPT' compared with OPT . Suppose OPT has T trades. Note that the number of trades in OPT' is least T .

Case 1: OPT' has at least $T + 1$ trades. Now in STR the top T (original and new) buyers and bottom T (original and new) sellers trade. The GFT from this is larger than the GFT from OPT which trades the top T original buyers and the bottom T original sellers.

Case 2: OPT' also has T trades. Our goal is to show that STR has the exact same T trades as OPT' . Thus the GFT of STR is the same as the GFT of OPT' , which is at least the GFT of OPT . By definition of STR, it suffices to show that $b(q(i_T)) \geq s(q(j_{T+1}))$, i.e. the T -th highest buyer value is least the $(T + 1)$ -th lowest seller cost.

First, we claim that $i_{n+c} \leq 2n + 2c$. Indeed, there are $n + c$ sellers in the augmented market. Thus, the $(n + c)$ -th highest value buyer must have index at most $2n + 2c$.

Next, recall that j_T^O is the index of the T -th lowest value original seller. We claim that $j_T^O \leq 2n + 2c$. For the sake of contradiction, suppose $j_T^O > 2n + 2c$. Recall that j_1^N is the index of the lowest-value new seller. Since $\pi \in \mathcal{E}'$, we have $S_{\text{NEW}}^\pi \cap [2n + 2c] \neq S_{\text{NEW}}^\pi$ and thus, $j_1^N > 2n + 2c$. In particular, $j_{T+1} \geq \min\{j_T^O, j_1^N\} > 2n + 2c \geq i_{n+c} \geq i_{T+1}$. Thus $q(i_{T+1}) > q(j_{T+1})$ and $b(q(i_{T+1})) > s(q(j_{T+1}))$. This implies that OPT' has at least $T + 1$ trades, a contradiction.

To finish the proof, we have $b(q(i_T)) \geq b(q(i_T^O)) \geq s(q(j_T^O)) \geq s(q(j_{T+1}))$. The first inequality uses $i_T \leq i_T^O$. The second inequality follows from the fact that OPT has T trades. The last inequality holds because $j_T^O \leq 2n + 2c < j_1^N$ and thus, $j_{T+1} \geq \min\{j_T^O, j_1^N\} = j_T^O$. We conclude that STR has the same T trades as OPT' . \square

B.3 Proof of Lemma 3.5

PROOF OF LEMMA 3.5. For every (\mathbf{q}, π) , let $\text{OPT}'(\mathbf{q}, \pi)$ be the GFT of the first-best matching in the augmented market. We clearly have $\text{OPT}(\mathbf{q}, \pi) \leq \text{OPT}'(\mathbf{q}, \pi)$. For each (\mathbf{q}, π) , we let $b^*(\mathbf{q}, \pi)$ (resp. $s^*(\mathbf{q}, \pi)$) denote the lowest value among buyers (resp. the largest value among sellers) traded in the augmented first-best matching. Let \mathcal{F} be the event that there is no trade in the augmented first-best matching and define $b^*(\mathbf{q}, \pi) = 0$ and $s^*(\mathbf{q}, \pi) = 0$ if there is no trade. Then by definition of the STR mechanism,

$$\begin{aligned} \text{OPT}(\mathbf{q}, \pi) - \text{STR}(\mathbf{q}, \pi) &\leq \text{OPT}'(\mathbf{q}, \pi) - \text{STR}(\mathbf{q}, \pi) \\ &\leq (b^*(\mathbf{q}, \pi) - s^*(\mathbf{q}, \pi)) \cdot \mathbb{1}[\pi \in \mathcal{F}]. \end{aligned}$$

We will show that

$$\mathbb{E} \left[b^*(\mathbf{q}, \pi) \mathbb{1} [\pi \in \mathcal{F}] \middle| \mathcal{E}_2 \right] \leq \mathbb{E}_i [b(q(i))]$$

where $i \sim I_1$ uniformly at random. A similar argument shows that

$$\mathbb{E} \left[s^*(\mathbf{q}, \pi) \mathbb{1} [\pi \in \mathcal{F}] \middle| \mathcal{E}_2 \right] \geq \mathbb{E}_j [s(q(j))]$$

where $j \sim J_1$ uniformly at random. To do so, we consider three cases: (i) where at least one new buyer is in I_1 , (ii) where no new buyers are in I_1 but at least one original buyer is in I_1 , and (iii) where no buyers (original or new) are in I_1 . In each of these cases, we prove that $\mathbb{E}_i [b(q(i))]$ is an upper bound on the value of lowest value traded buyer in the augmented market, in expectation.

Case 1: $|I_1 \cap B_{\text{NEW}}^\pi| \geq 1$. Let \mathcal{E}' be the event that $|I_1 \cap B_{\text{NEW}}^\pi| \geq 1$, i.e. at least one new buyer is in I_1 . The following lemma is similar to Lemma B.4 and follows from the definition of \mathcal{E}' and \mathcal{E}_2 .

Lemma B.5. $\mathcal{E}' \cap \mathcal{E}_2$ is swappable (see Definition B.3) in I_1 . Moreover, it is swappable in J_1 .

PROOF. Let π be any assignment in $\mathcal{E}' \cap \mathcal{E}_2 = \mathcal{E}' \cap \neg \mathcal{E}_1 \cap \{\hat{\pi} \in \Pi_{n,m,c} \mid S_{\text{NEW}}^{\hat{\pi}} \subseteq [2n + 2c]\}$ and π' be the assignment obtained by swapping any two labels in I_1 (or J_1). We notice that swapping the label for any two indices in I_1 (or J_1) will not change the value of $|I_1 \cap B_{\text{NEW}}^\pi|$, $|I_2 \cap B_{\text{OLD}}^\pi|$, $|J_1 \cap S_{\text{NEW}}^\pi|$, $|J_2 \cap S_{\text{OLD}}^\pi|$. Thus the new assignment π' is also in $\mathcal{E}' \cap \neg \mathcal{E}_1$. Moreover, since $I_1 \subseteq [2n + 2c]$, $S_{\text{NEW}}^\pi \subseteq [2n + 2c]$ implies that $S_{\text{NEW}}^{\pi'} \subseteq [2n + 2c]$. Thus $\pi \in \{\hat{\pi} \in \Pi_{n,m,c} \mid S_{\text{NEW}}^{\hat{\pi}} \subseteq [2n + 2c]\}$. \square

Consider the following random process of choosing an assignment π :

- (1) Choose an index i uniformly at random from I_1 and assign it to the “New Buyer” label.
- (2) Denote Π_i the set of valid assignments in set $\mathcal{E}' \cap \mathcal{E}_2$ such that i is assigned to the “New Buyer” label. Draw an assignment π uniformly at random from Π_i and assign the indices accordingly.

By Lemma B.5, we have that $|\Pi_{i'}| = |\Pi_{i''}|$ for any indices $i', i'' \in I_1$: For any assignment in $\Pi_{i'}$, we can swap the label for index i' and index i'' and generate an assignment in $\Pi_{i''}$ and vice versa. Moreover, for any valid assignment π , the number of “New Buyer” labels is c . So π is in c different Π_i ’s. Thus the above random process chooses the assignment π uniformly random from $\mathcal{E}' \cap \mathcal{E}_2$.

By Claim B.1, we note that for any realization of the above process, the new buyer with index i trades in the augmented first-best matching and therefore $\mathbb{1} [\pi \in \mathcal{F}] = 1$. Thus $b^*(\mathbf{q}, \pi)$, the lowest value among buyers traded in the augmented first-best matching, is upper bounded by $b(q(i))$. Thus $\mathbb{E}_\pi [b^*(\mathbf{q}, \pi) \mathbb{1} [\pi \in \mathcal{F}] \mid \mathcal{E}' \cap \mathcal{E}_2] \leq \mathbb{E}_i [b(q(i))]$, where i draws from I_1 uniformly at random.

Case 2: $|I_1 \cap B_{\text{NEW}}^\pi| = 0$ and $|I_1 \cap B_{\text{OLD}}^\pi| \geq 1$. Next, let \mathcal{E}'' be the event such that no new buyer is in I_1 and at least one old buyer is in I_1 , i.e. $|I_1 \cap B_{\text{OLD}}^\pi| \geq 1$ and $|I_1 \cap B_{\text{NEW}}^\pi| = 0$. One can easily verify that $\mathcal{E}'' \cap \mathcal{E}_2$ is also swappable in I_1 . And using a similar argument (by assigning index i to the “Old Buyer” label in the random process), we have $\mathbb{E}_\pi [b^*(\mathbf{q}, \pi) \mathbb{1} [\pi \in \mathcal{F}] \mid \mathcal{E}'' \cap \mathcal{E}_2] \leq \mathbb{E}_i [b(q(i))]$, where i draws from I_1 uniformly at random.

Case 3: $|I_1 \cap B_{\text{NEW}}^\pi| = 0$ and $|I_1 \cap B_{\text{OLD}}^\pi| = 0$. Finally, let $\mathcal{E}''' = \neg(\mathcal{E}' \cup \mathcal{E}'')$ be the event such that no buyer is in I_1 . Then for any $\pi \in \mathcal{E}'''$, $b^*(\mathbf{q}, \pi) \cdot \mathbb{1} [\pi \in \mathcal{F}] \leq b(q(\lceil n/10 \rceil)) \cdot \mathbb{1} [\pi \in \mathcal{F}]$. To see this, note that if there is no trade (i.e. $\pi \notin \mathcal{F}$) then both sides are equal to 0. On the other hand, if there is a trade (i.e. $\pi \in \mathcal{F}$) then all buyers have at most $b(q(\lceil n/10 \rceil))$ and thus, so does $b^*(\mathbf{q}, \pi)$. Thus $\mathbb{E}_\pi [b^*(\mathbf{q}, \pi) \mathbb{1} [\pi \in \mathcal{F}] \mid \mathcal{E}''' \cap \mathcal{E}_2] \leq b(q(\lceil n/10 \rceil)) \Pr[\mathcal{F}] \leq \mathbb{E}_i [b(q(i))]$, where i draws from I_1 uniformly at random.

Summarizing the three inequalities above, we have that

$$\mathbb{E}_\pi [b^*(\mathbf{q}, \pi) \mathbb{1} [\pi \in \mathcal{F}] \mid \mathcal{E}_2] \leq \mathbb{E}_i [b(q(i))],$$

where i draws from I_1 uniformly at random. An analogous argument gives that $\mathbb{E}_\pi [s^*(\mathbf{q}, \pi) \mathbb{1} [\pi \in \mathcal{F}] \mid \mathcal{E}_2] \geq \mathbb{E}_j [s(q(j))]$, where j draws from J_1 uniformly at random. Therefore,

$$\begin{aligned} & \mathbb{E}_\pi [\text{OPT}(\mathbf{q}, \pi) - \text{STR}(\mathbf{q}, \pi) \mid \mathcal{E}_2] \\ & \leq \mathbb{E}_\pi [(b^*(\mathbf{q}, \pi) - s^*(\mathbf{q}, \pi)) \cdot \mathbb{1} [\pi \in \mathcal{F}] \mid \mathcal{E}_2] \\ & \leq \mathbb{E}_{i,j} [b(q(i)) - s(q(j))], \end{aligned}$$

where $i \sim I_1, j \sim J_1$ uniformly at random. \square

B.4 Trade Reduction versus STR

As mentioned in Section 3.2, our argument crucially makes use of the fact that we use STR instead of the classic Trade Reduction mechanisms [25]. In particular, a key observation we use is that, if (i) the optimal allocation in the augmented market, OPT' is not the same as the optimal allocation in the original market OPT and (ii) the size of the optimal matching remains the same then STR and OPT' have the same GFT. This would not be true using McAfee’s Trade Reduction mechanism [25]. As an instructive example, consider the following scenario (we will assume the both sides have the same distribution so that values and quantiles are consistent). We have one original buyer with value 1, one original seller with value 0.9, one new buyer with value 0, and one new seller with value 0.8. In this example, the original first-best matching has size 1 and a GFT of 0.1. Once we add in the new buyers and sellers, the first-best matching remains at size 1 but the GFT is now 0.2. In STR, we use the second-lowest value seller to price the buyers and sellers, if possible. Here, this means a price of 0.9 is offered to buyer with value 1 and the seller with value 0.8 resulting in a trade. On the other hand, the Trade Reduction mechanism offers a price equal to the average of the next unmatched buyer and seller. This means a price of 0.45 is offered to the buyer and seller. The seller will not accept this price so the match is reduced resulting in zero trades.

It is not too difficult to extend the above example that show that \mathcal{E}_2 is not a necessary condition for OPT to outperform Trade Reduction. Concretely, suppose we have n original buyers with value 2, $n-1$ original sellers with value 1, one original seller with value $1+\varepsilon$, and $2c$ original buyers with value 0.9 for a total of $2n+2c$. Here, the first-best allocation trades the n value-2 buyers with all sellers, resulting in a GFT of $n+\varepsilon$. We then add c new buyers with value 0, $c-1$ new sellers with value 100, and one new seller with value 0.8. In particular, \mathcal{E}_2 does not happen since one of the new seller is outside the top $2n+2c$ values (and thus quantiles) in the augmented market. In the augmented first-best matching, the seller with value $1+\varepsilon$ would be removed from the matching and the new seller with value 0.8 would be added to the matching. STR would then offer a price of $1+\varepsilon$ which is accepted by the buyers with value

2 and the sellers with value at most 1. On the other hand, Trade Reduction would offer a price of $(0.9 + 1 + \varepsilon)/2 < 1$ if $\varepsilon < 0.1$. This price is not accepted by those sellers with value 1. Thus in Trade Reduction, $n - 1$ buyers with value 2 will trade with one (new) seller with value 0.8 and $n - 2$ (old) sellers with value 1. The GFT of Trade Reduction would be $2(n - 1) - (n - 2) - 0.8 = n - 0.8 < n$, which is worse than the original first-best GFT.

B.5 Example Where New Agents Can Trade but STR Loses a Trade

In this short section, we give an example where the new agents can trade but STR is still worse than OPT in the augmented market. Let $\varepsilon > 0$. There are 3 original buyers with values $b_1^O = 3, b_2^O = 2 + \varepsilon, b_3^O = 2$ and 3 original sellers all with value $s_1^O = s_2^O = s_3^O = 1$. The original optimal GFT is $4 + \varepsilon$. Now, suppose we add a new buyer with value $b^N = 2 + 3\varepsilon$ and a new seller with value $s^N = 2 + 2\varepsilon$. Note that b^N and s^N are eligible to trade with each other. The new optimal matching matches b_1^O, b^N, b_2^O with s_1^O, s_2^O, s_3^O . STR checks if s^N is able to price the buyers and the sellers; in this case it has higher value than b_2^O . Thus, STR removes b_2^O and s_3^O (say) and matches only b_1^O, b^N with s_1^O, s_2^O for a GFT of $3 + 3\varepsilon$. This is strictly worse than the original GFT if $\varepsilon < 1/2$. Note that a similar example could be possible even if there are multiple trades among the new agents.

REFERENCES

- [1] Moshe Babaioff, Yang Cai, Yannai A. Gonczarowski, and Mingfei Zhao. 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 (EC '18)*. Association for Computing Machinery, New York, NY, USA, 373. <https://doi.org/10.1145/3219166.3219203>
- [2] Moshe Babaioff, Kira Goldner, and Yannai A. Gonczarowski. 2020. Bulow-klemperer-style results for welfare maximization in two-sided markets. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '20)*. Society for Industrial and Applied Mathematics, USA, 2452–2471. <https://dl.acm.org/doi/10.5555/3381089.3381239>
- [3] Hedyeh Beyaghchi and S. Matthew Weinberg. 2019. Optimal (and benchmark-optimal) competition complexity for additive buyers over independent items. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 686–696. <https://dl.acm.org/doi/10.1145/3313276.3316405>
- [4] Liad Blumrosen and Shahar Dobzinski. 2021. (Almost) efficient mechanisms for bilateral trading. *Games and Economic Behavior* 130 (2021), 369–383. <https://www.sciencedirect.com/science/article/abs/pii/S0899825621001184>
- [5] Liad Blumrosen and Yehonatan Mizrahi. 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*. Springer, 400–413. https://dl.acm.org/doi/abs/10.1007/978-3-662-54110-4_28
- [6] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. 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. <https://dl.acm.org/doi/10.1145/3033274.3085148>
- [7] Jeremy I Bulow and Paul D Klemperer. 1994. Auctions vs. negotiations. <https://www.jstor.org/stable/2118262>
- [8] Linda Cai and Raghuvansh R Saxena. 2021. 99% revenue with constant enhanced competition. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, 224–241. <https://dl.acm.org/doi/10.1145/3465456.3467647>
- [9] Yang Cai, Kira Goldner, Steven Ma, and Mingfei Zhao. 2021. On multi-dimensional gains from trade maximization. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1079–1098. <https://dl.acm.org/doi/10.5555/3458064.3458131>
- [10] Yang Cai and Jinzhao Wu. 2023. On the Optimal Fixed-Price Mechanism in Bilateral Trade. (2023), 737–750. <https://doi.org/10.1145/3564246.3585171>
- [11] Riccardo Colini-Baldeschi, Paul Goldberg, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. 2017. Fixed price approximation of the optimal gain from trade. In *International Conference on Web and Internet Economics*. Springer, 146–160. https://dl.acm.org/doi/abs/10.1007/978-3-319-71924-5_11
- [12] Riccardo Colini-Baldeschi, Paul W Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, and Stefano Turchetta. 2020. Approximately efficient two-sided combinatorial auctions. *ACM Transactions on Economics and Computation (TEAC)* 8, 1 (2020), 1–29. <https://dl.acm.org/doi/10.1145/3381523>
- [13] Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. 2016. Approximately efficient double auctions with strong budget balance. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 1424–1443. <https://dl.acm.org/doi/10.5555/2884435.2884533>
- [14] Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. 2022. Approximately efficient bilateral trade. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, 718–721. <https://dl.acm.org/doi/10.1145/3519935.3520054>
- [15] Shaddin Dughmi, Tim Roughgarden, and Mukund Sundararajan. 2009. Revenue submodularity. In *Proceedings of the 10th ACM conference on Electronic commerce*, 243–252. <https://dl.acm.org/doi/10.1145/1566374.1566409>
- [16] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. 2014. Modularity and greed in double auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 241–258. <https://dl.acm.org/doi/10.1145/2600057.2602854>
- [17] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2017. The Competition Complexity of Auctions: A Bulow-Klemperer Result for Multi-Dimensional Bidders. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 343–343. <https://dl.acm.org/doi/10.1145/3033274.3085115>
- [18] Yunou Fei. 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*. Springer, 204–218. https://link.springer.com/chapter/10.1007/978-3-031-22832-2_12
- [19] Michal Feldman, Ophir Friedler, and Aviad Rubinstein. 2018. 99% revenue via enhanced competition. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 443–460. <https://dl.acm.org/doi/10.1145/3219166.3219202>
- [20] Hu Fu, Christopher Liaw, and Sikander Randhawa. 2019. The Vickrey auction with a single duplicate bidder approximates the optimal revenue. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 419–420. <https://dl.acm.org/doi/10.1145/3328526.3329597>
- [21] Jason D Hartline and Tim Roughgarden. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, 225–234. <https://dl.acm.org/doi/10.1145/1566374.1566407>
- [22] Zi Yang Kang, Francisco Pernice, and Jan Vondrák. 2022. Fixed-price approximations in bilateral trade. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2964–2985. <https://pubs.siam.org/doi/pdf/10.1137/1.9781611977073.115>
- [23] Siqi Liu and Christos-Alexandros Psomas. 2018. On the competition complexity of dynamic mechanism design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2008–2025. <https://dl.acm.org/doi/10.5555/3174304.3175436>
- [24] Jieming Mao, Renato Paes Leme, and Kangning Wang. 2022. Interactive Communication in Bilateral Trade. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik. <https://drops.dagstuhl.de/storage/00lipics/lipics-vol215-itcs2022/LIPIcs. ITCS.2022.105.LIPIcs. ITCS.2022.105.pdf>
- [25] R Preston McAfee. 1992. A dominant strategy double auction. *Journal of Economic Theory* 56, 2 (1992), 434–450. <https://www.sciencedirect.com/science/article/abs/pii/002205319090011U>
- [26] Roger B Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73. <https://www.jstor.org/stable/3689266>
- [27] Roger B Myerson and Mark A Satterthwaite. 1983. Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29, 2 (1983), 265–281. <https://www.sciencedirect.com/science/article/abs/pii/0022053183900480>
- [28] Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. 2012. Supply-limiting mechanisms. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 844–861. <https://dl.acm.org/doi/10.1145/2229012.2229077>
- [29] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. 2018. Double Auctions in Markets for Multiple Kinds of Goods. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*, Jérôme Lang (Ed.). ijcai.org, 489–497. <https://doi.org/10.24963/ijcai.2018/68>
- [30] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. 2018. MUDA: A Truthful Multi-Unit Double-Auction Mechanism. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, Sheila A. McIlraith and Kilian Q. Weinberger (Eds.). AAAI Press, 1193–1201. <https://www.aaai.org/ojs/index.php/AAAI/AAAI18/paper/view/16593>
- [31] Balasubramanian Sivan and Vasilis Syrgkanis. 2013. Vickrey auctions for irregular distributions. In *Web and Internet Economics: 9th International Conference, WINE 2013, Cambridge, MA, USA, December 11–14, 2013, Proceedings 9*. Springer, 422–435. https://dl.acm.org/doi/10.1007/978-3-642-45046-4_35
- [32] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16, 1 (1961), 8–37. <https://www.jstor.org/stable/2977633>

Received 11-NOV-2023; accepted 2024-02-11