

Strongly Budget Balanced Auctions for Multi-Sided Markets¹

Rica Gonen and Erel Segal-Halevi

Abstract

In two-sided markets, Myerson and Satterthwaite’s impossibility theorem states that one can not maximize the gain-from-trade while also satisfying truthfulness, individual-rationality and no deficit. Attempts have been made to circumvent Myerson and Satterthwaite’s result by attaining approximately-maximum gain-from-trade: the double-sided auctions of McAfee (1992) is truthful and has no deficit, and the one by Segal-Halevi et al. (2016) additionally has no surplus — it is strongly-budget-balanced. They consider two categories of agents — buyers and sellers, where each trade set is composed of a single buyer and a single seller.

The practical complexity of applications such as supply chain require one to look beyond two-sided markets. Common requirements are for: buyers trading with multiple sellers of different or identical items, buyers trading with sellers through transporters and mediators, and sellers trading with multiple buyers. We attempt to address these settings.

We generalize Segal-Halevi et al. (2016)’s strongly-budget-balanced double-sided auction setting to a multilateral market where each trade set is composed of any number of agent categories. Our generalization refines the notion of competition in multi-sided auctions by introducing the concepts of external competition and trade reduction. We also show an obviously-truthful implementation of our auction using multiple ascending prices.

Full version, including omitted proofs and simulation experiments, is available at <https://arxiv.org/abs/1911.08094>.

1 Introduction

Mechanism design for one-sided markets has been investigated for several decades in economics and in computer science. It aims to find an efficient (high social welfare) allocation of a set of items to a set of agents, while ensuring that truthfully reporting the input data is the best strategy for the agents. The Vickrey-Clarke-Groves (VCG) auction Vickrey (1961); Clarke (1971); Groves (1973) is a pillar of mechanism design. VCG auctions maximize the social welfare of the agents. They are dominant-strategy truthful (DST) — each agent’s dominant strategy is to truthfully report its preferences to the auction, regardless of what the other agents report. They can also be made individually rational (IR) — no agent loses from participating in the auction.

More recently, there has been increased attention on auctions for two-sided markets, in which the set of agents is partitioned into buyers and sellers. As opposed to the one-sided setting, where the auctioneer initially holds the items, in the two-sided setting the items are initially held by the set of sellers. The sellers express valuations for the items they hold, and are assumed to act rationally and strategically. Thus, the auctioneer is tasked with deciding which buyers and sellers should trade and with what prices.

The growing interest in two-sided markets can be attributed to various important applications. Examples range from selling display-advertising on ad exchange platforms, the US FCC spectrum license reallocation, and stock exchanges. However, little work has been done so far on the next level of generalization, i.e., multi-sided markets.

In two-sided markets, a further important requirement is strong budget-balance (SBB), which states that monetary transfers happen only among the agents in the market. This means that buyers and sellers are allowed to trade without leaving the auctioneer any share of their gains and without

¹This result was presented at AAAI2020 as an oral talk.

the auctioneer adding money into the market. A weaker version of SBB, often considered in the literature, is weak budget-balance (WBB). WBB only requires the auctioneer not to add money to the market. The problem with weak budget-balance is that the surplus of the auctioneer might consume most of the gain-from-trade, leaving little gain for the actual traders. This might drive traders away from the market.² Note that, in bilateral trade settings, VCG is usually not even WBB except in special cases Guo et al. (2013).

For double-sided auctions, the impossibility theorem of Myerson and Satterthwaite (1983) states that one can not maximize gain from trade (GFT, the difference between the total value of the sold items for the buyers and the total value of these items for the sellers) while also satisfying IR, DST, and no deficit.

The seminal double-auction mechanism of McAfee (1992) is DST, IR and WBB. It circumvents Myerson and Satterthwaite’s result by compromising on GFT: it may remove up to one deal from the optimal trade. In case a deal is removed, it is the one with the smallest GFT among the deals in the optimal trade; hence it attains at least $1 - 1/k$ of the optimal GFT, where k is the number of optimal deals. Thus, it is **asymptotically optimal** — its GFT approaches the optimum when $k \rightarrow \infty$.

Recently, Segal-Halevi et al. (2016) presented a SBB variant of McAfee’s mechanism, with similar GFT guarantees. Their mechanism may remove up to one *buyer* from the optimal trade, and it is the buyer with the lowest value among the buyers in the optimal trade. In case a buyer is removed, the remaining $k - 1$ buyers trade with $k - 1$ sellers selected at random from the k sellers in the optimal trade.

The complexity of practical requirements in areas such as supply chain require one to look beyond double-sided markets. As an example (Babaioff and Walsh, 2005), a market for lemonade may contain two kinds of sellers (lemon pickers and sugar producers), two kinds of buyers (juice drinkers and lemonade drinkers), and some intermediary agents (lemon squeezers, lemonade mixers, etc.) Our goal is to address such settings while keeping the strong budget balance requirement.

1.1 Our Contribution

Our contribution is twofold: First we generalize Segal-Halevi et al. (2016)’s SBB double auction to a multi-sided market where the trade set is composed of any number of agent types and any number of copies of any agent’s type. Our generalization refines the notion of competition in multi-sided auctions by introducing the concept of external competition — competition over who will act as a given participant’s trade partner(s) (complementarity).

The expanded notion of competition allows us to provide a simple well-performing procedure that generalizes McAfee (1992)’s trade reduction. The shift to thinking in terms of competition allows us to broadly address situations common to multi-sided auctions. These settings include trading entities that may be individuals or entire markets, transactions facilitated by zero or more intermediaries, and goods that can be exchanged individually or in bundles.

These settings also encompass many common commercial mechanisms including supply chains, distributed markets, security exchanges, and business to consumer auctions. Historically each of these settings has been considered unique and each presented the complex research problem of finding a suitable mechanism (see section 5 for details).

Second, in addition to the direct-revelation multi-sided auction, our result presents a multi-sided ascending-prices auction that implements the same outcome. In the theory of one-sided auctions, it is well-known that a second-price direct-revelation auction and an ascending-prices auction are strategically equivalent. In both auctions, the agent with the highest value wins and pays the second-

²The following “trick” can be used to convert any WBB auction to an SBB auction: before the auction starts, remove a random trader from the market; after the auction ends, give that trader all the surplus (if any). We do not support this trick since it might induce agents who have nothing to do with the auction (e.g. “sellers” with nothing to sell or “buyers” with no money) to come to the market, only because of the chance to win all the surplus. Like Colini-Baldeschi et al. (2017), we focus on *direct-trade auctions* — auctions that give/take money only to/from agents who actually participate in the trade.

highest value. However, an ascending-prices auction has the advantage that it is *obviously truthful* (see Li (2017) for formal definitions and proofs). The practical advantage of an obviously-truthful auction is that it is easier for people to understand that playing truthfully is best for them, even if they are not experts in game theory. This is particularly important when one deals with complex multi-sided markets with many entities.

1.2 Paper layout

Section 2 presents the formal definitions. Section 3 presents a special case of our extended multi-lateral auctions in which each trade requires exactly one agent of each category. Section 4 presents a more general case of our extended multilateral auctions in which each trade requires a fixed number of agents of each category, but this fixed number may be larger than 1. Section 5 compares our work to related work. The full version presents some simulation experiments evaluating the performance of our auctions. Section 6 concludes with some future work directions. An open-source implementation of our auctions, including example runs and experiments, is available at <https://github.com/erelsql/auctions>.

2 Preliminaries

2.1 Agents and categories

A *market* is defined by a set of *agents* grouped into different *categories*. N is the set of agents, G is the set of agent categories, and N_g is the set of agents in category $g \in G$. The categories are pairwise-disjoint, so $N = \sqcup_{g \in G} N_g$.

Each deal in the market requires a certain combination of agents. We call a subset of agents that can accomplish a single deal a *procurement-set* (PS). The *PS recipe* of the market is a vector of size $|G|$, denoted by $\mathbf{r} := (r_g)_{g \in G}$, where $r_g \in \mathbb{Z}_+$ for all $g \in G$. It describes the number of agents of each category that should be in each PS: each PS should contain r_1 agents of category 1, r_2 agents of category 2, and so on. As an example, the PS recipe of a standard two-sided market is $(1, 1)$, since there are two agent categories — buyers and sellers — and each PS should contain one buyer and one seller.

As another example, consider a market with three categories of agents — buyers, sellers and transporters, and PS recipe $(1, 2, 2)$. In such a market, each deal requires a buyer, two sellers and two transporters.

In general, one could think of markets with multiple PS recipes; however, in the present paper we restrict our attention to markets with a single PS recipe, denoted by \mathbf{r} .

Each agent $i \in N$ has a *value* $v_i \in \mathbb{R}$, which represents the monetary gain of an agent from participating in the trade. The value of an agent is the agent’s private information. It may be positive or negative. For example, in a two-sided market, the value of a buyer is typically positive while the value of a seller is typically negative. The agents are *quasi-linear in money*: the utility of agent i participating in some PS and paying p_i is $u_i := v_i - p_i$.

2.2 Trades and Gains

The *gain-from-trade* of a procurement-set S , denoted $GFT(S)$, is the sum of values of all agents in S :

$$GFT(S) := \sum_{i \in S} v_i.$$

In a standard two-sided market, the GFT of a PS with a buyer b and a seller s is $v_b - v_s$, since the seller’s value is $-v_s$.

Given a market (N, G, \mathbf{r}) , a *trade* is a collection of pairwise-disjoint procurement-sets. I.e, it is a collection of agent subsets, $S_1, \dots, S_k \subseteq N$, such that for each $j \in [k]$, the composition of agents in S_j corresponds to the recipe \mathbf{r} . The total GFT is the sum of the GFT of all procurement-sets participating in the trade:

$$\text{GFT}(S_1, \dots, S_k) := \sum_{j=1}^k \text{GFT}(S_j)$$

A trade is called *optimal* if its GFT is maximum over all possible trades.

2.3 Competition

Our direct-revelation auctions are based on the concept of *competition* between agents. Given a trade (S_1, \dots, S_k) , let $N_{rm} := N \setminus (S_1 \cup \dots \cup S_k)$ be the subset of agents who do not participate in the trade (the “remaining market”).

Consider a single PS S_j , and an agent $i \in S_j$ who belongs to category g , i.e, $i \in N_g$. Then, a subset of agents $T \subseteq N_{rm}$ is called an *external competition* for i if adding i to T yields a PS consistent with the recipe \mathbf{r} , with a positive GFT:

$$\text{GFT}(T \cup \{i\}) \geq 0$$

In a simple two-sided market, the external competition of a trading buyer is a non-trading seller whose value is sufficiently high such that, combining the trading buyer with the non-trading seller yields a pair with a GFT above 0.

In a three-sided market with buyers, sellers and mediators, with $\mathbf{r} = (1, 1, 1)$, an external competition of a trading buyer is a pair of a non-trading seller and a non-trading mediator, such that the GFT of the buyer+seller+mediator is at least 0.

3 One Agent Per Category

This section presents our two auctions for a special case in which the single PS recipe in the market is a vector of ones, so each PS must contain a single agent from each category.

Both auctions are parametrized by an ordering on the categories: each of the $|G|!$ possible orderings yields a different auction. The ordering should be fixed in advance and not depend on the agents’ values.

We present the auctions using a running example with three categories in the following order: buyers, sellers and mediators. The recipe is $(1, 1, 1)$. In each category there are five agents. The agents’ values are:

- Buyers: 17, 14, 13, 9, 6.
- Sellers: -1, -4, -5, -8, -11.
- Mediators: -1, -3, -4, -7, -10.

3.1 External-competition auction

The auction requests the agents to report their values, and then proceeds as follows.

Step 1: Optimal trade calculation. Order the agents in each category by descending order of their value. Combine the highest-value agents in each category to a PS. Combine the next-highest-value agents in each category into a PS. Keep constructing PS as long as the GFT of the constructed PS is positive. The resulting set of PS is the optimal trade. We denote by k the number of PS in the optimal trade.

In the running example, $k = 3$ and the optimal trade contains the following PS: $(17, -1, -1)$ with GFT 15, $(14, -4, -3)$ with GFT 7, and $(13, -5, -4)$ with GFT 4. The remaining market, denoted by N_{rm} , contains two buyers 9, 6, two sellers $-8, -11$ and two mediators $-7, -10$.

Step 2. Order the procurement-sets in the optimal trade by ascending GFT, such that, $\text{GFT}(S_1) \leq \dots \leq \text{GFT}(S_k)$. In the running example, S_1 is the PS $(13, -5, -4)$.

Step 3. Consider the agents in S_1 in the pre-determined order of categories. Initialize i to the first agent in S_1 by this ordering. In the running example, it is the buyer 13.

Step 4. Look for an external competition to i with a largest GFT. There are two cases.

Case 4a. No external competition for i is found. Then, i is removed from the trade (and added to N_{rm}), and we go back to step 4 with i being the next agent in S_1 .

In the running example, we consider first the buyer 13. The maximum GFT of a PS that contains this buyer and agents from N_{rm} is -2 , for the PS $(13, -8, -7)$. This GFT is negative so it is not considered an external competition. Hence, the buyer 13 is removed from trade.

Case 4b. An external competition for i is found; denote it by T_1 . From now on we call this agent i the *pivot agent* and its category the *pivot category*. Denote the pivot category by g_o . For each $g \neq g_o$, denote by v_g^T the value of the single agent in $T_1 \cap N_g$. Trade prices are calculated as follows:

- The price p_g for each agent in category $g \neq g_o$ is set to the value v_g^T .
- The price p_o for each agent in category g_o is set to: $p_o := -\sum_{g \neq g_o} v_g^T$.

In the running example, the next member of S_1 (after the buyer 13 is removed) is the seller -5 . The maximum GFT of a PS that contains this seller and agents from N_{rm} is $+1$, for the PS $(13, -5, -7)$; note that the removed buyer 13 participates in this PS. This GFT is positive so it is an external competition; the pivot category g_o is the sellers' category.

The prices are set to 13 for the buyers (like the buyer in T_1), -7 for the mediators (like the mediator in T_1), and $-6 = -(13 - 7)$ for the sellers. The final price-vector is thus $(13, -6, -7)$, i.e., all buyers pay 13, all sellers receive 6 and all mediators receive 7.

Step 5. Once the prices are calculated, the final trade is determined as follows:

- For each category, count the number of members remaining in the trade.
- In each category with the smallest count, all agents participate in the trade.
- In each category with a larger count, there is a lottery determining who will participate in the trade.

In the running example, there are two remaining buyers $(17, 14)$ all of whom trade at price 13; there are three remaining sellers $(-1, -4, -5)$ two of whom (selected at random) trade at price -6 ; similarly, there are three remaining mediators $(-1, -3, -4)$ two of whom trade at price -7 .

Note that selecting a different one of the 6 category-orders leads to a different outcome. A-priori, there is no reason to prefer one ordering over the other — our auction has the same desirable properties (proved below) for any ordering.

The SBBA auction of Segal-Halevi et al. (2016) is a special case of our auction, where the recipe is $(1, 1)$. Their two variants correspond to the two orderings — buyers-sellers or sellers-buyers.

3.2 Proof of correctness

First, note that there must be an agent $i \in S_1$ for whom an external competition exists. In the worst case, when only one last agent of S_1 remains in the trade, the other agents of S_1 (who were previously removed from trade) form an external competition for this agent. This is because their total GFT is $\text{GFT}(S_1)$, which is positive since S_1 is in the optimal trade.

Lemma 1. *For each category $g \in G$, denote by v_g^S the value of the single agent in $S_1 \cap N_g$. Then:*

$$\sum_{g \neq g_o} v_g^T \leq \sum_{g \neq g_o} v_g^S$$

Proof. For $g < g_o$, the agent in $S_1 \cap N_g$ had been removed from trade before the pivot was found, and was later used as an external competition for the pivot, so it is the same agent as in $T_1 \cap N_g$ and thus $v_g^T = v_g^S$.

For $g > g_o$, the agent in $S_1 \cap N_g$ had not been removed from trade, and thus, another (non-trading) agent was used as an external competition. Since the values of non-trading agents are smaller than that of trading agents, $v_g^T \leq v_g^S$. \square

Now we prove the properties of the auction.

Theorem 1. *The external-competition auction of Subsection 3.1 is strongly-budget-balanced, individually-rational and dominant-strategy truthful, and its gain-from-trade approaches the optimum when the optimal market size (k) approaches ∞ .*

Proof. Strong budget balance is obvious: the price p_o is calculated such that the sum of prices in each PS is 0.

Individual rationality: we prove that the price paid by each trading agent is at most the agent's reported value.

- Each trading agent in a category $g \neq g_o$ pays the value v_g^T of a non-trading agent in the same category g . The agents in each category are ordered by descending value, and the value of each trading agent is at least as large as the value of each non-trading agent in the same category, so it is at least v_g^T .
- Let v_o^S be the value of the pivot agent (who is an agent in S_1). By definition of external competition, the sum of values of agents in T_1 plus v_o^S is at least 0, so

$$\begin{aligned} v_o^S + \sum_{g \neq g_o} v_g^T &\geq 0 \\ \implies v_o^S &\geq - \sum_{g \neq g_o} v_g^T = p_o \end{aligned}$$

Since agents are ordered by descending value, the values of other trading agents are at least v_o^S which is at least p_o .

Truthfulness: By Myerson's theorem, it is sufficient to prove that the choice rule is monotone, and each trading agent pays his/her threshold value.

Monotonicity is obvious: an agent increasing his reported value (while other reports are fixed) is more likely to participate in the optimal trade, more likely to have an external competition, and thus more likely to remain in the trade.

To calculate the threshold value of an agent i from category g , we consider three cases, depending on the fixed ordering of the categories:

1. The category g comes before the pivot-category g_o . This means that an agent from g had been removed from S_1 before the pivot was found. All agents whose value is higher than v_g^T are in PS S_j for $j \geq 2$, they do not affect the auction in any way, and they remain in the trade. Any such agent whose value drops below v_g^T , replaces the v_g^T agent in the PS S_1 , and has no external competition, and so is removed from the trade. Therefore, v_g^T is a threshold-value for all agents of g , and indeed $p_g = v_g^T$.
2. $g = g_o$. All agents whose value is higher than v_o^S (the value of the pivot agent) are in PS S_j for $j \geq 2$, they do not affect the auction in any way, and they remain in the trade. Consider an agent of g_o whose value v_o drops below v_o^S but above p_o (recall that $v_o^S \geq p_o$). We claim that this agent remains in the trade. First, v_o is still in the optimal trade (it replaces the pivot agent in S_1), since:

$$\begin{aligned} v_o + \sum_{g \neq g_o} v_g^S &\geq v_o + \sum_{g \neq g_o} v_g^T \quad (\text{by Lemma 1}) \\ &= v_o - p_o \quad (\text{by definition of } p_o) \\ &\geq 0 \quad (\text{by assumption on } v_o), \end{aligned}$$

so the GFT of v_o plus the other agents in S_1 is still above 0. Second, T_1 is still an external competition for v_o , since:

$$v_o + \sum_{g \neq g_o} v_g^T = v_o - p_o \geq 0.$$

But, once v_o drops below p_o , the set T_1 is no longer an external competition, so the agent is removed from the trade. Hence, p_o is a threshold value for all agents of g_o .

3. The category g comes after the pivot-category g_o . This means that no agent from g had been removed from the trade before the pivot was found. As shown in Lemma 1, in this case $v_g^S \geq v_g^T$. All agents whose value is higher than v_g^S are in PS S_j for $j \geq 2$, they do not affect the auction in any way, and they remain in the trade.

Consider an agent of g whose value v_g drops below v_g^S but above v_g^T . We claim that this agent remains in the optimal trade (it replaces the agent v_g^S in S_1), since:

$$\begin{aligned} GFT(S_1) - v_g^S + v_g &\geq GFT(S_1) - v_g^S + v_g^T \\ &= v_o^S + \sum_{g \neq g_o} v_g^S - v_g^S + v_g^T \\ &\geq v_o^S + \sum_{g \neq g_o} v_g^T \quad (\text{by Lemma 1}) \\ &\geq 0 \quad (T_1 \text{ is external competition}) \end{aligned}$$

so the GFT of v_g plus the other agents in S_1 is above 0.

But, once v_g drops below v_g^T , it is replaced by the agent v_g^T in S_1 , and does not enter the optimal trade. Hence, v_g^T is a threshold-value for all agents of g , and $p_g = v_g^T$.

Gain-from-trade: For each $g \in G$ and $j \in \{1, \dots, k\}$, denote by v_g^j the value of the single agent in $N_g \cap S_j$. Then the optimal GFT is:

$$OPT = \sum_{g \in G} \sum_{j=1}^k v_g^j$$

If no traders are removed, then all these k PS are trading, and the GFT equals OPT. If some traders are removed, they are removed from S_1 which is the least profitable PS. In this case, $k - 1$ deals are made, where in each deal, the trader from each category g is:

- If g is before the pivot — one of the $k - 1$ high-value traders in g ;
- If g is the pivot or after the pivot — one of the k high-value traders in g , selected at random.

Hence, the expected GFT is at least:

$$\begin{aligned} & \sum_{g < g_o} \sum_{j=2}^k v_g^j + \sum_{g \geq g_o} \frac{k-1}{k} \sum_{j=1}^k v_g^j \\ & \geq \sum_{g \in G} \frac{k-1}{k} \sum_{j=1}^k v_g^j \\ & = \frac{k-1}{k} OPT. \end{aligned}$$

3.3 Ascending-prices auction

Our ascending-prices auction holds a price p_g for each category $g \in G$. All prices are initialized to $-\infty$, and initially all agents are in the trade (since every agent will be happy to pay $-\infty$). While the prices increase, each agent in category g with value v_g remains in the trade as long as $p_g < v_g$, and exits the trade when $p_g > v_g$ (since the prices increase monotonically, agents never return to the trade after exiting). When $p_g = v_g$, the agent is indifferent between trading and not trading; for simplicity, we assume that in this case the agent does not trade. Also, for simplicity we assume that the agents' valuations are *generic* in the sense that, for each category $g \in G$, all agents have different values. During the presentation of the ascending auction, we use the same running example as in Subsection 3.1.

Step 1: Initialization. For each category g , count the number of agents in N_g ; let n_{\min} be the size of the smallest category. For each category g with more than n_{\min} agents, increase the price p_g such that some agents leave the trade, until the number of remaining agents in all categories is n_{\min} .

In the running example, this step is not needed since initially there are 5 agents in each category.

Step 2. Loop over the categories in the pre-specified order. For each category g , increase p_g continuously until one of the following happens:

- an agent from category g exits the trade, or —
- the sum of prices increases to zero: $\sum_{g \in G} p_g = 0$.

In case (a), repeat the step with the next category (after the last category, return to the first one).

If a category becomes empty, the auction stops and there is no trade.

In case (b), stop and have the agents trade in the final prices: each agent in category $g \in G$ trades at price p_g . If, in some category, there are more remaining agents than in other categories, then a lottery is used to select who will trade.

In the running example, at the first round, the buyers' price increases to 6, the sellers' price increases to -11, the mediators' price increases to -10; after the first round, there are 4 agents remaining in each category, and the sum of prices is still negative, so we continue. At the second round, the prices increase to 9, -8, -7 and the sum is still negative. At the third round, the buyers' price increases to 13 and the sellers' price is increased towards -5, but when it hits -6, the sum of prices becomes 0 so the auction stops. The final trade is exactly the same as in the external-competition auction.

Theorem 2. *The ascending-prices auction of Subsection 3.3 is strongly-budget-balanced, individually-rational and obviously truthful, and its gain-from-trade approaches the optimum when the optimal market size (k) approaches ∞ .*

Proof. SBB and IR are immediate from the description.

As for obvious-truthfulness: Li (2017) defines a strategy S as obviously-dominant (for a given agent) if “for any deviating strategy T , starting from an earliest information set where S and T diverge, the best possible outcome from T is no better than the worst possible outcome from S ”. We show that, in the ascending-prices auction, for each agent i in category g , the strategy S of exiting when $p_g = v_i$ is obviously-dominant.

The worst outcome from S has a value of 0. We now show that, for any deviation T , the best possible outcome from T when S and T diverge has a value of at most 0. Indeed, if T is exiting too early (at some $v'_i < v_i$), then the point at which S and T diverge is when $p_g = v'_i$, and at that point the outcome from T has a value of 0. If T is exiting too late (at some $v''_i > v_i$), then the point at which S and T diverge is when $p_g = v_i$, and at that point all possible outcomes from T have a value of 0 or less.

We now analyze the gain-from-trade.

Let g_o be the category in which the protocol stops. Let n_o be the number of traders of this category that remain in the trade. Then, in each category $g < g_o$, there are $n_o - 1$ remaining traders, and in each category $g \geq g_o$, there are n_o remaining traders.

Recall that k is the number of deals in the optimal trade; we claim that $k = n_o$:

- First, suppose that the price p_g of each category $g \geq g_o$ is increased up to the value of the next agent in g (who did not exit the trade in the actual auction). Since the auction stopped when the sum of prices hit 0, the sum of prices after the increase is positive. Each price p_g equals the n_o -th highest value in category g . This means that there are at least n_o procurement-sets with a positive GFT, so $k \geq n_o$.
- Second, suppose that the price p_g of each category $g \leq g_o$ is decreased down to the value of the previous agent in g (who did exit the trade in the actual auction). Now the sum of prices is negative. Each price p_g equals the $(n_o + 1)$ -th highest value in category g . This means that there are not $(n_o + 1)$ procurement-sets with a positive GFT, so $k < n_o + 1$. Hence, $k = n_o$.

So at least $k - 1$ deals are done. From here, the proof is identical to the gain-from-trade proof in Theorem 1. \square

4 General Procurement-Set Recipes

This section extends the previous one by allowing the PS recipe to be an arbitrary vector of positive integers, rather than just a vector of ones. For each category g there is an integer $r_g \geq 1$, and every PS must contain exactly r_g traders from this category.

The external-competition auction can be extended to the setting of an arbitrary vector of positive integers. However, as the proof is somewhat involved, we choose to focus here on the ascending-prices auction extension.

We present the mechanisms using a running example in which there are two categories — buyers and sellers, and the recipe is $(1, 2)$, so that each PS should contain one buyer and two sellers. The market contains:

- Five buyers with values: 17, 14, 13, 9, 6.
- Nine sellers with values: -1, -2, -3, -4, -5, -7, -8, -10, -11.

Note that our setting of an arbitrary vector of positive integers might include an agent say a seller that has multiple items for sale. However the setting is a single parameter setting in which case that seller has identical value for each of his items.

Note that the optimal trade in this setting can be calculated just like in Section 3: the agents in each category are ordered by descending value, and then grouped greedily into procurement-sets. In the running example, the optimal trade contains three PS: $(17; -1, -2)$ with GFT 14, $(14; -3, -4)$ with GFT 7, and $(13; -5, -7)$ with GFT 1.

4.1 Ascending-prices auction

Similarly to subsection 3.3, the auction maintains a price p_g for each category g . All prices are initialized to $-\infty$, and initially all agents are in the trade.

Step 1: Initialization. For each category g , let $c_g := \text{floor}(|N_g|/r_g)$; this is the largest number of PS that can be composed of agents of category g . Let $c_{\min} := \min_{g \in G} c_g$; this is the largest number of PS that can be composed of the existing agents. For each category g for which $\text{floor}(|N_g|/r_g) > c_{\min}$, increase the price p_g such that some agents leave the trade, until the number of remaining agents in each category g decreases such that $\text{floor}(|N_g|/r_g) = c_{\min}$. Note that, when g is the category for which the minimum of c_g is attained, we already have $\text{floor}(|N_g|/r_g) = c_{\min}$ without any increase, so the price of this category (at least) remains $-\infty$, and the initial price-sum is negative.

Initialize $c := c_{\min}$. Informally, c is the number of procurement-sets that we aim to construct from the traders currently in the market. Note that initially, for every category g , we have $\text{floor}(|N_g|/r_g) = c_{\min}$ so $|N_g| \geq r_g \cdot c$.

In the running example, $c_{\text{buyers}} = \text{floor}(5/1) = 5$ and $c_{\text{sellers}} = \text{floor}(9/2) = 4$ and $c_{\min} = 4$. In the initialization step, p_{buyers} increases to 6 so that the low-value buyer leaves. The market now has 4 buyers and 9 sellers. The value of c is initially 4.

Step 2. Loop over the categories in the pre-specified order. For each category g , increase p_g until one of the following happens:

- The number of agents in g drops to $r_g \cdot c$, or —
- The *weighted* sum of prices increases to zero, where the weights are the recipe constants, i.e.: $\sum_{g \in G} r_g \cdot p_g = 0$.

In case (a), repeat the step with the next category in the pre-specified order. After the last category, set $c := c - 1$ and cycle back to the first category. If c drops to 0 then the auction ends and there is no trade.

In case (b), the auction terminates and the agents trade in the final prices. If, in some category, there are more remaining agents than needed to fill the procurement-sets, then a lottery is used to select who will trade.

In the running example, at the first round, p_{sellers} increases to -11 so that one seller leaves. Now there are 4 buyers and 8 sellers so exactly 4 PS are supported. However, the price-sum is $6 + 2 \cdot (-11) < 0$. Hence, we decrease c to 3 and continue.

In the second round, p_{buyers} increases to 9 such that one buyer leaves, and p_{sellers} increases to -8 such that two sellers leave. The weighted sum of prices is $9 + 2 \cdot (-8) < 0$, so we decrease c to 2 and continue.

In the third round, p_{buyers} increases to 13 such that one buyer leaves. We start increasing $p_{sellers}$ towards -5 such that two sellers would leave, but during the increase, the price hits -6.5 , and then the sum of prices is $13 + 2 \cdot (-6.5) = 0$, so the auction stops.

There are now two remaining buyers (17, 14) and five remaining sellers ($-1, -2, -3, -4, -5$). All remaining buyers trade and pay 13; 4 out of 5 remaining sellers are selected at random, they trade and receive 6.5.

Theorem 3. *The ascending-prices auction of Subsection 4.1 is strongly-budget-balanced, individually-rational and obviously truthful, and its gain-from-trade approaches the optimum when the optimal market size (k) approaches ∞ .*

Proof. SBB, IR and truthfulness are obvious — just as in Theorem 2. We now analyze the gain-from-trade. Let c_* be the final value of c . Note that the number of agents remaining in each category g is at least $r_g \cdot c_*$ and at most $r_g \cdot (c_* + 1)$.

We claim that $c_* \leq k \leq c_* + 1$:

- First, suppose that the price p_g of each category g is increased such that exactly $r_g \cdot c_*$ agents remain. Now the sum of prices is positive, and each price p_g equals the $(r_g \cdot c_* + 1)$ -th highest value in category g . This means that there are at least c_* procurement-sets with a positive GFT, so $k \geq c_*$.
- Second, suppose that the price p_g of each category g is decreased to its value at the end of round $c_* + 1$, such that exactly $r_g \cdot (c_* + 1)$ agents remain. Now the sum of prices is negative. Each price p_g equals the $(r_g \cdot c_* + r_g + 1)$ -th highest value in category g . The expression $(r_g \cdot c_* + r_g + 1)$ is at most $r_g \cdot (c_* + 2)$, with equality holding iff $r_g = 1$. This means that there are not $(c_* + 2)$ procurement-sets with a positive GFT, so $k \leq c_* + 1$.

Hence, in each category g , the number of agents is at least $r_g \cdot (k - 1)$ and at most $r_g \cdot (k + 1)$. At least $k - 1$ deals are done, and the participants are from the at most $k + 1$ highest sets of r_g traders in each category g . From here, the proof is similar to the gain-from-trade proof in Theorem 1. \square

5 Related Work

The literature on two-sided markets is large and we do not attempt to cover it here; see e.g. Brustle et al. (2017); Babaioff et al. (2018, 2020) for some recent references. Below we focus on auctions for markets with three or more sides.

Babaioff and Nisan (2004) handle a multi-sided market using multiple two-sided sub-markets, where each sub-market hosts an independent double-auction. Their example is the *lemonade industry*, which consists of lemon pickers, squeezers, and drinkers. In our auction, all three categories bid together in a single centralized auction with recipe $(1, 1, 1)$; in their setting, there are three different double-auctions, one each for pickers, squeezers and drinkers. The sub-markets are constructed so that the optimal number of deals is the same in all of them. So if the double-auction mechanisms make deals as a function of the optimal number only, their protocol is guaranteed to have a material balance. However, their protocol does not preserve SBB: in <https://github.com/erelsgl/auctions> we have a runnable example in which the same SBB double-auction is used in all sub-markets, but the combined outcome is not SBB. Babaioff and Nisan (2004) do present a SBB variant of their mechanism, but only in expectation, and it requires a prior distribution on the agents' valuations. In contrast, our mechanism is SBB with probability 1, and requires no prior.

Babaioff and Walsh (2005) extend the above work from a linear supply-chain to an arbitrary directed acyclic graph. For example, they consider a market in which pickers sell lemons to squeezers,

sugar-makers and squeezers sell to manufacturers, and manufacturers sell lemonade to drinkers. It still does not guarantee SBB.

Chen et al. (2005) consider a supply-chain auction with a sole buyer and single item-kind, but there are different producers in different supply-locations. The buyer needs a different quantity of the item in different demand-locations. The buyer conducts a reverse auction and has to pay, in addition to the cost of production, also the cost of transportation from the supply-locations to the demand-locations. They do not guarantee SBB. Gonen et al. (2007) generalized the above settings to a unified trade reduction procedure.

McAfee (2008) designs fixed price SBB double auction under the assumptions that the buyer's distribution dominates the seller's distribution or that there is exponential distribution. Our result does not assume knowledge of the distribution of participating categories. Additionally, we also allow for general number of categories as opposed to two.

Colini-Baldeschi et al. (2016); Colini-Baldeschi et al. (2017) also presents SBB auctions. Their auctions target double-sided and combinatorial markets respectively. However, their goal is to maximize social welfare as opposed to our goal which is maximizing gain from trade³. Thus their mechanisms are not asymptotically-optimal for gain from trade. They also require a prior on the agents' valuations.

Feldman et al. (2018); Feldman and Gonen (2018) presents a multilateral randomized market with buyers, mediators and sellers in the context of ad auctions. Their sellers are pre-associated with the mediators and they assume that mediators and buyers arrive over time in some uniform random order. Moreover, their trade matches are conducted in two stages: first the mediator trade with the buyer on behalf of his sellers and then the mediator transfers payments to his matched sellers. Our auction unites all three categories of buyer, seller and mediator actions into a single simultaneous trade step.

Some works attempt to handle mechanisms with buyers and sellers interacting through an intermediary or trader such as Feldman et al. (2010) and Blume et al. (2007). However, their design reduces to a one-sided auction or a two-sided auction.

The present work handles multiple categories of agents, but each agent is single-parametric. An orthogonal line of work Segal-Halevi et al. (2018a,b); Gerstgrasser et al. (2019); Gonen and Egri (2017) remains with two agent categories (buyers and sellers), but aims to handle multi-parametric agents. Another orthogonal line of work gets around Myerson –Satterthwaite in a different way: it relaxes truthfulness but keeps the maximum GFT. See e.g. Lubin et al. (2008).

6 Future Work

The results in this paper can be generalized in several ways.

1. Allow several recipes of procurement-sets. For example, suppose there are three categories: buyer, seller and purchaser. Each PS should contain either a buyer and a seller, or a purchaser and two sellers. In this case, even calculating the optimal trade, without any strategic considerations, may already be non-trivial.

2. Allow transaction costs. In general, each procurement-set may have a different cost-of-transaction, depending on the geographic locations of the agents in the PS, as well as other factors. Again, such transaction costs make the computation of the optimal trade difficult, even before strategic considerations, and even when all transaction costs are common knowledge. We have preliminary results showing that, without any restrictions on the transaction costs, there might be no auction

³When optimizing gain from trade we optimize the difference between the total value of the sold items for the buyers and the total value of these items for the sellers. When optimizing social welfare in a market we optimize the sum of the buying agents' valuations plus the sum of the unsold items' value held by selling agents at the end of trade. Despite their conceptual similarity, the two objectives are rather different in approximation. In many cases the social welfare approximation is close to the optimal social welfare solution; however, the gain from trade approximation may not be within any constant factor of the optimal gain from trade.

that satisfies all the desirable properties of Theorem 1. We believe that such an auction can be found given some natural restrictions on the transaction costs.

Acknowledgments

We are grateful to three referees of AAAI 2020 and three referees of COMSOC 2021 for their helpful comments.

References

- Moshe Babaioff, Yang Cai, Yannai A Gonczarowski, and Mingfei Zhao. 2018. The Best of Both Worlds: Asymptotically Efficient Mechanisms with a Guarantee on the Expected Gains-From-Trade. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM, 373–373. arXiv preprint 1802.08023.
- Moshe Babaioff, Kira Goldner, and Yannai A Gonczarowski. 2020. Bulow-Klemperer-Style Results for Welfare Maximization in Two-Sided Markets. In *Proceedings of SODA'20*. arXiv preprint arXiv:1903.06696.
- Moshe Babaioff and Noam Nisan. 2004. Concurrent Auctions Across the Supply Chain. *Journal of Artificial Intelligence Research (JAIR)* 21 (2004), 595–629. <https://doi.org/10.1613/jair.1316>
- Moshe Babaioff and William E. Walsh. 2005. Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. *Decision Support Systems* 39, 1 (March 2005), 123–149. <https://doi.org/10.1016/j.dss.2004.08.008>
- L. Blume, D. Easley, J. Kleinberg, and E. Tardos. 2007. Trading Networks with Price-Setting Agents. In *EC*. 13–16.
- Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. 2017. Approximating Gains from Trade in Two-sided Markets via Simple Mechanisms. (18 June 2017). arXiv:1706.04637 <http://arxiv.org/abs/1706.04637>
- Rachel R Chen, Robin O Roundy, Rachel Q Zhang, and Ganesh Janakiraman. 2005. Efficient auction mechanisms for supply chain procurement. *Management Science* 51, 3 (2005), 467–482.
- E. H. Clarke. 1971. Multipart pricing of public goods. *Public Choice* 2 (1971), 17–33.
- Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. 2016. Approximately efficient double auctions with strong budget balance. In *SODA*. 1424–1443.
- Riccardo Colini-Baldeschi, Paul W Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, and Stefano Turchetta. 2017. Approximately efficient two-sided combinatorial auctions. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 591–608.
- Jon Feldman, Vahab S. Mirrokni, S. Muthukrishnan, and Mallesh M. Pai. 2010. Auctions with intermediaries: extended abstract. In *ACM Conference on Electronic Commerce*. 23–32.
- Moran Feldman, Gonen Frim, and Rica Gonen. 2018. Multi-sided Advertising Markets: Dynamic Mechanisms and Incremental User Compensations. In *GameSec*. 227–247.
- Moran Feldman and Rica Gonen. 2018. Removal and Threshold Pricing: Truthful Two-Sided Markets with Multi-dimensional Participants. In *SAGT*. 163–175.
- Matthias Gerstgrasser, Paul W Goldberg, Bart de Keijzer, Philip Lazos, and Alexander Skopalik. 2019. Multi-unit bilateral trade. In *Proceedings of the AAAI'19*, Vol. 33. 1973–1980. arXiv preprint 1811.05130.
- Mira Gonen, Rica Gonen, and Pavlov Elan. 2007. Generalized trade reduction mechanisms. In *Proceedings of EC'07*. 20–29.

- Rica Gonen and Ozi Egri. 2017. DYCOM: A Dynamic Truthful Budget Balanced Double-sided Combinatorial Market. In *Proceedings of AAMAS'17*. 1556–1558.
- T. Groves. 1973. Incentives in teams. *Econometrica* 41 (1973), 617–631.
- Mingyu Guo, Evangelos Markakis, Krzysztof R. Apt, and Vincent Conitzer. 2013. Undominated Groves Mechanisms. *J. Artif. Intell. Res. (JAIR)* 46 (2013), 129–163.
- Shengwu Li. 2017. Obviously strategy-proof mechanisms. *American Economic Review* 107, 11 (2017), 3257–87.
- Benjamin Lubin, Adam I Juda, Ruggiero Cavallo, Sébastien Lahaie, Jeffrey Shneidman, and David C Parkes. 2008. ICE: An expressive iterative combinatorial exchange. *Journal of Artificial Intelligence Research* 33 (2008), 33–77.
- R. Preston McAfee. 1992. A dominant strategy double auction. *Journal of Economic Theory* 56, 2 (April 1992), 434–450.
- R. P. McAfee. 2008. The Gains from Trade Under Fixed Price Mechanisms. *Applied Economics Research Bulletin* 1 (2008).
- Roger B. Myerson and Mark A. Satterthwaite. 1983. Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29, 2 (April 1983), 265–281.
- Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. 2016. SBBA: A Strongly-Budget-Balanced Double-Auction Mechanism. In *Algorithmic Game Theory*, Martin Gairing and Rahul Savani (Eds.). Lecture Notes in Computer Science, Vol. 9928. Springer Berlin Heidelberg, 260–272. https://doi.org/10.1007/978-3-662-53354-3_21
- Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. 2018a. MUDA: A Truthful Multi-Unit Double-Auction Mechanism. In *Proceedings of AAAI'18*. AAAI Press. arXiv preprint 1712.06848.
- Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. 2018b. Double Auctions in Markets for Multiple Kinds of Goods. In *Proceedings of IJCAI'18*. AAAI Press. Previous name: "MIDA: A Multi Item-type Double-Auction Mechanism". arXiv preprint: 1604.06210.
- W. Vickrey. 1961. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance* 16 (1961), 8–37.

Rica Gonen
The Open University of Israel
Raanana, Israel
Email: ricagonen@gmail.com

Erel Segal-Halevi
Ariel University
Ariel, Israel
Email: erelsgl@gmail.com