Fast Convergence in the Double Oral Auction

Sepehr Assadi* Sanjeev Khanna* Yang Li* Rakesh Vohra†

Abstract

A classical trading experiment consists of a set of unit demand buyers and unit supply sellers with identical items. Each agent's value or opportunity cost for the item is their private information and preferences are quasi-linear. Trade between agents employs a double oral auction (DOA) in which both buyers and sellers call out bids or offers which an auctioneer recognizes. Transactions resulting from accepted bids and offers are recorded. This continues until there are no more acceptable bids or offers. Remarkably, the experiment consistently terminates in a Walrasian price. The main result of this paper is a mechanism in the spirit of the DOA that converges to a Walrasian equilibrium in a polynomial number of steps, thus providing a theoretical basis for the above-described empirical phenomenon. It is well-known that computation of a Walrasian equilibrium for this market corresponds to solving a maximum weight bipartite matching problem. The uncoordinated but rational responses of agents thus solve in a distributed fashion a maximum weight bipartite matching problem that is encoded by their private valuations. We show, furthermore, that every Walrasian equilibrium is reachable by some sequence of responses. This is in contrast to the well known auction algorithms for this problem which only allow one side to make offers and thus essentially choose an equilibrium that maximizes the surplus for the side making offers. Our results extend to the setting where not every agent pair is allowed to trade with each other.

*Department of Computer and Information Science, University of Pennsylvania.
Email: {sassadi,sanjeev,yangli2}@cis.upenn.edu. Supported in part by National Science Foundation grants CCF-1116961, CCF-1552909, and IIS-1447470.

†Department of Economics and the Department of Electrical and System Engineering.
Email: rvohra@seas.upenn.edu.

1 Introduction

Chamberlin reported on the results of a market experiment in which prices failed to converge to a Walrasian equilibrium [4]. Chamberlin's market was an instance of the assignment model with homogeneous goods. There is a set of unit demand buyers and a set of unit supply sellers, and all items are identical. Each agent's value or opportunity cost for the good is their private information and preferences are quasi-linear. Chamberlin concluded that his results showed competitive theory to be inadequate. Vernon Smith, in an instance of insomnia, recounted in [18] demurred:
"The thought occurred to me that the idea of doing an experiment was right, but what was wrong was that if you were going to show that competitive equilibrium was not realizable ... you should choose an institution of exchange that might be more favorable to yielding competitive equilibrium. Then when such an equilibrium failed to be approached, you would have a powerful result. This led to two ideas: (1) why not use the double oral auction procedure, used on the stock and commodity exchanges? (2) why not conduct the experiment in a sequence of trading days in which supply and demand were renewed to yield functions that were daily flows?"

Instead of Chamberlin's unstructured design, Smith used a double oral auction (DOA) scheme in which both buyers and sellers call out bids or offers which an auctioneer recognizes [17]. Transactions resulting from accepted bids and offers are recorded. This continues until there are no more acceptable bids or offers. At the conclusion of trading, the trades are erased, and the market reopens with valuations and opportunity costs unchanged. The only thing that has changed is that market participants have observed the outcomes of the previous days trading and may adjust their expectations accordingly. This procedure was iterated four or five times. Smith was astounded: "I am still recovering from the shock of the experimental results. The outcome was unbelievably consistent with competitive price theory." [18](p. 156)

As noted by Daniel Friedman [7], the results in [17], replicated many times, are something of a mystery. How is it that the agents in the DOA overcome the impediments of both *private information* and *strategic uncertainty* to arrive at the Walrasian equilibrium? A brief survey of the various (early) theoretical attempts to do so can be found in Chapter 1 of [7]. Friedman concluded his survey of the theoretical literature with a two-part conjecture. "First, that competitive (Walrasian) equilibrium coincides with ordinary (complete information) Nash Equilibrium (NE) in interesting environments for the DOA institution. Second, that the DOA promotes some plausible sort of learning process which eventually guides the *both clever and not-so-clever* traders to a behavior which constitutes an 'as-if' complete-information NE."

Over the years, the first part of Friedman's conjecture has been well studied (see, e.g., [6]; see also Section 3) but the second part of the conjecture is still left without a satisfying resolution. The focus of this paper is on the second part of Friedman's conjecture. More specifically, we design a mechanism which simulates the DOA, and prove that this mechanism always converges to a Walrasian equilibrium in polynomially many steps. Our mechanism captures the following four key properties of the DOA.

1. *Two-sided market:* Agents on either side of the market can make actions.
2. *Private information:* When making actions, agents have no other information besides their own valuations and the bids and offers submitted by others.
3. *Strategic uncertainty:* The agents have the freedom to choose their actions modulo mild rationality conditions.

4. *Arbitrary recognition:* The auctioneer (only) recognizes bids and offers in an *arbitrary* order.

Among these four properties, mechanisms that allow agents on either side to make actions (*two-sided market*) and/or limit the information each agent has (*private information*) have received more attention in the literature (see Section 1.1). However, very little is known for mechanisms that both work for strategically uncertain agents and recognize agents in an arbitrary order. Note that apart from resolving the second part of Friedman's conjecture, having a mechanism with these four properties itself is of great interest for multiple reasons. First, in reality, the agents are typically unwilling to share their private information to other agents or the auctioneer. Second, agents naturally prefer to act freely as oppose to being given a procedure and merely following it. Third, in large scale distributed settings, it is not always possible to find a real auctioneer who is trusted by every agent, and is capable of performing massive computation on the data collected from all agents. In the DOA (or in our mechanism) however, the auctioneer only recognizes actions in an arbitrary order, which can be replaced by any standard distributed token passing protocol, where an agent can take an action only when he is holding the token. In other words, our mechanism serves more like a platform (rather than a specific protocol) where rational agents always reach a Walrasian equilibrium no matter their actual strategy. To the best of our knowledge, no previous mechanism enables such a 'platform-like' feature. In the rest of this section, we summarize our results and discuss in more detail the four properties of the DOA in context of previous work.

1.1 Our Results and Related Work

We design a mechanism that simulates the DOA by simultaneously capturing two-sided market, private information, strategic uncertainty, and arbitrary recognition. More specifically, following the DOA, at each iteration of our mechanism, the auctioneer maintains a list of active price submission and a tentative assignment of buyers to sellers that 'clears' the market at the current prices (note that this can also be distributedly maintained by the agents themselves). Among the agents who wish to make or revise an earlier submission, an arbitrary one is recognized by the auctioneer and a new tentative assignment is formed. An agent can submit *any* price that strictly improves his payoff given the current submissions (rather than being forced to make a 'best' response, which is to submit the price that maximizes payoff). We show that as long as agents make myopically better responses, the market always converges to a Walrasian equilibrium in polynomial number of steps. Furthermore, *every* Walrasian equilibrium is the limit of some sequence of better responses. We should remark that the fact that an agent always improves his payoff does not imply that the total payoff of all agents always increases. For instance, a buyer can increase his payoff by submitting a higher price and 'stealing' the current match of some other buyer (whose payoff would drop).

To the best of our knowledge, no existing mechanism captures all four properties for the DOA that we proposed in this paper. For most of the early work on auction based algorithms (e.g., [1, 5, 6, 11, 16]), unlike the DOA, only one side of the market can make offers. By permitting only one side of the market to make offers, the auction methods essentially pick the Walrasian equilibrium (equilibria are not unique) that maximizes the total surplus of the side making the offers.

For two-sided auction based algorithms [2, 3], along with the 'learning' based algorithms studied more recently [9, 13], agents are required to follow a specific algorithm (or protocol) that determines their actions (and hence violates strategic uncertainty). For example, [3] requires that when an agent is activated, a buyer always matches to the 'best' seller and a seller always matches to the 'best' buyer (i.e., agents only make myopically *best* responses, which is not the case for the DOA). [9] has agents submit bids based on their current best alternative offer and prices are updated according to a common formula relying on knowledge of the agents opportunity costs and marginal values. [13], though not requiring agents to always make myopically best responses, has agents follow a specific (randomized) algorithm to submit conditional bids and choose matches. We should emphasize that agents acting based on some *random* process is different from agents being strategically uncertain. In particular, for the participants of the original DOA experiment (of [17]), there is no a priori reason to believe that they were following some specific random procedure during the experiment. On the contrary, as stated in Friedman's conjecture, there are *clever and not-so-clever* participants, and hence different agents could have completely different strategies and their strategies might even change when, for instance, seeing more agents matching with each other, or by observing the strategies of other agents. Therefore, analyzing a process where agents are strategically uncertain can be distinctly more complex than analyzing the case where agents behave in accordance with a well-defined stochastic process. In this paper, we consider an extremely general model of the agents: the agents are acting arbitrarily while only following some mild rationality conditions. Indeed, proving fast convergence (or even just convergence) for a mechanism with agents that are strategically uncertain is one of the main challenges of this work.

Arbitrary recognition is another critical challenge for designing our mechanism. For example, the work of [13, 14] deploys randomization in the process of recognizing agents. This is again in contrast to the original DOA experiment, since the auctioneer did not use a randomized procedure when recognizing actions, and it is unlikely that the participants *decide* to make an action following some random process (in fact, some participants might be more 'active' than others, which could lead to the 'quieter' participants barely getting *any* chance to make actions, as long as the 'active' agents are still making actions).

The classical work on the *stable matching* problem [8] serves as a very good illustration for the importance of arbitrary recognition. Knuth [10] proposed the following algorithm for finding a stable matching. Start with an arbitrary matching; if it is stable, stop; otherwise, pick a blocking pair and match them; repeat this process until a stable matching is found. Knuth showed that the algorithm could cycle if the blocking pair is picked *arbitrarily*. Later, [15] showed that picking the blocking pairs at random suffices to ensure that the algorithm eventually converges to a stable matching, which suggests that it is the arbitrary selection of blocking pairs that causes Knuth's algorithm to not converge.

The setting of Knuth's algorithm is very similar to the process of the DOA in the sense that in any step of the DOA, a temporary matching is maintained and agents can make actions to (possibly) change the current matching. But perhaps surprisingly, we show that arbitrary recognition does not cause the DOA to suffer from the same cycling problem as Knuth's algorithm. The main reason, or the main difference between the two models is that our assignment model involves both matching and prices, while Knuth's algorithm only involves matchings. As a consequence, in our mechanism, the preferences of the agents change over time (since an agent always favors the better price submission, the preferences could change when new prices are submitted). In the instance that leads Knuth's algorithm to cycle (see [10]), the fundamental cause is that the preferences of *all* agents form a cycle. However, in our mechanism, preferences (though changing) are always consistent for all agents.

Based on this observation, we establish the limit of the DOA by introducing a small friction into the market: restricting the set of agents on the other side that each agent can trade with[1]. We

[1] In Chamberlin's experiment, buyers and sellers had to seek each other out to determine prices. This search cost

show that in this case, there is an instance with a specific adversarial order of recognizing agents such that following this order, the preferences of the agents (over the entire order) form a cycle and the DOA may never converge. Finally, we complete the story by showing that if we change the mechanism to recognize agents randomly, with high probability, a Walrasian equilibrium will be reached in polynomial number of steps. This further emphasizes the distinction between random recognition and arbitrary recognition for DOA-like mechanisms.

Organization: The rest of the paper is organized as follows. In Section 2, we formally introduce the model of the market and develop some concepts and notation used throughout the paper. Section 3 establishes a connection between the stable states of the market and social welfare. Our main results are presented in Section 4. We describe our DOA style mechanism and show that in markets with no trading restrictions, it converges in a number of steps that is polynomially bounded in the number of agents. We then show that when each agent is restricted to trade only with an arbitrary subset of agents on the other side, the mechanism need not converge. A randomized variant of our mechanism is then presented to overcome this issue. Finally, we conclude with some directions for future work in Section 5.

2 Preliminaries

We will use the terms 'player' and 'agent' interchangeably throughout the paper. We use B to represent a buyer, S for a seller, and Z for either of them. Also, b is used as the bid submitted by a buyer and s as the offer from a seller.

Definition 1 (Market)**.** *A market is denoted by $G(\mathcal{B}, \mathcal{S}, E, val)$, where $\mathcal{B}$ and $\mathcal{S}$ are the sets of buyers and sellers, respectively. Each buyer $B \in \mathcal{B}$ is endowed with a valuation of the item, and each seller $S \in \mathcal{S}$ has an opportunity cost for the item. We slightly abuse the terminology and refer to both of these values as the valuation of the agent for the item. The valuation of any agent Z is chosen from range $[0, 1]$, and denoted by $val(Z)$. Finally, E is the set of undirected edges between $\mathcal{B}$ and $\mathcal{S}$, which determines the buyer-seller pairs that may trade.*

Let $m = |E|$ and $n = |\mathcal{B}| + |\mathcal{S}|$.

Definition 2 (Market State)**.** *The* state *of a market at time t is denoted $\mathcal{S}^t(P^t, \Pi^t)$ ($\mathcal{S}(P, \Pi)$ for short, if time is clear or not relevant), where P is a* price function *revealing the price submission of each player and Π is a* matching *between $\mathcal{B}$ and $\mathcal{S}$, indicating which players are currently paired. In other words, the bid (offer) of a buyer B (seller S) is $P(B)$ ($P(S)$), and B, S are paired in Π iff $(B, S) \in \Pi$. In addition, we denote a player $Z \in \Pi$ iff Z is matched with some other player in Π, and denote his match by $\Pi(Z)$.*

Furthermore, the state where each buyer submits a bid of 0, each seller submits an offer of 1, and no player is matched is called the zero-information state*.*

We use the term zero-information because no player reveals non-trivial information about his valuation in this state.

Definition 3 (Valid State)**.** *A state is called* valid *iff (a_1) the price submitted by each buyer (seller) is lower (higher) than his valuation, (a_2) two players are matched only when there is an edge between*

meant that each agent was not necessarily aware of all prices on the other side of the market.

them, and (a_3) *for any pair in the matching, the bid of the buyer is no smaller than the offer of the seller.*

In the following, we restrict attention to states that are valid.

Definition 4 (Utility)**.** *For a market* $G(\mathcal{B}, \mathcal{S}, E, val)$ *at state* $\mathcal{S}(P, \Pi)$*, the* utility *of a buyer is defined as* $val(B) - P(B)$*, if* B *receives an item, and zero otherwise. Similarly, the utility of a seller is defined as* $P(S) - val(S)$*, if* S *trades his item, and zero otherwise.*

Note that what we have called utility is also called surplus.

Definition 5 (Stable State)**.** *A* stable state *of a market* $G(\mathcal{B}, \mathcal{S}, E, val)$ *is a state* $\mathcal{S}(P, \Pi)$ *s.t.* (a_1) *for all* $(B, S) \in E$*,* $P(B) \leq P(S)$ (a_2) *if* $Z \notin \Pi$*, then* $P(Z) = val(Z)$*, and* (a_3) *if* $(B, S) \in \Pi$*, then* $P(B) = P(S)$*.*

Suppose $\mathcal{S}(P, \Pi)$ is not stable. Then, one of the following must be true.

1. There exists $(B, S) \in E$ such that $P(B) > P(S)$. Then, both B and S could strictly increase their utility by trading with each other using the average of their prices.

2. There exists $Z \notin \Pi$ such that $P(Z) \neq val(Z)$. This agent could raise his bid (if a buyer) or lower his offer (if a seller), without reducing his utility and having a better opportunity to trade.

3. There exists $(B, S) \in \Pi$ such that $P(B) > P(S)$ ($P(B) < P(S)$ results in an invalid state). One of the agents could do better by either raising his offer or lowering his bid.

Definition 6 (ε-Stable State)**.** *For any* $\varepsilon \geq 0$*, a state* $\mathcal{S}(P, \Pi)$ *of a market* $G(\mathcal{B}, \mathcal{S}, E, val)$ *is* ε*-stable iff* (a_1) *for any* $(B, S) \in E$*,* $P(B) - P(S) \leq \varepsilon$ (a_2) *if player* $Z \notin \Pi$*,* $P(Z) = val(Z)$*, and* (a_3) *if* $(B, S) \in \Pi$*,* $P(B) = P(S)$*.*

Note that the only difference between a stable state and an ε-stable state lies in the first property. At any ε-stable state, no matched player will have a move to increase his utility by more than ε.

Definition 7 (Social Welfare)**.** *For a market* $G(\mathcal{B}, \mathcal{S}, E, val)$ *with a matching* Π*, the social welfare (SW) of this matching is defined as the sum of the valuation of the matched buyers minus the total opportunity cost of the matched sellers. We denote by* $\mathtt{SW}_\Pi$ *the SW of matching* Π*.*

Definition 8 (ε-approximate SW)**.** *For any market, a matching* Π *is said to give an* ε*-approximate SW if* $\mathtt{SW}_\Pi \geq \mathtt{SW}_{\Pi^*} - n\varepsilon$ *for any* Π^* *that maximizes SW. In other words, on average, the social welfare collected from each player using* Π *is at most* ε *less than that collected using* Π^**.*

3 Stable State and Social Welfare

In this section we mainly establish the connection between stables states and social welfare in the market. We emphasize that most results in this section are well known in the literature and stated here for the sake of completeness.

The problem of finding a matching that maximizes SW can be formulated as a linear program (LP) (see [2] for example). For any edge $(B, S) \in E$, let $x_{B,S}$ be the variable indicating whether

(B, S) is selected in the matching or not, and define weight of the edge, $w_{B,S} = val(B) - val(S)$. Therefore, the LP (primal) and its dual can be defined as follows.

$$\begin{aligned} \max \sum_{(B,S)\in E} & w_{B,S} \cdot x_{B,S} \\ \text{s.t. } \forall B^* \in \mathcal{B}, & \sum_{(B^*,S)\in E} x_{B^*,S} \le 1 \\ \forall S^* \in \mathcal{S}, & \sum_{(B,S^*)\in E} x_{B,S^*} \le 1 \\ & x_{B,S} \ge 0 \end{aligned} \qquad \begin{aligned} \min & \sum_{B\in\mathcal{B}} y_B + \sum_{S\in\mathcal{S}} y_S \\ \text{s.t. } & \forall (B,S) \in E, y_B + y_S \ge w_{B,S} \\ & y_B, y_S \ge 0 \end{aligned}$$

In the following, we will refer to the above linear programs as 'primal' and 'dual', respectively. The dual variables y can be interpreted as the utilities that agents enjoy assuming every buyer gets an item and every seller sells the item. Since it only depends on the price function, we call this price-wise utility. The constraint $y_B + y_S \ge w_{B,S}$ essentially states that the sum of the utilities obtained by (B, S) must be at least as large as their gains from trade.

We use x^Π to denote the characteristic function of matching Π, i.e., $x^\Pi_{B,S} = 1$ iff $(B, S) \in \Pi$, and use y^P to denote the price-wise utility function of a price function P, i.e., $y^P_B = val(B) - P(B)$ and $y^P_S = P(S) - val(S)$. It is well known that SW is maximized at a Walrasian equilibrium (see [2]) and we state here a similar result for stable states (a simple proof can be found in Appendix A.1).

Theorem 3.1. *A state $\mathcal{S}(P, \Pi)$ is stable iff x^Π is an optimal solution for the primal and y^P is an optimal solution for the dual.*

Theorem 3.1 states that any stable state maximizes SW. In other words, a stable state is a Walrasian equilibrium of the market. Moreover, *any* pair of optimal primal and dual solutions can form a stable state. We now show that for a sufficiently small ε, an ε-stable state is almost as good as stable states in terms of achieving maximum SW. We defer the proof of the following theorem to Appendix A.2.

Theorem 3.2. *For any market $G(\mathcal{B}, \mathcal{S}, E, val)$, for any $\varepsilon > 0$, any ε-stable state realizes an ε-approximate SW. Moreover, if we define $\delta = \min\{|val(Z_1) - val(Z_2)| \mid Z_1, Z_2 \in \mathcal{B} \cup \mathcal{S}, val(Z_1) \neq val(Z_2)\}$, then for $0 \le \varepsilon < \delta/n$, any ε-stable state maximizes SW.*

We note that [6] also shows that a ε-stable state realizes an ε-approximate SW. However, the bound on ε given in Theorem 3.2 is new. In [2], using ε-complementary slackness, Bertsekas shows that for integer valuations, any ε-stable state achieves maximum SW if $\varepsilon < 1/n$. Therefore, for fractional valuations, by scaling valuations with a suitably large factor L, one can make the valuations integers, and deduce that $\varepsilon < 1/(nL)$ suffices for achieving maximum SW. Note that L is at least $1/\delta$ but can possibly be much larger.

We should point out that the bound $\varepsilon < \delta/n$ is not an immediate consequence of the fact that any matching in an ε-stable state is an ε-approximate SW, by arguing that the smallest non-zero difference in SW of two matchings is at least δ. Consider a market whose trading graph is a complete bipartite graph, with four players, where $val(B_1) = 0.1$, $val(S_1) = 0.05$ $val(B_2) = 0.2001$, $val(S_2) = 0.15$. The difference of valuation price between any two players is lower bounded by 0.05 ($\delta = 0.05$) but B_1, S_1 yields a SW of 0.05 and B_2, S_2 yields a SW of 0.0501 and the difference in SW could be made arbitrarily small.

Finally, it is worth mentioning that the fact that ε-stable state gives ε-approximate SW does have a corollary as follows, which is a weaker result compared to Theorem 3.2: If for any $(B,S) \in E$, $val(B) - val(S)$ is an integer multiple of δ, then for any $0 \leq \varepsilon < \delta/n$, an ε-stable state always maximizes SW.

4 Convergence to a Stable State

We establish our main results in this section. We will start by describing a mechanism in the spirit of DOA, and show that for any *well-behaved* stable state, there is a sequence of agent moves that leads to this state. When the trading graph is a complete bipartite graph, i.e, the case of the DOA expriment, we show that convergence to a stable state occurs in number of steps that is polynomially bounded in the number agents. However, convergence to a stable state is not guaranteed when the trading graph is an incomplete bipartite graph. We propose a natural randomized extension of our mechanism, and show that with high probability, the market will converge to a stable state in number of steps that is polynomially bounded in the number of agents.

4.1 The Main Mechanism

To describe our mechanism, we need the notion of an ε-*interested* player.

Definition 9 (ε-Interested Player)**.** *For a market at state $\mathcal{S}(P,\Pi)$ with any parameter $\varepsilon > 0$, a seller S is said to be* ε-interested *in his neighbor B iff either (a) $P(B) \geq P(S)$ and $S \notin \Pi$, or (b) $P(B) - P(S) \geq \varepsilon$ and $S \in \Pi$. The set of buyers interested in a seller S is defined analogously.*

When the parameter ε is clear from the context, we will simply refer to an ε-*interested* player as an interested player.

Mechanism 1. *(with input parameter $\varepsilon > 0$)*

- ***Activity Rule:*** *Among the unmatched buyers, any buyer that neither submits a new higher bid nor has a seller that is interested in him, is labeled as inactive. All other unmatched buyers are labeled as active. An active (inactive) seller is defined analogously. An inactive player changes his status iff some player on the other side matches with him.*[2]
- ***Minimum Increment:*** *Each submitted price must be an integer multiple of ε.* [3]
- ***Recognition:*** *Among all active players, an arbitrary one is recognized.*
- ***Matching:*** *After a buyer B is recognized, B will choose an interested seller to match with if one exists. If the offer of the seller is lower than the bid b, it is immediately raised to b. The seller action is defined analogously.*
- ***Tie Breaking:*** *When choosing a player on the other side to match to, an unmatched player is given priority (the* unmatched first *rule).*

[2] This is common for eliminating no trade equilibria.
[3] This is part of many experimental implementations of the DOA.

In each iteration, players are partitioned into two sets based on whether they are matched or not. The unmatched players are further partitioned into active players and inactive players. The only players with a myopic incentive to revise their submissions are those that are not matched.

Observe that since a buyer will never submit a bid higher than his valuation, and a seller will never make an offer below his own opportunity cost, by submitting only prices that are integer multiples of ε, an agent might not be able to submit his true valuation. However, since an agent can always submit a price at most ε away from the true valuation, if we pretend that the 'close to valuation' prices are true valuations, the maximum SW will decrease by at most $n\varepsilon$. By picking $\varepsilon' = \varepsilon/2$, if the market converges to an ε'-stable state, we still guarantee that the SW of the final state is at most $n\varepsilon$ away from the maximum SW.

When a buyer B chooses to increase his current bid: if s denotes the lowest offer in the neighborhood of B, and s' denotes the lowest offer of any unmatched seller in the neighborhood of B, then the new bid of B can be at most $\min\{s + \varepsilon, s'\}$. We refer to this as the *increment* rule. This may be viewed as a consequence of rationality – there is no incentive for a buyer to bid above the price needed to make a deal with some seller. A similar rule applies to sellers. With a slight abuse of the terminology, we call either rules increment rule. Notice, a player indifferent between submitting a new price and keeping his price unchanged will be assumed to break ties in favor of activity.

Note that the role of the auctioneer in Mechanism (1) is restricted to recognize agent actions, but never select actions for agents. In fact, the existence of an auctioneer is not even necessary for the mechanism to work. Minimum increment can be interpreted as setting the currency of the market to be ε. Arbitrary recognition can be achieved by a first come, first served principle. Activity rule and matching are both designed to ensure that players will keep making actions (submitting a new price or forming a valid match) if one exists.

We first prove some properties of Mechanism (1).

Claim 4.1. *For any market, if we use Mechanism (1) with any input parameter $\varepsilon > 0$ and start from any state that satisfies properties (a_1) and (a_3) of ε-stable states, any state reached satisfies properties (a_1) and (a_3) of ε-stable states.*

By the increment rule and matching rule respectively, the reached state satisfies properties (a_1) and (a_3) of ε-stable states.

If a state $\mathcal{S}(P, \Pi)$ satisfies $\forall (B, S) \in E, P(B) \leq P(S)$, then we call it a *valid starting state*. Note that a valid starting state satisfies properties (a_1) and (a_3) of ε-stable states (a valid Π matches a buyer to a seller only if the bid price of the buyer is at least the offer price of the seller). In the following, we only consider markets that begin with a valid starting state, and hence a matched player will never have a move to increase his utility by more than ε.

Claim 4.2. *For any market, if we use Mechanism (1) and begin with a valid starting state, then any final state of the market is ε-stable.*

Claim 4.1 ensures that the final state satisfies properties (a_1) and (a_3) of ε-stable states, and property (a_2) of ε-stable states holds because an unmatched buyer will always submit a new higher bid to avoid being inactive, unless he reaches his valuation. Same for the unmatched sellers.

Note that by Theorem 3.2, if a market converges to an ε-stable state, it always realizes ε-approximate SW.

Definition 10 (Well-behaved)**.** *A stable state $\mathcal{S}(P, \Pi)$, is* well-behaved *iff (a_1) for any $(B, S) \in E$, if $B \notin \Pi$ and $S \notin \Pi$, then $P(B) < P(S)$. An ε-stable state $\mathcal{S}(P, \Pi)$, is* well-behaved *iff not* *only property (a_1) is satisfied but also (a_2) for any $(B, S) \in E$, if either $B \notin \Pi$ or $S \notin \Pi$, then $P(B) \leq P(S)$.*

Note that the states ruled out by properties (a_1) and (a_2) of well-behaved states are the corner cases where a buyer-seller pair having the same valuation (thus having no contribution to SW) are not chosen in the matching, or players who can obtain utility at most ε stop attempting to match with others.

Theorem 4.3. *For any $\varepsilon > 0$, if we use Mechanism (1), and start from the zero-information state, any well-behaved ε-stable state can be reached via a sequence of at most n moves. Hence, any well-behaved stable state is also reachable.*

Proof. Given an ε-stable state $\mathcal{S}(P, \Pi)$, sort all pairs in Π in decreasing order of prices (arbitrarily break the ties), and denote the ordering as O. We propose a two-stage procedure: first stage handles the players in the matching and second stage deals with the remainders. Note that we only need to justify that the increment rule and unmatched first rule hold for every action.

In stage one, choose pairs of players following the order defined by O. For each pair (B, S), let the buyer submit $P(B)$, and then, let the seller submit $P(S) = P(B)$ and match with B. When B submits $P(B)$, no seller is submitting a price lower then $P(B)$, hence the increment rule is satisfied. The unmatched sellers are submitting 1, and hence either no one is interested in B or all of them are interested in B (if $P(B) = 1$, i.e., $P(B)$ is no less than the seller prices). In the later case, B can directly match with S.

For S, assume the highest bid he can see in his neighborhood is $P(B')$ submitted by buyer B'. By property (a_1) of ε-stable states, $P(B') \leq P(S)+\varepsilon = P(B)+\varepsilon$. Among the unmatched neighbors of S, B is the one submitting the highest price, and $P(S) = P(B) \geq \max\{P(B') - \varepsilon, P(B)\}$, the increment rule is satisfied. Since S matches with unmatched buyer B, the unmatched first rule is also satisfied.

In stage two, choose the unmatched players with an arbitrary order and let them submit their valuations. For any unmatched buyer B, by property (a_2) of well-behaved states, $P(S) \geq P(B)$ for any seller S visible to B, hence the increment rule is satisfied. In addition, for any unmatched seller S, by property (a_1) of well-behaved states, $P(B) < P(S)$, thus B cannot match with S. By analogy, any unmatched seller will also make a valid move and remain unmatched.

Thus, after exactly n steps the two stages end, and the market is in state $\mathcal{S}(P, \Pi)$. □

4.2 Complete Bipartite Graphs

We now prove that market with complete bipartite trading graph will always converge when using Mechanism (1).

Theorem 4.4. *For a market whose trading graph is a complete bipartite graph, if we use Mechanism (1) with any input parameter $\varepsilon > 0$, and begin with any valid starting state, then the market will converge to an ε-stable state after at most n^3/ε steps.*

We need the following lemma to prove Theorem 4.4.

Lemma 4.1. *For a market $G(\mathcal{B}, \mathcal{S}, E, val)$ whose trading graph is a complete bipartite graph, if we use Mechanism (1) with any input parameter $\varepsilon > 0$, then at any state $\mathcal{S}(P, \Pi)$ reached from a valid starting state, for any $(B, S) \in E$, if $P(B) > P(S)$, then both B and S are matched.*

Proof. Assume by contradiction that there exists some $(B, S) \in E$ with $P(B) > P(S)$ and, wlog, B being unmatched. Since in the starting state, $P(B) \leq P(S)$, let t be the first time that this happens. Therefore, at time $t-1$, either $P(B) \leq P(S)$ or B is matched. Note that since the prices are integer multiples of ε, a state with $P(B) > P(S)$ implies $P(B) - P(S) \geq \varepsilon$. On the other hand, since property (a_1) of ε-stable states always holds, $P(B) - P(S) \leq \varepsilon$. Thus $P(B) = P(S) + \varepsilon$ at time t.

If $P(B) \leq P(S)$ at time $t-1$, $P(B) > P(S)$ can only be a consequence of either B or S being recognized. If B is recognized and submits a bid of $P(S) + \varepsilon$, since S is interested in B, by the matching rule, B will be matched. If S is recognized and submits an offer of $P(B) - \varepsilon$, by the increment rule, B must be matched (otherwise S would not submit an offer lower than $P(B)$), a contradiction.

Assume that B was matched to some seller S' at time $t-1$. The only valid action at time $t-1$ that can make B unmatched is if some buyer B' overbids B and match with seller S'. If $S = S'$, then after the move, $P(B) < P(S)$, a contradiction. If $S \neq S'$, then this move will not change the bid of B or offer price of S, and hence, $P(B) = P(S) + \varepsilon$ in time $t-1$. Since the trading graph is a complete bipartite graph, S is a neighbor of B'. By the increment rule, B' can only submit a price at most equal to $P(S) + \varepsilon = P(B)$, thus B' is unable to overbid B, a contradiction. $\square$

Definition 11 (γ-feasible)**.** *A market state $\mathcal{S}(P, \Pi)$ is said to be γ-feasible iff there are exactly γ matches in Π.*

Proof of Theorem 4.4. Assume at any time t, the state $\mathcal{S}^t$ of the market is γ^t-feasible. Define the following potential function

$$\Phi_P = \sum_{S_i \in \mathcal{S}} P(S_i) + \sum_{B_i \in \mathcal{B}} (1 - P(B_i))$$

Note that the value of Φ_P is always an integer multiple of ε. We will first show that γ^t forms a non-decreasing sequence over time, and then argue that, for any γ, the market can stay in a γ-feasible state for a bounded number of steps. Specifically, we will show that, if γ does not change, Φ_P is a non-increasing function and can stay unchanged for at most γ steps. Since the maximum value of Φ_P is bounded by n, it follows that after at most $(\gamma n)/\varepsilon$ steps, the market moves from a γ-feasible state to a $(\gamma+1)$-feasible state (or converges).

We argue that γ^t forms a non-decreasing sequence over time. Since any recognized player is unmatched, if the action of an unmatched player Z results in a change in the matching, Z either matches with another unmatched player, or matches to a player that was already matched. In either case, the total number of matched pairs does not decrease.

Furthermore, we prove if γ does not change, then Φ_P is non-increasing. Moreover, the number of successive steps that Φ_P stay unchanged is at most γ.

To see that Φ_P is non-increasing, first note that Φ_P can increase only when either a buyer decreases his bid or a seller increases his offer. Assume an unmatched buyer B is recognized (seller case is analogous), and the price function before his move is P. To increase Φ_P, since B can only increase his bid, he must increase an offer by overbidding and matching with a seller S, resulting in the two of them submitting the same price b. The buyer bid increases by $b - P(B)$ and the seller offer increases by $b - P(S)$. Since B is unmatched, by Lemma 4.1, $P(B) \leq P(S)$, and hence Φ_P will not increase.

We now bound the maximum number of steps for which Φ_P could remain unchanged. A move from a buyer B that does not change Φ_P occurs only when B overbids a matched seller S, where the bid and the offer are equal both before and after the move. We call this a *no-change* buyer move. By analogy, a no-change seller move can be defined.

In the remainder of the proof, we first argue that a no-change buyer move can never be followed by a no-change seller move, and vice versa. After that, we prove the upper bound on the number of consecutive no-change moves to show that Φ_P will eventually decrease (by at least ε).

Assume at time t_1, a buyer B_{t_1} made a no-change move and matched with a seller S_{t_1}, who was originally paired with the buyer B'_{t_1}.[4] We prove that no seller can make a no-change move at time $t_1 + 1$. The case that a seller makes a no-change move first can be proved analogously. Suppose at time $t_1 + 1$, a seller S_{t_1+1} is recognized and decreases his offer by ε. Since B_{t_1} made a no-change move, we have

$$P^{t_1}(B'_{t_1}) = P^{t_1}(B_{t_1}) \tag{1}$$

Denote the lowest seller offer (highest buyer bid) at any time t by s^t (b^t). By Lemma 4.1, $P^{t_1}(B_{t_1}) \leq P^{t_1}(S)$ for any seller S, hence $P^{t_1}(B_{t_1}) \leq s^{t_1}$. Moreover, since $P^{t_1}(B_{t_1}) = P^{t_1}(S_{t_1}) \geq s^{t_1}$, we have

$$P^{t_1}(B_{t_1}) = s^{t_1} \tag{2}$$

In other words, a buyer can make a no-change move, only if his bid is equal to the lowest offer. Similarly, if S_{t_1+1} can make a no-change move at time $t_1 + 1$, his offer is equal to the highest bid. Since the highest bid at time t_1 (b^{t_1}) is at most $s^{t_1} + \varepsilon$ (property (a_1) of ε-stable states), after B_{t_1} submits a bid of $s^{t_1} + \varepsilon$, he will be submitting the highest bid at time $t_1 + 1$. Hence

$$P^{t_1+1}(S_{t_1+1}) = b^{t+1} = P^{t_1+1}(B_{t_1}) = P^{t_1+1}(B'_{t_1}) + \varepsilon \tag{3}$$

Therefore, at time $t_1 + 1$, after S_{t_1+1} decreases his offer by ε, the unmatched buyer B'_{t_1} is interested in S_{t_1+1}. By the unmatched first rule, S_{t_1+1} will match with an unmatched player, hence this cannot be a no-change move.

This proves that a no-change seller move can never occur after a no-change buyer move and vice versa. We now prove the upper bound on the number of consecutive no-change buyer moves.

For any sequence of consecutive no-change buyer moves, if there exists a time t_2 such that $s^{t_2} > s^{t_2-1}$, for any unmatched buyer B at time t_2, $P^{t_2}(B) \leq s^{t_2-1} < s^{t_2}$. By Equation (2), no buyer can make any more no-change move. Moreover, since any no-change buyer move will increase the submission of a matched seller who is submitting the lowest offer, after at most γ steps, the lowest offer must increase, implying that the length of the sequence is at most γ.

To conclude, the total number of steps that the market could stay in γ-feasible states is bounded by $(n/\varepsilon)\gamma$. As $\gamma \leq n$, the total number of steps before market converges is at most n^3/ε. □

4.3 General Bipartite Graphs

In this section, we study the convergence of markets with an arbitrary bipartite trading graph. Although by Theorem 4.3, using Mechanism (1), the market can reach any well-behaved ε-stable state, when the trading graph of a market can be an arbitrary bipartite graph, there is no guarantee that Mechanism (1) will actually converge.

[4] An action at time time t will take effect at the time $t + 1$, and P^t is the price function before any action is made at time t.

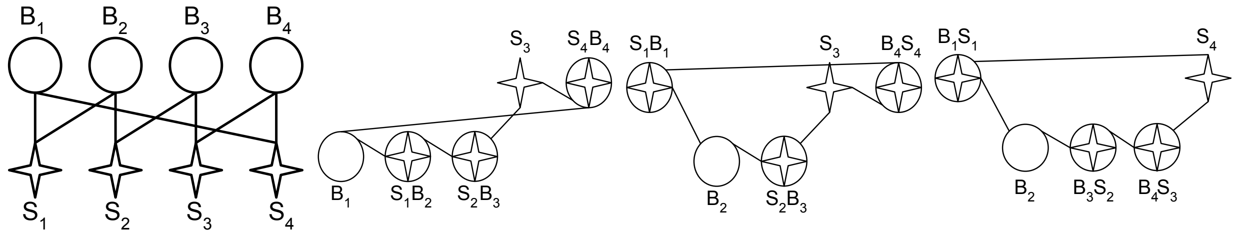

Figure 1: Unstable market with general trading graph and Mechanism (1)

Claim 4.5. *In a market whose trading graph is an arbitrary bipartite graph, Mechanism (1) may never converge.*

Consider the market shown in Figure 1. In this market, there are four buyers (B_1 to B_4) all with valuation 1 and four sellers (S_1 to S_4) all with opportunity cost 0. Moreover, the trading graph is a cycle of length 8, as illustrated by the first graph in Figure 1. Assume at some time t, the market enters the state illustrated by the second graph, where B_1, B_2, S_1, B_3, S_2 are submitting 5ε, S_3, B_4, S_4 are submitting 6ε, and pairs (B_2, S_1), (B_3, S_2) and (B_4, S_4) are matched.

At time $t+1$, since B_1 is unmatched, he can be recognized and submit 6ε. S_1 is the only interested seller, hence B_1, S_1 will match and the offer of S_1 increases to 6ε, which leads to the state shown in the third graph. Similarly, at time $t+2$, since S_3 is unmatched, he can be recognized and submit 5ε. B_4 is the only interested buyer, hence B_4, S_3 will match and bid of B_4 increases to 6ε, which leads to the state shown in the fourth graph.

Notice that the states at time t and $t+2$ are isomorphic. By shifting the indices and repeating above two steps, the market will never converge.

Observe that the cycle described in Claim 4.5 is caused by an adversarial coordination between the actions of various agents. To break this pathological coordination, we introduce Mechanism (2) which is a natural extension of Mechanism (1) that uses randomization. We first define this mechanism, and then prove that on any trading graph, with high probability, the mechanism leads to convergence in a number of steps that is polynomially bounded in the number of agents.

Mechanism 2. *(with input parameter $\varepsilon > 0$)*

- ***Activity Rule:*** *Among the unmatched buyers, any buyer that neither submits a new higher bid nor has a seller that is interested in him, is labeled as inactive. All other unmatched buyers are labeled as active. An active (inactive) seller is defined analogously. An inactive player changes his status iff some player on the other side matches with him.*

- ***Minimum Increment:*** *Each submitted price must be an integer multiple of ε.*

- ***Bounded Increment Rule:*** *In each step, a player is only allowed to change his price by ε.*

- ***Recognition:*** *Among all players who are active, one is recognized uniformly at random.*

- ***Matching:*** *After a player, say a buyer B, is recognized, if B does not submit a new price, then B will match to an interested seller if one exists. If the offer of the seller is lower than the bid b, it is immediately raised to b. The seller action is defined analogously.*

- ***Tie Breaking:*** *When choosing a player on the other side to match to, an unmatched player is given priority (unmatched first rule).*

Notice that we ask players to move cautiously through the bounded increment rule. Players can either change the price by ε or match with an interested seller, and always favor being active. Note that, any move in Mechanism (1) can be simulated by at most $(1/\varepsilon+1)$ moves in Mechanism (2) ($1/\varepsilon$ for submitting new price and 1 for forming a match). The following is an immediate consequence of results shown in Section 4.1.

Corollary 4.6. *For any market, if we use Mechanism (2), (i) starting from the zero-information state, any well-behaved ε-stable state can be reached in $n(1/\varepsilon + 1)$ steps, and (ii) beginning with a valid starting state, properties (a_1) and (a_3) of ε-stable states always hold, and the final state is ε-stable.*

We are now ready to prove our second main result, namely, for any trading graph, with high probability, Mechanism (2) converges to a ε-stable state in a number of steps that is polynomially bounded in the number of agents. We will utilize the following standard fact about random walk on a line (see [12], for instance).

Claim 4.7. *Consider a random walk on $\{0, 1, 2, ..., N\}$ such that for any $i \in [1, N]$, the random walk transition from state i to state $(i-1)$ happens with probability α, and for any $i \in [0, N-1]$, the random walk transition from state i to state $(i+1)$ happens with probability β, for some $\alpha + \beta = 1$. Then starting from any $i \in [0, N]$, with probability at least $1/2$, the random walk either reaches the state 0 or the state N, after $O(N^2)$ steps.*

Theorem 4.8. *For any market $G(\mathcal{B}, \mathcal{S}, E, val)$, if we use Mechanism (2) with any input parameter $\varepsilon > 0$, and begin with a valid starting state, the market will converge to an ε-stable state after at most $O((n^3/\varepsilon^2)\log n)$ steps with high probability.*

Proof. Let $u^t_{\mathcal{B}}$ and $u^t_{\mathcal{S}}$ denote the number of active buyers and sellers at time t, respectively, and let $u^t = u^t_{\mathcal{B}} + u^t_{\mathcal{S}}$. We will first show that $u^t_{\mathcal{B}}$ and $u^t_{\mathcal{S}}$ are both non-increasing functions of time and then argue that for any u, with high probability the market will remain in a state with u active players for a number of steps that is polynomially bounded in the number of players.

We first prove $u^t_{\mathcal{B}}$ and $u^t_{\mathcal{S}}$ are non-increasing. Note that the only move that can make a new player active is one where a player, say a buyer B, matches to a currently matched seller S. Let B' be the buyer that is currently matched to S. Then at time $t+1$, the buyer B moves out of the set of active players, while the buyer B' possibly joins the set of active players. Thus the number of active players remains unchanged. A similar argument applies to case when a seller is recognized and matches to a currently matched buyer.

In the remainder of the proof, we first show that if there exists an adjacent buyer-seller pair such that both players are unmatched and the buyer bid is not below the seller offer (we call such a pair to be an *active pair*), then after $O(n \log n)$ steps, with probability $1 - O(1/n^2)$, u^t will decrease. Next, in the absence of active pairs, we argue that either u^t decreases or an active pair appears in the market after $O((n/\varepsilon)^2 \log n)$ steps, with probability $1 - O(1/n^2)$. Note that if an active pair appears, by the same argument, after $O(n \log n)$ more steps, u^t will decrease with high probability. Since $u^t \le n$, we can conclude that the market converges in $O((n^3/\varepsilon^2)\log n)$ steps with high probability.

We first prove that, the existence of an active pair will lead to decrement of u^t. For any active pair (B, S). By the unmatched first rule, recognizing either B or S will increase the number of matches. Recognizing any other player who makes a move to match with B or S, will also increase the number of matches (note that only unmatched players will be recognized). Both cases decrease

u^t by 2. In other words, as long as u^t does not decrease, (B, S) will remain to be an active pair. Assume at time t_1, there is an active pair (B, S). Let Y_t be the random variable which is 1 iff B or S is recognized at time t. Hence, for any $t \geq t_1$ where $u^t = u^{t_1}$,

$$Pr(Y_t = 1) = \frac{2}{u^{t_1}} \geq \frac{2}{n}$$

It follows that after n steps from t_1 the probability that none of B or S is chosen is less than $1/2$. Therefore, after $2n \log n$ steps, with probability $1 - (1/n)^2$, either one of B and S has been recognized or u^t has already decreased. In either case, u^t decreases.

Next, in the absence of active pairs, we prove that after a bounded number of steps, either u^t decreases or an active pair appears. Consider the following potential function

$$\Phi = \sum_{B \in \mathcal{B}} P(B) + \sum_{S \in \mathcal{S}} P(S)$$

If there is no active pairs, by the design of the Mechanism (2), when recognized, any buyer will increase Φ by ε and any seller will decrease Φ by ε. Thus, at any time t with no active pairs, the probability that Φ increases by ε is $P^t_{\mathcal{B}} = u^t_{\mathcal{B}}/u^t$, and the probability that Φ decreases by ε is $P^t_{\mathcal{S}} = u^t_{\mathcal{S}}/u^t$ (note that $P^t_{\mathcal{B}}$ or $P^t_{\mathcal{S}}$ might be 0).

In the following, we will use Claim 4.7 to prove that as long as u^t does not change and no active pair appears, after a bounded number of steps, with high probability, Φ will have reached its upper or lower bound. If Φ reaches its upper bound then all buyers must be submitting their true valuations and all sellers must be submitting 1. Thus every unmatched buyer is inactive and u^t must have decreased. A similar situation also happens when Φ reaches its lower bound.

To use Claim 4.7, see that if u^t does not change and no active pair appears, $P^t_{\mathcal{B}}$ and $P^t_{\mathcal{S}}$ will also remain unchanged. During this time period, we can denote the probability of Φ increases by ε as $P_{\mathcal{B}}$ and the probability of Φ decreases by ε as $P_{\mathcal{S}}$. Let $\alpha = P_{\mathcal{S}}$, $\beta = P_{\mathcal{B}}$, and nodes be $\{0, \varepsilon, 2\varepsilon, \ldots, n\}$ (hence $N = n/\varepsilon$). Thus this is a random walk, and by Claim 4.7, after $O((n/\varepsilon)^2)$ steps, the probability of Φ reaches its upper bound n or lower bound 0 is at least $1/2$. Therefore, after $O((n/\varepsilon)^2 \log n)$ steps, with probability at least $1 - O(1/n^2)$, Φ reaches 0 or n.

To conclude, after $O((n/\varepsilon)^2 \log n)$ steps, u^t will decrease with probability at least $1 - O(1/n^2)$. As $u^t \leq n$, by union bound, the market will converge after $O((n^3/\varepsilon^2) \log n)$ steps with probability $1 - O(1/n)$.

□

5 Conclusions

We resolved the second part of Friedman's conjecture by designing a mechanism which simulates the DOA and proving that this mechanism always converges to a Walrasian equilibrium in polynomially many steps. Our mechanism captures four key properties of the DOA: agents on either side can make actions; agents only have limited information; agents can choose *any* better response (as opposed to the best response); and the submissions are recognized in an arbitrary order. An important aspect of our result is that, unlike previous models, *every* Walrasian equilibrium is reachable by some sequence of better responses.

For markets where only a restricted set of buyer-seller pairs are able to trade, we show that the DOA may never converge. However, if submissions are recognized randomly, and players only change their bids and offers by a small fixed amount, convergence is guaranteed. It is unclear that the latter condition is inherently necessary, and perhaps a convergence result can be established for a relaxed notion of bid and offer changes where players can make possibly large adjustments as long as they are consistent with the increment rule.

References

[1] Dimitri Bertsekas. A distributed algorithm for the assignment problem. *Lab. for Information and Decision Systems, Working Paper, M.I.T., Cambridge, MA*, 1979.

[2] Dimitri P Bertsekas. *Linear network optimization: algorithms and codes.* MIT Press, 1991.

[3] Dimitri P. Bertsekas and David A. Castañon. A forward/reverse auction algorithm for asymmetric assignment problems. *Computational Optimization and Applications*, 1(3):277–297, 1992.

[4] Edward H Chamberlin. An experimental imperfect market. *The Journal of Political Economy*, 56(2):95–108, 1948.

[5] Vincent P Crawford and Elsie Marie Knoer. Job matching with heterogeneous firms and workers. *Econometrica: Journal of the Econometric Society*, 1981.

[6] Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *The Journal of Political Economy*, pages 863–872, 1986.

[7] Daniel P Friedman and John Rust. *The double auction market: institutions, theories, and evidence*, volume 14. Westview Press, 1993.

[8] David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, pages 9–15, 1962.

[9] Yashodhan Kanoria, Mohsen Bayati, Christian Borgs, Jennifer Chayes, and Andrea Montanari. Fast convergence of natural bargaining dynamics in exchange networks. In *SODA*, 2011.

[10] D.E. Knuth. *Mariages stables et leurs relations avec d&autres problèmes combinatoires:*. Collection de la Chaire Aisenstadt. Presses de l'Université de Montréal, 1976.

[11] Harold W. Kuhn. The hungarian method for the assignment problem. In *50 Years of Integer Programming 1958-2008 - From the Early Years to the State-of-the-Art*, pages 29–47. 2010.

[12] Rajeev Motwani and Prabhakar Raghavan. *Randomized algorithms.* Chapman & Hall/CRC, 2010.

[13] Heinrich H. Nax, Bary S. R. Pradelski, and H. Peyton Young. Decentralized dynamics to optimal and stable states in the assignment game. In *CDC*, 2013.

[14] Bary S.R. Pradelski. Decentralized dynamics and fast convergence in the assignment game: Extended abstract. In *EC*, New York, NY, USA, 2015. ACM.

[15] Alvin E Roth and John H Vande Vate. Random paths to stability in two-sided matching. *Econometrica: Journal of the Econometric Society*, 1990.

[16] L.S. Shapley and M. Shubik. The assignment game i: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.

[17] Vernon L Smith. An experimental study of competitive market behavior. *The Journal of Political Economy*, 70(2):111–137, 1962.

[18] Vernon L Smith. *Papers in experimental economics*. Cambridge University Press, 1991.

A Omitted Details of Section 3

A.1 Proof of Theorem 3.1

Proof. To see the forward direction, assume $\mathcal{S}(P, \Pi)$ is a stable state. We first verify that x^Π and y^P are indeed feasible solutions. x^Π is clearly feasible since it is characteristic function of a valid matching. y^P preserves non-negativity constraint of dual, since no player could submit a price exceeding his valuation in P. Moreover, we can write $y^P_B + y^P_S = val(B) - P(B) + P(S) - val(S) = w_{B,S} + P(S) - P(B)$. By property (a_1) of stable states, $P(S) \geq P(B)$, hence $y^P_B + y^P_S \geq w_{B,S}$, preserving the dual constraint and implying that y^P is also feasible.

To prove optimality of x^Π and y^P, using weak duality, we only need to verify that value of primal is equal to value of dual.

$$\sum_{B \in \mathcal{B}} y_B + \sum_{S \in \mathcal{S}} y_S = \sum_{(B,S) \in \Pi} (val(B) - P(B) + P(S) - val(S)) + \sum_{Z \notin \Pi} y^P_Z \tag{4}$$

$$= \sum_{(B,S) \in \Pi} (val(B) - val(S)) = \sum_{(B,S) \in E} w_{B,S} \times x^\Pi_{B,S} \tag{5}$$

(1) to (2) uses properties (a_1) and (a_2) of stable states, $P(B) = P(S)$ for $(B, S) \in \Pi$, and $P(Z) = val(Z)$ for $Z \notin \Pi$. Thus x^Π and y^P are optimal solutions of primal and dual, respectively.

For the reverse direction, assume (x^Π, y^P) is a pair of optimal primal and dual solutions . Since y^P is a feasible solution, as we just stated, $y^P_B + y^P_S \geq w_{B,S}$ will give us $P(S) \geq P(B)$, thus property (a_1) of stable states holds. For properties (a_2) and (a_3) of stable states, since

$$\sum_{B \in \mathcal{B}} y_B + \sum_{S \in \mathcal{S}} y_S = \sum_{(B,S) \in E} w_{B,S} \times x^\Pi_{B,S}$$

$$\sum_{(B,S) \in \Pi} (val(B) - P(B) + P(S) - val(S)) + \sum_{Z \notin \Pi} y^P_Z = \sum_{(B,S) \in \Pi} val(B) - val(S)$$

$$\sum_{(B,S) \in \Pi} (P(S) - P(B)) + \sum_{Z \notin \Pi} y^P_Z = 0$$

Since $P(S) \geq P(B)$ and also y^P is a non-negative vector, both terms in the last expression must be zero, which implies that properties (a_2) and (a_3) of stable states also hold. Therefore $\mathcal{S}(P, \Pi)$ is a stable state. $\square$

A.2 Proof of Theorem 3.2

To simplify the notation, we treat an agent who is unmatched as being matched with themselves. To this end, for each buyer we introduce a *dummy* seller with an opportunity cost equal to his valuation, similarly for each seller. An agent matched with their dummy counterpart is interpreted as being unmatched. We denote the dummy seller of buyer B as $\bar{S}_B$ and the dummy buyer of seller S as $\bar{B}_S$.

Proof. We define the following ε-primal and ε-dual pair.

$$\begin{aligned} &\max \sum_{(B,S)\in E} (w_{B,S} - \varepsilon) x_{B,S} && \min \sum_{B\in\mathcal{B}} y_B + \sum_{S\in\mathcal{S}} y_S \\ &\text{s.t. } \forall B^* \in \mathcal{B}, \sum_{(B^*,S)\in E} x_{B^*,S} \le 1 && \text{s.t. } \forall (B,S)\in E, y_B + y_S \ge (w_{B,S} - \varepsilon) \\ &\forall S^* \in \mathcal{S}, \sum_{(B,S^*)\in E} x_{B,S^*} \le 1 && y_B, y_S \ge 0 \\ &x_{B,S} \ge 0 \end{aligned}$$

Given a ε-stable state $\mathcal{S}(P,\Pi)$, since property (a_1) of ε-stable states is equivalent to $val(B) - P(B) + P(S) - val(S) \ge w_{B,S} - \varepsilon$, y^P is a feasible solution of ε-dual (non-negativity constrains hold since P is valid price function). By properties (a_2) and (a_3) of ε-stable states,

$$\begin{aligned} \sum_{B\in\mathcal{B}} y_B^P + \sum_{S\in\mathcal{S}} y_S^P &= \sum_{(B,S)\in\Pi} (val(B) - P(B) + P(S) - val(S)) + \sum_{Z\notin\Pi} y_Z^P \\ &= \sum_{(B,S)\in\Pi} (val(B) - val(S)) = \mathtt{SW}_\Pi \end{aligned}$$

On the other hand, if we take a matching Π^* that maximizes SW, and plug x^{Π^*} into the ε-primal, we have

$$\sum_{(B,S)\in E} (w_{B,S} - \varepsilon) x_{B,S} = \sum_{(B,S)\in E} w_{B,S} x_{B,S} - \sum_{(B,S)\in E} \varepsilon x_{B,S} \ge \mathtt{SW}_{\Pi^*} - n\varepsilon$$

The last inequality comes from the fact that n is an upper bound on the number of possible pairs (i.e., number of possible 1's in $x_{B,S}$) for any matching. By the weak duality, any value of ε-primal is less than or equal to any value of ε-dual, thus $\mathtt{SW}_\Pi \ge \mathtt{SW}_{\Pi^*} - n\varepsilon$.

We now proceed to prove the condition for an ε-stable state to maximize SW. Fix a matching Π^* that maximizes SW. Construct graph $G'(V', E')$ with $V' = \mathcal{B} \cup \mathcal{S}$ and $E' = \{(B,S) \mid (B,S) \in \Pi \vee (B,S) \in \Pi^*\}$. As any player can be matched with at most one other player in each matching, the degree of each node in G' is at most two. Consequently, the connected components of G' could only be cycles or paths. Note that such cycles and paths are formed by the different pairs of the two matchings. We now prove that for any of those cycles or paths, the local SW of the two matchings are the same.

For any cycle $B_0, S_0, B_1, S_1, \ldots, B_k, S_k, B_0$, pair (B_i, S_i) belongs to one matching while pair (B_{i+1}, S_i) belongs to the other one. If we only consider these players, every buyer gets an item and every seller sells the item, thus the SW of both matchings are the same.

For any path $Z_0, Z_1, Z_2, Z_3, \ldots, Z_k$, wlog, we can assume $val(Z_0) \ge val(Z_k)$. If Z_0 is a seller, add his dummy buyer to the left of the path. If Z_k is a buyer, add his dummy seller to the right of the path. Therefore, the path starts with a buyer and ends with a seller. We can denote the path as $B_0, S_0, B_1, S_1, \ldots, B_k, S_k$.

For the same reason as cycle case, the players in the middle contributes same amount of SW to both matchings, thus the difference of SW is $val(B_0) - val(S_k)$. Since Π^* is a matching that maximizes SW, it must be the case that $(B_i, S_i) \in \Pi^*$ and $(B_{i+1}, S_i) \in \Pi$.

If the difference of SW is 0, then we are done. Suppose not, then $val(B_0) - val(S_k) \ge \delta$. By properties (a_1) and (a_3) of ε-stable states,

$$P(B_{i+1}) = P(S_i) \ge P(B_i) - \varepsilon \Rightarrow P(B_0) \le P(B_k) + k\varepsilon$$

We now have

$$val(B_0) - val(S_k) = P(B_0) - P(S_k) \tag{6}$$
$$\leq P(B_k) + k\varepsilon - P(S_k) \leq (k+1)\varepsilon \leq n\varepsilon < \delta \tag{7}$$

where (3) is because both B_0 and S_k are matched in Π^* but not in Π, implying that their submitted prices are equal to their valuation.

Thus on one side we have $val(B_0) - val(S_k) \geq \delta$, and from the other side by inequality (7) above we have $val(B_0) - val(S_k) < \delta$, a contradiction. It concludes that all such cycles and paths generate the same SW for both matchings and thus Π also maximizes SW. $\square$