

Double Auctions in Markets for Multiple Kinds of Goods

Erel Segal-Halevi¹, Avinatan Hassidim,² Yonatan Aumann²

¹ Ariel University, Ariel 40700, Israel

² Bar-Ilan University, Ramat-Gan 52900, Israel

erelsgl@gmail.com, avinatanh@gmail.com, yaumann@gmail.com

Abstract

Motivated by applications such as stock exchanges and spectrum auctions, there is a growing interest in mechanisms for arranging trade in two-sided markets. Existing mechanisms are either not truthful, or do not guarantee an asymptotically-optimal gain-from-trade, or rely on a prior on the traders’ valuations, or operate in limited settings such as a single kind of good. We extend the random market-halving technique used in earlier works to markets with multiple kinds of goods, where traders have gross-substitute valuations. We present MIDA: a Multi Item-kind Double-Auction mechanism. It is prior-free, truthful, strongly-budget-balanced, and guarantees near-optimal gain from trade when market sizes of all goods grow to ∞ at a similar rate.

1 Introduction

The recent auction for radio spectrum reallocation [Leyton-Brown *et al.*, 2017] has incited a surge of interest in *double auction* mechanisms — mechanisms for arranging trade in two-sided markets. Such markets differ from one-sided markets in that the valuations of both buyers and sellers are their private information, and both sides might act strategically.

An important requirement from a double-auction is *efficiency*, which is measured by its *gain-from-trade* (GFT) — the total value gained by the buyers minus the total value contributed by the sellers. As an example, consider a stock market with two sellers: Alice holds a unit of stock x which she values as 41 and Bob holds a unit of stock y which he values as 46. There is a single buyer, George, who wants a unit of a single stock which can be either x or y ; he values x as 47 and y as 48. Then, if George buys x from Alice the GFT is $47 - 41 = 6$ while if he buys y from Bob the GFT is $48 - 46 = 2$. Therefore, an optimal double-auction mechanism would make the former trade. If the mechanism makes the wrong trade, it attains only $1/3$ of the optimal GFT.¹

A common double-auction mechanism is the *Walrasian mechanism* [Rustichini *et al.*, 1994; Babaioff *et al.*, 2014].

¹Note that the *social welfare* — the sum of agents’ valuations — is much easier to approximate than the GFT. E.g, in the above example, the wrong deal attains $\frac{48+41}{47+46} > 95\%$ of the optimum.

It computes an *equilibrium price-vector* — a price for each good, at which the supply of each good equals its demand: the total number of units that sellers want to sell at this price equals the total number of units that buyers want to buy at this price. This mechanism attains the maximum GFT [Nisan *et al.*, 2007, Th. 11.13]. Unfortunately, it is not *truthful* — traders may gain by reporting false valuations.

In fact, in a two-sided market, *any* truthful mechanism that attains the maximum GFT cannot be *budget-balanced* (BB) — the market-maker has to subsidize the trade [Myerson and Satterthwaite, 1983]. Therefore, it is interesting to develop double-auction mechanisms that are truthful, BB and attain an *approximately* maximal GFT. We define the *competitive ratio* of a mechanism as the minimum ratio (over all utility profiles) of its GFT divided by the optimal GFT. The first truthful and BB approximation mechanism was presented in the seminal work of McAfee [1992]: Its competitive-ratio is $1 - 1/k$, where k is the number of units traded in the optimal situation. (i.e, its GFT is always at least $1 - 1/k$ of the maximum GFT). Thus, McAfee’s mechanism is *asymptotically optimal* — when the market-size k grows to infinity, the GFT approaches the maximum. Another advantage of McAfee’s mechanism is that it is *prior-free* (PF) — it does not require any probabilistic knowledge on traders’ valuations; it works even for adversarial valuations. Its main drawback is that it works only in a *single-good* market when each trader can trade a *single unit*.

Currently, as far as we know, no truthful mechanism attains asymptotic optimality in a two-sided market with multiple kinds of goods (see related work in §1.3). Recently, Segal-Halevi *et al.* [2018a] presented a truthful, BB, PF and asymptotically-optimal mechanism for a *single-good* market that allows *many units* per trader. It is called *MUDA* and is based on a random-halving scheme:

Split the market into two sub-markets, left and right, by sending each trader to each side with probability $1/2$, independently of the others. Then, in each sub-market:

1. Calculate a Walrasian equilibrium price (p^R at the right, p^L at the left).
2. Let the traders trade at the price from the other market (p^L at the right, p^R at the left).

The competitive ratio of MUDA is $1 - O(M\sqrt{\frac{\ln k}{k}})$, where k is again the total number of units traded in the optimal situation and M is the maximum number of units per trader.

The goal of this paper is to break the single-good barrier and attain truthful, BB, PF and asymptotic optimality in a multi-good market. We show that it is possible to extend the random-halving scheme of MUDA to multi-good markets. We call the extended version MIDA — Multiple Item-kind Double Auction. However, the extension is not trivial and requires several additional assumptions, presented below.

1.1 Truthfulness in a Multi-Good Market

First, since *step 1* of MUDA calculates Walrasian equilibria, extending it requires that a Walrasian equilibrium exists. Gul and Stacchetti [1999] prove that a sufficient condition for the existence of a Walrasian equilibrium is that all traders' valuations satisfy a condition called *gross-substitute (GS)*.

The GS condition means that, if an agent wants to buy a good x and the price of another good y increases, the agent still wants to buy x (so the demand for good x weakly increases). A typical example of a GS market is a market for used cars: usually, people consider cars of different models to be substitutes, so when the price of one model increases, the demand for other models weakly increases.

Second, *step 2* of MUDA lets traders trade in a price that is not an equilibrium price at their own market, so it must handle excess demand and supply. In a single-good market this is easy since there is excess only in one side of the market, so it can be solved using randomization. In a multi-good market this is harder since in each side we can have excess demand in some goods and excess supply in other goods.

Currently we know to solve this problem for markets in which only one of the two sides can trade multiple goods. For concreteness, we assume that the buyers can trade multiple goods while each seller sells multiple units of a single good.

Under these two assumptions, MIDA works as follows.

Halve the market randomly (like MUDA). In each half:

1. Calculate a Walrasian equilibrium price-vector — a price for each good-kind (\mathbf{p}^R at the right, \mathbf{p}^L at the left).

2. Let the traders trade at the price-vector from the other market (\mathbf{p}^L at the right, \mathbf{p}^R at the left) as follows:

- For each good-kind x , put all the sellers of x in a queue, ordered randomly. For each x -seller, calculate the number of x -units that maximize his gain at the market price-vector; sum all these numbers to get the aggregate supply of x .
- Put all buyers in a single queue, ordered randomly. For each buyer i in the buyers' queue, calculate the bundle that maximizes his net gain at the market price-vector, among all bundles that are contained in the aggregate supply. Let the buyer buy this bundle from the first sellers in the sellers' queues. Whenever a first seller exhausts his optimal supply, he leaves the market and the next seller in this queue becomes first.

Theorem 1. *Suppose each seller sells multiple units of a single good, and all traders have gross-substitute valuations. Then, MIDA is strongly budget-balanced (= has no deficit and no surplus) and truthful. Moreover, these properties hold ex-post — for every outcome of the randomization.*

Proof. Strong budget-balance is obvious since all monetary transfers are between buyers and sellers. We now prove that this mechanism is truthful for every randomization outcome.

The buyers effectively play random-serial-dictatorship, which is known to be truthful. Whenever it's a buyer's turn to buy a bundle, the mechanism picks for him the best possible bundle given the available supply.

The sellers have GS valuations. Gul and Stacchetti [1999] prove that GS valuations are *submodular*. This means that a seller gains the highest utility from selling the first unit, a lower utility from selling the second unit, etc. Hence, for each seller there is an optimal number of units to sell in the given market-prices, and it is always optimal for a seller to sell as many units as possible up to the optimal number, then leave. This is exactly what the mechanism does for them. □

Remark. (a) MIDA is truthful only for single-good sellers. Consider a seller who can sell multiple goods. For example, an additive seller who, in the market prices, gains 5 from selling x and 3 from selling y (and 8 from selling both). MIDA has to let such a seller participate simultaneously in both seller-queues. Suppose there is a single unit-demand buyer who gains 6 from buying x and 8 from buying y . If the seller reports truthfully then the buyer buys y . But the seller can strategize by claiming that he loses from selling y (i.e., standing only in the queue of x -sellers); then the buyer will be forced to buy x and the seller will gain more.

(b) MIDA is truthful only for submodular sellers. Consider a non-submodular seller. For example, a seller who values a single unit as 10 and two units as 40. If the market price is 25 per unit, then the seller gains -5 from selling a single unit and +10 from selling two units. Suppose the first buyer in the buyers' queue wants a single unit. The best response of the seller depends on the preferences of the following buyers: if another buyer wants a unit then it is best to truthfully report a supply of 2, but if no other buyer wants a unit then it is best to untruthfully report a supply of 0.

1.2 Asymptotic Optimality in a Multi-Good Market

Asymptotic optimality means that, when the market is sufficiently large, the GFT approaches its maximum value. We have to define what is meant by "large market". The definition implied by McAfee [1992] is that k — the number of units traded in the optimal situation — approaches infinity. This definition cannot be used as-is in a multi-good market. As an extreme example, suppose the market has two independent goods x and y , where most GFT comes from trading a single unit of x , but the market-size of y approaches infinity. It is clear that the competitive ratio will not approach 1 in this case. Therefore, we extend the "large market" definition as follows. Let G be the number of different good-kinds. For any good-kind $x \in \{1, \dots, G\}$, let k_x be the total number of

units of x traded in the optimal situation. A large market is a market where for all x , $k_x \rightarrow \infty$. To ensure that not all these units come from few large traders, we also assume that the number of units traded by a single person is bounded by a constant M (a similar assumption is used in MUDA).

Surprisingly we found out that, even in a large market, the random halving technique might attain sub-optimal GFT. This can happen when the market-size in one good is much larger than the market-size in another good. Appendix A shows a concrete example where $G = 2$ and $M = 1$ but the GFT is at most $63/64$ of the maximum even if $\forall x : k_x \rightarrow \infty$.

Intuitively, The random halving process creates an imbalance in the number of traders in the two market-halves. Standard tail bounds tell us that, with high probability, the imbalance is in the order of square-root of the number of traders sampled. With one good, this imbalance is negligible as $\sqrt{k}/k \rightarrow 0$ when $k \rightarrow \infty$. However, if we have two goods x and y and $\sqrt{k_x} > k_y$, then the sample deviation in x might destroy all the trade in y .

The example in Appendix A is interesting, since the random-halving technique has been used successfully in various domains [Goldberg *et al.*, 2001; 2006; Devanur *et al.*, 2015; Balcan *et al.*, 2008; 2007; Devanur and Hayes, 2009], and one could think it is an omnipotent technique. Our example shows the limitations of random-halving and illustrates the inherent difficulties in multi-good markets.

In contrast to this negative example, we prove that MIDA is asymptotically-optimal when the market sizes in all goods go to infinity in a similar rate. We define $k_{\min} := \min_x k_x$ and $k_{\max} := \max_x k_x$ and $c := k_{\max}/k_{\min}$ = the largest ratio between the market-sizes of different goods. Then:

Theorem 2. *Suppose all traders have GS valuations. Then the expected competitive ratio of MIDA is at least:*

$$1 - 18 \cdot M^{2G+2} \cdot c \cdot \frac{\ln(10Gk_{\max})}{\sqrt{k_{\max}}} - o(1/k_{\max})$$

where G is the number of good-kinds, M the maximum number of units of a single good traded by a single person, k_{\max} the largest market-size of a single good, and $c = k_{\max}/k_{\min}$ = the largest ratio between market-sizes.

Note that previously it was not known if asymptotic optimality is possible even with two goods ($G = 2$). Therefore, it is encouraging that we can get asymptotic optimality for a constant number of goods.

The assumption that c is bounded means that, as the population grows, the market-sizes in the different good-kinds grow at a similar rate. It is reasonable if we believe that the ratio between market-sizes depends on some characteristics of the population that remain approximately constant when the population grows. This is analogous to the *thickness* assumption introduced e.g. by Kojima and Pathak [2009] in the context of two-sided matching.

It is possible to replace c with a different parameter — h — the ratio of the largest to the smallest marginal gain from a deal in a single item. We assume that the traders' valuations are normalized such that, whenever a seller sells an item to a buyer, the marginal GFT contributed by this deal is between 1 and h . This assumption comes from the random-halving

literature (e.g. Goldberg *et al.* [2001]). There, it is common to assume that the valuation of each buyer with positive valuation is bounded in $[1, h]$, so that each buyer contributes at most h to the total welfare. Our definition is the natural generalization of this assumption to a two-sided market.

Theorem 3. *Suppose all traders have GS valuations. Then the expected competitive ratio of MIDA is at least:*

$$1 - 18 \cdot M^{2G+2} \cdot G \cdot h \cdot \frac{\ln(10Gk_{\max})}{\sqrt{k_{\max}}} - o(1/k_{\max})$$

where G is the number of good-kinds, M is the maximum number of units of a single good traded by a single person, and the marginal gain from each unit traded is in $[1, h]$.

So when M and G are bounded, and either c or h (or both) are bounded, and $k_{\max} \rightarrow \infty$, The GFT of MIDA approaches its maximum value.

Note that MIDA does not have to know the parameters used in Theorems 2 and 3 (M, G, c, h, k_{\max} , etc.); they are used only in the analysis. The mechanism itself is prior-free.

Note also that Theorems 2 and 3 do not need the assumption that each seller sells a single good-kind. This assumption is needed only for the truthfulness (see Theorem 1 and the following remark).

Paper layout. Related work is surveyed below. Formal model is presented in Section 2. Theorems 2 and 3 are proved in Section 3. Section 4 discusses the limitations of our results and presents directions for future work.

1.3 Related Work

Mechanisms for multi-unit double auction are surveyed in Segal-Halevi *et al.* [2018a]. Below we survey more advanced mechanisms that handle multi-good markets.

1. Feng *et al.* [2012] present TAHES — Truthful double-Auction for HETerogeneous Spectrum. The mechanism is inspired by the problem of re-allocating radio frequency-spectrum from primary owners to secondary users. In their setting, each seller (primary owner) has a single unit of a unique good-kind, each buyer (secondary user) has a different valuation for each kind, and each buyer wants a single unit. This is similar to a special case of our setting, in which all buyers have unit-demand valuations. The authors do not analyze the competitive ratio of their mechanism.

2. Blumrosen and Dobzinski [2014] present two *Combinatorial Reallocation* mechanisms for multi-good markets. The first one attains a competitive ratio of $1/[8 \cdot \Theta(\log M)]$ where M is the maximum number of items per single seller, but it is not prior-free since it needs to know the median value of the initial endowment of each seller. The second one is prior-free and attains a competitive ratio of $1/48$. Both competitive ratios do not approach 1.

3. Gonen *et al.* [2007] present a general scheme for converting a truthful mechanism with deficit to a truthful mechanism without deficit. Their scheme can handle combinatorial double auctions, but only when traders are single-valued, e.g. when each trader has a set of desired bundles and values all these bundles the same.

4. Chu and Shen [2006] present a double-auction mechanism that can handle multiple goods. However, its truthfulness depends on substitutability conditions between buyers to

buyers, sellers to sellers, and buyers to sellers. It is not clear whether these conditions hold in our setting. Moreover, they do not analyze the competitive ratio in a multi-good market.

5. Colini-Baldeschi *et al.* [2017b] present a combinatorial double-auction mechanism in a Bayesian setting (not prior-free). In contrast to our work, they approximate the social welfare rather than the gain-from-trade. Note that *any* mechanism that attains a fraction α of the optimal GFT also attains a fraction of at least α of the optimal social welfare [Brustle *et al.*, 2017; Colini-Baldeschi *et al.*, 2017a], but the opposite is not true. Moreover, their competitive ratios (1/4, 1/6 or 1/16) do not approach 1.

6. Gonen and Egri [2017] present a combinatorial dynamic double-auction mechanism that relies on partial prior information — the maximum and minimum values of the traders. They, too, approximate the social welfare rather than the GFT: they show by simulations that their mechanism attains about 0.5 of the maximum social welfare.

7. Hirai and Sato [2017] present a double-auction mechanism for budgeted buyers, that is Pareto-optimal and truthful for the buyers, but not truthful for the sellers.

2 Model

Traders and Valuations. We consider a market in which some traders, the “sellers”, are endowed with discrete goods, and other traders, the “buyers”, are endowed with unlimited money. There are G different kinds of goods and each trader can hold at most M units of each good, for some constants G and M . A *bundle* is a vector $X \in \{0, \dots, M\}^G$. The empty bundle is denoted by $\emptyset = (0, \dots, 0)$. A *price-vector* is a vector $\mathbf{p} = (p_1, \dots, p_G)$. The price of a bundle X is $\mathbf{p} \cdot X$.

Each trader i has a value function v_i on bundles, normalized such that $v_i(\emptyset) = 0$. All traders’ utilities are quasi-linear in money, so the net gain of a buyer i from buying a bundle X at price \mathbf{p} is $u_i(X, \mathbf{p}) = v_i(X) - \mathbf{p} \cdot X$.

For each seller j with initial endowment E_j , we define a *sale-value function* v'_j on negative bundles: $v'_j(-X) := v_j(E_j) - v_j(E_j - X)$. Hence the net gain of seller j from selling a bundle X at price-vector \mathbf{p} is $u_j(-X, \mathbf{p}) = \mathbf{p} \cdot X - v'_j(-X)$. Note that the gain from no trade is 0 both for a buyer and a seller: $\forall i, \mathbf{p} : u_i(\emptyset, \mathbf{p}) = 0$.

We present the gain of a trader as a sum of *marginal gains*, in which there is a term for each unit traded, ordered from good 1 to good G . For example, for a buyer who buys two units of good 1 and one unit of good 2, the net gain is:

$$\begin{aligned} u_i([2, 1], \mathbf{p}) = & u_i([1, 0], \mathbf{p}) + \\ & [u_i([2, 0], \mathbf{p}) - u_i([1, 0], \mathbf{p})] + \\ & [u_i([2, 1], \mathbf{p}) - u_i([2, 0], \mathbf{p})]. \end{aligned}$$

We present each trader as a set of at most MG *virtual traders*, each of whom trades a single unit and enjoys the marginal utility of that unit (a similar idea was used by Chawla *et al.* [2010]). So the above buyer represents three virtual buyers: the first buys good 1 and gains $u_i([1, 0], \mathbf{p})$, the second buys good 1 and gains $u_i([2, 0], \mathbf{p}) - u_i([1, 0], \mathbf{p})$, and the third buys good 2 and gains $u_i([2, 1], \mathbf{p}) - u_i([2, 0], \mathbf{p})$.

Demand and Supply. We assume that traders’ valuations are *generic*, i.e. values of different bundles and of different traders are different and linearly-independent over the integers (no linear combination with integer coefficients equals zero). This assumption guarantees that, given a price-vector \mathbf{p}^R determined in the right market, every trader i in the left market has a unique bundle that maximizes $u_i(\cdot, \mathbf{p})$, and vice versa.² We call this maximizing bundle the *demand* of i at \mathbf{p} . If i is a seller then his demand is a weakly-negative vector; we call the absolute value of this vector the *supply* of i at \mathbf{p} .

Gross Substitutes. We assume that all traders’ valuations are *gross-substitute* (GS); see Kelso and Crawford [1982] and [Gul and Stacchetti, 1999] for formal definitions. Intuitively, GS means that there are no complementarities between different items, so that when the price of a single unit of a single good increases, the demands for the other units and the other goods do not decrease. For example, suppose at price-vector \mathbf{p} the demand of some buyer is $[2, 1]$, e.g. two units of good x and one unit of good y . Suppose that the price of the first unit of x increases while the prices of the other units remain fixed. GS means that the buyer will still want to buy at least one unit of x and one unit of y .

When all traders have GS valuations, a Walrasian equilibrium exists and can be found efficiently using an ascending auction [Gul and Stacchetti, 2000; Ben-Zwi *et al.*, 2013]. Intuitively, the GS property guarantees that the auctioneer can increase the prices of over-demanded items, without creating a demand-shortage in other items, until the market converges to an equilibrium.

GS valuations are always *submodular*; see Gul and Stacchetti [1999] for a formal definition and a proof. Intuitively, submodularity means that, if a trader loses some deals, its marginal gain from the other deals weakly increases. For example, if the buyer from above arrives at the head of the buyers’ queue when only one unit of good 1 is available for sale, then its gain from this unit will be $u_i([1, 0], \mathbf{p})$ which is at least as large as $u_i([2, 0], \mathbf{p}) - u_i([1, 0], \mathbf{p})$, and its additional gain from the unit of good 2 will be $u_i([1, 1], \mathbf{p}) - u_i([1, 0], \mathbf{p})$ which is at least as large as $u_i([2, 1], \mathbf{p}) - u_i([2, 0], \mathbf{p})$.

Mechanisms. A *mechanism* is a (randomized) function that takes the traders’ valuations and returns (1) a price-vector \mathbf{p} , (2) for each buyer(seller) i , the bundle X_i to buy(sell) at \mathbf{p} . A mechanism is *materially-balanced* if for each good, the number of units bought equals the number of units sold: $\sum_{i \in \text{Buyers}} X_i = \sum_{j \in \text{Sellers}} X_j$.

A mechanism is *ex-post truthful*, if a trader can never increase his gain by pretending to have different valuations, even when the trader knows in advance the outcomes of all randomizations done by the mechanism.

Given a trading scenario in which each trader i buys/sells a bundle X_i , the *gain-from-trade* (GFT) is the sum of gains of

²The genericity assumption can be dropped by using centralized tie-breaking, i.e. whenever a trader has two or more demands, the market-manager may select one of them in a way that maximizes GFT. See Hsu *et al.* [2016] for other ways to handle ties in markets.

all traders: $GFT := \sum_{i \in \text{AllTraders}} u_i(X_i)$.³

In a materially-balanced setting, the GFT does not depend on the price-vector, since the prices are canceled in the summation; money is only transferred between buyers and sellers.

3 Competitive-Ratio Analysis

For truthfulness we had to assume that, in one side of the market, each trader specializes in a single good. In the present section we do not need this assumption — both buyers and sellers can trade multiple goods and are entirely symmetric. We first analyze the optimal trade, then the right sub-market and finally the left sub-market.

3.1 Optimal Trade

By the gross-substitute assumption, there exists a Walrasian equilibrium in the global population, and it attains the maximum GFT (see Section 2). By the genericity assumption, the equilibrium allocation is unique. We call it the “optimal trade” and denote any equilibrium price-vector by \mathbf{p}^O .

For any good $x \in \{1, \dots, G\}$, denote by B_{x^*} the set of virtual buyers and by S_{x^*} the set of virtual sellers participating in the optimal trade. By definition of equilibrium, the numbers of virtual-traders in both groups are the same; this is the number we denoted by k_x :

$$\forall x \in \{1, \dots, G\} : |B_{x^*}| = |S_{x^*}| = k_x \quad (1)$$

We call these virtual traders the *efficient traders*. We make the pessimistic assumption that all GFT in the sub-markets comes only from these efficient traders. Therefore, the GFT of our mechanism depends on the numbers of efficient traders that trade in each sub-market.

The reduction in GFT has two reasons: one is the *sampling error* — efficient buyers and sellers land in different sub-markets, so they do not meet and cannot trade. This error is easy to bound using standard tail inequalities. The second reason is the *pricing error* — the prices at the sub-market might be too high or too low, which might create imbalance in the demand and supply. Analyzing this error requires careful analysis of the equilibrium in the optimal situation vs. the equilibrium in each sub-market.

3.2 Right Sub-Market

In the right, we calculate an equilibrium price-vector \mathbf{p}^R . For each good $x \in \{1, \dots, G\}$, define the sets of virtual-traders:

- B_{x-} is the set of efficient virtual buyers of x (members of B_{x^*}) who do not demand x at \mathbf{p}^R .
- S_{x-} is the set of efficient virtual-sellers (members of S_{x^*}) who do not supply x at \mathbf{p}^R .
- B_{+x} is the set of inefficient virtual-buyers (not members of B_{x^*}) who demand x at \mathbf{p}^R .
- S_{+x} is the set of inefficient virtual-sellers (not members of S_{x^*}) who supply x at \mathbf{p}^R .

³In some papers, the gain-for-trade is termed *trader surplus* or *price improvement* [Chakraborty et al., 2015].

We denote $D_{+x} := B_{+x} \cup S_{x-}$ and $D_{x-} := B_{x-} \cup S_{+x}$. These sets D_{+x}, D_{x-} represent the pricing error — they are responsible for the loss of welfare due to efficient traders quitting or inefficient traders competing with the efficient ones. The goal of our analysis below is to bound the sizes of the sets D_{x-} and D_{+x} .

For any set T of traders, denote by T^R the subset of T that is sampled to the right market and by T^L the subset of T sampled to the left market. By definition of the equilibrium price-vector \mathbf{p}^R , for every good x :

$$|B_{x^*}^R| - |B_{x-}^R| + |B_{+x}^R| = |S_{x^*}^R| - |S_{x-}^R| + |S_{+x}^R| \quad (2)$$

In order to relate (1) and (2), we have to relate $B_{x^*}^R, S_{x^*}^R$ to B_{x^*}, S_{x^*} . This is be done using the following lemma.

Lemma. For every set T of virtual-traders and integer $q > 0$

$$\text{w.p. } 1 - 2 \exp\left(\frac{-2q^2}{M^2|T|}\right) : \left| |T^R| - \frac{|T|}{2} \right| < q \quad (3)$$

(“w.p. x ” is a shorthand to “with probability at least x ”).

The lemma is proved using Hoeffding’s inequality. The proof is standard and we omit it.

We apply this lemma $2G$ times, to B_{x^*} and S_{x^*} for each $x \in \{1, \dots, G\}$, and combine the outcomes using the union bound. This gives, $\forall q > 0$:

$$\text{w.p. } 1 - 4G \exp\left(\frac{-2q^2}{M^2k_{\max}}\right) : \quad (4)$$

$$\forall x : \left| |B_{x^*}^R| - k_x/2 \right| < q \text{ and } \left| |S_{x^*}^R| - k_x/2 \right| < q$$

Combining equations 1, 2 and 4 gives, $\forall q > 0$:

$$\text{w.p. } 1 - 4G \exp\left(\frac{-2q^2}{M^2k_{\max}}\right) : \quad (5)$$

$$\forall x : \left| |B_{x-}^R \cup S_{+x}^R| - |B_{+x}^R \cup S_{x-}^R| \right| < 2q$$

$$\implies \left| |D_{x-}^R| - |D_{+x}^R| \right| < 2q$$

Our goal now is to get a bound on $|D_{+x}|, |D_{x-}|$. We proceed in several steps.

Step A: From Bound on Difference to Bound on Sets.

Below we fix some $q > 0$ and assume that the suffix of inequality (5) holds. This gives us an upper bound on $\left| |D_{x-}^R| - |D_{+x}^R| \right|$. We need upper bounds on $|D_{x-}^R|$ and $|D_{+x}^R|$. To get it, we use a property of GS valuations called *Downward Demand Flow (DDF)* [Segal-Halevi et al., 2016]. Consider two price-vectors \mathbf{p}^O and \mathbf{p}^R , and let $\Delta := \mathbf{p}^R - \mathbf{p}^O$ = the vector of price-increases. Intuitively, DDF means that a trader’s demand goes down the price-change ladder — a trader switches from demanding good x at price \mathbf{p}^O to demanding good y at price \mathbf{p}^R , only if $\Delta_x > \Delta_y$.

Lemma (DDF Lemma, Segal-Halevi et al. [2016]). Every trader with gross-substitute valuations has the following property. Whenever the price-vector changes from \mathbf{p}^O to \mathbf{p}^R , the trader’s demand changes in the following way.

(1) If the trader stopped demanding some good x with $\Delta_x \leq 0$, then he started demanding a good y with $\Delta_y < \Delta_x$.

(2) If the trader started demanding some good x with $\Delta_x \geq 0$, then he stopped demanding a good y with $\Delta_y > \Delta_x$.

By combining the DDF lemma and (5), we can prove:

$$\text{w.p. } 1 - 4G \exp\left(\frac{-2q^2}{M^2 k_{\max}}\right) : \quad (6)$$

$$\forall x \in \{1, \dots, G\} : |D_{x-}^R| < 2^G \cdot q \ \& \ |D_{+x}^R| < 2^G \cdot q$$

The proof is in the full version [Segal-Halevi *et al.*, 2018b].

Step B: From a Bound on Sampled Sets to a Bound on Their Parent Sets. Our goal is now to derive from (6) an upper bound on $|D_{x-}|$ and $|D_{+x}|$ — the number of virtual traders that change their demand or supply due to the pricing error. These two sets are entirely analogous; we focus on D_{x-} . First, we derive from (3) the following, by substituting $q \rightarrow |T|/4$:

$$\forall \text{ set } T : \text{w.p. } 1 - 2 \exp\left(\frac{-|T|}{8M^2}\right) : |T| < 4|T^R| \quad (7)$$

We would like to apply (7) to D_{x-} and conclude that $|D_{x-}| < 2^G 4q$. But we cannot do so directly since D_{x-} is a random-set — it depends on the random-sampling through \mathbf{p}^R — while T in (7) must be a deterministic set — independent of the random-sampling. Our solution is to use the union bound. Let E_q be the set of all possible values of D_{x-} with cardinality $2^G 4q$ (so E_q is a set of sets of virtual traders). We apply (7) simultaneously to all sets in E_q using the union bound. We get, for every integer q :

$$\text{w.p. } 1 - 2|E_q| \exp\left(\frac{-2^G 4q}{8M^2}\right) : \forall T \in E_q : 2^G q < |T^R|$$

The same arguments are true for D_{+x} and for every good $x \in \{1, \dots, G\}$. Combining this with (6) gives, for every q :

$$\text{w.p. } 1 - 4G \exp\left(\frac{-2q^2}{M^2 k_{\max}}\right) - 4G \exp\left(\frac{-\ln|E_q| \cdot 2^G 4q}{8M^2}\right) : \quad (8)$$

$$\forall x \in \{1, \dots, G\} : |D_{x-}| < 2^G 4q \ \& \ |D_{+x}| < 2^G 4q$$

Step C: Bounding the Number of Possible Sets. Our goal now is to find an upper bound on $|E_q|$. In fact, we will bound $|E|$ — the number of possible values of D_{x-} , regardless of their cardinality. Initially, we bound $|E|$ as a function of n — the total number of real traders that come to the auction. Then, we replace n by a function of k_{\max} .

We present the set D_{x-} as a function of simpler sets. For every two bundles Y and Z , denote by $R_{Y \succ Z}$ the set of traders who, at price-vector \mathbf{p}^R , prefer bundle Y to bundle Z . $R_{Y \succ Z}$ is a *one dimensional* random set: it contains all the traders i for whom:

$$v_i(Y) - v_i(Z) > \mathbf{p}^R \cdot Y - \mathbf{p}^R \cdot Z$$

Therefore, for every integer $t \in \{1, \dots, n\}$, there is exactly one possible value of $R_{Y \succ Z}$ with cardinality t — it is the set of t traders for whom the difference $v_i(Y) - v_i(Z)$ is the highest among all n traders. Therefore the random set $R_{Y \succ Z}$ has n possible values. The total number of sets $R_{Y \succ Z}$ is at most the number of bundle-pairs, which is at most $(M^G)^2 = M^{2G}$. Hence there are at most $n^{M^{2G}}$ ways to choose these

sets. Once all the sets $R_{Y \succ Z}$ are determined, the set D_{x-} is completely determined too. Therefore:

$$|E| < n^{M^{2G}} \quad (9)$$

We now want to replace n by an expression that depends only on k_{\max} . The reason is that n — the total number of traders in the market — might be very large: in theory, every person on Earth might come to the auction as a “trader” without buying or selling anything. Therefore, we cannot hope to get any meaningful approximation factor based on n . Indeed, the approximation factors in the prior-free double-auction literature since McAfee [1992] are given as a function of the number of optimal deals (k), which is a better indicator of the actual market size.

Define a *relevant trader* as a trader who trades at least one unit of at least one good in at least one price-vector — \mathbf{p}^O or \mathbf{p}^R . Our arguments for bounding $|E|$ clearly hold when we use the number of relevant traders instead of n . We will now prove that, with high probability, the number of relevant traders is at most n' , where $n' := 10Gk_{\max}$. Then we will use this bound in (9) and get that $|E| < (n')^{M^{2G}}$.

Lemma. *w.p. $1 - 2 \exp\left(\frac{-2G^2 k_{\max} \ln(8Gk_{\max})}{M^2}\right)$, the number of relevant traders is at most $10Gk_{\max}$.*

Proof. [Proof sketch] By definition, the number of traders who want to trade at \mathbf{p}^O is at most $2Gk_{\max}$. It remains to count the traders who trade nothing at \mathbf{p}^O , but want to trade something at \mathbf{p}^R . Let D_0 be the set of these traders.

Using the DDF lemma, we can prove that their number in the right market is at most the number of traders who want to trade at \mathbf{p}^O , so $|D_0^R| \leq 2Gk_{\max}$. From this we can infer that:

$$\text{w.p. } 1 - 2 \exp\left(\frac{-2G^2 k_{\max} \ln(8Gk_{\max})}{M^2}\right) : |D_0| < 8Gk_{\max}$$

so the total number of relevant traders is at most $10Gk_{\max} = n'$. The proof details are deferred to the full version [Segal-Halevi *et al.*, 2018b]. \square

Substituting $n \rightarrow n'$ in (9) and (9) in (8) gives the failure probability:

$$4G \exp\left(\frac{-2q^2}{M^2 k_{\max}}\right) + 4G \exp\left(M^{2G} \ln n' - \frac{2^G q}{2M^2}\right) + 2 \exp\left(\frac{-2G^2 k_{\max} \ln(8Gk_{\max})}{M^2}\right)$$

Denote this probability by $\mathcal{P}\mathcal{F}(q)$. So by (8), w.p. $1 - \mathcal{P}\mathcal{F}(q)$:

$$\forall x \in \{1, \dots, G\} : |D_{x-}| < 2^G 4q \ \& \ |D_{+x}| < 2^G 4q$$

3.3 Left Sub-Market

For every good x , the set of efficient virtual-buyers in the left market is $B_{x*}^L = B_{x*} \setminus B_{x*}^R$. All these virtual-buyers want to trade at price \mathbf{p}^O ; all the virtual-buyers who are not in B_{x-} want to trade at price \mathbf{p}^R too. Therefore, the set of efficient buyers who are sampled to the left market and want to trade at the price-vector posted for the left market is: $B_{x*} \setminus B_{x*}^R \setminus B_{x-}$. Some of these buyers might be “wasted” by trading

with inefficient virtual-sellers; the set of inefficient virtual-sellers who want to trade x at price \mathbf{p}^R is S_{+x} . All in all, the number of efficient virtual-buyers in the left submarket who want to trade at price \mathbf{p}^R and are not wasted on inefficient sellers is at least:

$$\begin{aligned}\mathcal{B}_x^L &:= |B_{x*}| - |B_{x*}^R| - |B_{x-}| - |S_{+x}| \\ &= k_x - |B_{x*}^R| - |D_{x-}|\end{aligned}$$

Similarly, the number of efficient virtual-sellers who want to trade at \mathbf{p}^R and are not wasted on inefficient buyers is at least:

$$\begin{aligned}\mathcal{S}_x^L &:= |S_{x*}| - |S_{x*}^R| - |S_{x-}| - |B_{+x}| \\ &= k_x - |S_{x*}^R| - |D_{+x}|\end{aligned}$$

By inequalities (4) and (8) from the previous subsection, we have for every integer q , with probability $1 - \mathcal{P}\mathcal{F}(q)$:

$$\begin{aligned}\mathcal{B}_x^L &\geq k_x/2 - q - 2^G 4q \\ \mathcal{S}_x^L &\geq k_x/2 - q - 2^G 4q\end{aligned}$$

Thus, in the left market, at least $k_x/2 - q - 2^G 4q$ efficient deals in good x are done. The same analysis is true in the right market. Therefore, the total number of efficient deals in x done in all submarkets is at least $k_x - 2q - 2^G 8q \geq k_x - 2^G 9q$.

From here we proceed to prove Theorems 2, 3.

Approximation Based on c . By the random ordering of the traders in each market, the $k_x - 2^G 9q$ efficient deals executed in the sub-markets are chosen at random from the set of k_x efficient deals. The same is true for every good-kind x . We recall that, by submodularity, if a virtual-trader loses a deal, then the marginal-gain of other virtual-traders of the same real-trader do not decrease. Hence, with probability $1 - \mathcal{P}\mathcal{F}(q)$, the expected competitive ratio is at least $1 - \frac{2^G 9q}{k_{\min}}$. Hence, with probability 1, the expected competitive ratio is at least $1 - \frac{2^G 9q}{k_{\min}} - \mathcal{P}\mathcal{F}(q) =$

$$\begin{aligned}1 - \frac{2^G 9q}{k_{\min}} - 4G \exp\left(\frac{-2q^2}{M^2 k_{\max}}\right) \\ - 4G \exp\left(M^{2G} \ln n' - \frac{2^G q}{2M^2}\right) - o\left(\frac{1}{k_{\max}}\right)\end{aligned}$$

All our analysis so far holds simultaneously for every integer q . Now we substitute $q = \frac{M^{2G+2}}{2^{G-1}} \cdot \ln n \cdot \sqrt{k_{\max}}$. The competitive ratio becomes:

$$\begin{aligned}1 - 18M^{2G+2} \frac{\sqrt{k_{\max}} \ln n'}{k_{\min}} - 4G \exp\left(-2 \frac{M^{4G+2}}{2^{G-1}} \ln^2 n'\right) \\ - 4G \exp\left(-M^{2G} \ln n' [\sqrt{k_{\max}} - 1]\right) - o\left(\frac{1}{k_{\max}}\right)\end{aligned}$$

We now substitute $n' \rightarrow 10Gk_{\max}$ and $k_{\min} \rightarrow k_{\max}/c$. The latter three expressions are clearly in $o(1/k_{\max})$. Therefore MIDA's competitive ratio is at least as claimed in Theorem 2:

$$1 - 18 \cdot M^{2G+2} \cdot c \cdot \frac{\ln(10Gk_{\max})}{\sqrt{k_{\max}}} - o(1/k_{\max})$$

Approximation Based on h . Recall that in each good-kind x , w.p. $1 - \mathcal{P}\mathcal{F}(q)$, at most $2^G 9q$ deals are lost. By assumption, each deal contributes to the GFT at most h . Therefore, the total loss of GFT in all goods together is at most $2^G 9qGh$. On the other hand, by assumption, each deal contributes to the GFT at least 1. Therefore, the optimal GFT in all goods together is at least k_{\max} . Hence, w.p. $1 - \mathcal{P}\mathcal{F}(q)$, the relative loss of GFT is at most $\frac{2^G 9qGh}{k_{\max}}$. Hence, w.p. 1, the expected competitive ratio is at least $1 - \frac{2^G 9qGh}{k_{\max}} - \mathcal{P}\mathcal{F}(q)$. This is exactly the same expression as in the previous paragraph, except that the $\frac{1}{k_{\min}}$ in the denominator is replaced by $\frac{Gh}{k_{\max}}$. Therefore, by the same analysis as above we get that the expected competitive ratio is at least as claimed in Theorem 3:

$$1 - 18 \cdot M^{2G+2} \cdot G \cdot h \cdot \frac{\ln(10Gk_{\max})}{\sqrt{k_{\max}}} - o(1/k_{\max})$$

4 Limitations and Future Work

Our goal in this work was to prove that truthful asymptotic optimality is theoretically possible even with multiple kinds of goods. While we proved this possibility, our work has several limitations that should be handled in future work.

(1) MIDA requires that agents have gross-substitute valuations. We do not know how to handle agents with more general valuations, e.g. submodular. The main problem is that a Walrasian equilibrium might not exist, so we cannot use prices from one market to control trade in the other market.

(2) MIDA's truthfulness (Theorem 1) requires that in one side of the market (e.g. the seller side), each trader specializes in a single good-kind. We do not know how to handle markets where both sides trade multiple good-kinds. The impossibility result in Segal-Halevi and Hassidim [2017] implies that this might be a hard problem even when the prices are fixed exogenously.

(3) MIDA's asymptotic optimality (Theorems 2,3) requires that either the parameter c or the parameter h (or both) are bounded. Even with these assumptions, the competitive ratio guarantee becomes positive only for a very large market. E.g. with $M = 2, G = 2, c = 1$, where the market-size in both goods is k , Theorem 2 guarantees $1 - 1152 \cdot \frac{\ln(20k)}{\sqrt{k}}$, which becomes positive only for $k \approx 750,000,000$.

Often, the real-life performance of a mechanism is much better than its theoretical worst-case guarantee. This was shown empirically for MUDA [Segal-Halevi *et al.*, 2018a] and we believe this should be the case for MIDA as well. However, currently it is very difficult to find data on multi-good two-sided markets, since even the user-interface of such markets (e.g. in stock exchanges) usually accepts only single-good inputs. We hope that, once multi-good double-auction mechanisms become available, two-sided markets will allow multi-good inputs, which will then enable empirical research.

A Failure of Random Halving

This Appendix shows that the common random-market-halving technique might fail to attain an asymptotically-optimal gain-from-trade. In the example, there is a two-good market where the numbers of traders are parametrized

| Name | Value x | Value y | # Total | Optimal behavior | # Right | # Left | Behavior In \mathbf{p}^R : $1 = p_y^R \leq p_x^R \leq 6$ |
|------------------|-----------------------------------|-------------------------------------|----------|------------------|-----------------|-----------------|--|
| B_{yy} | 0 | 9 | $2k^4$ | Buy y | $k^4 - d_1$ | $k^4 + d_1$ | Buy y |
| B_{xy} | 9 | 9 | $2k - 1$ | Buy x | $k - 1 - d_2$ | $k + d_2$ | Buy y |
| B_{xx} | k^{100} | 0 | 1 | Buy x | 0 | 1 | Buy x |
| S_{yy} | - | $1 + i/k^5$ $i = 1, \dots, 2k^4$ | $2k^4$ | Sell y | $k^4 + k + d_3$ | $k^4 - k - d_3$ | $k^4 + k - 1 - d_1 - d_2$ sell y ; 1 is indifferent; all others do nothing. |
| $S_{x\emptyset}$ | $6 + i/k^5$ $i = 1, \dots, 2k$ | - | $2k$ | Sell x | $k + d_4$ | $k - d_4$ | At most 1 is indifferent; all others do nothing. |

Table 1: Data for negative example.

by k . With constant probability (independent of k), random-halving loses almost all the gain-from-trade even when $k \rightarrow \infty$. Hence, the competitive ratio does not approach 1 even when the market-sizes in both goods approach ∞ .

We consider a market for medications. There are two goods (medications) denoted y and x :

- Good y is a medication for headache. It is produced by a group S_{yy} of $2k^4$ sellers who value it as ≈ 1 . More precisely, the valuation of seller i in this subset is $1 + i/k^5$, for $i \in \{1, \dots, |S_{yy}|\}$.
- Good x is a strong medication that can cure both headache and a much more lethal disease. It is produced by a group $S_{x\emptyset}$ of $2k$ sellers, each of whom produces a single unit and values it as ≈ 6 : the valuation of seller i in this subset is $6 + i/k^5$, for $i \in \{1, \dots, |S_{x\emptyset}|\}$.

The groups of sellers are disjoint so all sellers are single-good and single-unit. There are three groups of buyers (patients).

- B_{yy} is a large group of headache patients, each of whom wants a unit of y and values it as 9.
- B_{xx} contains a single patient who has the lethal disease and must get x ; he values it very highly — k^{100} .
- B_{xy} is a group of patients with a mild disease that can be cured by both x and y ; each of them wants a single unit of either x or y and values it as 9.

The valuations and numbers are summarized in Table 1. The first three columns show the trader sets and their valuations. The next two columns show the total number of traders in each group, and what they do in an optimal equilibrium. All sellers sell and all buyers buy; to balance the number of deals in x , all the B_{xy} buyers buy x . The number of efficient deals in *both* good-kinds is an increasing function of k .

The next two columns show the number of traders from each group sampled to each sub-market. Here, the d_j are random variables and the only thing we require about them is that $d_j \geq 0$. The probability of getting such a sample approaches a constant independent of k :

- The probability that in B_{yy} at least half the members fall in the left market ($d_1 \geq 0$) is $1/2$; similarly for B_{xy} , $S_{x\emptyset}$.
- The probability that the member of B_{xx} falls in the left market is $1/2$.

- The probability that in S_{yy} at least $k^4 + k$ members fall in the right market is less than $1/2$; however, because $k \ll \sqrt{2k^4}$, this probability converges to $1/2$ as $k \rightarrow \infty$.

Additionally, we require that lowest-valued seller in $S_{x\emptyset}$ — the seller with value $6 + 1/k^5$ — falls in the right market. The probability for this is $1/2$. All in all, the probability of having such a sample converges to $1/64$ as $k \rightarrow \infty$.

We now calculate the equilibrium price-vector \mathbf{p}^R .

- The number of y -sellers is higher than the total number of potential x -buyers. Therefore, p_y^R must be near 1, so that the higher-valued sellers from S_{yy} do not sell.
- If p_x^R is below 6 then no sellers from $S_{x\emptyset}$ want to sell x ; if p_x^R is at least 6 then no buyers from B_{xy} want to buy x (they prefer y). In both cases, no units of x are traded in the right market. Therefore p_x^R is at most the value of the lowest-valued seller in the right, which is $6 + 1/k^5$ (in fact, p_x^R can be any number between p_y^R and $6 + 1/k^5$).

When \mathbf{p}^R is applied in the left market, the price of x is too low for the $S_{x\emptyset}$ sellers so they do not sell at all. The supply of x is 0 and the B_{xx} patient cannot buy the strong medication. Almost all welfare is lost. Since this happens with probability $\approx 1/64$, the competitive ratio is at most $\approx 63/64$ — it does not approach 1.

This disaster *cannot* happen in any of the following cases.

- It cannot happen when the ratio between the market-sizes of the two goods is bounded by some constant c . In this case, the probability that the sampling-error in S_{yy} is of the same order-of-magnitude as B_{xy} approaches 0 as $k \rightarrow \infty$. This is consistent with our Theorem 2.
- It cannot happen when the welfare contributed by a single agent is bounded by some constant h . In this case, the relative loss of welfare in the sampling scenario above is $h/\Omega(k)$, which approaches 0 as $k \rightarrow \infty$. This is consistent with our Theorem 3.
- It cannot happen with single good-kind, even when most welfare comes from a single buyer. In this case, as long as k (the number of efficient deals) is sufficiently large, there will be sufficiently many sellers willing to sell the item in a price that the important buyer is willing to pay. This is consistent with Segal-Halevi *et al.* [2018a].

References

- [Babaioff *et al.*, 2014] Moshe Babaioff, Brendan Lucier, Noam Nisan, and Renato Paes Leme. On the efficiency of the Walrasian mechanism. In *Proc. EC '14*, pages 783–800, 2014.
- [Balcan *et al.*, 2007] Maria-Florina Balcan, Nikhil Devanur, Jason D. Hartline, and Kunal Talwar. Random Sampling Auctions for Limited Supply. Technical report, Carnegie Mellon University, 2007.
- [Balcan *et al.*, 2008] Maria-Florina Balcan, Avrim Blum, Jason D. Hartline, and Yishay Mansour. Reducing mechanism design to algorithm design via machine learning. *J Comput. Sys. Sci.*, 74(8):1245–1270, 2008.
- [Ben-Zwi *et al.*, 2013] Oren Ben-Zwi, Ron Lavi, and Ilan Newman. Ascending auctions and Walrasian equilibrium, July 2013.
- [Blumrosen and Dobzinski, 2014] Liad Blumrosen and Shahar Dobzinski. Reallocation Mechanisms. In *Proc. EC*, page 617. ACM, 2014.
- [Brustle *et al.*, 2017] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proc. EC '17*, 2017.
- [Chakraborty *et al.*, 2015] Mithun Chakraborty, Sanmay Das, and Justin Peabody. Price evolution in a continuous double auction prediction market with a scoring-rule based market maker. In *AAAI*, pages 835–841, 2015.
- [Chawla *et al.*, 2010] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter Mechanism Design and Sequential Posted Pricing. In *Proc. STOC*, pages 311–320, 2010.
- [Chu and Shen, 2006] Leon Y. Chu and Zuo-Jun M. Shen. Agent Competition Double-Auction Mechanism. *Management Science*, 52(8):1215–1222, 2006.
- [Colini-Baldeschi *et al.*, 2017a] 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.
- [Colini-Baldeschi *et al.*, 2017b] Riccardo Colini-Baldeschi, Paul W Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, and Stefano Turchetta. Approximately efficient two-sided combinatorial auctions. In *Proc. EC '17*, pages 591–608, 2017.
- [Devanur and Hayes, 2009] Nikhil R. Devanur and Thomas P. Hayes. The Adwords Problem: Online Keyword Matching with Budgeted Bidders Under Random Permutations. In *Proc. EC*, pages 71–78, 2009.
- [Devanur *et al.*, 2015] Nikhil R. Devanur, Jason D. Hartline, and Qiqi Yan. Envy freedom and prior-free mechanism design. *JET*, 156:103–143, 2015.
- [Feng *et al.*, 2012] Xiaojun Feng, Yanjiao Chen, Jin Zhang, Qian Zhang, and Bo Li. TAHES: Truthful double Auction for Heterogeneous Spectrums. In *Proc. INFOCOM*, pages 3076–3080. IEEE, 2012.
- [Goldberg *et al.*, 2001] Andrew V. Goldberg, Jason D. Hartline, and Andrew Wright. Competitive Auctions and Digital Goods. In *Proc. SODA '01*, 2001.
- [Goldberg *et al.*, 2006] Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, Michael Saks, and Andrew Wright. Competitive auctions. *GEB*, 55:242–269, 2006.
- [Gonen and Egri, 2017] Rica Gonen and Ozi Egri. Dycom: A dynamic truthful budget balanced double-sided combinatorial market. In *EXPLORE 2017 workshop, part of AAMAS'17*, pages 1556–1558, 2017.
- [Gonen *et al.*, 2007] Mira Gonen, Rica Gonen, and Elan Pavlov. Generalized Trade Reduction Mechanisms. In *Proc. EC*, pages 20–29, 2007.
- [Gul and Stacchetti, 1999] Faruk Gul and Ennio Stacchetti. Walrasian Equilibrium with Gross Substitutes. *JET*, 87(1):95–124, 1999.
- [Gul and Stacchetti, 2000] Faruk Gul and Ennio Stacchetti. The English Auction with Differentiated Commodities. *JET*, 92(1):66–95, 2000.
- [Hirai and Sato, 2017] Hiroshi Hirai and Ryosuke Sato. Polyhedral clinching auctions for two-sided markets. *arXiv preprint arXiv:1708.04881*, 2017.
- [Hsu *et al.*, 2016] Justin Hsu, Jamie Morgenstern, Ryan Rogers, Aaron Roth, and Rakesh Vohra. Do prices coordinate markets? *SIGecom Exch.*, 15(1):84–88, 2016.
- [Kelso and Crawford, 1982] Alexander S. Kelso and Vincent P. Crawford. Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica*, 50(6):1483–1504, November 1982.
- [Kojima and Pathak, 2009] Fuhito Kojima and Parag A. Pathak. Incentives and Stability in Large Two-Sided Matching Markets. *The American Economic Review*, 99(3):608–627, 2009.
- [Leyton-Brown *et al.*, 2017] Kevin Leyton-Brown, Paul Milgrom, and Ilya Segal. Economics and computer science of a radio spectrum reallocation. *Proceedings of the National Academy of Sciences*, 114(28):7202–7209, 2017.
- [McAfee, 1992] R. Preston McAfee. A dominant strategy double auction. *JET*, 56(2):434–450, 1992.
- [Myerson and Satterthwaite, 1983] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *JET*, 29(2):265–281, 1983.
- [Nisan *et al.*, 2007] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [Rustichini *et al.*, 1994] Aldo Rustichini, Mark A Satterthwaite, and Steven R Williams. Convergence to efficiency in a simple market with incomplete information. *Econometrica*, pages 1041–1063, 1994.
- [Segal-Halevi and Hassidim, 2017] Erel Segal-Halevi and Avinatan Hassidim. Truthful Bilateral Trade is Impossible even with Fixed Prices. 2017. *arXiv preprint 1711.08057*.
- [Segal-Halevi *et al.*, 2016] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. Demand-flow of agents with gross-substitute valuations. *Operations Research Letters*, 44:757–760, 2016.
- [Segal-Halevi *et al.*, 2018a] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. MUDA: A Truthful Multi-Unit Double-Auction Mechanism. In *Proceedings of AAAI'18*, February 2018a. *arXiv preprint 1712.06848*.
- [Segal-Halevi *et al.*, 2018b] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. Double Auctions in Markets for Multiple Kinds of Goods. In *Proceedings of IJCAI'18*, July 2018b. Previous name: "MIDA: A Multi Item-type Double-Auction Mechanism". *arXiv preprint: 1604.06210*.