

On Multi-Dimensional Gains from Trade Maximization

Yang Cai^Y
Yale University, USA
yang.cai@yale.edu

Kira Goldner^Z
Columbia University, USA
kgoldner@cs.columbia.edu

Steven Ma
Yale University, USA
steven.ma@yale.edu

Mingfei Zhao
Yale University, USA
mingfei.zhao@yale.edu

Abstract

We study gains from trade in multi-dimensional two-sided markets. Specifically, we focus on a setting with n heterogeneous items, where each item is owned by a different seller i , and there is a constrained-additive buyer with feasibility constraint F . Multi-dimensional settings in one-sided markets, e.g. where a seller owns multiple heterogeneous items but also is the mechanism designer, are well-understood. In addition, single-dimensional settings in two-sided markets, e.g. where a buyer and seller each seek or own a single item, are also well-understood. Multi-dimensional two-sided markets, however, encapsulate the major challenges of both lines of work: optimizing the sale of heterogeneous items, ensuring incentive-compatibility among both sides of the market, and enforcing budget balance. We present, to the best of our knowledge, the first worst-case approximation guarantee for gains from trade in a multi-dimensional two-sided market.

Our first result provides an $O(\log(1/r))$ -approximation to the first-best gains from trade for a broad class of downward-closed feasibility constraints (such as matroid, matching, knapsack, or the intersection of these). Here r is the minimum probability over all items that a buyer's value for the item exceeds the seller's cost. Our second result removes the dependence on r and provides an unconditional $O(\log n)$ -approximation to the second-best gains from trade. We extend both results for a general constrained-additive buyer, losing another $O(\log n)$ -factor en-route. The first result is achieved using a fixed posted price mechanism, and the analysis involves a novel application of the prophet inequality or a new concentration inequality. Our second result follows from a stitching lemma that allows us to upper bound the second-best gains from trade by the first-best gains

from trade from the “likely to trade” items (items with trade probability at least $1/n$) and the optimal profit from selling the “unlikely to trade” items. We can obtain an $O(\log n)$ -approximation to the first term by invoking our $O(\log(1/r))$ -approximation on the “likely to trade” items. We introduce a generalization of the fixed posted price mechanism—seller adjusted posted price—to obtain an $O(\log n)$ -approximation to the optimal profit for the “unlikely to trade” items. Unlike fixed posted price mechanisms, not all seller adjusted posted price mechanisms are incentive compatible and budget balanced. We develop a new argument based on “allocation coupling” to show the seller adjusted posted price mechanism used in our approximation is indeed budget balanced and incentive-compatible.

1 Introduction

Two-sided markets are ubiquitous in today's economy: take for example the New York Stock Exchange, online ad exchange platforms (e.g., Google's Doubleclick, Microsoft's AdECN, etc.), crowdsourcing platforms, FCC's spectrum auctions, or sharing economy platforms such as Uber, Lyft, and Airbnb. Yet mechanism design for such two-sided markets, where both the buyer(s) and seller(s) are strategic, is known to be substantially harder than for one-sided markets, i.e. auctions where the seller designs the mechanism. The additional challenges stem from the following requirements: (1) now the allocation rule must satisfy incentive-compatibility for both sides of the market; and (2) the buyer and seller payments must satisfy budget balance, that is, the mechanism must not run a deficit. The limitations of these constraints are best illustrated by the seminal impossibility result of Myerson and Satterthwaite [30]. They show that even in the simplest possible two-sided market—bilateral trade, when one seller is selling a single item to a buyer—no Bayesian incentive compatible (BIC), individually rational (IR), and budget balanced (BB) mechanism can achieve the first-best efficiency: the maximum efficiency achievable without any of the previous constraints. The second-best ef-

[—]The arXiv version of this paper can be found at <https://arxiv.org/abs/2007.13934>.

^YSupported by a Sloan Foundation Research Fellowship and the NSF Award CCF-1942583 (CAREER).

^ZSupported by NSF Award DMS-1903037 and a Columbia Data Science Institute postdoctoral fellowship.

efficiency is the maximum efficiency achievable by any BIC, IR, and BB mechanism.

Despite the additional challenges, significant progress has very recently been made in understanding single-dimensional two-sided markets [1, 2, 5, 6, 7, 15]. Yet, in reality, many two-sided markets involve agents with multi-dimensional preferences. For example, a customer searching for a place to stay on Airbnb typically values a listing based on its location, number of rooms, amenities, reviews, and more. For one-sided markets, multi-dimensional mechanism design has been the core of algorithmic mechanism design in the past decade. See [13, 17] and the references therein for more details. Our goal in this paper is to study efficiency maximization in multi-dimensional two-sided markets.

There are two ways to measure efficiency in two-sided markets. One is the standard notion of welfare. The other is the gains from trade (GFT), which is the welfare of the final allocation minus the total cost of the sellers. Intuitively, the GFT captures how much more welfare the mechanism brings to the market. Clearly, the two measures are equivalent if efficiency is maximized. However, approximating the GFT is much more challenging than welfare. For example, if the buyer's value is 10 and the seller's cost is 9, not trading the item is a $9=10$ -approximation to the optimal welfare but a 0-approximation to the optimal GFT. Obviously, any good approximation to the optimal GFT immediately gives a good approximation to the optimal welfare, but the opposite direction is rarely true.

Several results show that generalizations of posted price mechanisms can achieve a constant fraction of the first-best welfare in fairly general multi-dimensional two-sided markets [5, 14, 16, 19]. However, GFT maximization in multi-dimensional settings has remained elusive. We present, to the best of our knowledge, the first worst-case approximation guarantee for GFT in a multi-dimensional two-sided market. We focus on a setting with n heterogeneous items, where each item is owned by a different seller i , and there is a constrained-additive buyer with feasibility constraint F . The Airbnb example is a special case of our setting, where the customer is a unit-demand buyer, and there are n hosts, each listing a property. We further assume that the prior distributions of the buyer's valuations and sellers' costs are public knowledge and independent; the realized valuations and costs are private.

Recall that in one-sided markets, maximizing revenue for even a single (non-constrained) additive buyer is far more challenging than for single-dimensional buyers, both optimally and approximately [4, 18, 22, 26, 29]. Maximizing GFT in two-sided markets suffers from this curse of dimensionality as well. As with revenue, single-dimensional settings can leverage an analog to Myerson's virtual value theory by using the optimal dual variables, as shown in [7], but this does not extend to multiple dimensions. Note also that

while Colini-Baldeschi et al. [16] are able to extend an $O(1)$ -approximation to welfare to a two-sided market with XOS buyers and additive sellers, their mechanism gives no guarantee for GFT.

The first main result is a distribution-parameterized approximation to the first-best GFT.

RESULT 1. There is a fixed posted price mechanism whose GFT is an $O(\frac{1}{\log(1-r)})$ -approximation to the first-best GFT when the buyer's feasibility constraint F is $(;)$ -selectable (Definition 3.3), and an $O(\log(n)\log(1-r))$ -approximation for a general constrained-additive buyer. r is a distributional parameter: the minimum trade probability over all items. We define the trade probability of item i as the probability that the buyer's value for i exceeds the seller's cost.

The notion of $(;)$ -selectability is introduced by Feldman et al. [21] as a sufficient condition for prophet-inequality-type online algorithms to exist. Many familiar feasibility constraints such as matroid, matching, knapsack, and the compositions of each, are known to be $(;)$ -selectable with constant α and [21], so our result provides an $O(\log(1-r))$ -approximation for all of these environments. We will give the formal definition of $(;)$ -selectability in Section 3.4. The main takeaway from Result 1 is that there is an $O(\log(1-r))$ -approximation to the first-best GFT for the feasibility constraints stated above. Next we introduce the class of fixed posted price mechanisms.

Fixed Posted Price (FPP): In a fixed posted price mechanism, there is a collection of fixed prices $(b_i^B; s_i^S)_{i \in [n]}$, where b_i^B, s_i^S for each item i . Let R be the set of sellers that are willing to sell their item at price s_i^S . The buyer can purchase any item i in R at price b_i^B . Trade only occurs when the buyer wants to buy the item and the seller is willing to sell it.

Our result is a generalization of the result by Colini-Baldeschi et al. [15], where they provide the same approximation using a fixed posted price mechanism for bilateral trade. Importantly, our approximation ratio has the optimal dependence on r up to a constant factor. Example 1 (adapted from an example by Blumrosen and Dobzinski [5]) in Appendix A shows that, for any $r > 0$, there is an instance of our problem with minimum trade probability r such that no fixed posted price mechanism can achieve more than a $\frac{c}{\log(1-r)}$ -fraction of even the second-best GFT for some absolute constant c . In our fixed posted price mechanism, we allow b_i^B to be strictly greater than s_i^S . This is crucial for our analysis, but makes the mechanism only ex-post weakly budget balanced. We leave it as an interesting open question as to whether our approximation ratio can be achieved by an ex-post strongly budget balanced fixed posted price mechanism.

When the trade probability of each item is not too low,

our first result provides a good approximation to the first-best GFT using a simple fixed posted price mechanism. However, r can be arbitrarily small in the worst-case, making our approximation too large to be useful. Is it possible to produce an unconditional worst-case approximation guarantee? We provide an affirmative answer to this question with an unconditional $O(\log n)$ -approximation to the second-best GFT.

RESULT 2. There is a dominant strategy incentive compatible (DSIC), ex-post IR, and BB mechanism whose GFT is at least $\frac{1}{2^{\log n}}$ -fraction of the second-best GFT when the buyer's feasibility constraint F is $(;)$ -selectable, and at least $\frac{1}{2^{\log(n)}}$ -fraction of the second-best GFT when the buyer is general constrained-additive.

As we show in Example 1, no fixed posted price mechanism can provide such a guarantee. We develop two new mechanisms. The first one is a multi-dimensional extension of the "Generalized Buyer Offering Mechanism" by Brustle et al. [7]. We provide a full description of the mechanism in Section 4.2. The second mechanism is a generalization of the fixed posted price mechanism that we call the Seller Adjusted Posted Price Mechanism.

Seller Adjusted Posted Price (SAPP): The sellers report their costs s . The mechanism maps the cost profile to a collection of posted prices $f_i(s)g_{i2[n]}$ for the buyer. The buyer can purchase at most one item, and pays price $f_i(s)$ if she buys item i . An item trades if the buyer decides to purchase that item.

The main advantage of using a SAPP mechanism is that it provides the flexibility to set prices based on the sellers' costs, which allows a SAPP mechanism to achieve GFT that can be unboundedly higher than the GFT attainable by even the best fixed posted price mechanism (see Example 2). Example 3 in Appendix A shows that the class of SAPP mechanisms is necessary to obtain any finite approximation ratio to the second-best: both the best FPP mechanisms and the "Generalized Buyer Offering Mechanism" [7] have an unbounded gap compared to the second-best GFT, even in the bilateral trade setting.

An astute reader may have already realized that the payments to the sellers are not yet defined in the SAPP mechanism. This is because the allocation rule of a SAPP mechanism is not necessarily monotone in the sellers' costs if the mappings $f_i(\cdot)g_{i2[n]}$ are not chosen carefully. Interestingly, we show that if the mappings $f_i(\cdot)g_{i2[n]}$ satisfy a strong type of monotonicity that we call bi-monotonicity (Definition 4.1), then the allocation rule is indeed monotone in each seller's reported cost. Since the sellers are single-dimensional, we can apply Myerson's payment identity to design an incentive compatible payment rule. The final property we need to establish is budget balance, which turns out to be the major technical challenge for us. We provide more

details and intuition about our solution to this challenge in the discussion of the techniques.

In Section 5, we draw a connection between a lower bound to our analysis and one of the major open problems in single dimensional two-sided markets. We prove a reduction from approximating the first-best GFT in the unit-demand setting to bounding the gap between the first-best and second-best GFT in a related single-dimensional setting (Theorem 5.1). If in the latter market, the gap between first-best and second-best GFT is at most c , then our mechanism is a $2c$ -approximation to the first-best GFT in the former market.

1.1 Our Approach and Techniques

1. $\log(1=r)$ -Approximation (Section 3): Our starting point is similar to Colini-Baldeschi et al. [15]. We first argue that the probability space of each item i can be partitioned into $O(\log(1=r))$ events $f_{E_{ij}}g_{j2[\log(2=r)]}$, such that in each event E_{ij} , the median of the buyer's value b_i for item i dominates the median of the i -th seller's cost s_i . The first-best GFT is upper bounded by the sum of the contribution to GFT from each of these events. In bilateral trade, simply setting the posted price to be the median of the buyer's value is sufficient to obtain $1=2$ of the optimal GFT from E_{ij} as shown by McAfee [27]. The $\log(1=r)$ -approximation by Colini-Baldeschi et al. [15] essentially follows from this argument.

To illustrate the added difficulty from multiple items, it suffices to consider a unit-demand buyer. Setting the posted price on each item to be the median of the buyer's value does not provide a good approximation, because the buyer will purchase the item that gives her the highest surplus, which could be very different from the item that generates the most GFT. Similar scenarios are not uncommon in multi-dimensional auction design, and prophet inequalities [23, 24] have been proven to be effective in addressing similar challenges. The main barrier for applying the prophet inequality to two-sided markets is choosing the appropriate random variable as the reward for the prophet/gambler. It is not obvious how to choose a random variable that will translate to a two-sided market mechanism, and in fact, for some choices, no translation between the thresholding policy for the gambler and a two-sided market mechanism is possible.¹ Our key insight is to replace event E_{ij} with a related but different event \bar{E}_{ij} where there is a fixed number ij such that s_i and b_i are always separated by ij ($s_i < b_i$). We further show that the GFT contribution

¹For example, one can choose the GFT from the i^{th} item ($b_i - s_i$)⁺ as the reward of the i^{th} round, but no fixed posted price mechanism corresponds to the policy that only accepts items whose GFT is above a certain threshold. Indeed, no BIC, IR, and BB mechanism can implement a thresholding policy with threshold 0 due to the impossibility result by Myerson and Satterthwaite [30].

from event \bar{E}_{ij} is at least half of the GFT contribution from E_{ij} . Importantly, the GFT contributed by item i in event \bar{E}_{ij} : $(b_i - s_i)^+ = (b_i - i_j)^+ 1[s_i - i_j] + (i_j - s_i)^+ 1[b_i - i_j]$. Note that if we replace E_{ij} with \bar{E}_{ij} , the LHS can exceed the RHS when $i_j > b_i > s_i$. The decomposition of $(b_i - s_i)^+$ using i_j is critical for us to apply the prophet inequality. We can now choose the reward for the gambler to be $v_i = (i_j - s_i)^+ 1[b_i - i_j]$, and the thresholding policy with a threshold T can be implemented with a posted price mechanism where the price for the buyer is i_j and the price for the seller is $i_j - T$.³

When the buyer's feasibility constraint is general downward-closed, the only known prophet inequalities are due to Rubinstein [32] and are $O(\log n)$ -competitive. Unfortunately, the prophet inequalities in [32] are highly adaptive, and thus cannot translate into prices for a single buyer. Further, an almost matching lower bound of $O(\log n = \log \log n)$ is shown by Babaioff et al. [3], precluding much possible improvement for this approach. Instead, we use a constrained fixed posted price mechanism that forces the buyer to buy at least h items (at their posted prices) if she wants to buy any; otherwise, she must leave with nothing. We divide the same variables v_i into $O(\log n)$ buckets based on their contribution to seller surplus. Within each bucket k , all variables v_i lie in $[L_k; 2L_k]$ for some L_k . We prove a concentration inequality for the maximum size of a feasible and affordable set. It guarantees that with constant probability, the buyer will be willing to purchase at least h items (for an appropriate choice of h), generating sufficient GFT.

2. Benchmark of the Second-Best GFT (Section 4.1):

As our goal is to obtain a benchmark of the second-best GFT that is unconditional, the benchmark from the previous (distribution-parameterized) result cannot be used here. We derive a novel benchmark in two steps. Step (i): we create two imaginary one-sided markets: the super seller auction and the super buyer procurement auction. We show that the second-best GFT of the two-sided market is upper bounded by the optimal profit from the super seller auction and the optimal buyer utility from the super buyer procurement auction. Step (ii): we provide an extension of the marginal mechanism lemma [22, 9] to the optimal profit. We show that the optimal profit for selling all items in $[n]$ is upper bounded by the first-best GFT from items in T and the optimal profit for selling items in $[n] \setminus T$, where T is an arbitrary subset of $[n]$. Our key insight is to choose T to be the "likely to trade" items, which are the ones with trade probability at least $1/n$, and apply the marginal mechanism lemma. This partition allows us to use our first result to provide an $O(\log n)$ -approximation to the first-best GFT of the "likely to trade" items using a fixed posted price mechanism. Moreover, we

prove that the optimal buyer utility from the super buyer procurement auction is upper bounded by the GFT of an extension of the "generalized buyer offering mechanism" [7]. Finally, we provide an $O(\log n)$ -approximation to the optimal profit for selling the "unlikely to trade" items using a SAPP mechanism. Note that the approximation crucially relies on the fact that in expectation at most one item can trade among the "unlikely to trade" items.

3. Budget Balance of Seller Adjusted Posted Price Mechanisms (Section 4.3): As mentioned earlier, we restrict our attention to bi-monotonic mappings from cost profiles to buyer prices $f_i(\cdot)g_{i \in [n]}$ to guarantee incentive-compatibility. However, budget balance does not follow from bi-monotonic mappings. We extend the definition of bi-monotonicity to allocation rules and show that all bi-monotonic allocation rules can be transformed into a DSIC, IR, and BB SAPP mechanism. In our proof of the budget balance property, we identify an auxiliary allocation rule q , which may not be implementable by a BB mechanism. We then show that the allocation rule of our SAPP mechanism is "coupled" with q . In particular, our allocation probability is always between $q=4$ and $q=2$. The upper bound $q=2$ allows us to upper bound the payment to the seller, and the lower bound $q=4$ allows us to lower bound the payment we collect from the buyer. Surprisingly, we can prove that the upper bound of the payment to the seller is no more than the lower bound of the buyer's payment. We suspect this type of allocation coupling argument may also be useful in other problems.

1.2 Related Work

Gains from Trade: The main related works are on worst-case GFT approximation. Blumrosen and Mizrahi [6] guarantee an ϵ -approximation to the first-best GFT in the setting of bilateral trade—one buyer, one seller, one item—when the buyer's distribution satisfies the monotone hazard rate condition. Brustle et al. [7] study the more general double auction setting: there are many buyers and sellers, but the goods are identical, and each buyer and seller is unit-demand or unit-supply respectively. In addition, any downward-closed feasibility constraint over the buyer-seller pairs that can trade is allowed. They use the better of a "seller-offering" or "buyer-offering" mechanism to achieve a 2-approximation to the second-best GFT, for general buyers' and sellers' distributions. Colini-Baldeschi et al. [15] show that a simple fixed price mechanism obtains an $O(\frac{1}{r})$ -approximation to GFT in the bilateral trade and double auction settings, but a more careful setting of the fixed price gives an $O(\log \frac{1}{r})$ -approximation for bilateral trade. Our setting is the first multi-dimensional setting with a worst-case approximation guarantee, and we match the $O(\log \frac{1}{r})$ -approximation of [15] while providing an unconditional $O(\log n)$ -approximation.

² $x^+ = \max\{x, 0\}$.

³ A similar fixed posted price mechanism can take care of $(b_i - i_j)^+ 1[s_i - i_j]$.

Other lines of work provide (1) 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 [28, 36, 35], (2) dual asymptotic and worst-case guarantees for double auctions and matching markets [1], and (3) Bulow-Klemperer-style guarantees of the number of additional buyers (or sellers) needed in double auctions in order for the GFT of the new setting running a simple mechanism to beat the first-best GFT of the original setting [2].

Multi-Dimensional Revenue: In the setting where one seller owns all of the items, has no cost for the items, and is the mechanism designer, much more is known. However, even when selling to a single additive bidder (e.g. with no feasibility constraints), posted prices can achieve at best an $O(\log n)$ -approximation [22, 25]. In order to obtain a constant-factor approximation for an additive buyer, Babaioff et al. [4] use the better of posted prices and posting a price on the grand bundle, and a variation works for a single subadditive (which includes constrained-additive) buyer as well [33]. However, in a two-sided market where items are owned by separate sellers, it is not clear how to implement bundling in an incentive-compatible way. The mechanisms used to obtain constant-approximations for multiple constrained-additive, XOS, or subadditive buyers [12, 10] are only more complex.

Welfare in Two-Sided Markets: Colini-Baldeschi et al. [14] consider welfare maximization in the double auction setting with matroid feasibility constraints. They generalize sequential posted price mechanisms (SPMs) to the two-sided market setting, guaranteeing a constant-factor approximation to welfare. The mechanism posts prices for each buyer-seller combination (not just for each item), visits the buyers and sellers simultaneously in the given order, and advances on either side when the price is rejected. Trade occurs when both sides accept the trade. Follow up work of Colini-Baldeschi et al. [16] generalizes the idea to the setting where buyers are XOS and sellers are additive. Here, there is a posted price for each item, but only “high welfare” items are considered. The buyers visit and pick out the bundles they want among the high welfare items. Then sellers are given the opportunity to sell their entire bundle of items demanded by the buyers (but not any subset), and they are skipped with some probability. Like the previous work, this mechanism is ex-post IR, DSIC, and strongly BB (buyer payments equal seller payments). As only “high welfare” items are considered, it is possible for their mechanism to not trade any item when the minimum trade probability r is a constant.

Blumrosen and Dobzinski [5] give an IR, BIC, and strongly BB mechanism for bilateral trade that obtains in expectation a constant-fraction of the optimal welfare. Dütting et al. [19] study welfare maximization in the prior-free set-

ting and present DSIC, IR, and weakly BB (buyer payments exceed seller payments) mechanisms for double auctions with feasibility constraints on either side.

2 Preliminaries

Two-sided Markets. We focus on two-sided markets between a single buyer and n unit-supply sellers. Every seller i sells a heterogeneous item. For simplicity we denote the item sold by seller i as item i . Each seller i has cost s_i for producing item i , where s_i is drawn independently from distribution D_i^S . The buyer has value b_i for every item i where b_i is drawn independently from distribution D_i^B . D_i^S and D_i^B are public knowledge. Let $D^B = \prod_{i=1}^n D_i^B$ be the distribution of the buyer’s value profile and $D^S = \prod_{i=1}^n D_i^S$ be the distribution of the cost profile for all sellers. Let $b = (b_1; \dots; b_n)$ and $s = (s_1; \dots; s_n)$ denote the value (or cost) profile for the buyer and all sellers. For notational convenience, for every i we denote b_{-i} (or s_{-i}) to be the value (or cost) profile without item i . For every i , F_i ; f_i (or G_i ; g_i) denote the cumulative distribution function and density function of D_i^B (or D_i^S). Throughout the paper we assume that all distributions are continuous over their support, and so the inverse cumulative functions F_i^{-1} and G_i^{-1} exist.⁴

Throughout this paper, we assume that the buyer has a constrained-additive valuation over the items, which means that the buyer is additive over the items, but is only allowed to take a feasible set of items with respect to a downward-closed⁵ constraint $F \subseteq 2^{[n]}$. Formally, for every b and $S \subseteq [n]$, the buyer’s value for a set of items S is: $v(b; S) = \max_{T \subseteq F; T \subseteq S} \sum_{i \in T} b_i$.

Mechanism and Constraints. Any mechanism in the two-sided market defined above is specified by the tuple $(x; p^B; p^S)$ where x is the allocation rule of the mechanism and p^B, p^S are the payment rules. For every profile $(b; s)$ and every i , $x_i(b; s)$ is the probability that the buyer trades with seller i under profile $(b; s)$. $p^B(b; s)$ is the payment from the buyer and $p_i^S(b; s)$ is the gains for (or payment to) seller i . All agents in the market have linear utility functions.⁶ We call the mechanism ex-ante Strongly Budget Balanced (SBB) or Weakly Budget Balanced (WBB) if the buyer’s expected payment equals, or is greater than, the sum of all

⁴Any discrete distribution can be made continuous by replacing each point mass a with a uniform distribution on $[a - \epsilon; a + \epsilon]$, for arbitrarily small ϵ . Thus our result applies to discrete distributions as well by losing arbitrarily small GFT.

⁵ $F \subseteq 2^{[n]}$ is downward-closed if for every $S \subseteq F$, we have $S^0 \subseteq F$; $S^0 = \{i \in [n] : \exists S \subseteq F, i \in S\}$.

⁶Without loss of generality we can assume that the mechanism will only allow the buyer to trade with a (possibly randomized) set S of sellers where $S \subseteq F$. For any trading set T , let S denote the utility-maximizing feasible subset, $S = \arg \max_{S \subseteq F; S \subseteq T} \sum_{i \in S} b_i$. If we only allow the buyer to trade with the sellers in S instead of all of T , the gains from trade of the mechanism will not decrease.

sellers' expected gains, respectively, over the randomness of the mechanism and the profiles of all agents. We call the mechanism ex-post SBB (or ex-post WBB) if this property holds for every agent's profile. The definition of incentive compatibility and individual rationality are as follows.

BIC: For every agent, reporting her true value (or cost) maximizes her expected utility over the profiles of other agents.

DSIC: For every agent, reporting her true value (or cost) maximizes her expected utility, no matter what other agents report.

(Bayesian) IR: For every agent, reporting her true value (or cost) derives non-negative utility over the profiles of other agents.

Ex-post IR: For every agent, reporting her true value (or cost) derives non-negative utility, no matter what other agents report.

Gains from Trade. We aim to maximize the Gains from Trade (GFT), i.e. the gains of social welfare induced by the mechanism. Formally, given a mechanism $M = (x; p^B; p^S)$, the expected GFT of M is

$$GFT(M) = E_{b, D^B; s, D^S} \left[\sum_{i=1}^n x_i(b; s) (b_i - s_i) \right];$$

We use SB-GFT to denote the optimal GFT attainable by any BIC, IR, ex-ante WBB mechanism (also known as the "second-best" mechanism). Let FB-GFT denote the maximum possible gains of social welfare among all feasible allocations (known as the "first-best"). Formally

$$FB-GFT = E_{b, D^B; s, D^S} \max_{S \subseteq F} \sum_{i \in S} (b_i - s_i);$$

In Section 3, the distribution-parameterized approximation uses the parameter r , the minimum probability over all items i that the buyer's value for item i is at least seller i 's cost. Formally, for every item $i \in [n]$, let $r_i = \Pr_{b_i \in D^B; s_i \in D^S} [b_i \geq s_i]$ denote the probability that the buyer's value for item i exceeds seller i 's cost. Without loss of generality, assume that $r_i > 0$ for all $i \in [n]$.⁷ Let $r = \min_{i \in [n]} r_i > 0$.

3 A Distribution-Parameterized Approximation

In this section, we present an $O(\frac{\log(1+r)}{r})$ -approximation to FB-GFT when the buyer's feasibility constraint F is $(;)$ -selectable, and an $O(\log(n) \log(\frac{1}{r}))$ -approximation for a general constrained-additive buyer. In Section 3.1, we show that FB-GFT can be bounded by the sum of four

⁷If $r_i = 0$ the mechanism should never trade between the buyer and seller i , and so it can remove seller i from the market. This will not decrease the GFT of the mechanism as $b_i < s_i$ with probability 1.

separate terms. In Section 3.2 we show that two of the terms ("buyer surplus") are relatively easy to bound using fixed posted price (FPP) mechanisms with the same prices posted on both sides. In Section 3.3, we consider the special case of a unit-demand buyer and bound the other two terms ("seller surplus") using FPP mechanisms combined with the prophet inequality. In Section 3.4, we introduce the concept of selectability [21] and bound the seller surplus for any selectable feasibility constraint by using a constrained FPP mechanism. In Section 3.5, we present our result for a general constrained-additive buyer.

3.1 Upper Bound of FB-GFT. For every i , let $\bar{F}_i = 1 - F_i$ denote the complementary CDF of b_i . Let x_i and y_i be the $\frac{r_i}{2}$ -quantile of the buyer's and seller's distribution for item i , respectively. Formally, $x_i = \bar{F}_i^{-1}(\frac{r_i}{2})$; $y_i = G_i^{-1}(\frac{r_i}{2})$. We first prove that $x_i \geq y_i$.

LEMMA 3.1. For every $i \in [n]$, $x_i \geq y_i$.

Proof. Note that for every $i \in [n]$, $b_i < x_i \wedge s_i > y_i$ implies that $b_i < s_i$. We have

$$\Pr_{b_i \in D^B; s_i \in D^S} [b_i < s_i] \geq \Pr_{b_i \in D^B; s_i \in D^S} [b_i < x_i \wedge s_i > y_i] = (1 - \Pr_{b_i \in D^B} [b_i < x_i]) \Pr_{s_i \in D^S} [s_i > y_i] = (1 - \frac{r_i}{2}) \Pr_{s_i \in D^S} [s_i > y_i];$$

Suppose $x_i < y_i$. Then $(1 - \frac{r_i}{2}) \Pr_{s_i \in D^S} [s_i > x_i] > (1 - \frac{r_i}{2}) \Pr_{s_i \in D^S} [s_i > y_i] = (1 - \frac{r_i}{2}) \Pr_{s_i \in D^S} [s_i > x_i]$. This is a contradiction. Thus $x_i \geq y_i$. \square

In the following upper bound, we will separate the probability space for each item i into $2d \log(2/r)$ events, and then divide the GFT into buyer surplus and seller surplus terms according to the cutoff for each event. For every $b; s$, define the feasible set that maximizes the GFT as $S(b; s) = \arg \max_{S \subseteq F} \sum_{k \in S} (b_k - s_k)$, and break ties arbitrarily. Observe the following upper bound for the first-best GFT:

$$FB-GFT = E \left[\sum_{i \in [n]} \max_{b; s \in F} (b_i - s_i)^+ \right] \\ = E_{b; s} \left[\sum_{i \in [n]} (b_i - s_i) \mathbb{1}[i \in S(b; s)] \mathbb{1}[b_i - s_i \wedge s_i < x_i] \right] \\ + E_{b; s} \left[\sum_{i \in [n]} (b_i - s_i) \mathbb{1}[i \in S(b; s)] \mathbb{1}[b_i - s_i \wedge y_i \leq s_i] \right];$$

where the inequality holds because $x_i \geq y_i$ for all i . We refer to the two terms of RHS as ① and ② accordingly. We first consider term ①. For every $i \in [n]$; $j \in \{1, 2, \dots, d\}$; $\log(\frac{2}{r})e$, let $i_{j-1} = \bar{F}_i^{-1}(\frac{1}{2^{j-1}})$. Let $E_{i,j}$ be the event that $\bar{F}_i^{-1}(\frac{1}{2^{j-1}}) \leq s_i < \bar{F}_i^{-1}(\frac{1}{2^j}) \wedge b_i < \bar{F}_i^{-1}(\frac{1}{2^{j-1}})$. Then we have

$$\textcircled{1} = E_{b; s} \left[\sum_{j=1}^{d \log(\frac{2}{r})e} \sum_{i \in [n]} (b_i - s_i)^+ \mathbb{1}[i \in S(b; s) \wedge E_{i,j}] \right].$$

As discussed in Section 1.1, in order to bound the benchmark with fixed posted price mechanisms, we will consider a more restrictive event \bar{E}_{ij} and show that the GFT contribution from event \bar{E}_{ij} is at least half of the GFT contribution from E_{ij} .

LEMMA 3.2. For every i, j , let \bar{E}_{ij} be the event that $F_i^{-1}(\frac{1}{2}) \leq s_i \leq F_i^{-1}(\frac{1}{2}) \wedge F_j^{-1}(\frac{1}{2}) \leq s_j \leq F_j^{-1}(\frac{1}{2})$. Then the following inequality holds for every $j = 1, \dots, d \log(2=r)e$:

$$E_{b,s} [\sum_{i \in S(b;s)} (b_i - s_i)^+ \mathbb{1}[E_{ij}]] \geq \frac{1}{2} E_{b,s} [\sum_{i \in S(b;s)} (b_i - s_i)^+ \mathbb{1}[\bar{E}_{ij}]]$$

Moreover,

$$\textcircled{1} \quad 2 \sum_{j=1}^{d \log(2=r)e} E \max_{b,s} \sum_{i \in S} (b_i - s_i)^+ \mathbb{1}[s_i \leq p_i] \leq \sum_{j=1}^{d \log(2=r)e} E \max_{b,s} \sum_{i \in S} (b_i - s_i)^+ \mathbb{1}[b_i \leq p_i]$$

We will refer to the two terms of RHS as $\textcircled{3}$ and $\textcircled{4}$ accordingly.

Readers may notice that $\Pr[\bar{E}_{ij}] = \frac{1}{2} \Pr[E_{ij}]$. However, this alone does not prove the first statement of Lemma 3.2, since both the indicator $\mathbb{1}[E_{ij}]$ and the contributed GFT $(b_i - s_i)^+ \mathbb{1}[i \in S(b;s)]$ depend on the realization of b_i, s_i . In Lemma 3.2 we show that the two random variables are positively correlated with respect to b_i , which allows us to prove the first statement. The second statement follows from the fact that $(b_i - s_i)^+ \leq (b_i - p_i)^+ + (p_i - s_i)^+$ for every b_i, s_i , and that $S(b;s) \subseteq S(b;p_i)$ for every b,s .

In Lemma 3.3, we bound term $\textcircled{2}$ in a similar way. The proof of Lemmas 3.2 and 3.3 can be found in the full version of the paper.

LEMMA 3.3. For every $i \in [n]$ and $j = 1, \dots, d \log(2=r)e$, let $\bar{G}_{ij} = G_i^{-1}(\frac{1}{2})$. Then

$$\textcircled{2} \quad 2 \sum_{j=1}^{d \log(2=r)e} E \max_{b,s} \sum_{i \in S} (b_i - s_i)^+ \mathbb{1}[s_i \leq \bar{G}_{ij}] \leq \sum_{j=1}^{d \log(2=r)e} E \max_{b,s} \sum_{i \in S} (b_i - s_i)^+ \mathbb{1}[b_i \leq \bar{G}_{ij}]$$

We will refer to the two terms of RHS as $\textcircled{5}$ and $\textcircled{6}$ accordingly.

We refer to terms $\textcircled{3}$ and $\textcircled{5}$ as buyer surplus, and $\textcircled{4}$ and $\textcircled{6}$ as seller surplus. In the rest of this section we will bound each term separately.

3.2 Bounding Buyer Surplus. We bound terms $\textcircled{3}$ and $\textcircled{5}$ using fixed posted price mechanisms. Let GFT_{FPP} denote the optimal GFT among all fixed posted price mechanisms. Recall that our market is not symmetric: a single multi-dimensional buyer with a feasibility constraint faces multiple single-dimensional sellers. As a result, even for the general constrained-additive buyer, bounding buyer surplus is fairly straightforward using fixed price mechanisms that set $p_i = b_i$ (or $p_i = s_i$) for each term.

LEMMA 3.4. For any $p_i \in [0, 1]$,

$$E \max_{b,s} \sum_{i \in S} (b_i - p_i)^+ \mathbb{1}[s_i \leq p_i] \leq \text{GFT}_{\text{FPP}}$$

Thus both $\textcircled{3}$ and $\textcircled{5}$ are upper bounded by $O(\log(\frac{1}{r})) \text{GFT}_{\text{FPP}}$.

Proof. Consider the fixed posted price mechanism M with $p_i = b_i$. For every s , let $A(s) = \{i \in [n] \mid s_i \leq p_i\}$ be the set of available items. Then the buyer will choose the best set $S \subseteq A(s); S \subseteq F$ that maximizes $\sum_{i \in S} (b_i - p_i)^+$ (and not buy any item if $b_i - p_i \leq 0$ for all i). Thus the gains from trade $\sum_{i \in S} (b_i - s_i)$ is at least $\sum_{i \in S} (b_i - p_i)$. We have

$$\text{GFT}(M) = E_{b,s} \max_{S \subseteq A(s); S \subseteq F} \sum_{i \in S} (b_i - p_i)^+ = E_{b,s} \max_{S \subseteq F} \sum_{i \in S} (b_i - p_i)^+ \mathbb{1}[s_i \leq p_i]$$

To bound terms $\textcircled{3}$ and $\textcircled{5}$, just apply the above inequality with $p_i = b_i$ (or $p_i = s_i$). \square

3.3 Bounding Seller Surplus for One Unit-Demand Buyer. In the remainder of this section, we will bound the seller surplus terms $\textcircled{4}$ and $\textcircled{6}$. As a warm-up, we first focus on the case where the buyer is unit-demand, i.e. the buyer is only interested in at most one item. Here, the prophet inequality suffices for our bound.

LEMMA 3.5. When the buyer is unit-demand, for any $p_i \in [0, 1]$, we have

$$E_{b,s} [\max_{i \in S} (b_i - s_i)^+ \mathbb{1}[b_i \leq p_i]] \leq 2 \text{GFT}_{\text{FPP}; j=1}$$

Hence terms $\textcircled{4}$ and $\textcircled{6}$ are both upper-bounded by $O(\log(\frac{1}{r})) \text{GFT}_{\text{FPP}}$.

Proof. For every i , let $v_i = (b_i - s_i)^+ \mathbb{1}[b_i \leq p_i]$ be a random variable that depends on b_i and s_i . Let $v = (v_i)_{i \in [n]}$. Let V_i be the distribution of v_i where $b_i \in D_i; s_i \in \mathcal{D}_i$, and $V = \prod_{i=1}^n V_i$ be the distribution of v .

Then the LHS of the inequality in the Lemma statement is equal to $E_{v,v} [\max_i v_i]$.

Consider any threshold $\epsilon > 0$. Observe that $v_i \leq p_i$ if and only if $b_i \leq p_i - s_i$. Consider the fixed posted price mechanism M with $b^i = p_i$ and $s^i = p_i - \epsilon$ for every $i \in [n]$. Whenever the buyer purchases some item i , we must have $b_i \leq p_i$ (the buyer buys) and $s_i \leq p_i - \epsilon$ (the seller sells), and the contributed GFT satisfies $b_i - s_i \geq \epsilon$. In addition, the buyer will purchase some item if and only if there exists some i such that $v_i \geq \epsilon$. Therefore we can apply the prophet inequality [23, 24, 34] with threshold $\epsilon = \frac{1}{2} E_{v,v} [\max_i v_i]$ to ensure that the GFT of mechanism M is at least $\frac{1}{2} E_{v,v} [\max_i v_i]$. \square

3.4 Bounding Seller Surplus with Selectability. In this subsection we bound terms (4) and (6) for a more general class of constraints F using a variant of a fixed posted price (FPP) mechanism which we call constrained FPP. In the variant, the mechanism determines a (randomized) subconstraint $F^0 \subseteq F$ upfront. Then the buyer is only allowed to take a feasible set in F^0 (among all items that the sellers agree to sell at prices $\{s_{i \in [n]}\}$) and pays the price b^i for each item she takes.⁸ Let GFT_{CFPP} denote the optimal GFT among all constrained FPP mechanisms.⁹ Since all of the posted prices as well as the subconstraint are independent from the agents' reported profiles, the mechanism is DSIC and ex-post IR. The mechanism is also ex-post WBB since $b_i \geq s_i$ for all $i \in [n]$.

To present our result, we introduce a concept for downward-closed constraints called $(;)$ -selectability [21]. Feldman et al. introduce $(;)$ -selectability in the study of Online Contention Resolution Schemes (OCRS) [21]. An OCRS is an algorithm defined for the following online selection problem: There is a ground set I , and the elements are revealed one by one, with item i active with probability x_i independent of the other items. The algorithm is only allowed to accept active elements and has to irrevocably make a decision whether to accept an element before the next one is revealed. Moreover, the algorithm can only accept a set of elements subject to a feasibility constraint F . We use the vector x to denote active probabilities for the elements and $R(x)$ to denote the random set of active elements.

DEFINITION 3.1. (RELAXATION) We say that a polytope $P \subseteq [0; 1]^{[n]}$ is a relaxation of P_F if it contains the same $f_0; 1g$ -points, i.e., $P \setminus \{f_0; 1g\} = P_F \setminus \{f_0; 1g\}$.

⁸Throughout this paper, we assume for simplicity that the buyer will purchase item i when $b_i = p_i$ as long as the bundle remains feasible after including i . Without this tie-breaking rule, one can simply decrease the posted price for each item by an arbitrarily small value ϵ , and the loss of GFT will be arbitrarily small.

⁹Note that FPP is a subclass of constrained FPP, and therefore $GFT_{FPP} \leq GFT_{CFPP}$.

DEFINITION 3.2. An Online Contention Resolution Scheme (OCRS) for a polytope $P \subseteq [0; 1]^{[n]}$ and feasibility constraint F is an online algorithm that selects a feasible and active set $S \subseteq R(x)$ and $S \subseteq F$ for any $x \in P$. A greedy OCRS greedily decides whether or not to select an element in each iteration: given the vector $x \in P$, it first determines a sub-constraint $F_{;x} \subseteq F$. When element i is revealed, it accepts the element if and only if i is active and $S \subseteq F_{;x}$, where S is the set of elements accepted so far. In most cases, we choose P to be P_F , the convex hull of all characteristic vectors of feasible sets in F : $P_F = \text{conv}(\{1_S \mid S \subseteq F\})$.

DEFINITION 3.3. ((;)-SELECTABILITY [21]) For any $\epsilon \in (0; 1)$, a greedy OCRS for P and F is $(;)$ -selectable if for every $x \in P$ and $i \in I$,

$$\Pr[S \subseteq F_{;x}; |S \cap R(x)| \geq \epsilon] \geq \epsilon$$

The probability is taken over the randomness of $R(x)$ and the subconstraint $F_{;x}$. We slightly abuse notation and say that F is $(;)$ -selectable if there exists a $(;)$ -selectable greedy OCRS for P_F and F .

The following lemma is adapted from [21] and connects $(;)$ -selectability to constrained FPP mechanisms. Once again, the OCRS gives us both a GFT guarantee and a mechanism: variables v_i correspond to the bound on seller surplus, buyer item prices are $p_i g_{i \in [n]}$, seller prices are $p_i - \epsilon g_{i \in [n]}$, and the subconstraint is suggested by the OCRS. The proof of Lemma 3.6 can be found in the full version of the paper.

LEMMA 3.6. Suppose there exists a $(;)$ -selectable greedy OCRS for the polytope P_F , for some $\epsilon \in (0; 1)$. Fix any $f_{p_i g_{i \in [n]}} \in \mathbb{R}_+^n$. For every $i \in [n]$, let $v_i = (p_i - s_i)^+ + 1[b_i \leq p_i]$. For any $q \in P_F$ that satisfies $q_i \leq \Pr_{b_i, s_i}[b_i - p_i > s_i] + \epsilon$, let $\epsilon_i = p_i - G_i$ ($q_i = \Pr[b_i \leq p_i]$).¹⁰ We have

$$\mathbb{E}_{b_i, s_i} [v_i - 1[v_i \leq \epsilon_i]] \geq GFT_{CFPP}$$

Moreover, there exists a choice of q such that

$$\mathbb{E}_{b_i, s_i} \max_{S \subseteq F} \sum_{i \in S} (p_i - s_i)^+ + 1[b_i \leq p_i] g_i \leq \mathbb{E}_{b_i, s_i} [v_i - 1[v_i \leq \epsilon_i]] + GFT_{CFPP} + \epsilon$$

For each j in the summation, choose p_j from Lemma 3.6 to be ϵ_j (or ϵ_j).^c Then both terms (4) and (6) are bounded by $\log(1/\epsilon) GFT_{CFPP}$. Theorem 3.1 then follows directly from Lemmas 3.2, 3.3, 3.4, and 3.6.

Feldman et al. [21] show that many natural constraints—including matroids, matchings, knapsack, and their

¹⁰When $q_i \leq \Pr_{b_i, s_i}[b_i - p_i > s_i]$, $q_i = \Pr[b_i \leq p_i] - 1$. Thus ϵ_i is well-defined.

compositions—are $(;)$ -selectable for some constants α and β . For all of these, Theorem 3.1 implies that GFT_{CFPP} is an $O(\log(1+r))$ -approximation to FB-GFT. We refer the readers to the full version of our paper for more details.

THEOREM 3.1. Suppose the buyer’s feasibility constraint F is $(;)$ -selectable for some $\alpha; \beta \in (0; 1)$. Then $FB\text{-}GFT \in O(\frac{\log(1+r)}{\alpha\beta}) GFT_{CFPP}$.

3.5 General Constrained-Additive Buyer. In this section, we consider the case of a general constrained-additive buyer, and prove an $O(\log(n) \log(1+r))$ approximation to FB-GFT using constrained FPP mechanisms. Note that Lemmas 3.2, 3.3, and 3.4 still hold in this setting. It is sufficient to bound the seller surplus term with GFT_{CFPP} .

Throughout this section, we will use the following variant of FPP mechanisms: Other than posted prices, the mechanism also determines an integer $h > 0$ upfront. The buyer can purchase any set of items of size at least h by paying the posted prices for each item in the set; otherwise, she leaves with nothing. This is a subclass of constrained FPP, with subconstraint $F^0 = \{S \subseteq [n] : |S| \geq h\}$.¹¹

LEMMA 3.7. For any $f, p_i, g_{i \in [n]} \in \mathbb{R}_+^n$,

$$A = E_{b; s} \max_{T \subseteq [n]} \sum_{i \in T} f(p_i - s_i)^+ \mathbb{1}[b_i \geq p_i] g_i$$

$$O(\log(n)) GFT_{CFPP}$$

Hence terms ④ and ⑥ are both upper-bounded by $O(\log(n) \log(\frac{1}{\alpha\beta})) GFT_{CFPP}$.

For every $i \in [n]$, again construct random variables $v_i = (p_i - s_i)^+ \mathbb{1}[b_i \geq p_i]$. The main issue here is that in an FPP mechanism, say with posted prices $p_i = \hat{p}_i$, the buyer will pick the maximum weight feasible set (among all items that sellers are willing to sell) according to weight $b_i - p_i$ (her utility). However, this might be far from the set used in the benchmark, i.e. the maximum weight feasible set according to weight $p_i - s_i$. In the previous section (when the constraint F had selectability), by setting different prices for both sides and adding a more restrictive constraint, we guaranteed that if both the buyer and seller accept the posted prices for some item, then the buyer would purchase this item with at least constant probability. For general downward-closed F , it is unclear how to achieve this property with a constrained FPP mechanism.

For every $b; s$, let $T(b; s) = \arg\max_{T \subseteq [n]} \sum_{i \in T} v_i$ be the optimal set used in the benchmark. We divide A into three terms according to the value of v_i when i is in this optimal set: $v_i < A/2n$, $v_i \in [A/2n; 2nA]$ and $v_i > 2nA$. Denote the three terms $A_S; A_M; A_L$ accordingly. First we notice that the contribution of A_S is at most a constant fraction

¹¹If $h = 1$, the mechanism becomes a standard FPP mechanism without any subconstraint F^0 .

of A , as $A_S = E_{b; s} \sum_{i \in T(b; s)} v_i \mathbb{1}[v_i < \frac{A}{2n}] \leq E_{b; s} \frac{A}{2n} \leq \frac{A}{2}$.

For A_L , in Lemma 3.9, we first prove that $\Pr_{b_i; s_i} [b_i - p_i \wedge p_i - s_i \geq 2nA] \leq \frac{1}{2n}$ holds for every i . This implies that in a standard FPP mechanism (where $h = 1$) with $p_i = p_i; s_i = (p_i - 2nA)^+$ for all i , the buyer purchases each item i with probability at least $\frac{1}{2}$ if both the buyer and seller i accept the posted prices. First, we will need Lemma 3.8.

LEMMA 3.8. Given any constrained FPP mechanism M with posted prices $f_i^B, g_{i \in [n]}^S$ and $h = 1$, suppose $\Pr_{b_i; s_i} [b_i - p_i \wedge p_i - s_i \geq 2nA] \leq \frac{1}{2}$. Then

$$GFT(M) \geq \frac{1}{2} \sum_{i \in [n]} E_{b_i; s_i} (b_i - s_i) \mathbb{1}[b_i \geq p_i \wedge p_i \geq s_i]:$$

Proof. For any item i , the buyer will purchase item i if both of the following events happen:

1. $b_i \geq p_i$ and $s_i \leq p_i$;
2. For all items $k \neq i$, either $s_k > p_k$ or $b_k < p_k$.

By the union bound, the second event happens with probability at least $1 - \prod_{k \neq i} \Pr[b_k \geq p_k \wedge s_k \leq p_k] \geq \frac{1}{2}$. Since both events are independent, we have $GFT(M) \geq \frac{1}{2} \sum_{i \in [n]} E_{b_i; s_i} (b_i - s_i) \mathbb{1}[b_i \geq p_i \wedge s_i \leq p_i]$. \square

LEMMA 3.9. $A_L = E_{b; s} [\sum_{i \in [n]} v_i \mathbb{1}[v_i \in [A/2n; 2nA]]] \leq 2 GFT_{FPP}$.

Proof. Consider the FPP mechanism with $p_i = p_i; s_i = (p_i - 2nA)^+$ for all i (and $h = 1$).

Note that for every $i \in [n]$, it must hold that $\Pr_{b_i; s_i} [b_i - p_i \wedge p_i - s_i \geq 2nA] \leq \frac{1}{2n}$. In fact,

$$A_S = E_{b; s} [\sum_{i \in [n]} (p_i - s_i)^+ \mathbb{1}[b_i \geq p_i]]$$

$$\leq 2nA \sum_{b; s} \Pr [b_i - p_i \wedge p_i - s_i \geq 2nA] \leq \frac{1}{2}$$

Thus by Lemma 3.8 and the fact that $b_i - s_i \leq p_i - s_i$ when $b_i \geq p_i$, we have $GFT_{FPP} \geq \frac{1}{2} \sum_{i \in [n]} E [(p_i - s_i) \mathbb{1}[b_i \geq p_i \wedge p_i - s_i \geq 2nA]] \geq \frac{1}{2} A_L$. \square

In Lemma 3.10 we bound A_M , which is the primary challenge for this approximation.

LEMMA 3.10. $A_M = O(\log(n)) GFT_{CFPP}$.

Proof. We further divide the interval $[A/2n; 2nA]$ into $O(\log(n))$ buckets, where in each bucket k , v_i falls in the range $[L_k; 2L_k]$ for some L_k . Formally, for any $k \in \{1; 2; \dots; d\}$ let $L_k = \frac{A}{2n} 2^{k-1}$. We have $A_M = \sum_{k=1}^d \sum_{i \in [n]} E [v_i \mathbb{1}[v_i \in [L_k; 2L_k]]]$:

In the rest of the proof, we will show that for any k , there exists some constant $c > 0$ such that

$$\mathbb{E}_{b,s} \sum_{i \in [n]} v_i \mathbb{1}[i \in T(b; s) \wedge v_i \geq L_k; 2L_k] \geq c \text{GFT}_{\text{CFPP}}; i \quad \#$$

Fix any k . For every $i \in [n]$, let $t_i^{(k)} = \frac{v_i - L_k}{2L_k} \mathbb{1}[L_k \leq v_i \leq 2L_k]$. This is a random variable in $[\frac{1}{2}, 1]$. Note that all random variables $t = \{t_i^{(k)}\}_{i \in [n]}$ are independent. Let $Z(t) = \max_{T \subseteq [n]} \sum_{i \in T} t_i^{(k)}$. Then the contribution to A_L from values in this range is bounded by the expectation of the random variable $Z(t)$:

$$\mathbb{E}_{b,s} \sum_{i \in [n]} v_i \mathbb{1}[i \in T(b; s) \wedge v_i \geq L_k; 2L_k] \geq \sum_{i \in [n]} \max_{S \subseteq [n]} \sum_{i \in S} v_i \mathbb{1}[L_k \leq v_i \leq 2L_k] = 2L_k \mathbb{E}[Z(t)]; i \in S \quad \#$$

Now consider the constrained FPP mechanism with $b_i = p_i$ and $s_i = (p_i - L_k)^+$ for every i (the threshold h is determined later). Then in the mechanism, whenever the buyer purchases an item, the contributed GFT is at least L_k . Thus it is sufficient to show that the expected size of the purchasing set is at least a constant factor of $\mathbb{E}[Z(t)]$. Note that $Z(t)$ is a random variable on t , which is the maximum weight feasible set over n independent random variables in $[0; 1]$. In Lemma 3.11, we prove that $Z(t)$ concentrates near its mean. The proof is postponed to Section 3.5.1.

LEMMA 3.11. For any $c \in (0; 1)$,

$$\Pr_t[Z(t) \leq c \mathbb{E}[Z(t)]] \leq \frac{(1 - c)^2}{1 + \frac{(1 - c)^2}{2 \mathbb{E}[Z(t)]}}$$

We first suppose that $\mathbb{E}[Z(t)] \geq \frac{1}{4}$. By applying Lemma 3.11 with $c = \frac{1}{2}$, we get

$$\Pr_t[Z(t) \geq \frac{\mathbb{E}[Z(t)]}{2}] \geq \frac{1}{12}$$

Let $h = \max\{\frac{\mathbb{E}[Z(t)]}{2}, 1\}$. In mechanism M_k , note that for every i , $t_i^{(k)} > 0$ implies that item i is on the market and that the buyer can afford it. With probability at least $\frac{1}{12}$, $Z(t) \geq h$, which implies that the item set $\{i \in [n] : \arg\max_{S \subseteq [n]} \sum_{i \in S} t_i^{(k)} \wedge t_i^{(k)} > 0\}$ is a feasible set of size at least h . (Recall that all $t_i^{(k)}$ are in $[\frac{1}{2}, 1]$). In this scenario, the buyer will purchase a set of items of size at least h . For every item i she purchases, the contributed GFT is $b_i - s_i = L_k$. Thus, $\text{GFT}(M_k) \geq \frac{1}{12} h L_k$. Readers who are familiar with mechanism design may notice that the role of the size threshold h is similar to an “entry fee” in the

posted price mechanism in auctions [4, 8, 10, 12, 33, 37], though the buyer doesn’t have to pay extra money to attend the auction. It guarantees that the buyer will purchase at least h items when she can afford it, as otherwise she gets no utility.

When $\mathbb{E}[Z(t)] < \frac{1}{4}$, we have $h = \frac{\mathbb{E}[Z(t)]}{4}$. Thus

$$\mathbb{E}_{b,s} \sum_{i \in [n]} v_i \mathbb{1}[i \in T(b; s) \wedge v_i \geq L_k; 2L_k] \geq \frac{1}{4} \mathbb{E}[Z(t)] \geq \frac{1}{16} \mathbb{E}[Z(t)] \geq \frac{1}{16} \text{GFT}_{\text{CFPP}}; i \quad \#$$

Now we consider the case where $\mathbb{E}[Z(t)] < \frac{1}{4}$. For every i , let $q_i = \Pr_i[t_i^{(k)} > 0] = \Pr_{b_i, s_i}[b_i - p_i \wedge p_i - s_i \geq L_k; 2L_k]$. Then it holds that $\Pr[\sum_{i \in [n]} t_i^{(k)} = 0] = \prod_{i=1}^n (1 - q_i) > \frac{1}{2}$: This is because if there exists i such that $t_i^{(k)} > 0$, then $Z(t) = \max_{T \subseteq [n]} \sum_{i \in T} t_i^{(k)} \geq \frac{1}{2}$ as $t_i^{(k)} \in [\frac{1}{2}, 1]$ for every i . Thus if $\Pr[\sum_{i \in [n]} t_i^{(k)} = 0] > \frac{1}{2}$, then $\mathbb{E}[Z(t)] < \frac{1}{4}$, which leads to a contradiction.

Consider the constrained FPP mechanism M with $b_i = p_i$, $s_i = (p_i - L_k)^+$, and $h = 1$. For every i , define event $E_i = \{t_j^{(k)} > 0 \wedge t_j^{(k)} = 0; \sum_{j \in [n]} t_j^{(k)} = 1\}$. Note that $t_i^{(k)} > 0$ implies that seller i accepts price s_i and the buyer can afford item i . Under event E_i , there is at least one item on the market that the buyer can afford, i.e. item i . Thus the buyer must purchase some item j on the market that she can afford (possibly item i). For this item j , we have $b_j - s_j \geq p_j - s_j \geq L_k$. Since all E_i s are disjoint events, we have $\text{GFT}(M) \geq \sum_{i \in [n]} \Pr[E_i] L_k = \sum_{i \in [n]} q_i \prod_{j \neq i} (1 - q_j) L_k \geq \sum_{i \in [n]} q_i \prod_{j \neq i} (1 - q_j) > \frac{1}{2} L_k$; where the equality uses the fact that all $t^{(k)}$ s are independent. On the other hand, since $t_i^{(k)} \leq 1$ for any i , $\mathbb{E}[Z(t)] \leq \sum_{i \in [n]} \Pr_i[t_i^{(k)} > 0] = \sum_{i \in [n]} q_i$. Thus

$$\mathbb{E}_{b,s} \sum_{i \in [n]} v_i \mathbb{1}[i \in T(b; s) \wedge v_i \geq L_k; 2L_k] \geq \frac{1}{2} L_k \mathbb{E}[Z(t)] \geq \frac{1}{4} \text{GFT}_{\text{CFPP}}; i \quad \#$$

Summing the inequality over all k finishes the proof.

Proof of Lemma 3.7: By Lemmas 3.9, 3.10, and the fact that $\mathbb{E}[A_S] \geq \frac{1}{2} \mathbb{E}[A]$, we have that

$$A \geq 2(A_M + A_L) - O(\log(n)) \text{GFT}_{\text{CFPP}}; 2$$

Theorem 3.2 summarizes our result for a general constrained-additive buyer. It directly follows from Lemmas 3.2, 3.3, 3.4, and 3.7.

THEOREM 3.2. For any downward-closed constraint F , $\text{FB-GFT} \geq O(\log(n) \log(\frac{1}{\epsilon})) \text{GFT}_{\text{CFPP}}$.

3.5.1 Proof of Lemma 3.11 We recall the statement of Lemma 3.11: For any $c \in (0, 1)$,

$$\Pr_t[Z(t) \leq c E[Z(t)]] \leq \frac{(1-c)^2}{1 + \frac{(1-c)^2}{2 E[Z(t)]}}$$

Recall that $Z(t) = \max_{i \in T} t_i^{(k)}$. In the proof we will omit the superscript k as it is fixed. The random seed t is also omitted if clear from context.

LEMMA 3.12. (Paley-Zygmund Inequality [31]) For any random variable $Z \geq 0$ with finite variance, for any $c \in (0, 1)$,

$$\Pr[Z \geq c E[Z]] \geq (1-c)^2 \frac{E[Z]^2}{\text{Var}[Z] + E[Z]^2}$$

To use Lemma 3.12, we only need to show an upper bound on $\text{Var}[Z(t)]$.

LEMMA 3.13. $\text{Var}[Z(t)] \leq E[Z(t)]$.

Proof. By the Efron-Stein Inequality [20],

$$\begin{aligned} \text{Var}[Z(t)] &\leq \frac{1}{2} \sum_i E_{t_i, t_{-i}} [(Z(t_i; t_{-i}) - Z(t_{-i}; t_i))^2] \\ &= \frac{1}{2} \sum_i \text{Var}[Z(t) | t_{-i}] \end{aligned}$$

Here t_{-i} shares the same distribution with t_i (a fresh sample). Note that for every fixed t_{-i} , $\text{Var}_{t_i}[Z(t_i; t_{-i})] \leq E_{t_i} [(Z(t_i; t_{-i}) - a)^2]$ for any constant $a \geq 0$. For every i , let $Z_i(t_{-i}) = \max_{j \in T} t_j$, which only depends on t_{-i} . We have $\text{Var}[Z(t) | t_{-i}] \leq \text{Var}[Z_i(t_{-i})]$ and $E[(Z(t) - Z_i(t_{-i}))^2] \leq E[(Z_i(t_{-i}) - Z_i(t_{-i}))^2]$; where the last inequality follows from the fact that $Z_i(t_{-i}) \leq Z(t) \leq Z_i(t_{-i}) + 1$, as every random variable $t_j \in [0, 1]$.

Now fix any t . Let $T = \arg\max_{T \subseteq [n]} \sum_{i \in T} t_i$. Then for every i , by the definition of Z_i , $Z_i(t_{-i}) = \max_{j \in T} t_j$. Thus $\sum_{i \in T} Z_i(t_{-i}) = \sum_{i \in T} \max_{j \in T} t_j = (n-1) \sum_{j \in T} t_j = (n-1) Z(t)$. Hence,

$$\text{Var}[Z(t)] \leq \frac{1}{2} \sum_i E[(Z(t) - Z_i(t_{-i}))^2] \leq \frac{1}{2} E[Z(t)]$$

□

4 An Unconditional Approximation for a Single Constrained-Additive Buyer

In this section, we prove Theorem 4.1, an unconditional $O(\log n)$ -approximation when the buyer's feasibility constraint is selectable, and an unconditional $O(\log^2(n))$ -approximation for a general constrained-additive buyer—without dependence on distributional parameters. The result combines the $\log(1/r)$ -approximation and a novel mechanism—the seller adjusted posted price mechanism.

THEOREM 4.1. Suppose the buyer's feasibility constraint F is (ϵ, δ) -selectable for some $\epsilon \in (0, 1)$. Then there exists a DSIC, ex-post IR, ex-ante WBB mechanism M such that $\text{SB-GFT}(M) \geq O(\log n) \cdot \text{GFT}(M)$. Moreover, for a general constrained-additive buyer, there exists a DSIC, ex-post IR, ex-ante WBB mechanism M such that $\text{SB-GFT}(M) \geq O(\log^2(n)) \cdot \text{GFT}(M)$.

4.1 An Upper Bound of the Second-Best GFT. Formally, we use SB-GFT to denote the optimal GFT attainable by any BIC, IR, ex-ante WBB mechanism. Notice that the GFT of any two-sided market mechanism can be broken down into the buyer's expected utility of this mechanism, plus the sum of all sellers' expected utilities (or profit), plus the difference between buyer's and sellers' expected payment. We show that the SB-GFT is upper bounded by the sum of the designers' utilities in two related one-sided markets: the super seller auction and the super buyer procurement auction.

Super Seller Auction. Consider a one-sided market, where the designer is the super seller who owns all the items, and replaces all the original sellers. The buyer is the same as in our two-sided market setting. The super seller designs a mechanism to sell the items to the buyer. The main difference between the super seller auction and the original two-sided market is that the mechanism only needs to be BIC and IR for the buyer and does not have any incentive compatibility constraints for the super seller. We use OPT-S to denote the maximum profit (revenue minus her cost) achievable by any BIC and IR mechanism in the super seller auction.

To avoid ambiguity in further proofs, for every subset $T \subseteq [n]$ and downward-closed feasibility constraint J with respect to T , we let $\text{OPT-S}(T; J)$ denote the optimal profit in the following super seller auction: the super seller owns the set of items in T and has cost $s_i \leq D_i$ for every item $i \in T$. The buyer has value $b_i \leq D_i$ for every item $i \in T$ and is additive subject to constraint J . We slightly abuse notation and write $\text{OPT-S}(T; \text{ADD})$ if the buyer is additive ($J = 2^T$) and $\text{OPT-S}(T; \text{UD})$ if the buyer is unit-demand ($J = \{i\} : i \in T$). Clearly, $\text{OPT-S} = \text{OPT-S}([n]; F)$.

Super Buyer Procurement Auction. Similarly, let the super buyer procurement auction be the one-sided market where the super buyer (same as the real buyer) designs the mechanism to procure items from the sellers. Here the mechanism only needs to be BIC and IR for all of the sellers, but not the buyer. We use OPT-B to denote the maximum utility (value minus payment) of the super buyer attainable by any BIC and IR mechanism in the super buyer procurement auction.

First, we extend the upper bound of Brustle et al. [7] to our multi-dimensional setting. We prove that the GFT of any IR, BIC, ex-ante WBB mechanism $M = (x; p^B; p^S)$ is

upper bounded by $\text{OPT-S} + \text{OPT-B}$.

LEMMA 4.1. [7] $\text{SB-GFT} \leq \text{OPT-S} + \text{OPT-B}$:

Proof. Take any BIC, IR, ex-ante WBB mechanism $M = (x; p^B; p^S)$. Since every seller i is BIC and IR, we have for any $s_i; s_i^0$, $E_{b_i; s_i} [p_i^S(b; s) - s_i x_i(b; s)] \geq \max_{s_i} E_{b_i; s_i} [p_i^S(b; s_i^0; s_i)] - s_i x_i(b; s_i^0; s_i); 0$:

Observe that $M^0 = (x; p^S)$ is a valid super buyer procurement auction. The above inequalities are exactly the BIC and IR constraints for seller i . Thus M^0 is BIC and IR. Similarly, $M^{00} = (x; p^B)$ is BIC and IR, so it is a valid super seller auction. Since M is ex-ante WBB, $E_{b_i; s_i} [p_i^B(b; s) - p_i^S(b; s)] \geq 0$. Thus we have $\text{GFT}(M) = E_{b_i; s_i} [p_i^B(b; s) - p_i^S(b; s)] \geq E_{b_i; s_i} [p_i^B(b; s) - p_i^S(b; s)] \geq \text{OPT-S} + \text{OPT-B}$: Taking M to be the GFT-maximizing mechanism completes the proof. \square

Next we prove an analog of the “Marginal Mechanism Lemma” [9, 22] for the optimal profit. Namely, let $(T; R)$ be a partition of the items in $[n]$. Then the optimal profit in a super seller auction with items in $[n]$ is upper bounded by the first-best GFT for items in T plus the optimal profit in a super seller auction with items in R .

LEMMA 4.2. (MARGINAL MECHANISM FOR PROFIT)

For any subset $T \subseteq [n]$, we let $F = f_{S \setminus T} : S \subseteq [n] \rightarrow \mathbb{R}^n$ denote the restriction of F to T . We use $\text{FB-GFT}(T; F_T)$ to denote the first-best GFT obtainable between sellers in T and the F_T -constrained additive buyer, that is,

$$\text{FB-GFT}(T; F_T) = E_{b_T; s_T} [\max_{S \subseteq [n]} \sum_{i \in S} (b_i - s_i)^+];$$

where $b_T = f_{b_i} g_{i \in T}$, $s_T = f_{s_i} g_{i \in T}$. Let $(R; T)$ be any partition of the items in $[n]$. Then

$$\text{OPT-S}([n]; F) \leq \text{OPT-S}(R; F_R) + \text{FB-GFT}(T; F_T);$$

Proof. Consider the optimal BIC and IR mechanism $M = (x; p)$ in the super seller auction with item set $[n]$. We will construct a BIC and IR mechanism $M^0 = (x^0; p^0)$ in the super seller auction with item set R as follows. The mechanism only sells items in R using the same allocation x . The payment for the buyer is defined as the payment p in M minus the buyer’s expected total value for all items in T . Formally, for every $b_R = f_{b_j} g_{j \in R}$, $s_R = f_{s_j} g_{j \in R}$ and $i \in R$, let

$$x_i^0(b_R; s_R) = E_{b_T; s_T} [x_i(b; s)]$$

$$p^0(b_R; s_R) = E_{b_T; s_T} [p(b; s) - \sum_{j \in T} b_j x_j(b; s)] :$$

Notice that in M^0 , the expected utility of the buyer with type b_R when reporting b_R^0 is $E_{s_R} [\sum_{i \in T} b_i x_i(b_R^0; s_R)] - p^0(b_R^0; s_R) = E_{b_T; s_T} [\sum_{i \in [n]} b_i x_i(b_R^0; b_T; s) - p(b_R^0; b_T; s)];$

Since M is BIC and IR, M^0 is also BIC and IR. Thus $\text{OPT-S}([n]; F) = E_{b; s} [p(b; s) - \sum_{i \in [n]} s_i x_i(b; s)] = E_{b_R; s_R} [p^0(b_R; s_R) - \sum_{i \in R} s_i x_i^0(b_R; s_R)] + E_{b; s} [\sum_{i \in T} (b_i - s_i) x_i(b; s)] \leq \text{OPT-S}(R; F_R) + \text{FB-GFT}(T; F_T) : \square$

We partition the items into the set of “likely to trade” items, that is, items with trade probability $r_i = \Pr_{b_i; s_i} [b_i \geq s_i]$, $1 \leq i \leq n$, and the “unlikely to trade” items. We can bound the OPT-S by the first-best GFT of the “likely to trade” items and the optimal profit of the super seller auction with the “unlikely to trade” items. We can further replace the first-best GFT of the “likely to trade” by $O(\log n)$ GFT_{CFPP} according to Theorem 3.1 or by $O(\log \{n\}) \text{GFT}_{\text{CFPP}}$ according to Theorem 3.2 depending on the buyer’s feasibility constraint. Formally,

LEMMA 4.3. Define $H = \{i \in [n] : r_i \geq \frac{1}{n} g\}$ and $L = [n] \setminus H = \{i \in [n] : r_i < \frac{1}{n} g\}$. Suppose the buyer’s feasibility constraint F is $(;)$ -selectable for some $; \geq 2 (0; 1)$. Then

SB-GFT is upper bounded by

$$\text{OPT-B} + \text{OPT-S}(L; F_L) + \text{FB-GFT}(H; F_H) \leq \text{OPT-B} + \text{OPT-S}(L; F_L) + O\left(\frac{\log n}{;}\right) \text{GFT}_{\text{CFPP}};$$

For a general constrained-additive buyer, SB-GFT is upper bounded by

$$\text{OPT-B} + \text{OPT-S}(L; F_L) + O(\log^2(n)) \text{GFT}_{\text{CFPP}};$$

Proof. The first inequality follows from Lemma 4.1 and 4.2. Since F is $(;)$ -selectable, F_H is also $(;)$ -selectable. We derive the second inequality by applying Theorem 3.1 on the items in H . For a general constrained-additive buyer, we derive the inequality by applying Theorem 3.2 on the items in H . \square

It is well known that in multi-item auctions, the revenue of selling the items separately is an $O(\log n)$ -approximation to the optimal revenue when there is a single additive buyer [25]. Cai and Zhao [11] provide an extension of this $O(\log n)$ -approximation to profit maximization. We build on this in Section 4.4 to upper bound the $\text{OPT-S}(L; F_L)$ term, where with jLj items, we get a $\log(jLj)$ factor (Lemma 4.10).

All together, this gives the following upper bound on the second-best GFT.

LEMMA 4.4. (UPPER BOUND ON SECOND-BEST GFT) Define $H = \{i \in [n] : r_i \geq \frac{1}{n} g\}$ and $L = [n] \setminus H = \{i \in [n] : r_i < \frac{1}{n} g\}$. Suppose the buyer’s feasibility constraint F is $(;)$ -selectable for some $; \geq 2 (0; 1)$. Then

$$\text{SB-GFT} \leq \text{OPT-B} + O\left(\frac{\log n}{;}\right) \text{GFT}_{\text{CFPP}} + O(\log(jLj)) \sum_{i \in L} E_{b_i; s_i} (b_i - s_i)^+ ;$$

For a general constrained-additive buyer, the $O \log n$ factor above becomes $O \log^2(n)$.

Next, Section 4.2 gives details on constructing a mechanism for a two-sided market whose GFT is at least OPT-B. In Section 4.3, we show how to use a generalization of posted price mechanisms to approximate the second term in the upper bound by the GFT of the Seller Adjusted Posted Price mechanism. The approximation heavily relies on the fact that in expectation, only one item can trade, so it is crucial that L only contains the “unlikely to trade” items.

4.2 Bounding the Optimal Buyer Utility in the Super Buyer Procurement Auction. In this section, we construct a two-sided market to bound OPT-B for any constrained additive buyer.

LEMMA 4.5. Consider the mechanism $M = (x; p^B; p^S)$ where for every item i , buyer profile b , and seller profile s , $x_i(b; s) = 1[b_i \geq \tilde{v}_i(s_i) - p_i^S(b; s)]$. Here $\tilde{v}_i(s_i)$ is Myerson’s ironed virtual value function¹² for seller i ’s distribution D_i^S . For every seller i , since $\tilde{v}_i(s_i)$ is non-decreasing in s_i , $x_i(b; s)$ is non-increasing in s_i . Define $p_i^S(b; s)$ as the threshold payment for seller i , i.e., the largest cost s_i such that $x_i(b; s_i; s_{-i}) = 1$. Define the buyer’s payment $p^B(b; s) = \sum_i x_i(b; s) \tilde{v}_i(s_i)$. M is DSIC, ex-post IR, ex-ante SBB¹³ and

$$\text{GFT}(M) \geq \text{OPT-B} = E \left[\max_{b; s} \sum_{i \in S} (b_i - e(s_i))^+ \right]$$

Proof. Since the seller’s allocation rule is monotone and we use the threshold payment, M is DSIC and ex-post IR for each seller.

Note that for any seller profile s , when the buyer has true type b , her expected utility from reporting b^0 is $\sum_i x_i(b^0; s) (b_i - e(s_i))$. According to the definition of x , the buyer’s utility is maximized when $b^C = b$. Hence, M is DSIC for the buyer. Moreover we have ex-post IR, as the buyer’s expected utility when reporting truthfully is $\max_{b; s} \sum_{i \in S} (b_i - e(s_i))^+ \geq 0$.

It only remains to prove that the mechanism is ex-ante SBB and to lower bound its GFT. By Myerson’s lemma, for every b we have

$$E_s \sum_i p_i^S(b; s) = E_s \sum_i x_i(b; s) e(s_i) = E_s [p^B(b; s)]$$

Thus the mechanism is ex-ante SBB.

¹²The seller’s unironed virtual value function is $\tilde{v}_i(s_i) = s_i + \frac{G_i(s_i)}{g_i(s_i)}$

¹³One can make the mechanism IR and ex-post SBB by defining $p^B(b; s) = \sum_i p_i^S(b; s)$. The mechanism is still DSIC for all sellers. It is only BIC for the buyer, as the sellers’ gains only equal the virtual welfare when taking expectation over sellers’ profile.

Why is $\text{OPT-B} = E_{b; s} [\max_{i \in S} (b_i - e(s_i))^+]$? Notice that only the sellers are strategic in a super buyer procurement auction, and their types are all single-dimensional. One can apply the standard Myersonian analysis to the super buyer procurement auction and show that the optimal buyer utility is exactly $E_{b; s} [\max_{i \in S} (b_i - e(s_i))^+]$.

Note that the buyer’s expected utility in M is exactly OPT-B. As M is an ex-ante SBB mechanism, the expected GFT of M is equal to the buyer’s expected utility plus the sum of all sellers’ expected utility, and the latter is non-negative since M is ex-post IR for every seller. \square

4.3 The Seller Adjusted Posted Price Mechanism. In this section, we introduce a new mechanism—the Seller Adjusted Posted Price (SAPP) Mechanism. We define an adjusted price mechanism to first elicit each seller’s cost s_i , and then produce posted prices $f_i(s)_{i \in [n]}$ as a function of the reported profile s ; thus the mechanism is a collection of posted prices depending on the reported seller cost profile. The items are offered to the buyer at each posted price $f_i(s)$, with the buyer only allowed to purchase at most one item by paying the posted price. See Mechanism 1 for a complete description of the SAPP mechanism. We show that for a properly selected mapping $f_i(s)_{i \in [n]}$, the SAPP mechanism is DSIC, ex-post IR, and ex-ante WBB. Moreover, its GFT is at least $E_{b; s} [\sum_{i \in L} (b_i - s_i)^+]$.

Since the posted prices depend on the reported seller cost profile, we need to be careful to ensure that there is no incentive for any seller to misreport the cost. We identify a sufficient condition for the posted prices, called bi-monotonicity, to make sure the corresponding mechanism is DSIC and ex-post IR.

DEFINITION 4.1. (BI-MONOTONIC PRICES) We say the posted prices $f_i(s)_{i \in [n]}$ are bi-monotonic, if (i) $f_i(s) \geq s_i$ for all seller profile s and seller i ; (ii) $f_i(s)$ is non-decreasing in s_i and non-increasing in s_j for all $j = i$.

In Lemma 4.6, we prove that bi-monotonic posted prices induce a monotone allocation rule for every seller, enabling threshold payments [29, 30]. Formally, for every seller i let $\mathbb{1}_i(b; s)$ denote the probability that the buyer trades with seller i under profile $(b; s)$. This is either 0 or 1 since all $f_i(s)$ are fixed values when s is fixed. If $\mathbb{1}_i(b; s) = 1$, $p_i^S(b; s)$ is defined as the maximum value s_i^0 such that $\mathbb{1}_i(b; s_i^0; s_{-i}) = 1$. Otherwise $p_i^S(b; s) = 0$. This makes the SAPP mechanism DSIC and ex-post IR.

LEMMA 4.6. Let M be an SAPP mechanism with bi-monotonic posted prices $f_i(s)_{i \in [n]}$. Then the allocation of the mechanism $\mathbb{1}_i(b; s)$ is non-increasing in s_i for all sellers i , and M is DSIC and ex-post IR for the buyer and the sellers.

Proof. Notice that for every type b , the buyer chooses the

item that maximizes $b_i - p_i(s)$ (and does not choose any item if she cannot afford any of the items). For every i , by bi-monotonicity, when s_i decreases, $b_i - p_i(s)$ does not decrease while $b_j - p_j(s)$ does not increase for all $j \neq i$. Thus if the buyer chooses item i under the original s_i , she must continue to choose item i for smaller reports s_i^c . Thus $\mathfrak{X}_i(b; s)$ is non-increasing in s_i . Since every seller receives the threshold payment, she is DSIC and ex-post IR. As the buyer simply faces a posted price mechanism, the mechanism is DSIC and ex-post IR for the buyer. \square

Mechanism 1 Seller Adjusted Posted Price Mechanism

Require: $\delta_i \in [0, 1]$, function $p_i(\cdot)$ that maps each seller cost profile to a price for item i . Input $(b; s)$.

- 1: Given the sellers' reported cost profile s , offer each item i to the buyer at price $p_i(s)$.
 - 2: The buyer is allowed to purchase at most one item by paying the corresponding posted price.
 - 3: If no item is picked, then no trade happens and payment is 0 for every agent. Otherwise, if the buyer chooses item i , she receives item i and pays $p_i(s)$. Seller i sells her item and receives threshold payment.
-

The main challenge we face here is establishing the budget balance condition. Unfortunately, having bi-monotonic posted prices is not sufficient. Consider the $n = 1$ case: the posted price $p(s) = s$ is trivially bi-monotonic. Clearly, the corresponding SAPP mechanism achieves FB-GFT. However, due to the impossibility result by Myerson and Satterthwaite [30], no BIC, IR, and ex-ante WBB mechanism can always achieve FB-GFT, so the SAPP mechanism must sometimes violate the budget balance constraint. In Lemma 4.7, we show that even though bi-monotonic posted prices do not imply budget balance, there is indeed a wide range of bi-monotonic posted prices that induce budget balanced SAPP mechanisms. Our budget balance proof heavily relies on an allocation coupling argument (Lemma 4.8) that simultaneously provides a lower bound on the buyer's payment, as well as an upper bound on the payment to the seller.

LEMMA 4.7. Let $x = f x_i(b; s) g_{i \in [n]}$ be an arbitrary allocation rule that satisfies (i) the buyer never purchases more than one item in expectation under each profile $(b; s)$, i.e., $\sum_{i \in [n]} x_i(b; s) \leq 1$, and (ii) for every buyer type b and seller i , $x_i(b; s)$ is non-increasing in s_i , and non-decreasing in s_j for all $j \neq i$. We define $q_i(s) = E_b[x_i(b; s) \mathbb{1}[\tilde{v}_i(b_i) \geq s_i]]$, where $\tilde{v}_i(b_i)$ is Myerson's ironed virtual value for D_i^B , and $p_i(s) = F_i^{-1}(1 - \frac{q_i(s)}{2})$. The posted prices

$f_i(s) g_{i \in [n]}$ are bi-monotonic, and the corresponding SAPP mechanism M is DSIC, ex-post IR, and ex-ante WBB. Moreover, $GFT(M) \geq \frac{1}{2} E_{b; s} [\sum_{i \in [n]} (\tilde{v}_i(b_i) - s_i) x_i(b; s)]$.

Proof. It is not hard to verify that $f_i(s) = F_i^{-1}(1 - \frac{q_i(s)}{2}) g_{i \in [n]}$ is bi-monotonic. Now we proceed to prove that the SAPP mechanism M is ex-ante WBB. We require the following lemma.

LEMMA 4.8. For every seller i and every seller profile s , $\mathfrak{X}_i(s) \geq \frac{q_i(s) + q_i(s)^2}{4}; \frac{q_i(s)}{2}$.

Proof. Note that the buyer will purchase item i if both of the following conditions are satisfied:

1. The buyer can afford item i , i.e., $b_i - p_i(s)$.
2. The buyer cannot afford any other items, i.e., $b_j < p_j(s); \forall j \neq i$.

By choice of $p_i(s)$, the first event happens with probability $\Pr[b_i - p_i(s) \geq 0] = q_i(s)$.

Note that $\sum_{i \in [n]} q_i(s) = E_b[\sum_{i \in [n]} x_i(b; s)] \leq 1$. For each $j \neq i$, $\Pr[b_j < p_j(s)] = 1 - \frac{q_j(s)}{2}$. Thus $\Pr_{j \neq i} [b_j < p_j(s)] \geq \prod_{j \neq i} (1 - \frac{q_j(s)}{2}) \geq \frac{1}{2} (1 - \sum_{j \neq i} \frac{q_j(s)}{2})$. The second event happens with probability $\prod_{j \neq i} (1 - \frac{q_j(s)}{2}) \geq \frac{1}{2} (1 - \sum_{j \neq i} \frac{q_j(s)}{2})$. The equality holds when one out of the $n - 1$ $q_j(s)$'s equals 1 and the rest are all equal to 0. Notice that the two events are independent, so we have the upper and lower bound on $\mathfrak{X}_i(s)$. \square

We return to the proof of Lemma 4.7. For easy reference, we list our notation again:

$x = f x_i(b; s) g_{i \in [n]}$ is an arbitrary allocation.

$\mathfrak{X}_i(b; s)$ is the probability that item i trades in M under profile $(b; s)$.

$\mathfrak{X}_i(s) = E_b[\mathfrak{X}_i(b; s)]$ is the probability that item i trades over the randomness of buyer valuations, i.e. the interim trade probability.

$q_i(s) = E_b[x_i(b; s) \mathbb{1}[\tilde{v}_i(b_i) \geq s_i]]$ is the probability that item i trades in allocation x and the buyer's ironed virtual value for item i is above the seller's cost.

$p_i(s) = F_i^{-1}(1 - \frac{q_i(s)}{2})$ is the buyer's posted price set such that $\Pr[b_i - p_i(s) \geq 0] = q_i(s)$.

Fix any seller profile s . For simplicity, we slightly abuse notation and use $\mathfrak{X}_i(z)$ and $q_i(z)$ to denote $\mathfrak{X}_i(z; s_{-i})$ and $q_i(z; s_{-i})$. The expected payment from the buyer under cost profile s is $\sum_{i \in [n]} \mathfrak{X}_i(s_i) p_i(s)$. For every seller i , denote $p_i^S(s) = E_b[p_i^S(b; s)]$ as her expected payment under cost profile s .

Note that for every $b; s$, the threshold payment $p_i^S(b; s)$ can be rewritten as the quantity $\int_{s_i}^{\tilde{v}_i(b_i)} \mathfrak{X}_i(b; t; s_{-i}) dt + s_i \mathfrak{X}_i(b; s_i; s_{-i})$: When $\mathfrak{X}_i(b; s) = 0$, then $\mathfrak{X}_i(b; t; s_{-i})$ for

all $t \leq s_i$ since $\chi_i(b; s)$ is non-increasing in s_i . Thus the above quantity is 0. When $\chi_i(b; s) = 1$, let s_i^0 be the maximum value such that $\chi_i(b; s_i^0; s_i) = 1$. Then the above quantity is equal to $\int_{s_i^0}^{s_i} \int_{s_i}^1 dt + s_i = s_i^0 = p_i^s(b; s)$. Thus $p_i^s(s) = E_b[p_i^s(b; s)] = \int_{s_i}^1 \int_{s_i}^1 \chi_i(z; s_i) dz + s_i \chi_i(s_i; s_i)$: We will show that $p_i^s(s) \leq \chi_i(s_i) + i(s)$. By definition,

$$\begin{aligned} p_i^s(s) &= \int_{s_i}^1 \chi_i(z) dz + s_i \chi_i(s_i) \\ &= \int_{s_i}^{s_i} \int_{s_i}^1 1[\chi_i(z) \leq t] dt dz + s_i \chi_i(s_i) \\ &= \int_{s_i}^0 \int_{s_i}^1 1[\chi_i(z) \leq t] dt dz + s_i \chi_i(s_i) \\ &= \int_{s_i}^0 \int_{s_i}^1 1[\chi_i(z) \leq t] dz dt + s_i \chi_i(s_i) \end{aligned}$$

The second inequality follows from $\chi_i(z) = 1[\chi_i(z) \leq t]$; $8z$. The third inequality is due to $1[\chi_i(z) \leq t] = 0$; $8z \leq s_i$ and $t > \chi_i(s_i)$. The last equality follows from Fubini's Theorem, as the integral is finite due to the monotonicity of $\chi_i(\cdot)$.

Moreover, since $\chi_i(\cdot)$ is non-increasing, for every $z \leq s_i$; $t \leq \chi_i(s_i)$, we have $\chi_i(z) \geq \chi_i(s_i)$. Thus $\int_{s_i}^0 \int_{s_i}^1 1[\chi_i(z) \leq t] dz dt = \int_{s_i}^0 \int_{s_i}^1 1 dz dt = s_i \chi_i(s_i)$:

Combining the two equations, we have

$$\begin{aligned} p_i^s(s) &= \int_{s_i}^0 \int_{s_i}^1 1[\chi_i(z) \leq t] dz dt \\ &\quad + \int_{s_i}^0 \int_{s_i}^1 1[q_i(z) \leq 2t] dz dt \\ &\quad + \int_{s_i}^0 \int_{s_i}^1 1[\Pr[\hat{e}_i(b_i) \leq z] \leq 2t] dz dt \end{aligned}$$

The first inequality follows from Lemma 4.8 and the second inequality follows from the definition of $q_i(\cdot)$. For every t , we prove that $\int_{s_i}^0 \int_{s_i}^1 1[\Pr_b[\hat{e}_i(b_i) \leq z] \leq 2t] dz dt \leq \int_{s_i}^0 \int_{s_i}^1 1[\Pr_b[\hat{e}_i(b_i) \leq z] \leq 2t] dz dt$ for any $\epsilon > 0$. In fact, let $z = \hat{e}_i(F^{-1}(1 - 2t + \epsilon))$. For every $z > z$, $\Pr[\hat{e}_i(b_i) \leq z] \leq \Pr[\hat{e}_i(b_i) > z] = \Pr[b_i > F^{-1}(1 - 2t + \epsilon)] = 2t - \epsilon$. So $1[\Pr[\hat{e}_i(b_i) \leq z] \leq 2t] = 0$ for all $z > z$.

Therefore, for any $\epsilon > 0$, we have the following. In the second line, we change the variable by denoting $y =$

$$\begin{aligned} &F_i^{-1}(1 - 2t + \epsilon) \\ &= \int_{F_i^{-1}(1 - 2\chi_i(s_i) + \epsilon)}^s \hat{e}_i(F_i^{-1}(1 - 2t + \epsilon)) dt \\ &= \int_{F_i^{-1}(1 - 2\chi_i(s_i) + \epsilon)}^s \hat{e}_i(y) dy \\ &= \int_{F_i^{-1}(1 - 2\chi_i(s_i) + \epsilon)}^s \hat{e}_i(y) f_i(y) dy \\ &= \frac{1}{2} F_i^{-1}(1 - 2\chi_i(s_i) + \epsilon) [1 - F_i(F_i^{-1}(1 - 2\chi_i(s_i) + \epsilon))] \\ &= F_i^{-1}(1 - 2\chi_i(s_i) + \epsilon) (\chi_i(s_i) - \epsilon) \\ &= \chi_i(s_i) - F_i^{-1}(1 - 2\chi_i(s_i) + \epsilon) \end{aligned}$$

If $q_i(s_i) = 0$, then $\chi_i(s_i) - F_i^{-1}(1 - 2\chi_i(s_i) + \epsilon) = 0 = \chi_i(s_i) - i(s)$. Otherwise, choose ϵ to be any number in $(0; \frac{q_i(s_i)^2}{4})$. Then, according to Lemma 4.8 and our choice of ϵ ,

$$1 - 2\chi_i(s_i) + \epsilon \leq \frac{q_i(s_i)}{2} - \frac{q_i(s_i)^2}{4} \leq 1 - \frac{q_i(s_i)}{2}$$

Hence, $F_i^{-1}(1 - 2\chi_i(s_i) + \epsilon) < i(s)$. Thus $p_i(s) \leq \chi_i(s) + i(s)$ for every i and s , which implies that $E_s[p_i(s) \chi_i(s)] \leq E_s[p_i^s(s; s_i)]$. Hence M is ex-ante WBB.

We now need to lower bound the GFT from mechanism M .

$$\begin{aligned} \text{GFT}(M) &= E_{b;s} \sum_i (b_i - s_i) \chi_i(b; s) \\ &= E_s \sum_i (i(s) - s_i) \chi_i(s) \\ &= \frac{1}{2} E_{b;s} \sum_i (\hat{e}_i(b_i) - s_i) [1 - b_i F_i^{-1}(1 - q_i(s))] \\ &= \frac{1}{4} E_{b;s} \sum_i (\hat{e}_i(b_i) - s_i) x_i(b; s) [1 - \chi_i(b_i) - s_i] \\ &= \frac{1}{4} E_{b;s} \sum_i (\hat{e}_i(b_i) - s_i) x_i(b; s) \end{aligned}$$

Here the second inequality uses the definition of $i(s)$, $q_i(s)$ and Lemma 4.8. The third inequality follows from

Myerson's lemma. The second-to-last inequality uses the fact that $E_{b_i}[\hat{e}_i(b_i) - 1] \leq \frac{1}{2} \frac{q_i(s_i)}{f_i(b_i)}$. $E_{b_i}[\hat{e}_i(b_i) - x_i(b; s)] \geq 1 - \frac{1}{2} \frac{q_i(s_i)}{f_i(b_i)}$ holds for every s and i . This is because the right hand side $\frac{1}{2} E_{b_i}[\hat{e}_i(b_i) - x_i(b; s)] \geq 1 - \frac{1}{2} \frac{q_i(s_i)}{f_i(b_i)}$ can be viewed as the expectation of $\hat{e}_i(b_i)$ on an event of b_i with a total probability mass $E_{b_i}[\frac{1}{2} E_{b_i}[\hat{e}_i(b_i) - x_i(b; s) - 1]] = \frac{q_i(s_i)}{2}$; while the left hand side is the maximum expectation of $\hat{e}_i(b_i)$ on any event of b_i with total probability mass $\frac{q_i(s_i)}{2}$, as $\hat{e}_i(b_i)$ is non-decreasing on b_i . \square

Lemma 4.9 shows how to choose an allocation rule x so that the induced SAPP mechanism (using Lemma 4.7) has GFT at least

$$\frac{1}{4e} \sum_{i \in L} E_{b_i; s_i} [(\hat{e}_i(b_i) - s_i)^+].$$

Note that the existence of such an x heavily relies on the fact that in expectation there is only one item that can trade among the "unlikely to trade" items.

LEMMA 4.9. We let $GFT_{SAPP}(S)$ denote the optimal GFT attainable by any DSIC, ex-post IR, and ex-ante WBB SAPP mechanism over items in S for any subset $S \subseteq [n]$. $GFT_{SAPP}(L) \geq \frac{1}{4e} \sum_{i \in L} E_{b_i; s_i} [(\hat{e}_i(b_i) - s_i)^+]$.

Proof. Let $b_L = (b_i)_{i \in L}$ and $s_L = (s_i)_{i \in L}$. For every $i \in L$, define the event that only i is tradeable:

$$A_i = \{b_L; s_L : b_i \leq s_i \wedge b_j > s_j; \forall j \in L, j \neq i\}$$

We consider the following allocation rule:

$$x_i(b_L; s_L) = \begin{cases} 1[\hat{e}_i(b_i) \geq s_i] & ; \text{ if } (b; s) \in A_i \\ 0 & ; \text{ otherwise} \end{cases}$$

Notice that $(b_L; s_L) \in A_i$ implies that $(b_L; s_i^0; (s_j)_{j \in L, j \neq i}) \in A_i$ for any $s_i^0 \leq s_i$. Thus, $x_i(b_L; s_L)$ is non-increasing in s_i . Similarly, it is easy to verify that $x_i(b_L; s_L)$ is non-decreasing in all s_j where $j \in L, j \neq i$. Furthermore, $\sum_{i \in L} x_i(b_L; s_L) \leq 1$ for all $b_L; s_L$. If we choose the posted prices according to Lemma 4.7, then the corresponding mechanism has GFT that is at least $\frac{1}{4} E_{b; s} [\sum_{i \in L} (\hat{e}_i(b_i) - s_i) x_i(b; s)]$.

Moreover, by the definition of $x_i(b; s)$,

$$\begin{aligned} & \sum_{b; s} E_{b; s} \sum_{i \in L} (\hat{e}_i(b_i) - s_i) x_i(b; s) \\ &= \sum_{b; s} E_{b; s} [\sum_{i \in L} (\hat{e}_i(b_i) - s_i)^+] \Pr[b_j < s_j; \forall j \in L, j \neq i] \\ & \leq \sum_{b; s} E_{b; s} [\sum_{i \in L} (\hat{e}_i(b_i) - s_i)^+] (1 - \frac{1}{n})^{j \in L, j \neq i} \\ & \leq \frac{1}{e} \sum_{b; s} E_{b; s} [\sum_{i \in L} (\hat{e}_i(b_i) - s_i)^+] \end{aligned}$$

The first inequality holds because for each item $j \in L$, $\Pr_{b_j; s_j}[b_j < s_j] \geq 1 - \frac{1}{n}$. Hence,

$$GFT_{SAPP}(L) \geq \frac{1}{4e} \sum_{i \in L} E_{b_i; s_i} [(\hat{e}_i(b_i) - s_i)^+]$$

\square

4.4 Bounding the Optimal Profit from the Unlikely to Trade Items. In this section, we provide an upper bound for the optimal super seller profit from items in L . It is well known that in multi-item auctions the revenue of selling the items separately is a $O(\log n)$ -approximation to the optimal revenue when there is a single additive buyer [25]. Cai and Zhao [11] provide an extension of this $O(\log n)$ -approximation to profit maximization. Combining this approximation with some basic observations based on the Cai-Devanur-Weinberg duality framework [8], we derive the following upper bound of $OPT-S(L; F_L)$.

LEMMA 4.10. $OPT-S(L; F^L) \leq O(\log(jL)) \sum_{b; s} E_{b; s_i} [(\hat{e}_i(b_i) - s_i)^+]$: Here $\hat{e}_i(b_i)$ is Myerson's ironed virtual value function¹⁴ for the buyer's distribution for item i , D_i^B .

To bound $OPT-S(L; F^L)$, we need the following result from [11]. It provides a benchmark of the optimal profit using the Cai-Devanur-Weinberg duality framework [8]: The profit of any BIC, IR mechanism is upper bounded by the buyer's virtual welfare with respect to some virtual value function, minus the sellers' total cost for the same allocation.

A sketch of the framework is as follows: We first formulate the profit maximization problem as an LP. Then we Lagrangify the BIC and IR constraints to get a partial Lagrangian dual of the LP. Since the buyer's payment is unconstrained in the partial Lagrangian, one can argue that the corresponding dual variables must form a flow in order for the benchmark to be finite. By weak duality, any choice of the dual variables/flow derives a benchmark for the optimal profit. In [11], they also construct a canonical flow and prove that there exists a BIC and IR mechanism whose profit is within a constant factor times the benchmark w.r.t. the flow for any single constrained-additive buyer.

LEMMA 4.11. [11] For any $T \subseteq [n]$ and feasibility constraint J with respect to T , consider the super seller auction with item set T and a J -constrained buyer. Any flow τ induces a finite benchmark for the optimal profit, that is,

$$OPT-S(T; J) \leq \max_{x \geq 0} \sum_{b; s} E_{b; s} \sum_{i \in T} x_i(b; s) (\tau_i(b) - s_i)$$

¹⁴The buyer's ironed virtual value function is $\hat{e}_i(b_i) = b_i - \frac{1 - F(b_i)}{f_i(b_i)}$. These values are averaged to make the function monotonic in quantile space, which creates $\hat{e}_i(b_i)$.

where $\tau(b) = b_i - \frac{1}{f_i(b_i)} P_{j_i}(b^0; b)(b^0 - b_i)$ can be viewed as buyer i 's virtual value function, and P_j is the set of all feasible allocation rules. More specifically, $\tau(b; b)^0$ is the Lagrangian multiplier for the BIC/IR constraint that states that when the buyer has true type b , she does not want to misreport b^0 . The equality sign is achieved when the optimal dual is chosen.

Next, we show that $\text{OPT-S}(L; F_L)$ is no more than $\text{OPT-S}(L; \text{ADD})$ using Lemma 4.11.

LEMMA 4.12. $\text{OPT-S}(L; F_L) \leq \text{OPT-S}(L; \text{ADD})$:

Proof. Let $\hat{\lambda}$ be the optimal dual in Lemma 4.11 when the buyer is additive without any feasibility constraint, and $\hat{v}_i(b)$ be the induced virtual value function. We have that

$$\begin{aligned} \text{OPT-S}(L; F_L) &= \max_{x \in \mathcal{P}_{F_L}} \sum_{i \in [n]} \mathbb{E}_{b_i, s_i} x_i(b; s) (\hat{v}_i(b) - s_i) \\ &= \max_{x_i(b; s) \in [0, 1]} \sum_{i \in [n]} \mathbb{E}_{b_i, s_i} x_i(b; s) (\hat{v}_i(b) - s_i) \\ &= \text{OPT-S}(L; \text{ADD}); \end{aligned}$$

□

Cai and Zhao [11] also give a logarithmic upper bound of the optimal profit for a single additive buyer, using the sum of optimal profit for each individual item.

LEMMA 4.13. [11]

$$\begin{aligned} \text{OPT-S}(L; \text{ADD}) &\leq \log(jLj) \sum_{i \in [n]} \text{OPT-S}(f_i) \\ &= \log(jLj) \sum_{i \in [n]} \mathbb{E}_{b_i, s_i} [(\hat{v}_i(b_i) - s_i)^+]; \end{aligned}$$

Together, Lemmas 4.12 and 4.13 conclude the proof of Lemma 4.10:

$$\text{OPT-S}(L; F_L) \leq \log(jLj) \sum_{i \in [n]} \mathbb{E}_{b_i, s_i} (\hat{v}_i(b_i) - s_i)^+$$

Proof of Lemma 4.4: It directly follows from Lemmas 4.3 and 4.10. 2

Proof of Theorem 4.1: The theorem follows directly from Lemmas 4.4, 4.5, 4.7, and 4.9. 2

5 Lower Bounds and the First-Best–Second-Best Gap

In the unconditional approximation results stated in Section 4, we compare the GFT of our mechanism to SB-GFT. Readers may be interested in whether our mechanism is also an approximation to FB-GFT. In fact, this question is related to one of the major open problems in two-sided markets: How large is the gap between the second-best and the

first-best GFT? In this section, we consider a unit-demand buyer and present a reduction from achieving a FB-GFT of approximation in our multi-dimensional setting to the open problem regarding the gap in single-dimensional two-sided markets.

Matching Markets. This setting has a two-sided market with n buyers, n sellers, and n identical items. Each seller owns one item and each buyer is interested in buying at most one item. Thus the value (or cost) for every agent is a scalar. Here we consider a special case where for every $i \in [n]$, buyer i and seller i can only trade with each other, and at most one pair of agents in the market can trade. This is bilateral trade when $n = 1$.

THEOREM 5.1. Suppose the buyer is unit-demand in the multi-dimensional setting, and define FB-GFT; OPT-B, GFT_{SAPP} as in the previous section. Also consider the following matching market with n buyers and n sellers: for every $i \in [n]$, buyer i has value drawn from D_i^B and seller i has cost drawn from D_i^S . Let $\text{FB-GFT}^{\text{SD}} = \mathbb{E}_{b_i, s_i} [\max_i (b_i - s_i)]$ be the first-best GFT of the matching market defined above (which is the same as FB-GFT in the multi-dimensional unit-demand setting) and $\text{SB-GFT}^{\text{SD}}$ be the second-best GFT. For any $c > 1$, suppose $\text{SB-GFT}^{\text{SD}} \geq \frac{1}{c} \text{FB-GFT}^{\text{SD}}$, then

$$\max\{\text{OPT-B}; \text{GFT}_{\text{SAPP}}\} \geq \frac{1}{2c} \text{FB-GFT};$$

Proof. We construct the following allocation rule $x = \{x_i(b; s)\}_{i \in [n]}$. For every i and $b; s$, let

$$x_i(b; s) = 1 [i = \arg \max_k (\hat{v}_k(b_k) - s_k) \wedge \hat{v}_i(b_i) - s_i];$$

Then x satisfies both properties in Lemma 4.7. By Lemma 4.7, $\text{GFT}_{\text{SAPP}} = \mathbb{E}_{b_i, s_i} [\sum_i (\hat{v}_i(b_i) - s_i) x_i(b; s)] = \mathbb{E}_{b_i, s_i} [\max_i (\hat{v}_i(b_i) - s_i)^+]$:

Moreover, using the upper-bound on $\text{SB-GFT}^{\text{SD}}$ given by Brustle et al. [7], we have that $\text{SB-GFT}^{\text{SD}} \leq \mathbb{E}_{b_i, s_i} [\max_i (\hat{v}_i(b_i) - s_i)^+] + \mathbb{E}_{b_i, s_i} [\max_i (b_i - \tilde{v}_i(s_i))^+]$. Thus by Lemma 4.5, we have

$$\begin{aligned} \max\{\text{OPT-B}; \text{GFT}_{\text{SAPP}}\} &\geq \frac{1}{2} \mathbb{E}_{b_i, s_i} \max_i (\hat{v}_i(b_i) - s_i)^+ + \mathbb{E}_{b_i, s_i} \max_i (b_i - \tilde{v}_i(s_i))^+ \\ &\geq \frac{1}{2} \text{SB-GFT}^{\text{SD}} + \frac{1}{2c} \text{FB-GFT}^{\text{SD}}; \end{aligned}$$

□

The main takeaway of Theorem 5.1 is that, if the largest gap between $\text{FB-GFT}^{\text{SD}}$ and $\text{SB-GFT}^{\text{SD}}$ is at most (i.e. a constant) c for matching markets, then our mechanism is a $2c$ -approximation to FB-GFT. Note that if the buyer is additive, such a reduction clearly exists: In the additive case,

items can be treated separately without impacting the IC constraint. Then performing a Buyer (or Seller) Offering mechanism¹⁵ for every item separately obtains GFT at least SB-GFT^{SD} [7], thus approximating FB-GFT by the assumption. Theorem 5.1 shows that for a unit-demand buyer, a similar reduction also exists using the SAPP mechanism.

On the other hand, finding a lower bound for our result (compared to SB-GFT) is at least as hard as finding a lower bound for the approximation ratio w.r.t. FB-GFT, and thus is at least as hard as finding an instance in the matching market that separates FB-GFT^{SD} from SB-GFT^{SD} — a problem that has long remained open. Indeed, even in bilateral trade, deciding whether the gap is finite or not is still open.

Acknowledgements

The authors would like to thank Anna Karlin for helpful discussions in the early stages of the paper.

A Examples

Tight Example of the $\log(\frac{1}{r})$ -Approximation. Consider the case when $n = 1$ (bilateral trade). We introduce an example provided by Blumrosen and Dobzinski [5]. They prove that in this example, no fixed posted price mechanisms can achieve an approximation ratio better than $(\log(1=r))$ compared to the first-best GFT. In addition, we will verify that the statement also holds for the second-best GFT for the same example. It implies that our $\log(\frac{1}{r})$ -approximation is tight even compared to the second-best GFT.

EXAMPLE 1. (EXAMPLE IN BILATERAL TRADING [5]) For any $t > 0$, consider a buyer and a seller with values on the support $[0; t]$. Let $b \sim \frac{1}{e-t} e^{-\frac{b}{e-t}}$. Let $F(b) = (1 - e^{-\frac{b}{e-t}})$ with $f(b) = e^{-\frac{b}{e-t}}$ and $G(s) = (e^{\frac{s}{e-t}} - e^{-\frac{s}{e-t}})$ with $g(s) = e^{\frac{s}{e-t}}$. Then $r = \frac{1}{e-t}$. $\Pr[b \leq s] = \int_0^s \frac{1}{e-t} e^{-\frac{b}{e-t}} e^{\frac{s}{e-t}} ds = \frac{1}{e-t} (e^{\frac{s}{e-t}} - e^{-\frac{s}{e-t}})$. Thus $\text{FB-GFT} = \int_0^t \int_0^b (b-s) e^{-\frac{b}{e-t}} e^{\frac{s}{e-t}} ds db = \frac{2}{e-t} (t - \frac{t+2}{e-t})$. In any fixed posted price mechanism, note that the mechanism always achieves a larger GFT by choosing the same price for both agents. The gains from trade from posting at price p is $\text{GFT}(p) = \int_0^p \int_0^p (b-s) e^{-\frac{b}{e-t}} e^{\frac{s}{e-t}} ds db = \frac{2}{e-t} (p - \frac{p+2}{e-t})$. When t is sufficiently large, FB-GFT is about $2 \frac{t-2}{e-t}$ while $\text{GFT}(p)$ is at most $2 \frac{p-2}{e-t}$, as $\frac{p+2}{e-t}$ is negligible. Thus $\text{GFT}(p) = O(1=t)\text{FB-GFT}$. On the other hand, $r = (\frac{1}{e-t})$ for large t , $\log(\frac{1}{r}) = t$. Thus $\text{GFT}(p) = O(1=\log(\frac{1}{r})) \cdot \text{FB-GFT}$.

We now verify that $\text{GFT}(p) = O(1=\log(\frac{1}{r})) \text{SB-GFT}$ for any $p \in [0; t]$ and sufficiently large t . By [7],

$$\text{SB-GFT} = \mathbb{E}_{b,s} [(b-s) \mathbb{1}[\tilde{r}(b) \leq s]]$$

For the above distribution, $\tilde{r}(b) = b \frac{1 - F(b)}{f(b)} = b \frac{1 - e^{-\frac{b}{e-t}}}{e^{-\frac{b}{e-t}}}$ is monotonic increasing in b . Thus $\tilde{r}(b) = \tilde{r}(b)$.

$$\begin{aligned} & \frac{1}{2} \mathbb{E}_{b,s} [(b-s) \mathbb{1}[\tilde{r}(b) \leq s]] \\ &= \frac{1}{2} \int_0^t \int_0^b (b-s) e^{-\frac{b}{e-t}} ds db \\ &= \frac{1}{2} \int_0^t \int_0^b (b-s) e^{-\frac{b}{e-t}} ds db \\ &= e^{-t} \int_0^t \int_0^b (b-s) e^{-\frac{b}{e-t}} ds db \quad (k = b-s) \\ &= e^{-t} \int_0^t \int_0^1 (e^{-k} (k+1) \frac{b}{e-t}) ds \\ &= e^{-t} \int_0^t \frac{2}{e-t} e^{-\frac{b}{e-t}} e^{-b} (b+1) ds \\ &= e^{-t} \left(\frac{2t}{e-t} + \frac{t+2}{e-t} \right) \end{aligned}$$

Thus when t is sufficiently large, $\text{SB-GFT} = \frac{t}{e-t}$ and we have $\text{GFT}(p) = O(1=t)\text{SB-GFT} = O(1=\log(\frac{1}{r})) \text{SB-GFT}$.

EXAMPLE 2. (GFT_{SAPP} VS. GFT_{FPP}) For any fixed n , consider the following instance for an additive buyer. D_1^B and D_1^S are distributions in Example 1 for some sufficiently large t . Pick any $C > 0$. For every $i = 2; \dots; n$, D_i^B is a degenerate distribution at C , i.e. the value is C with probability 1. Distribution D_i^S takes value $C + \frac{1}{2n}$ with probability $\frac{1}{2n}$ and C with probability $\frac{1}{2n}$, for some small $\frac{1}{2n} > 0$. As shown in Example 2, when t is large, $r_1 = (\frac{t-2}{e-t}) < \frac{1}{n}$. For $i \geq 2$, $r_i = \frac{1}{2n}$. Thus all items are “unlikely to trade” items ($r_i < \frac{1}{n}$).

Note that for $i \geq 2$, b_i is always no more than s_i . By Lemma 4.9, $\text{GFT}_{\text{SAPP}} = \mathbb{E}_{b_1; s_1} [(r_1(b_1) - s_1)^+] = \mathbb{E}_{b_1; s_1} [(r_1(b_1) - s_1)^+] = \mathbb{E}_{b_1; s_1} [(r_1(b_1) - s_1)^+] = \frac{2}{e-t} (t - \frac{t+2}{e-t})$. In Example 1, when t is sufficiently large, $\text{GFT}_{\text{FPP}} = O(1=t)$. Thus $\text{GFT}_{\text{FPP}} = O(1=t)$, which can be arbitrarily far from GFT_{SAPP} as t goes to infinity.

Dependence on r is Necessary. We show that the dependence on $r = \min_i r_i$ is necessary for the approximation result of fixed posted price mechanisms. More formally, suppose fixed posted price mechanisms achieve an approximation ratio of $f(r_1; \dots; r_n)$, for some n -ary function f .

¹⁵In bilateral trade, a Buyer Offering mechanism lets the buyer choose a take-it-or-leave-it price for the seller according to her value. And in the Seller Offering mechanism, the seller is asked to pick the price for the buyer.

We will show that $f(r_1; \dots; r_n) = (\log(1=r))$. Consider the instance shown in Example 2. Clearly $FB-GFT = E[(b_1 - s_1)^+]$. Since all items other than item 1 always contribute 0 gains from trade, no fixed posted price mechanism can achieve better than $(\log(1=r_1))$ -approximation to the first-best. Thus $f(r_1; \dots; r_n) = (\log(1=r_1))$. Similarly we have $f(r_1; \dots; r_n) = (\log(1=r_i))$ for all $i = 1; \dots; n$. Thus $f(r_1; \dots; r_n) = (\log(1=r))$.

SAPP Mechanism is Necessary. We provide the following example (Example 3) to show that the class of SAPP mechanisms defined in Mechanism 4.3 is necessary to obtain any finite approximation ratio to SB-GFT. More specifically, we show that in bilateral trading, the best FPP mechanism and the mechanism used in Lemma 4.5 can both be arbitrarily far from SB-GFT.

EXAMPLE 3. For every positive integer $m \geq 2$, consider the bilateral trading instance where the seller's and buyer's (discrete) distributions are shown in the following tables. In the table, $g(s)$ (or $f(b)$) represents the density of the corresponding value in the support.

s	0	$2^m - 1$...	$2^m - 2^k$...	$2^m - 1$
$g(s)$	$\frac{1}{2^m}$	$\frac{1}{2^m}$...	$\frac{1}{2^{k+1}}$...	$\frac{1}{2}(s)$
	0	2^m	...	2^m	...	2^m

Table 1: Seller's Distribution

b	$2^m - 2^L$...	$2^m - 2^k$...	$2^m - 1$
$f(b)$	p_L	...	p_k	...	p_0

Table 2: Buyer's Distribution

For the seller's distribution, one can verify that the virtual value (s) is 0 if $s = 0$ and 2^m elsewhere.¹⁶

For the buyer's distribution, choose $L = \lceil \log(m)c \rceil$. Then define the sequence $\{p_k\}_{k=0}^L$ as follows: Construct the sequence $\{q_k\}_{k=0}^L$ with $q_0 = 1$, $q_1 = \frac{1}{m-1}$, and for every $k = 2; \dots; L$, $q_k = \frac{m-k+2}{m-k} q_{k-1}$. Then for every $k = 0; \dots; L$ define $p_k = q_k \prod_{j=0}^{L-k} q_j$ for every k .

By induction, we have $\prod_{j=0}^k q_j = q_{k+1} (m - k - 1)$. Thus

$$(A.1) \quad p_k = p_{k+1} (m - k - 1);$$

LEMMA A.1. For any sufficiently large integer m , let M be the mechanism used in Lemma 4.5.¹⁷ Then in Example 3 we have

$$\max_{GFT_{FPP}; GFT(M)} g \leq O\left(\frac{1}{\log(m)}\right) SB-GFT;$$

The proof of Lemma A.1 can be found in the full version of the paper.

References

- [1] Moshe Babaioff, Yang Cai, Yannai A. Gonczarowski, and Mingfei Zhao. 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, Ithaca, NY, USA, June 18-22, 2018, page 373, 2018.
- [2] Moshe Babaioff, Kira Goldner, and Yannai A. Gonczarowski. Bulow-klemperer-style results for welfare maximization in two-sided markets. In Shuchi Chawla, editor, Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020, pages 2452–2471. SIAM, 2020.
- [3] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms, pages 434–443, 2007.
- [4] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2014.
- [5] Liad Blumrosen and Shahar Dobzinski. (almost) efficient mechanisms for bilateral trading. CoRR, abs/1604.04876, 2016.
- [6] Liad Blumrosen and Yehonatan Mizrahi. Approximating gains-from-trade in bilateral trading. In Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings, pages 400–413, 2016.
- [7] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017, pages 589–590, 2017.
- [8] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In the 48th Annual ACM Symposium on Theory of Computing (STOC), 2016.
- [9] Yang Cai and Zhiyi Huang. Simple and Nearly Optimal Multi-Item Auctions. In the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2013.

¹⁶For discrete distributions, the virtual value for the seller's distribution is defined as $(s) = s + \frac{t < s}{g(t)(s - s^0)}$, where s^0 is the largest type in the support that is smaller than s .

¹⁷In bilateral trading, the mechanism is essentially the Buyer Offering mechanism [7]: The buyer picks a take-it or leave-it price according to her value and the seller can choose whether to sell at that price.

- [10] Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, pages 170–183, 2017.
- [11] Yang Cai and Mingfei Zhao. Simple mechanisms for profit maximization in multi-item auctions. In Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019., pages 217–236, 2019.
- [12] Shuchi Chawla and J. Benjamin Miller. Mechanism design for subadditive agents via an ex-ante relaxation. In Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016, pages 579–596, 2016.
- [13] Shuchi Chawla and Balasubramanian Sivan. Bayesian algorithmic mechanism design. ACM SIGecom Exchanges, 13(1):5–49, 2014.
- [14] Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016, pages 1424–1443, 2016.
- [15] Riccardo Colini-Baldeschi, Paul Goldberg, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Fixed price approximability of the optimal gain from trade. In International Conference on Web and Internet Economics, pages 146–160. Springer, 2017.
- [16] Riccardo Colini-Baldeschi, Paul W Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, and Stefano Turchetta. Approximately efficient two-sided combinatorial auctions. ACM Transactions on Economics and Computation (TEAC), 8(1):1–29, 2020.
- [17] Constantinos Daskalakis. Multi-item auctions defying intuition? ACM SIGecom Exchanges, 14(1):41–75, 2015.
- [18] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. In Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015, pages 449–450, 2015.
- [19] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. Modularity and greed in double auctions. In ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014, pages 241–258, 2014.
- [20] Bradley Efron and Charles Stein. The jackknife estimate of variance. The Annals of Statistics, pages 586–596, 1981.
- [21] Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016, pages 1014–1033, 2016.
- [22] Sergiu Hart and Noam Nisan. Approximate Revenue Maximization with Multiple Items. In the 13th ACM Conference on Electronic Commerce (EC), 2012.
- [23] Robert Kleinberg and S. Matthew Weinberg. Matroid Prophet Inequalities. In the 44th Annual ACM Symposium on Theory of Computing (STOC), 2012.
- [24] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. Advances in Prob, 4(197-266):1–5, 1978.
- [25] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. Proceedings of the National Academy of Sciences, 110(28):11232–11237, 2013.
- [26] A. M. Manelli and D. R. Vincent. Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly. Journal of Economic Theory, 137(1):153–185, 2007.
- [27] Preston R McAfee. The gains from trade under fixed price mechanisms. Applied Economics Research Bulletin, 1(1):1–10, 2008.
- [28] R Preston McAfee. A dominant strategy double auction. Journal of economic Theory, 56(2):434–450, 1992.
- [29] Roger B. Myerson. Optimal Auction Design. Mathematics of Operations Research, 6(1):58–73, 1981.
- [30] Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. Journal of economic theory, 29(2):265–281, 1983.
- [31] REAC Paley and Antoni Zygmund. A note on analytic functions in the unit circle. In Mathematical Proceedings of the Cambridge Philosophical Society, volume 28, pages 266–272. Cambridge University Press, 1932.
- [32] Aviad Rubinfeld. Beyond matroids: secretary problem and prophet inequality with general constraints. In Daniel Wicks and Yishay Mansour, editors, Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016, pages 324–332. ACM, 2016.
- [33] Aviad Rubinfeld and S. Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015, pages 377–394, 2015.
- [34] Ester Samuel-Cahn et al. Comparison of threshold stop rules and maximum for independent nonnegative random variables. the Annals of Probability, 12(4):1213–1216, 1984.
- [35] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. Double auctions in markets for multiple kinds of goods. In Jérôme Lang, editor, Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, pages 489–497. ijcai.org, 2018.
- [36] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. MUDA: A truthful multi-unit double-auction mechanism. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), pages 1193–1201. AAAI Press, 2018.
- [37] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In SODA, 2015.