

# MUDA: A Truthful Multi-Unit Double-Auction Mechanism

Erel Segal-Halevi  
Ariel University  
Ariel, Israel 40700

Avinatan Hassidim, Yonatan Aumann  
Bar-Ilan University  
Ramat-Gan, Israel 52900

## Abstract

In a seminal paper, McAfee (1992) presented a truthful mechanism for double auctions, attaining asymptotically-optimal gain-from-trade without any prior information on the valuations of the traders. McAfee’s mechanism handles single-parametric agents, allowing each seller to sell a single unit and each buyer to buy a single unit. This paper presents a double-auction mechanism that handles multi-parametric agents and allows multiple units per trader, as long as the valuation functions of all traders have decreasing marginal returns. The mechanism is prior-free, ex-post individually-rational, dominant-strategy truthful and strongly-budget-balanced. Its gain-from-trade approaches the optimum when the market size is sufficiently large.

## 1 Introduction

In a two-sided market, there are several sellers who hold items for sale and several buyers who consider buying these items. Examples are stock exchanges, used-car markets, emission trading markets (Godby 1999; Sturm 2008), Internet advertisements (Feldman and Gonen 2016) and markets for spectrum reallocation (Leyton-Brown, Milgrom, and Segal 2017). Each trader has a different valuation to each bundle of items. In contrast to a one-sided market, here the valuations of both the buyers and the sellers are their private information, and both sides might act strategically. A *double auction* is a mechanism for organizing a two-sided market — deciding who will buy, who will sell and at what prices.

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, in a used-car market with a single buyer and a single seller holding a single car, if the seller values the car as  $v_s$  and the buyer as  $v_b > v_s$ , then the potential GFT is  $v_b - v_s$ .

The most commonly used double-auction mechanism is the *Walrasian mechanism* (Rustichini, Satterthwaite, and Williams 1994; Babaioff et al. 2014). It computes an *equilibrium price* — a price at which the supply equals the demand: the total number of units that sellers are interested to sell at this price equals the total number of units that buyers are interested to buy at this price. In a single-good market, an equi-

librium price exists whenever the agents have *decreasing-marginal-returns* (DMR) — the marginal utility for an agent from having one more unit is weakly-decreasing in his current number of units (Gul and Stacchetti 2000). Moreover, by the First Welfare Theorem, this mechanism attains the maximum GFT (Nisan et al. 2007, Theorem 11.13). Unfortunately, this mechanism is not *incentive-compatible* (IC) — agents have an incentive to misreport their valuations in order to manipulate the price.

The problem of designing an IC double-auction has already been considered by Vickrey. In his seminal paper (Vickrey 1961) he described a variant of his famous second-price auction for a two-sided market. Like its one-sided variant, it is dominant-strategy incentive-compatible (DSIC, aka *truthful*) — it is a weakly-dominant strategy for each agent to reveal its true valuation function. But unlike the one-sided variant it is not *budget-balanced* (BB) — it has a deficit — it requires the market-maker to bring money from home. While it is possible to charge entrance-fees to cover this cost, this makes the mechanism not *individually rational* (IR) — some traders might lose from participating.

Myerson and Satterthwaite (1983) proved that, in a two-sided market, *any* mechanism that is IR, BB and IC cannot be efficient. Intuitively, the reason is that it is impossible to truthfully determine prices for trading. Consider any mechanism that charges the buyer  $p_b$  and pays the seller  $p_s$ . If  $p_s < v_b$ , then the seller is incentivized to bid  $(p_s + v_b)/2$  to force the price up (the mechanism has to do the deal since it is still efficient). Similarly, if  $p_b > v_s$ , the buyer is incentivized to force the price down. Setting  $p_s = v_b$  and  $p_b = v_s$  leads to a deficit. One way out is to determine take-it-or-leave-it prices independently of the traders’ valuations, but this might result in a total loss of GFT.

This impossibility result initiated a search for double-auction mechanisms that are IC, IR and 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 approximation mechanism was presented by McAfee McAfee (1992) for the case when each seller has a single unit and each buyer wants a single unit. McAfee’s mechanism is truthful — it is a weakly-dominant strategy for each agent to report his true valuation for the item. 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 efficient* — when the market-size  $k$  grows to infinity, the GFT approaches the maximum. In addition to being truthful, IR, BB and asymptotically-efficient, McAfee’s mechanism has a fifth nice property — it is *prior-free* (PF). This means that it does not require any statistical information about the traders’ valuations; it works well even for worst-case (adversarial) valuations.

The main drawback of McAfee’s mechanism is that it works only for single-unit traders. Another potential drawback is that it is only *weakly-budget-balanced* (WBB) — the market-maker need not bring money, but may the have to take money. Moreover, in some cases the market-maker consumes almost all the GFT, leaving only little GFT to the traders (Segal-Halevi, Hassidim, and Aumann 2016b). This may be acceptable when the market-maker is a monopolist (e.g. a government), but might be problematic when the market-maker faces competition, e.g. in stock exchange platforms, since a low GFT for the traders might drive them away to the competitors.<sup>1</sup>

This paper presents MUDA — a double-auction mechanism for traders that buy and sell multiple units. The only assumption required is that all traders have decreasing-marginal-returns; this is the same assumption that guarantees the existence of market-equilibrium.

The main idea of MUDA is to calculate equilibrium prices using *random halving*. The general scheme is presented below; it is explained in more detail in Section 3.

**MUDA** (general scheme): Split the market to 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 market-equilibrium price ( $p^R$  at the right,  $p^L$  at the left).
2. Let the agents trade at the price from the other market ( $p^L$  at the right,  $p^R$  at the left).

The random-sampling technique was found useful in one-sided markets (Goldberg, Hartline, and Wright 2001; Goldberg et al. 2006; Devanur, Hartline, and Yan 2015; Balcan et al. 2008; 2007) and matching markets (Devanur and Hayes 2009). However, applying it to a two-sided market poses a new challenge: *material balance* — the number of units bought and sold must be exactly equal. This is in contrast to a one-sided market, where the seller may leave some units unsold. Random-sampling in two-sided markets was used only with one-parametric agents

<sup>1</sup>Technically it is possible to convert a WBB mechanism to a SBB one: randomly pick an agent before the trade, disallow him to participate in the trade, and give him all the surplus after the trade. However, this might induce a lot of people without any interest in the trade (e.g. “sellers” with no goods or “buyers” with no money) to participate in the auction, in the hope of winning the lottery. We believe this goes against the idea of truthfulness, and might have adverse effects on the market.

— agents whose valuations are characterized by a single signal Baliga and Vohra; Kojima and Yamashita (2003; 2014).

With random-sampling, the price at each sub-market is determined exogenously based on the other sub-market, so the agents cannot manipulate it by reporting strategically. However, because of the material-balance requirement, the traders have another reason to report strategically — they might try to manipulate the *quantity* of trade. Thus, the multi-unit two-sided setting is more difficult than both the single-unit two-sided setting of McAfee and the multi-unit one-sided setting common in the random-sampling literature. To illustrate this added difficulty, we prove:

**Theorem 1.** Suppose a single seller holds  $M$  units of a good and a single buyer considers buying at most  $M$  units of this good and the price per unit of the good is determined exogenously. Then, the expected competitive ratio of every DSIC and IR mechanism, whether deterministic or randomized, is:

- at most  $1/M$  for agents with general valuations, and —
- at most  $1/H_M$  for agents with DMR valuations, where  $H_M$  is the  $M$ -th harmonic number ( $H_M \approx \ln M + \frac{1}{M}$ ).

Both these upper bounds are tight.

The proof is given in Segal-Halevi and Hassidim (2017). Theorem 1 can be seen as a dual to the Myerson–Satterthwaite impossibility theorem. In their setting, the quantity is determined exogenously and the traders might manipulate the price; in our setting, the price is determined exogenously and the traders might manipulate the quantity.

We present two variants of MUDA that overcome this impossibility when the market is sufficiently large. We call them *Lottery-MUDA* and *Vickrey-MUDA*.

Step 1 is the same in both variants. Step 2 in both variants starts by calculating the aggregate demand and aggregate supply in each sub-market, at the prices calculated in the other sub-market. Usually, the aggregate demand will be larger or smaller than the aggregate supply, so the sub-market will have a *long side* and a *short side* (e.g. if there is more demand than supply then the buyers are the long side and the sellers are the short side). The short side can always trade their optimum quantity; i.e. if the buyers are short, we can let each buyer buy the number of units that maximizes his utility given the price. The two variants of MUDA differ in the way they handle the long side: in *Lottery-MUDA* the traders in the long side are selected using a random permutation that ignores their values, while in *Vickrey-MUDA* the high-value traders in the long side are selected. In *Vickrey-MUDA*, each selected trader pays the market-maker a positive trading-fee calculated as in a Vickrey auction; this fee is separate from the money transferred among the traders (the inspiration to this idea came from a recent working paper by Loertscher and Mezzetti (2016)).

**Theorem 2.** Both variants of MUDA are prior-free, dominant-strategy incentive-compatible, individually-rational and budget-balanced (no deficit). In addition, *Lottery-MUDA* is strongly-budget-balanced (no surplus).

*Proof.* Prior-freeness holds by design: MUDA does not use any statistical information on the valuations.

DSIC and IR will be proved in Section 4 after presenting the mechanism details.

Lottery-MUDA is strongly-budget-balanced (SBB) since all money goes from buyers to sellers; the market-maker neither brings nor takes any money. Vickrey-MUDA is only *weakly budget-balanced (WBB)* — the market-maker does not lose money, but may make profit from trading-fees.  $\square$

We emphasize that the properties are true *ex-post* — for every outcome of the randomization in the mechanism.

We distinguish between the *total-GFT*, which includes the gain of both the agents and the market-maker, and the *agents-GFT*, which includes only the gain of the agents. Since Lottery-MUDA is SBB, the agents-GFT equals the total-GFT. In Vickrey-MUDA, the total-GFT is always higher since the most profitable deals are selected. However, the agents-GFT might be much lower. In a particular example shown in the arXiv version, the total-GFT is high but the agents-GFT is near 0. Thus, Vickrey-MUDA may be preferred when the market-maker is a monopolist (e.g. a government), while Lottery-MUDA may be preferred when the market-maker faces competition (e.g. a stock-exchange).

The competitive-ratio of MUDA depends on  $k$  — the number of units traded in the optimal situation (the  $k$  used by McAfee 1992), and  $M$  — the maximum number of units offered(demanded) by a single seller(buyer).  $M$  represents the market concentration — how much market power is held by a single trader.

**Theorem 3.** The expected total-GFT of both Lottery-MUDA and Vickrey-MUDA is at least a fraction  $1 - O(M\sqrt{\frac{\ln k}{k}})$  of the maximum total-GFT.

The proof is given in Section 5.

While Vickrey-MUDA attains more total-GFT than Lottery-MUDA, their asymptotic behavior is similar — both of them approach the optimal total-GFT when the market size ( $k$ ) is large, as long as the market-concentration factor (represented by  $M$ ) is kept constant. In contrast, if  $M$  is very large ( $M = k$ ) the market effectively has a single buyer and a single seller, and the impossibility result of (Myerson and Satterthwaite 1983) implies that no positive approximation of the GFT is possible. The arXiv version shows an example in which the *agents-GFT* of Vickrey-MUDA is close to 0 (the agents-GFT of Lottery-MUDA always equals its total-GFT).

While we focus on approximating the maximum GFT (buyers’ values minus trading sellers’ values), other mechanisms in the literature approximate the maximum *social welfare* (buyers’ values plus non-trading sellers’ values). But *any* mechanism that attains a fraction  $\alpha$  of the optimal GFT also attains a fraction of at least  $\alpha$  of the optimal social welfare (Brustle et al. 2017). Hence, MUDA is asymptotically-optimal with respect to the social-welfare too.

The lower bound of Theorem 3 depends on  $k$  — the optimal trade size. Ideally, we would like a bound that depends on the number of traders that come to the market (say,  $n$ ). However, we cannot attain such bound theoretically in a worst-case analysis, even for the social-welfare. As an example, consider a single-good single-unit market having  $n$

sellers with value 2,  $n - 1$  buyers with value 1 and one buyer with value 1000000. Here  $k = 1$  as there is only one relevant trade. The competitive ratio of any mechanism depends only on the probability with which it performs this single trade; this probability does not change even when  $n \rightarrow \infty$ .

We complement our worst-case analysis that depends on  $k$  with simulations of both variants of MUDA on agents drawn from both synthetic and realistic distributions. The simulations show that, when valuations are random (and not worst-case), the competitive-ratio of MUDA increases with the number of traders. These are presented in Section 6.

## 1.1 Related Work

Most existing mechanisms for multi-unit double-auction are not truthful, e.g. Plott and Gray (1990).

The research on truthful double auction has made many advancements since McAfee (1992). There are variants of McAfee’s mechanism for maximizing the auctioneer’s surplus (Deshmukh et al. 2002), handling spatially-distributed markets (Babaioff, Nisan, and Pavlov 2004), transaction costs (Chu and Shen 2006), supply chains (Babaioff and Walsh 2006), constraints on the set of traders that can trade simultaneously (Yao, Lu, and Jiang 2011; Dütting et al. 2017; Brustle et al. 2017), online arrival of buyers (Wang et al. 2010) and strong budget-balance (Colini-Baldeschi et al. 2016). All these advancements are for single-unit agents.

Other double-auction mechanisms either assume that the agents’ valuations are additive (Huang, Scheller-Wolf, and Sycara 2002; Xu, Jin, and Li 2010; Feldman and Gonen 2016; Goel et al. 2016; Hirai and Sato 2017) or assume that their valuations are represented by a single parameter Gonen, Gonen, and Pavlov; Kojima and Yamashita (2007; 2014). The DMR valuations handled by MUDA are multi-parametric and include additive valuations as a special case.

Several recent double-auction mechanisms work in a Bayesian setting — they assume that the traders’ valuations are drawn from some probability distribution that is public knowledge. Such knowledge allows the mechanism designer to attain approximate efficiency without relying on the agents’ reports. Examples are Yoon; Colini-Baldeschi et al. (2008; 2017). Some other mechanisms require partial prior knowledge on the valuations, such as their median (Blumrosen and Dobzinski 2014) or their maximum and minimum value Gonen and Egri (2017). (Baliga and Vohra 2003) assume that the agents’ valuations are drawn from some *unknown* distribution (they handle single-unit traders). In contrast, MUDA needs no prior information on the agents’ valuations and does not even assume that they are drawn from some distribution.

We are aware of two truthful mechanisms that handle multi-parametric agents with DMR valuations in a prior-free way. The first is by Blumrosen and Dobzinski (2014): their competitive ratio is 1/48 — it is not asymptotically-efficient. The second is by Loertscher and Mezzetti (2016), which is being developed simultaneously to our work. Their mechanism (that does not use market-halving) is asymptotically-efficient when the valuations of the traders are drawn from probability distributions satisfying certain conditions. In

contrast, our convergence theorem does not assume that agents' valuations are drawn from a distribution at all.

## 2 Model

**Agents and valuations** We consider a market for a single good. Some agents, the “sellers”, are endowed with at most  $M$  units of that good, and other agents, the “buyers”, are endowed with an unlimited supply of money. Each agent  $j$  has a valuation-function  $v_j$  that returns, for every integer  $t$  between 0 and  $M$ , the agent's value for owning  $t$  units. The valuations are normalized such that  $v_i(0) = 0$ .

All agents have *decreasing marginal returns* (DMR). Formally,  $v_j(t+1) - v_j(t) \leq v_j(t) - v_j(t-1)$ , for every agent  $j$  and  $t \in \{1, \dots, M-1\}$ .

We will often represent a multi-unit agent as  $M$  single-unit *virtual agents*; the value of virtual-agent  $t$  of agent  $j$  is the agent's marginal value for having the  $t$ -th unit:  $v_{j,t} := v_j(t) - v_j(t-1)$  for  $t \in \{1, \dots, M\}$  (a similar idea was used by Chawla et al. (2010)).

We assume that the agents' valuations are *generic*, i.e., all marginal values of different traders are different and linearly-independent over the integers (no linear combination with integer coefficients equals zero). This assumption can be dropped if we use centralized tie-breaking; see Hsu et al. (2016) for other ways to handle ties in markets.

Given a price  $p$  per unit of the good, the *gain* of a buyer  $i$  from buying  $t$  units is  $Gain_i(t, p) = v_i(t) - t \cdot p$ , and the gain of a seller  $j$  from selling  $t$  units is  $Gain_j(t, p) = t \cdot p + v_j(M_j - t) - v_j(M_j)$ , where  $M_j$  is the number of units initially held by seller  $j$ ,  $M_j \leq M$ .

A *mechanism* is a (randomized) function that takes the agents' valuations and returns (1) a trading-price  $p$ , (2) for each buyer(seller)  $i$ , the amount of units  $t_i$  he should buy(sell) at price  $p$ , and possibly a trading-fee  $f_i$  paid to the market-maker. A mechanism is *materially-balanced* if the number of units bought equals the number of units sold:  $\sum_{i \in \text{Buyers}} t_i = \sum_{j \in \text{Sellers}} t_j$ . It is *budget-balanced* (BB) if it is materially-balanced and  $\sum_{i \in \text{Traders}} f_i \geq 0$ . It is *strongly-budget-balanced* (SBB) if  $\sum_{i \in \text{Traders}} f_i = 0$ .

A mechanism is *individually-rational* (IR) if every agent  $i$  has a weakly-positive gain:  $Gain_i(t_i, p) - f_i \geq 0$  with probability 1. It is *DSIC* (= *truthful*) if an agent can never increase his gain by pretending to have different valuations.

The *demand* of buyer  $i$  is the number of units that maximizes the gain:  $Demand_i(p) = \arg \max_{t \in \{0, \dots, M\}} Gain_i(t, p)$ . The genericity assumption implies that the maximum is unique. When  $i$  has DMR,  $Demand_i(p)$  equals the number of virtual-buyers  $i, t$  with  $v_{i,t} > p$ .<sup>2</sup>

The *aggregate-demand* at price  $p$  is the sum of demands of all buyers. Equivalently, this is the number of virtual-

<sup>2</sup>This is true only for a DMR agent. For example, suppose buyer  $i$  values one unit as 3 and two units as 4. Then  $v_{i,1} = 3$  and  $v_{i,2} = 4$ . If the price is 2, then there is one virtual buyer with value above the price, and indeed the agent's demand is 1. However, if the buyer values one unit as 1 and two units as 4, then  $v_{i,1} = 1$  and  $v_{i,2} = 3$ . If the price is 2, then there is still one virtual buyer with value above the price, but the agent's demand is 0.

buyers with  $v_{j,t} > p$ . The *supply* and *aggregate-supply* are defined analogously.

The *total-gain-from-trade* is the sum of gains of all buyers and sellers:  $total\ GFT := \sum_{i \in \text{Traders}} Gain_i(t_i)$ . Note that in a materially-balanced mechanism, the total-GFT does not depend on the trading-price. The *agents-GFT* is the total GFT minus the total fees,  $agents\ GFT := total\ GFT - \sum_{i \in \text{Traders}} f_i$ .

## 3 Mechanism Details

This section explains the details of the MUDA mechanism presented in the introduction. The steps are done in each sub-market separately. For convenience we describe step 1 in the right market and step 2 in the left market; the execution in the opposite direction is entirely analogous.

**Step 1: Price calculation.** We calculate a price that is an *equilibrium price* at the right market — a price  $p^R$  for which  $Demand^R(p^R) = Supply^R(p^R)$ . Such a price exists even in more general (multi-good) settings. It can be found, for example, by simulating an English auction (Gul and Stacchetti 2000), or by binary search.

**Step 2: Posted-price trade.** For each buyer  $i$  in the left market, calculate  $Demand_i(p^R)$ . For each seller  $j$  in the left market, calculate  $Supply_j(p^R)$ . Let  $Demand^L$  be the sum of demands and  $Supply^L$  the sum of supplies. If  $Demand^L = Supply^L$ , then we can let the traders trade freely at price  $p^R$  and the market will clear. Usually, however, we will not be so lucky: there will be either excess demand ( $Demand^L > Supply^L$ ) or excess supply ( $Demand^L < Supply^L$ ). These two cases are handled analogously; henceforth we describe how to handle excess supply. First, we ask each buyer to pay in advance for the optimal number of units he wants to buy at price  $p^R$ , so we have money for  $Demand^L$  units. Since  $Demand^L < Supply^L$ , we will have more than enough units to give to all these buyers for their money. The problem is how to select the sellers that will supply these  $Demand^L$  units. We present two solutions.

(a) *Lottery*: Order the sellers randomly; let each seller in turn sell at price  $p^R$  as many units as he likes, while there is money (i.e., while at most  $Demand^L$  units are sold).

(b) *Vickrey-style auction*: Order the *virtual* sellers in increasing order of their value. From the  $Supply^L$  virtual sellers whose value is below  $p^R$ , pick the  $Demand^L$  virtual-sellers with the lowest values, and have each of them sell an item at price  $p^R$ .

The Vickrey-style auction is followed by a 4th step: **trading-fees**. For each seller  $j$ , let  $k_j$  be the number of virtual-sellers of this seller that are picked. Note that  $k_j \leq Supply_j(p^R)$ . The fee paid by seller  $j$  is determined by the potential gains of the virtual-sellers that are “pushed” out of the market because of seller  $j$ . Specifically, consider the set of virtual-sellers who want to trade in the left-market, *except* the  $k_j$  virtual-sellers of  $j$ . From this set, pick the (at most)  $Demand^L$  low-value ones. In this set, there are (at

most)  $k_j$  virtual-sellers that do not trade when seller  $j$  is present. Seller  $j$  pays to the market-maker, the gain (price minus value) of these virtual sellers. Note that the trading-fee is positive if  $k_j > 0$  and zero if  $k_j = 0$ .

**Example 1.** Suppose  $p^R = 50$ . The virtual-buyers in the left market have values 100, 90, 80, 60, 40, 20, so the aggregate demand is 4. There are two sellers: Alice’s values are 10, 20, 40, 60, 70 and Bob’s values are 15, 25, 35, 45, 65, so the aggregate supply is 7 and there is excess supply. The four high-value virtual-buyers (with values 100,90,80,60) each pays 50 and is guaranteed a unit.

In Lottery-MUDA, the two sellers are ordered randomly; if Alice is first then she sells 3 units and gains  $40 + 30 + 10 = 80$ , and Bob sells 1 unit and gains 35, so the GFT is 115. If Bob is first then he sells 4 units and gains  $35 + 25 + 15 + 5 = 80$ , and Alice sells nothing. The price is 50 per unit, all money collected from the buyers is given to the sellers, and no money is given to the market-maker.

In Vickrey-MUDA, the four low-value virtual-sellers are picked: they are the virtual-sellers with values 10,15,20,25. So each each real seller sells two units for 50 per unit. At this point Alice’s gain is  $(50 - 10) + (50 - 20) = 70$  and Bob’s gain is  $(50 - 15) + (50 - 25) = 60$ , so the total-GFT is 130. Now, Alice pays a fee of 20, since her presence prevents Bob from selling two units worth for him 35 and 45, so her net gain is 50. Bob pays 10, since his presence prevents Alice from selling a unit worth for her 40 (the unit worth 60 would not have been sold anyway), so his net gain is 50 too, and the agents-GFT is 100. The market-maker gains 30.  $\square$

## 4 Strategic Properties of MUDA

In both variants of MUDA, the traders cannot affect their trading-price. Moreover, in both variants, the traders in the short side of each sub-market trade the amount of units that maximizes their gain given the price, so for them the mechanism is obviously IR and DSIC. As for the long-side traders:

(a) in Lottery-MUDA they play random serial dictatorship: the first agents in the line trade their optimal quantity given the price and the last agents in the line cannot trade at all. There is at most a single agent, in the middle of the line, who trades less than his optimal quantity. Because of the DMR assumption, it is always optimal to trade as many units as possible up to the optimal quantity (since the highest gain comes from trading the first units). Therefore, the mechanism is IR and DSIC for all traders.

(b) In Vickrey-MUDA, long-side virtual-traders trade only if they have positive gain. In this case, a trader with  $k_j$  participating virtual-traders pays a fee that is equal to the gain of at most  $k_j$  non-participating virtual-traders. Since the mechanism always selects the virtual-traders with the highest gain, the total fee paid by any trader is lower than his gain, so the net gain remains positive and the mechanism is IR. The mechanism is DSIC since it is effectively a multi-unit Vickrey-auction with a reserve-price. It is known that such a mechanism is DSIC; we omit the proof.

## 5 Competitive-Ratio Analysis

In this section we prove Theorem 3. In fact, we prove a more general claim that depends on a third parameter,  $m$  — the minimum number of units offered(demanded) by a single seller(buyer) at same value. We assume that the virtual-traders of each trader  $i$  are divided to groups of size at least  $m$ , such that the values in each group are the same. So each agent  $i$  values  $m_{i,1}$  units as  $v_{i,1}$ ,  $m_{i,2}$  units as  $v_{i,2}$ , etc., with  $m_{i,l} \geq m$  for all  $l$  and  $\sum_l m_{i,l} \leq M$ . Theorem 3 follows from the expression at the end of this section by setting  $m = 1$ .

We first analyze the optimal trade, then the right sub-market and finally the left sub-market.

### 5.1 Optimal trade

In the optimal trade, there is a set  $B_*$  of virtual-buyers who buy goods from a set  $S_*$  of virtual-sellers. By material-balance the numbers of virtual-agents in both groups are the same; this is the number we denoted by  $k$ :

$$|B_*| = |S_*| = k \quad (1)$$

We call these  $k$  buyers and sellers the *efficient traders*. We make the pessimistic assumption that all GFT in the sub-markets comes 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 bounds. The second reason is the *pricing error* — the price 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.

### 5.2 Right sub-market

In the right sub-market, MUDA calculates an equilibrium price  $p^R$ . We define four sets of virtual-traders:

- $B_-$  is the set of efficient virtual-buyers (members of  $B_*$ ) whose value is below  $p^R$  (so they won’t buy at price  $p^R$ ).
- $S_-$  is the set of efficient virtual-sellers (members of  $S_*$ ) whose value is above  $p^R$  (so they won’t sell at price  $p^R$ ).
- $B_+$  is the set of inefficient virtual-buyers whose value is above  $p^R$  (so they want to buy at price  $p^R$ ).
- $S_+$  is the set of inefficient virtual-sellers whose value is below  $p^R$  (so they want to sell at price  $p^R$ ).

These sets represent the pricing error, so we want to upper-bound their sizes.

For any set  $T$  of agents, 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  $p^R$ :

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

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

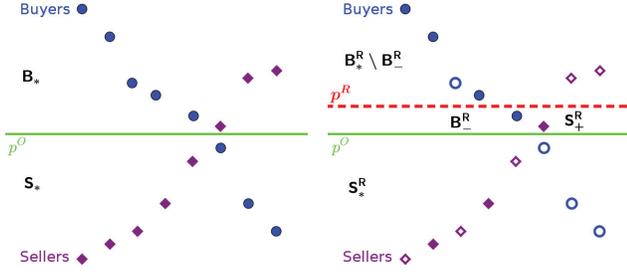


Figure 1: Price-distortion due to sampling. Each ball is the value of a single-unit buyer and each square is the value of a single-unit seller. *Left*: the optimal situation where there are 5 deals: the buyers above the line  $p^O$  trade with the sellers below the line. *Right*: the right market: full markers represent the traders sampled to the right and empty markers represent the traders sampled to the left. The equilibrium price  $p^R$  is higher than  $p^O$ , so some efficient buyers quit ( $B_-$ ) and some inefficient sellers join ( $S_+$ ).

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

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

(“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 twice, to  $B_*$  and  $S_*$ , and combine the outcomes using the union bound. This gives,  $\forall q > 0$ :

$$\text{w.p. } 1 - 4e^{-\frac{2mq^2}{M^2k}} : \left| |B_*^R| - \frac{k}{2} \right| < q \text{ and } \left| |S_*^R| - \frac{k}{2} \right| < q \quad (4)$$

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

$$\text{w.p. } 1 - 4e^{-\frac{2mq^2}{M^2k}} : \left| |B_-^R \cup S_+^R| - |B_+^R \cup S_-^R| \right| < 2q \quad (5)$$

Of the two sets in the left-hand side, at least one must be empty: if  $p^R$  is too high (relative to some optimal equilibrium price  $p^O$ ) then efficient buyers quit and inefficient sellers join, but no inefficient buyers join and no efficient sellers quit, so  $B_+^R = S_-^R = \emptyset$ . This situation is illustrated in Figure 1. Analogously, if  $p^R$  is too low then  $B_-^R = S_+^R = \emptyset$ .

From now on we assume that the situation is like in Figure 1 (the other situation is analogous). So (5) implies:

$$\forall q : \text{w.p. } 1 - 4e^{-\frac{2mq^2}{M^2k}} : |B_-^R| < 2q \text{ and } |S_+^R| < 2q \quad (6)$$

Our goal is now to derive an upper bound on  $B_-$  and  $S_+$ . They are entirely analogous; we focus on  $B_-$ . Note that we cannot apply (3) directly to  $B_-$ , since  $B_-$  is a random-set — it depends on the random-sampling through  $p^R$ . (3) does not apply to sets  $T$  that depend on the random-sampling; as an illustration, suppose the set  $T$  is selected such that it contains

only virtual-agents from the right market. Then  $T^R = T$  so (3) is obviously not satisfied. Fortunately,  $B_-$  is a special random-set — it is *one-dimensional*: for every integer  $t > 0$ , it has only a single possible value with cardinality  $t$ , which is the set of  $t$  virtual-buyers with the lowest value in  $B_*$  (see Figure 1, where  $t = 1$ ). Denote these sets by  $B_{-,t}$ . For every  $t$ ,  $B_{-,t}$  is independent of the random-sampling, so it is eligible to (3). Substituting there  $q \rightarrow 2t - 2q$  gives:

$$\forall q, t : \text{w.p. } 1 - 2e^{-\frac{2m(2t-2q)^2}{M^2 \cdot 4t}} : \left| |B_{-,4t}^R| - 2t \right| < 2t - 2q \quad (7)$$

$$\implies |B_{-,4t}^R| > 2q$$

If  $|B_-^R| < 2q$  and  $|B_{-,4t}^R| > 2q$ , then necessarily  $|B_-| < 4t$ . So by combining (6) and (7) with the union bound we get,  $\forall q, t$ :

$$\text{w.p. } 1 - 4e^{-\frac{2m(t-q)^2}{M^2t}} - 4e^{-\frac{2mq^2}{M^2k}} : |B_-| < 4t \text{ \& } |S_+| < 4t \quad (8)$$

To simplify the expression we choose  $t = 2q$ ; assuming  $2q < k$ , this implies  $\frac{-2m(t-q)^2}{M^2t} < \frac{-2mq^2}{M^2k}$ , so (8) simplifies to:

$$\forall q < k/2 : \text{w.p. } 1 - 8e^{-\frac{2mq^2}{M^2k}} : |B_-| < 8q \text{ \& } |S_+| < 8q \quad (9)$$

### 5.3 Left sub-market

Denote by  $B^L(S^L)$  the set of efficient buyers(sellers) who want to buy(sell) in the left sub-market at price  $p^R$ :

$$B^L := B_*^L \setminus B_-^L = B_* \setminus B_*^R \setminus B_-^L$$

$$S^L := S_*^L \setminus S_-^L = S_*^L = S_* \setminus S_*^R$$

(Recall that we assume the case in which  $S_-$  and  $B_+$  are empty; the other case is analogous).

Using (4) and (9) with the same  $q$  gives, for every  $q < \frac{k}{2}$ :

$$\text{w.p. } 1 - 8e^{-\frac{2mq^2}{M^2k}} : |B^L| > \frac{k}{2} - 9q \text{ and } |S^L| > \frac{k}{2} - q$$

In Vickrey-MUDA, the most efficient traders in each side trade with each other. Therefore, the mechanism makes at least the  $\frac{k}{2} - 9q$  most efficient deals in the left submarket. Similar considerations are true in the right submarket. All in all, Vickrey-MUDA does at least the  $k - 18q$  most efficient out of the  $k$  efficient deals. Therefore, w.p.  $1 - 8e^{-\frac{2mq^2}{M^2k}}$ , the competitive ratio is at least  $1 - 18\frac{q}{k}$ .

In Lottery-MUDA, the efficient sellers in  $S_*^L$  have to compete with the inefficient sellers in  $S_+^L$  in the random lottery. The expected number of efficient deals carried out is thus:

$$\frac{|B^L|}{|S_+^L| + |S_*^L|} \geq \frac{k/2 - 9q}{k/2 + 9q} \geq \frac{k}{2} - 18q$$

Therefore, w.p.  $1 - 8e^{-\frac{2mq^2}{M^2k}}$ , the *expected* competitive ratio is at least  $1 - 36\frac{q}{k}$ . So w.p. 1, the expected competitive ratio is at least  $1 - 8e^{-\frac{2mq^2}{M^2k}} - 36\frac{q}{k}$ .

All our analysis so far holds simultaneously for every  $q < \frac{k}{2}$ . Now, we select  $q$  to maximize the expected competitive ratio. With some tedious calculations we find that, when  $k$  is sufficiently large, we can select  $q$  such that the competitive ratio is at least:

$$1 - \frac{36M}{2\sqrt{mk}} \left( 1 + \sqrt{\ln \left( \frac{16^2 mk}{36^2 M^2} \right)} \right) \\ = 1 - O \left( M \sqrt{\frac{\ln mk}{mk}} \right)$$

The analysis for Vickrey-MUDA is obtained by replacing 36 with 18; the asymptotic behavior is the same.

## 6 Simulations

To complement our theoretic analysis, we simulated MUDA on traders with valuations sampled both from a synthetic distribution and an empirical distribution based on real stock-exchange data.

### 6.1 Uniform distribution

In the first experiment, for each agent, we sampled  $M/m$  values from a uniform distribution with support  $[V - A, V + A]$ . We considered each of these values as the marginal-value of  $m$  virtual-traders, so that each agent has  $M$  virtual-traders. We ordered the values in decreasing order to get DMR valuations. In the experiments, we took  $m = 100$  and  $V = 500$  and varied the noise-amplitude  $A$  between 50 and 450. Here we show the results for  $A = 250$ ; varying  $A$  did not have much effect on the results. We repeated each experiment 100 times and averaged the results.

In the first sub-experiment, we kept  $M$  constant (at 100,000) and varied the number of real traders between 0 and 1000. The results are shown in Figure 2/Left. The competitive ratios of all variants of MUDA increase towards 1 when the number of traders grows.

In the second sub-experiment, we kept constant the total number of units held by all sellers. We increased  $M$  from 100 to  $10^8$ , and decreased the number of traders accordingly so that the total number of units remains constant. The results are shown in Figure 2/Right. The competitive ratios decrease when  $M$  increases and the same number of units are concentrated in the hand of fewer traders. When  $M$  is very large we effectively have one buyer and one seller, and then the impossibility result of (Myerson and Satterthwaite 1983) implies that no positive competitive ratio is possible.

Both plots show that the GFT of Lottery-MUDA is approximately middle-way between the total-GFT and the agents-GFT of Vickrey-MUDA. This fact has a practical implication regarding the choice of mechanism: if the market is managed by a monopolist (e.g. the government), then Vickrey-MUDA is better since its total-GFT is higher; but if the market competes with other markets (e.g. stock trading platforms), then Lottery-MUDA is better since its agents-GFT is higher.

We found that, when the number of units per trader is fixed, the performance is determined by the number of

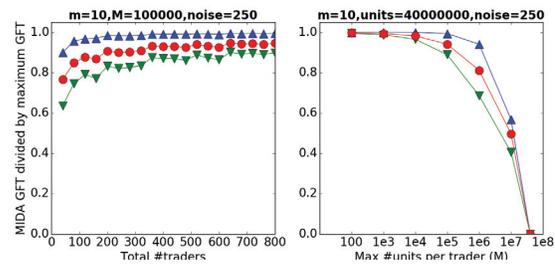


Figure 2: Competitive ratio of MUDA when traders’ valuations are drawn from a uniform distribution. Legend:

- \* Downward-wedges = agents-GFT of Vickrey-MUDA;
- \* Discs = Lottery-MUDA;
- \* Upward-wedges = total-GFT of Vickrey-MUDA.

**Left:** maximum number of units per trader ( $M$ ) is fixed, and number of traders (and total number of units) increases.

\* Agents-GFT of Vickrey-MUDA ranges from 0.58 to 0.86.

\* GFT of Lottery-MUDA ranges from 0.74 to 0.92.

\* Total-GFT of Vickrey-MUDA ranges from 0.897 to 0.991.

**Right:** total number of units is fixed, and maximum number of units per trader increases (so the total number of traders decreases).

traders in the market. Therefore, the plot at the left of Figure 2 remains the same even when we set e.g.  $m = 1$  and  $M = 10$ . As for the plot at the right of Figure 2, if we set the total number of units to e.g. 1000, then the plot approaches zero already when  $M = 1e3$ , since in this case there is a single buyer and a single seller.

### 6.2 Empirical stock-exchange distribution

In the second experiment, we used the TORQ database (Hasbrouck 1992; Lee and Radhakrishna 2000). It contains buy and sell orders given for a sample of 144 NYSE stocks for the three months 1990-11 through 1991-01. In the NYSE, each day before the continuous trade begins, there is a phase of “start-of-day auctions”. For each stock, a separate multi-unit double-auction is conducted in the following way. All buy and sell orders given for that stock before the start of day are collected. An equilibrium price is calculated. All buy-orders above the price and all sell-orders below the price are executed. As explained in the introduction, this mechanism is efficient but not truthful. Therefore, the orders may not represent the true values of the traders. While there are econometric methods for estimating the true values from the reported values, these are beyond the scope of the present paper, since we do not need to know the true value of every trader — all we need is a distribution of values. Our results are valid as long as the empirical distribution of reported values resembles the distribution of values in the real world.

In the present paper we assume that the empirical distribution of the orders is similar to the empirical distribution of the true values.

For the experiment, we considered only the start-of-day orders. These orders are given in the format (Symbol, Date, Order date, Side, Price, Quantity) where Symbol is the three-

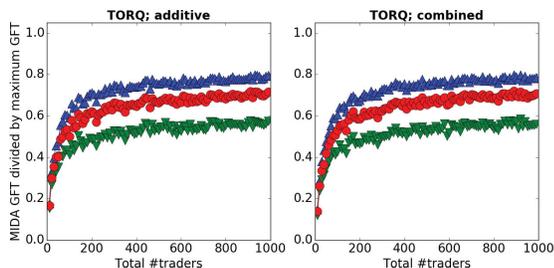


Figure 3: Competitive ratio of MUDA when traders’ valuations are drawn from an empirical distribution based on NYSE orders. Discs = Lottery-MUDA; upward/downward-wedges = total/agents-GFT of Vickrey-MUDA. At the left, each order is assumed to belong to a different trader (so all traders are additive). At the right, orders from the same order-date are merged to a single trader.

letter acronym of the stock; Date is the day of the auction. Order date is the day the order was made, which may be earlier than Date; Side is BUY or SELL; and Price and Quantity are the actual bid. The dataset contains 144 symbols, 7914 different symbol-date combinations and  $\approx 210000$  orders.

The dataset does not contain the identities of the traders. We made two experiments: in the first we treated each order as a separate trader, so that all traders are additive (each trader values a certain quantity of units for the same price-per-unit). In the second experiment, we simulated traders with non-additive valuations by merging bids based on the order-date. I.e, we heuristically assumed that two bids with the same symbol, date, order-date and side belong to the same trader. We ordered the bids of the same trader in descending order to create DMR valuations, similarly to the uniform experiment. In average, each merged bidder contained 10 different orders. The number of units per order ranges between 100 and 99000, so  $m = 100$ . The number of units per trader, after combining orders from the same order-date, ranges between 5000 and 12000000 with median 148000, so  $M = 12000000$ .

For each symbol, we created a collection of all traders in all days to represent the empirical distribution of this symbol. Then, in each experiment we sampled  $n$  traders, where  $n = 10, 20, \dots, 990$ . All in all we simulated  $99 \cdot 144$  auctions and averaged the 144 auctions for each  $n$ . The results are shown in Figure 3.

### 6.3 Discussion

It is interesting that the plots in Figure 3 look almost the same, i.e, the non-additivity had little effect on the results. The reason might be that stock-traders have many different bids with very similar prices, so that they are almost additive.

The competitive ratio in the TORQ experiment is somewhat lower than in the uniform case. We attribute this to the large variation in the number of units: while some traders hold only 5000 units, others hold as many as 12 million.

Source code for reproducing the experiments is available at <https://github.com/erelsgl/economics>.

## 7 Future Work

It is interesting whether MUDA can be extended to agents with general valuations (not DMR). This poses two challenges. First, in step 1, a price-equilibrium might not exist (the existence of a price-equilibrium is guaranteed only when the agents have DMR). Second, in step 2, material balance might fail. For example, if there is one buyer who only wants to buy 2 units and one seller who only wants to sell 3 units, then with DMR we can assume that both agree to trade 2 units, but without DMR this assumption might be false — the seller might have negative utility from selling 3 units.

Another challenge is handling double auctions with multiple types of items. Preliminary results in this direction were presented by Segal-Halevi, Hassidim, and Aumann (2016a).

We will also be happy to compare MUDA to other double-auction mechanisms. In particular, an interesting practical question is when would MUDA’s truthfulness actually matter — how does its GFT compare to the GFT in Nash equilibrium of the standard (non-truthful) market mechanism? As far as we know, the Nash-equilibrium GFT of the non-truthful market mechanism (i.e, its “price-of-anarchy”) is still an open question. The closest reference we found is by Babaioff et al. (2014), who consider the price-of-anarchy of the market mechanism in a single-sided market. We are not aware of analogous results for double-sided markets.

But non-truthful market mechanisms have disadvantages beyond the price-of-anarchy. First, the players need to spend resources to obtain information about other players and calculate their best response. Second, the players may not even play an equilibrium. It is possible that players do not know the correct state of the market and hence play best-response to a fictitious state. This can bring social welfare to zero. Consider any mechanism that sets a single price clearing the market. For every active buyer(seller), it is a best response to bid below(above) the market price in order to push the price down(up). If all agents play these best responses, the outcome is that nothing gets sold. Such market failures are prevented by the truthfulness of MUDA.

## 8 Acknowledgments

Assaf Romm helped us accessing and analyzing the TORQ dataset and provided many helpful comments. Ron Adin, Simcha Haber, Ron Peretz, Tom van der Zanden, Jack D’Aurizio, Andre Nicolas, Brian Tung, Robert Israel, Clement C. and C. Rose helped us with the probability calculations. The participants of the game-theory seminar in Bar-Ilan University, computational economics and economic theory seminars in the Hebrew university of Jerusalem, algorithms seminar in Tel-Aviv university and computer science seminar in Ariel university provided helpful comments on various aspects of the paper. The anonymous reviewers in Theoretical Economics, SODA 2016, EC 2016, SODA 2017, EC 2017, and, of course, AAAI 2018, gave us detailed and helpful feedback.

Erel was supported by ISF grant 1083/13, Doctoral Fellowships of Excellence Program and Mordecai and Monique Katz Graduate Fellowship Program at Bar-Ilan University. Avinatan Hassidim is supported by ISF grant 1394/16.

## References

- Babaioff, M., and Walsh, W. 2006. Incentive Compatible Supply Chain Auctions. In *Multiagent based Supply Chain Management*, volume 28. Springer, 315–350.
- Babaioff, M.; Lucier, B.; Nisan, N.; and Paes Leme, R. 2014. On the efficiency of the walrasian mechanism. In *Proc. EC '14*, 783–800.
- Babaioff, M.; Nisan, N.; and Pavlov, E. 2004. Mechanisms for a Spatially Distributed Market. In *Proc. EC*, 9–20.
- Balcan, M.-F.; Devanur, N.; Hartline, J. D.; and Talwar, K. 2007. Random Sampling Auctions for Limited Supply. Technical report, Carnegie Mellon University.
- Balcan, M.-F.; Blum, A.; Hartline, J. D.; and Mansour, Y. 2008. Reducing mechanism design to algorithm design via machine learning. *J Comput. Sys. Sci.* 74(8):1245–1270.
- Baliga, S., and Vohra, R. 2003. Market Research and Market Design. *Advances in Theoretical Economics* 3(1).
- Blumrosen, L., and Dobzinski, S. 2014. Reallocation Mechanisms. In *Proc. EC*, 617. ACM.
- Brustle, J.; Cai, Y.; Wu, F.; and Zhao, M. 2017. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proc. EC '17*.
- Chawla, S.; Hartline, J. D.; Malec, D. L.; and Sivan, B. 2010. Multi-parameter Mechanism Design and Sequential Posted Pricing. In *Proc. STOC*, 311–320.
- Chu, L. Y., and Shen, Z.-J. M. 2006. Agent Competition Double-Auction Mechanism. *Management Science* 52(8):1215–1222.
- Colini-Baldeschi, R.; de Keijzer, B.; Leonardi, S.; and Turchetta, S. 2016. Approximately Efficient Double Auctions with Strong Budget Balance. In *Proc. SODA '16*.
- Colini-Baldeschi, R.; Goldberg, P. W.; de Keijzer, B.; Leonardi, S.; Roughgarden, T.; and Turchetta, S. 2017. Approximately efficient two-sided combinatorial auctions. In *Proc. EC '17*, 591–608.
- Deshmukh, K.; Goldberg, A. V.; Hartline, J. D.; and Karlin, A. R. 2002. Truthful and Competitive Double Auctions. In *Proc. ESA*, volume 2461 of *Lecture Notes in Computer Science*. Springer, 361–373.
- Devanur, N. R., and Hayes, T. P. 2009. The Adwords Problem: Online Keyword Matching with Budgeted Bidders Under Random Permutations. In *Proc. EC*, 71–78.
- Devanur, N. R.; Hartline, J. D.; and Yan, Q. 2015. Envy freedom and prior-free mechanism design. *JET* 156:103–143.
- Dütting, P.; Talgam-Cohen, I.; Roughgarden, T.; et al. 2017. Modularity and greed in double auctions. *Games and Economic Behavior* 105(C):59–83.
- Feldman, M., and Gonen, R. 2016. Double-sided markets with strategic multi-dimensional players. *arXiv preprint arXiv:1603.08717*.
- Godby, R. 1999. Market power in emission permit double auctions. *Research in experimental economics* 7:121–162.
- Goel, G.; Leonardi, S.; Mirrokni, V.; Nikzad, A.; and Paes-Leme, R. 2016. Reservation exchange markets for internet advertising. In *LIPICs*, volume 55.
- Goldberg, A. V.; Hartline, J. D.; Karlin, A. R.; Saks, M.; and Wright, A. 2006. Competitive auctions. *GEB* 55:242–269.
- Goldberg, A. V.; Hartline, J. D.; and Wright, A. 2001. Competitive Auctions and Digital Goods. In *Proc. SODA '01*.
- Gonen, R., and Egri, O. 2017. Dycom: A dynamic truthful budget balanced double-sided combinatorial market. In *Proc. AAMAS '17*, 1556–1558.
- Gonen, M.; Gonen, R.; and Pavlov, E. 2007. Generalized Trade Reduction Mechanisms. In *Proc. EC*, 20–29.
- Gul, F., and Stacchetti, E. 2000. The English Auction with Differentiated Commodities. *JET* 92(1):66–95.
- Hasbrouck, J. 1992. Using the torq database. *New York Stock Exchange* 11.
- Hirai, H., and Sato, R. 2017. Polyhedral clinching auctions for two-sided markets. *arXiv preprint arXiv:1708.04881*.
- Hsu, J.; Morgenstern, J.; Rogers, R.; Roth, A.; and Vohra, R. 2016. Do prices coordinate markets? *SIGecom Exch.* 15(1):84–88.
- Huang, P.; Scheller-Wolf, A.; and Sycara, K. 2002. Design of a multi-unit double auction e-market. *Computational Intelligence* 18(4):596–617.
- Kojima, F., and Yamashita, T. 2014. Double auction with interdependent values: Incentives and efficiency. *Th.Econ.*
- Lee, C. M., and Radhakrishna, B. 2000. Inferring investor behavior: Evidence from torq data. *Journal of Financial Markets* 3(2):83–111.
- Leyton-Brown, K.; Milgrom, P.; and Segal, I. 2017. Economics and computer science of a radio spectrum reallocation. *Proceedings of the National Academy of Sciences* 114(28):7202–7209.
- Loertscher, S., and Mezzetti, C. 2016. Dominant strategy, double clock auctions with estimation-based tâtonnement. Technical report. <https://sites.google.com/site/clamezzetti/work-papers>.
- McAfee, R. P. 1992. A dominant strategy double auction. *JET* 56(2):434–450.
- Myerson, R. B., and Satterthwaite, M. A. 1983. Efficient mechanisms for bilateral trading. *JET* 29(2):265–281.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V., eds. 2007. *Algorithmic Game Theory*. Cambridge University Press.
- Plott, C. R., and Gray, P. 1990. The multiple unit double auction. *Journal of Economic Behavior & Organization* 13(2):245–258.
- Rustichini, A.; Satterthwaite, M. A.; and Williams, S. R. 1994. Convergence to efficiency in a simple market with incomplete information. *Econometrica* 1041–1063.
- Segal-Halevi, E., and Hassidim, A. 2017. Truthful Bilateral Trade is Impossible even with Fixed Prices. *arXiv preprint*.
- Segal-Halevi, E.; Hassidim, A.; and Aumann, Y. 2016a. Mida: A multi item-type double-auction mechanism. *arXiv preprint 1604.06210*.
- Segal-Halevi, E.; Hassidim, A.; and Aumann, Y. 2016b. Sbba: A strongly-budget-balanced double-auction mechanism. In *Proc. SAGT '16*, 260–272.
- Sturm, B. 2008. Double auction experiments and their relevance for emissions trading. In *Emissions Trading*. Springer, 49–68.
- Vickrey, W. 1961. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J Finance* 16(1):8–37.
- Wang, S.; Xu, P.; Xu, X.; Tang, S.; Li, X.; and Liu, X. 2010. TODA: Truthful Online Double Auction for Spectrum Allocation in Wireless Networks. In *New Frontiers in Dynamic Spectrum*, 1–10. IEEE.
- Xu, H.; Jin, J.; and Li, B. 2010. A Secondary Market for Spectrum. In *Proc. INFOCOM*, 1–5. IEEE.
- Yao, E.; Lu, L.; and Jiang, W. 2011. An efficient truthful double spectrum auction design for dynamic spectrum access. In *Proc. CROWNCOM*, 181–185. IEEE.
- Yoon, K. 2008. The participatory vickrey–clarke–groves mechanism. *JME* 44(3):324–336.