Equilibrium Selection in the Stable Marriage Problem: Experimental Evidence

Marco Castillo* Ahrash Dianat**

Abstract

We present experimental evidence on equilibrium selection in the stable marriage problem. By automating the side of the market that has a dominant strategy, we induce a coordination game with two symmetric and Pareto-ranked equilibria: an equilibrium in "truncation" strategies and an equilibrium in "permutation" strategies. This construction allows us to apply the equilibrium selection concepts of payoff-dominance and risk-dominance in the context of two-sided matching. By keeping subjects' ordinal preferences fixed while changing their cardinal representation, our experimental treatments vary the risk-dominant equilibrium prediction. We observe several regularities in the patterns of equilibrium play. First, both equilibrium strategies are played more often when they are risk-dominant. Second, the unique payoff-dominant strategy is played more often in later rounds of the experiment. Our results also provide support for the empirical relevance of truncation strategies in centralized matching clearinghouses.

JEL codes: C72, C78, D47
Keywords: stable matching; equilibrium selection; risk-dominance

* Department of Economics, Texas A&M University, marco.castillo@tamu.edu
** Division of the Humanities and Social Sciences, California Institute of Technology, adianat@caltech.edu

This research was supported by the International Foundation for Research in Experimental Economics (IFREE).

1 Introduction

In economics and game theory, the *stable marriage problem* refers to the question of how to find a stable matching between two finite and disjoint sets of agents, given a preference ordering for each agent.[1] The theoretical structure of this problem is rich and well-studied.[2] Moreover, in recent decades the basic theory has been extended and successfully applied to the design of matching markets in the real world.[3] These markets are typically organized as centralized clearinghouses, in which participants report rank-order lists of their preferences and then a particular algorithm uses the reported lists to calculate the final outcome (i.e., who is matched with whom).

Both theory and practice suggest the use of matching clearinghouses that produce stable outcomes.[4] The emphasis on stability introduces the question of *which* stable matching to implement when there are multiple candidates for consideration. It is well-known that, in markets where all agents have identical preferences, there is a unique stable outcome characterized by positive assortative matching. However, this is an unreasonable assumption in most real-world settings. We often wish to model heterogeneity in preferences, which necessarily introduces the possibility of multiple stable matchings.

Moreover, there can be significant welfare implications associated with the selection of stable matchings. When attention is confined to stable outcomes, the interests of the two sides of the market are opposed in a fundamental sense: the best stable matching for one side of the market is the worst stable matching for the other side of the market.[5] This makes selection among stable matchings an important and relevant consideration for policymakers, who may have reasons to favor the welfare of one side of the market over another when designing matching institutions. An example of this is provided by the history of the National Resident Matching Program (NRMP), the entry-level labor market for American physicians. In May 1997, the NRMP unanimously voted to alter the algorithm that was being used over concerns that the original design unduly favored hospitals at the expense of students.[6]

Under the assumption of truthful preference reporting, matching clearinghouses used in practice implement an extremal stable outcome (i.e., the most preferred stable outcome for one side of the market). However, in markets with more than one stable matching, truth-telling is not a dominant strategy for all agents. If agents strategically misrepresent their preferences, then it is less clear

[1] A matching is said to be stable if no agent prefers remaining unmatched to her current allocation and no pair of agents both prefer each other to their current allocations.

[2] For a survey of the theoretical results, see Roth and Sotomayor (1992).

[3] For an application to school choice programs, see Abdulkadiroğlu and Sönmez (2003). For an application to the medical residency matching program, see Roth and Peranson (1999).

[4] Indeed, Roth (1991) provides evidence that centralized markets producing unstable outcomes often perform no better than the decentralized markets they replace.

[5] This is a consequence of the fact that the set of stable matchings has a lattice structure.

[6] Specifically, the NRMP switched from a version of the hospital-proposing deferred acceptance algorithm to a version of the student-proposing deferred acceptance algorithm.

which stable matching is implemented or even whether the final matching is stable. To gain insight on selection among stable outcomes, we address the related question of equilibrium selection in the *induced preference-revelation game*.[7] There is a close connection between the set of stable matchings and the set of Nash equilibria: any Nash equilibrium of the induced preference-revelation game produces an outcome that is stable with respect to the true preferences (Roth, 1984) and any stable outcome with respect to the true preferences can be achieved through a Nash equilibrium profile of reported preferences (Gale and Sotomayor, 1985).[8]

A popular stable matching mechanism used in field applications is the deferred acceptance algorithm, originally introduced by Gale and Shapley (1962). The deferred acceptance algorithm divides the market into "proposers" and "responders" and then implements a chain of offers and acceptance/rejection decisions that generate the proposer-optimal stable matching (with respect to the reported preferences). In this environment, it is well-known that the proposers have a dominant strategy of truth-telling (Dubins and Freedman, 1981). However, in markets with more than one stable matching, at least one responder will have an incentive to misrepresent her preferences (Gale and Sotomayor, 1985). Fixing the behavior of the proposers, we identify a coordination game among the responders with respect to the manner in which they misrepresent their preferences.

Specifically, the coordination game has at least two symmetric and Pareto-ranked equilibria in pure strategies: an equilibrium in "truncation" strategies (i.e., removing less preferred match partners from the tail end of a preference list) and an equilibrium in "permutation" strategies (i.e., switching the order of match partners in a preference list). Although the truncation equilibrium yields a higher payoff for all responders, the presence of strategic uncertainty makes truncation strategies less appealing. Intuitively, truncation generates a trade-off between the likelihood of matching and the quality of match partner (conditional on matching). Thus, an agent who plays a truncation strategy opens herself up to the possibility of remaining unmatched for some profile of other agents' preference reports.[9] This insight allows us to apply equilibrium selection concepts from a class of coordination games called *stag hunt* games, which feature a similar tension between profitability and safety. In particular, we will use the Harsanyi and Selten (1988) characterization of *risk-dominance* to capture the intuition of one equilibrium being more or less "risky" than another equilibrium.

However, standard equilibrium selection arguments may not have full explanatory power in this environment for several reasons. First, truncation behavior by one individual creates positive

[7] The induced preference-revelation game is defined in Section 2.

[8] These equivalences do not necessarily hold for the more general case of many-to-one matching (often referred to as the "college admissions problem"). For a discussion of the differences between models of one-to-one matching and many-to-one matching, see Roth (1985).

[9] Even after removing this aspect of strategic uncertainty, there is also the risk of "over-truncation" (i.e., playing the wrong kind of truncation strategy). However, this risk is not present in our experimental set-up. In a related experiment, Castillo and Dianat (2016) present evidence that laboratory subjects are less likely to truncate their preferences when the possibility of over-truncation exists.

spillovers for other individuals on that side of the market.[10] This implies that asymmetric equilibria also exist in which a subset of agents play truncation strategies and the remaining agents report their true preferences, effectively free-riding on others' truncation behavior. Second, the permutation strategy we identify is weakly dominated by truth-telling.[11,12] In fact, truth-telling is the unique *protective* strategy for all agents in the deferred acceptance mechanism (Barberà and Dutta, 1995).[13] In our context, this is because truth-telling accomplishes two tasks: (1) it secures against the worst possible outcome (i.e., remaining single) and (2) it leads to the best possible outcome for some profile of other agents' preference reports. Thus, while no stable matching mechanism admits truth-telling as a dominant strategy for all agents (Roth, 1982), truthful behavior might carry natural appeal as a heuristic.

We study this coordination game in the lab, using a simple setting that allows us to isolate the strategic features of interest. Each experimental market consists of four participants: two firms (automated roles) and two workers (subject roles). Subjects play 20 rounds of the induced preference-revelation game based on the firm-proposing deferred acceptance algorithm, with random and anonymous re-pairing across rounds. We use an ordinal constellation of preferences in which each subject has two stable match partners. However, these ordinal preferences can be represented by different cardinal utilities. Indeed, we have full freedom in choosing the payoff difference between two ordered alternatives. In our experiment, we choose cardinal representations such that our treatments vary whether the criterion of risk-dominance selects the truncation or the permutation equilibrium. When remaining unmatched is particularly costly (to be precisely defined later), then permutation is risk-dominant. When an agent has a strong intensity of preference for her first choice partner (to be precisely defined later), then truncation is risk-dominant.

We now preview our main results. Overall, we find that truthful preference reporting is more common than either equilibrium strategy. However, we observe several regularities in the patterns of equilibrium play. First, both truncation and permutation strategies are played significantly more often when they are the risk-dominant equilibrium prediction. This result is more prominent in later rounds of the experiment and is robust to whether the treatment effect is measured at the subject or the session level. Second, in both treatments, truncation strategies are played significantly more often and permutation strategies are played significantly less often in later rounds of the experiment. This suggests that the salience of payoff-dominance as a selection principle increases with subject

[10]This result has been proven for the complete-information case by Ashlagi and Klijn (2012). A similar result has been shown to hold when comparing Bayesian-Nash equilibria of the incomplete-information version of the game Coles and Shorrer (2014).

[11]In fact, any permutation strategy that does not list an agent's true first choice at the head of her list is weakly dominated (Roth and Sotomayor, 1992).

[12]The "reasonableness" of equilibria in weakly dominated strategies remains a question of interest. Although arguments from evolutionary game theory eliminate strictly dominated strategies from consideration, they need not eliminate weakly dominated strategies entirely (Samuelson, 1998).

[13]A protective strategy is a refinement of a maximin strategy.

experience.

To the best of our knowledge, we are the first to identify the conditions under which preference misrepresentation in the stable marriage problem can be modeled as a coordination game and to apply standard equilibrium selection arguments to this context. As such, our paper naturally fits into several different strands of literature. First, there is a growing body of work on the performance of matching mechanisms in the lab.[14] Our work is most closely related to Castillo and Dianat (2016), which is an experimental investigation of truncation behavior in what approximates a decision-theoretic setting. To that end, Castillo and Dianat (2016) employ a restricted strategy space that yields a unique equilibrium outcome. The current design, on the other hand, introduces multiple Pareto-ranked equilibria and allows us to better understand the conditions that favor the implementation of one equilibrium over another. There are also experimental studies of deferred acceptance mechanisms that investigate whether intensity of preference has implications for strategic behavior. For instance, Echenique et al. (2016) report that the cardinal representation of subjects' preferences has a significant effect on the stability of final outcomes, with instability more likely to arise in the presence of weak incentives. However, Klijn et al. (2013) find that subject behavior in the Gale-Shapley mechanism is fairly robust to changes in cardinal preferences.

Second, there is a large experimental literature on behavior in coordination games.[15] Many of these studies focus on stag hunt games, which are a class of coordination games that feature a tension between payoff-dominance and risk-dominance. Existing work provides mixed results on the merits of different selection principles in this context. For instance, Cooper et al. (1990) find evidence for Nash equilibrium play, but not for payoff-dominance as the relevant selection criterion among the set of Nash equilibria. Battalio et al. (2001) construct a class of stag hunt games in which the basin of attraction of the risk-dominant equilibrium is fixed, but in which the optimization premium (i.e., the pecuniary incentive to play a best-response strategy) is allowed to vary. In their experiment, subjects learn to gravitate toward the risk-dominant action more effectively in treatments with a larger optimization premium. In a sequence of coordination games, Schmidt et al. (2003) find that changes in the level of risk-dominance (but not payoff-dominance) have a significant effect on subject behavior. Rankin et al. (2000), on the other hand, report on a series of laboratory stag hunt games in which cosmetic details of the game (e.g., action labels, player labels) are constantly perturbed. They find that subject behavior is consistent with payoff-dominance rather than risk-dominance.

Finally, there is a literature that attempts to appropriately generalize and extend the concept of risk-dominance to other strategic environments. In particular, our work can be viewed as method-

[14]See, for instance, Ding and Schotter (2016), Chen and Sönmez (2006), Fragiadakis and Troyan (2015), Castillo and Dianat (2016), Echenique et al. (2016), Featherstone and Mayefsky (2015), Featherstone and Niederle (2014), Harrison and McCabe (1989), Pais and Pintér (2008), and Klijn et al. (2013).

[15]For a survey of this literature, see Camerer (2003).

ologically related to Bó and Fréchette (2011), which applies risk-dominance as a selection criterion in the context of the infinitely-repeated prisoner's dilemma. In their construction, risk-dominance remains a pairwise comparison among Nash equilibria: instead of applying the concept to the entire strategy space, they focus on a simplified version of the game that consists of only two equilibrium strategies. In contrast, Morris et al. (1995) introduce a more stringent notion (*p-dominance*) for symmetric, many-action coordination games in which candidate action pairs are compared against all other actions (not just equilibrium actions).

The rest of the paper is organized as follows. Section 2 introduces the necessary theoretical background, Section 3 presents our experimental design, Section 4 presents our experimental results, and Section 5 discusses broader implications and concludes.

2 Theoretical Background

A *cardinal marriage market* is denoted by the triplet (M, W, u), where M is the set of men, W is the set of women, and $u = (u_i)_{i \in M \cup W}$ is a profile of von Neumann-Morgenstern utility functions.[16] The preferences of woman $w \in W$ are represented by the function $u_w : M \cup \{w\} \to \mathbb{R}$, where $u_w(w)$ is the utility she derives from remaining single. We assume that each function u_w is one-to-one, such that it induces a strict preference ordering P_w on the set $M \cup \{w\}$. We will refer to P_w as the *true preference list* of woman w. The preferences of the men are defined similarly and we let $P = (P_i)_{i \in M \cup W}$ denote the profile of agents' true preference lists. A matching is a function $\mu : M \cup W \to M \cup W$ such that: (1) for any $m \in M$, $\mu(m) \in W \cup \{m\}$, (2) for any $w \in W$, $\mu(w) \in M \cup \{w\}$, and (3) for any $m \in M$, $w \in W$, $\mu(m) = w$ if and only if $\mu(w) = m$.

Man m is said to be *acceptable* to woman w if she prefers him to remaining single (i.e., $u_w(m) > u_w(w)$). A matching μ is *individually rational* if every individual is either matched to an acceptable partner or remains unmatched. A pair of agents (m, w) is said to *block* a matching μ if they are not matched to one another at μ but they prefer each other to their assignments at μ (i.e., $u_w(m) > u_w(\mu(w))$ and $u_m(w) > u_m(\mu(m))$). A matching μ is *stable* if it is individually rational and not blocked by any pair of agents. A man m and a woman w are said to be *achievable* for each other in a marriage market (M, W, u) if they are matched to each other at some stable matching. A stable matching is called *W-optimal* (denoted μ_W) if each woman is matched to her most preferred achievable man. Similarly, a stable matching is called *M-optimal* (denoted μ_M) if each man is matched to his most preferred achievable woman.

We make two simplifying assumptions. First, we assume that there are an equal number of men and women in the market. Second, we assume that all men are acceptable to all women and

[16] Although the classical results we present are usually framed in terms of ordinal preferences, the solution concept of risk-dominance is inherently cardinal. Thus, we assume cardinal preferences throughout the analysis.

vice versa. These restrictions ensure that all agents are matched at all stable matchings.[17] For convenience, we will also suppress the outside option of remaining single from an agent's preference list. Thus, P_w is a strict preference ordering on the set M and P_m is a strict preference ordering on the set W.

Let $\mathcal{M}$ denote the set of all possible matchings, $\mathcal{Q}$ denote the set of all possible preference profiles, and $\mathcal{Q}_i$ denote the set of all possible preference lists for agent $i \in M \cup W$. Let μ, Q, and Q_i denote arbitrary elements of the sets $\mathcal{M}$, $\mathcal{Q}$, and $\mathcal{Q}_i$, respectively. A mechanism is a function $\phi : \mathcal{Q} \to \mathcal{M}$ that assigns a matching to each preference profile. A mechanism ϕ that for each preference profile Q produces a matching $\phi(Q)$ that is stable with respect to Q is called a stable mechanism. If $\phi(Q)$ is the M-optimal stable matching with respect to Q, then ϕ is called the M-optimal stable mechanism. We denote the M-optimal stable mechanism by ϕ_M.

The M-optimal stable mechanism can be modeled as a non-cooperative game in which the strategy space is the set of all possible ordinal preference lists. In this preference-revelation game, it is well-known that the men have a dominant strategy of truth-telling (Dubins and Freedman, 1981). In markets with more than one stable matching, however, at least one woman will have an incentive to misrepresent her preferences to improve her match outcome (Gale and Sotomayor, 1985). Our goal is to characterize the different equilibria that can arise in this environment. To simplify our analysis, we define an *induced preference-revelation game* based on the M-optimal stable mechanism:[18]

Definition 1. *Consider a cardinal marriage market (M, W, u) in which $P = (P_i)_{i \in M \cup W}$ denotes the profile of agents' true preference lists and $Q = (Q_i)_{i \in M \cup W}$ denotes the profile of agents' reported preference lists. The* ***induced preference-revelation game*** *based on the M-optimal stable mechanism ϕ_M is the preference-revelation game in which the men are restricted to truth-telling. That is, for any profile of women's reports $(Q_i)_{i \in W}$, the induced preference-revelation game produces the matching $\phi_M((P_i)_{i \in M}, (Q_i)_{i \in W})$.*

In other words, in the induced preference-revelation game based on the M-optimal stable mechanism, the men are constrained to play their dominant strategy of truth-telling while the women are free to report any preference ordering.[19]

We will find it useful to define two types of misrepresentation strategies for the women:

[17] To see this, suppose an agent is unmatched at a stable matching μ. Since $|M| = |W|$, there exists an unmatched agent on both sides of the market. Denote this pair of unmatched agents by (m, w). Since $u_w(m) > u_w(w)$ and $u_m(w) > u_m(m)$, the pair (m, w) blocks the matching μ.

[18] This terminology originally appears in Echenique et al. (2016).

[19] The assumption that men play their dominant strategy in the M-optimal stable mechanism is not entirely innocuous. While the M-optimal stable mechanism is strategy-proof for men, Ashlagi and Gonczarowski (2016) show that it is not obviously strategy-proof in the sense of Li (2016). Furthermore, empirical studies by Rees-Jones (2016) and Hassidim et al. (2015) find that a small fraction of participants fail to play their dominant strategy in strategy-proof matching mechanisms.

Definition 2. *A* ***truncation*** *of a preference list P_w containing k acceptable men is a list Q_w containing $k' < k$ acceptable men such that the k' elements of Q_w are the first k' elements of P_w, in the same order.*[20]

Definition 3. *A* ***permutation*** *of a preference list P_w is a list $Q_w \neq P_w$ that is not a truncation of P_w.*[21]

In other words, a truncation involves misrepresenting preferences by removing acceptable partners from the tail end of a preference list, while a permutation involves misrepresenting preferences by switching the order of acceptable partners in a preference list (regardless of the length of the list).

We now state and prove some basic results that are relevant to our experimental design. Throughout, we let μ_M and μ_W denote the M-optimal and W-optimal stable matchings with respect to the true preferences P.

Proposition 1. *Consider a cardinal marriage market (M, W, u) in which all agents have more than one achievable partner. In the induced preference-revelation game based on the M-optimal stable mechanism ϕ_M, there is a payoff-dominant equilibrium in which all women play truncation strategies.*

Proof. Let T denote the profile of reported preferences in which each woman w truncates her preference list by removing all men ranked below $\mu_W(w)$. By Theorem 4.17 of Roth and Sotomayor (1992), T is a Nash equilibrium and it produces the matching μ_W. Suppose another Nash equilibrium Q payoff-dominates T. Let μ denote the matching that is produced by Q. Since Q is a Nash equilibrium, we know by Theorem 4.16 of Roth and Sotomayor (1992) that the matching μ is also stable with respect to the true preferences P. Since Q payoff-dominates T, we know that $u_w(\mu(w)) > u_w(\mu_W(w))$ for all $w \in W$. We have arrived at a contradiction, since μ_W is the W-optimal stable matching with respect to P. Thus, there is no other Nash equilibrium that payoff-dominates T. We conclude that T is payoff-dominant. □

We will refer to the equilibrium in which all women play truncation strategies as the "symmetric" truncation equilibrium. However, one woman's truncation decision creates positive spillovers for other women in the market (Ashlagi and Klijn, 2012; Coles and Shorrer, 2014). This implies that asymmetric equilibria also exist in which a subset of women truncates its preferences and the remaining women report their true preferences, effectively free-riding on others' truncation behavior.

[20] This definition is taken from Roth and Rothblum (1999). However, it has been slightly modified such that truthful preference revelation is no longer an "edge case" of a truncation strategy.

[21] The term "dropping strategy" is often used to refer to the act of removing an acceptable partner from the middle of a preference list (rather than from the tail end of a preference list). According to our definitions, a dropping strategy would be classified as a permutation strategy.

We now construct a "symmetric" permutation equilibrium in which all women play permutation strategies.

Proposition 2. *Consider a cardinal marriage market (M, W, u) in which all agents have more than one achievable partner. In the induced preference-revelation game based on the M-optimal stable mechanism ϕ_M, there is a payoff-dominated equilibrium in which all women play permutation strategies.*

Proof. Let Q denote the profile of reported preferences in which each woman w reports a preference list Q_w that ranks $\mu_M(w)$ in the first position (regardless of the length of the list). Each preference list Q_w is clearly a permutation since $\mu_M(w)$ is not at the head of any woman's true preference list.[22] It is straightforward to see that Q produces the matching μ_M.

We argue that the profile of reported preferences Q constitutes a Nash equilibrium.[23] To see this, suppose that Q is not a Nash equilibrium. Then, there exists some woman w who can deviate and report a preference list Q'_w, which leads to a new profile of reported preferences $Q' = (Q_{-w}, Q'_w)$ and a new matching μ' such that $u_w(\mu'(w)) > u_w(\mu_M(w))$. Let $m = \mu'(w)$. Then man m must have been matched to a woman he prefers to w at μ_M, otherwise (m, w) would have blocked the matching μ_M under the true preferences P. But now man m and woman $\mu_M(m)$ block the matching μ' under the reported preferences Q', which is a contradiction. Therefore, Q is a Nash equilibrium. Furthermore, Q is payoff-dominated by the truncation equilibrium constructed in Proposition 1. □

Although all women prefer the truncation equilibrium to the permutation equilibrium, truncation behavior introduces the possibility of remaining unmatched for some profiles of other agents' reported preferences. This exposure to the worst possible outcome is not present for the permutation strategy that we identify. To see this more clearly, it is instructive to consider the steps of the man-proposing deferred acceptance algorithm. A consequential truncation (i.e., a truncation that affects the outcome) requires a woman to reject a proposal from an achievable man. This rejection frees the man to make other proposals, which may cause a chain of further rejections (by the women). If the truncating woman does not receive new proposals from this rejection chain, then she will remain single. The permutation strategy in Proposition 2 is inherently not consequential. It produces the same outcome as truth-telling: it merely prioritizes the least preferred achievable man by elevating him to the top of a woman's preference ordering.

[22] If $\mu_M(w)$ were at the head of any woman's true preference list, then this contradicts the assumption that all women have more than one achievable partner.

[23] The proof of this claim closely follows the proof of Theorem 4.15 in Roth and Sotomayor (1992). The only difference is that we allow the preference lists in Q to be of any length.

$$P_{f_1} = w_1, w_2 \quad P_{w_1} = f_2, f_1$$
$$P_{f_2} = w_2, w_1 \quad P_{w_2} = f_1, f_2$$

Table 1: The ordinal preferences used in the experiment. The firm-optimal stable matching is shown in red and the worker-optimal stable matching is shown in blue.

	Truth	*Truncate*	*Permute*
Truth	v_2, v_2	v_1, v_1	v_2, v_2
Truncate	v_1, v_1	v_1, v_1	v_3, v_2
Permute	v_2, v_2	v_2, v_3	v_2, v_2

Table 2: Normal-form representation of the induced preference-revelation game ($v_1 > v_2 > v_3$). The Nash equilibrium that implements the firm-optimal stable matching is shown in red. The Nash equilibria that implement the worker-optimal stable matching are shown in blue.

3 Experimental Design

In the experiment, the two sides of the market are referred to as "firms" and "workers." Each experimental market consists of four participants: two firms and two workers. The firms are automated to play their dominant strategy of truth-telling. Subjects are assigned to the role of workers and they are free to report any preference ordering. In each session, subjects play 20 rounds of the induced preference-revelation game based on the firm-proposing deferred acceptance algorithm. Subjects are randomly and anonymously re-grouped in markets (i.e., pairs) at the start of each round.

Table 1 shows the ordinal preferences used across all 20 rounds of the experiment. With this constellation of preferences, there are two disjoint stable matchings (i.e., each agent has two achievable partners). For convenience, we let v_1, v_2 and v_3 denote an agent's utilities from matching with her most preferred partner, her least preferred partner, and remaining single, respectively. Table 2 depicts the normal-form representation of the induced preference-revelation game. It should be noted that *Permute* combines two pure strategies that are strategically equivalent.[24]

We now apply the Harsanyi and Selten (1988) criterion of risk-dominance to this strategic environment. Risk-dominance concerns the pairwise comparison of Nash equilibria. In 2×2 symmetric coordination games, an equilibrium is risk-dominant if its equilibrium action is a best-response to the belief that assigns equal probabilities to both actions of the opponent.[25] However, generalizing the concept of risk-dominance presents complications.[26] In the spirit of Bó and Fréchette (2011), we restrict attention to two focal equilibrium actions: truncation and permutation.[27] We will say

[24] For instance, consider the situation facing worker w_1 with true preference list $P_{w_1} = f_2, f_1$. Both permutation strategies ($Q_{w_1} = f_1, f_2$ and $Q'_{w_1} = f_1$) generate the same payoff for all preference reports by the other player.

[25] This definition is equivalent to the following characterization: an equilibrium is risk-dominant if it has the larger basin of attraction under either the best-reply or the replicator dynamics (Samuelson, 1998).

[26] For instance, the binary relation imposed by risk-dominance can fail transitivity. Morris et al. (1995) provide an example of a 3×3 game with three strict Nash equilibria in which the risk-dominance relationship is cyclical.

[27] Bó and Fréchette (2011) use risk-dominance to compare the "grim trigger" strategy and the "always defect"

	Treatment	
	PDOM	TDOM
v_2	15	5
Truth	protective	protective
Truncation	Payoff dominant (PD)	PD and RD
Permutation	Risk dominant (RD)	neither PD nor RD

Table 3: Our experimental treatments vary the risk-dominant equilibrium prediction.

that one of the two equilibrium actions is risk-dominant if it is the best response to the other player choosing between truncation and permutation with equal probabilities. This formulation captures both the logic and the danger of truncation. When remaining unmatched is particularly costly (i.e., $v_2 - v_3 > v_1 - v_2$), then permutation is risk-dominant. When an agent has a strong intensity of preference for her first choice partner (i.e., $v_1 - v_2 > v_2 - v_3$), then truncation is risk-dominant.

For all agents, however, truth-telling constitutes the unique *protective* strategy in stable matching mechanisms (Barberà and Dutta, 1995).[28] This is because truth-telling is the only strategy that accomplishes the following two tasks: (1) it secures against the worst possible outcome (i.e., remaining single) and (2) it leads to the best possible outcome for some profile of other agents' preference reports.

We review the key strategic features of this game:[29]

1. $(Truth, Truth)$ is not an equilibrium.

2. There are two symmetric equilibria: $(Truncate, Truncate)$ and $(Permute, Permute)$.

3. There are asymmetric equilibria in which one agent truncates her preferences and the other agent reports her true preferences.

4. Any equilibrium in truncation strategies is payoff-dominant.

5. If $v_2 - v_3 > v_1 - v_2$, then permutation is risk-dominant.

6. If $v_1 - v_2 > v_2 - v_3$, then truncation is risk-dominant.

7. $Truth$ is the unique protective strategy.

Table 3 summarizes our experimental design. In our experiment, we fix the payoff from matching with the most preferred partner ($v_1 = 20$) and from the outside option of remaining single ($v_3 = 0$).

strategy in the context of the infinitely-repeated prisoner's dilemma.

[28]A protective strategy is a refinement of a maximin strategy. Notice that while both $Truth$ and $Permute$ are maximin strategies, $Truth$ weakly dominates $Permute$.

[29]Recall that these features are not a function of the ordinal preferences we are using in the experiment. Rather, these features exist in any marriage market where agents have multiple achievable partners (i.e., where there are disjoint stable matchings).

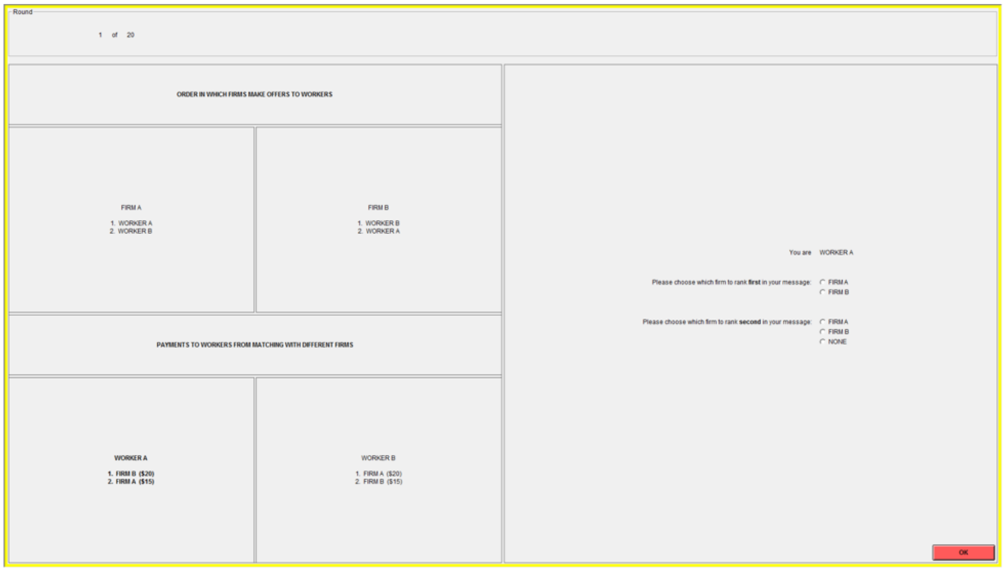

Figure 1: The experimental interface for the PDOM treatment.

Our treatments vary the risk-dominant equilibrium prediction by manipulating the payoff from matching with the least preferred partner ($0 < v_2 < 20$). When $v_2 = 15$, permutation is risk-dominant (PDOM treatment). When $v_2 = 5$, truncation is risk-dominant (TDOM treatment).

To secure comprehension, subjects are required to complete a demonstration of the deferred acceptance algorithm (with a hypothetical set of reported preferences) and correctly answer a series of questions. In each round of the experiment, subjects observe the preferences of all market participants and they are then asked to report a preference ordering. At the end of a round, subjects receive feedback about the identity of their match partner (i.e., FIRM A or FIRM B) and their payoff in that particular round. The experimental interface is shown in Figure 1 and the experimental instructions for both treatments are included at the end of the paper.[30]

4 Experimental Results

The experimental sessions were run at the Experimental Social Science Laboratory (ESSL) at UC Irvine. A total of 120 subjects participated in the experiment (PDOM treatment: 64 subjects, TDOM treatment: 56 subjects). Each experimental session lasted approximately 60 minutes. One of the 20 experimental rounds was randomly selected for subject payment. Average subject earnings were $21.75 (including a $7 show-up payment). The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007). We now present our main results.

[30]To reduce experimenter demand effects, the terminology of preferences is never used in the experiment. A subject's true preference list is referred to as a "list of payments" and a subject's reported preference list is referred to as a "message."

	Round 1		Round 20	
Strategy	PDOM	TDOM	PDOM	TDOM
Truth	39 (61%)	39 (70%)	41 (64%)	26 (46%)
Truncation	3 (5%)	4 (7%)	16 (25%)	30 (54%)
Permutation	22 (34%)	13 (23%)	7 (11%)	0 (0%)
Total	64 (100%)	56 (100%)	64 (100%)	56 (100%)

Table 4: Treatment by Round 1 and Round 20 Choices.
PDOM: permutation is risk-dominant. TDOM: truncation is risk-dominant.

Result 1. *Initial behavior is statistically different from random behavior.*

Although subjects choose among four pure strategies, both permutation strategies are strategically equivalent. Thus, we combine them under the heading of a single strategy for the remainder of the analysis. Table 4 presents the Round 1 and Round 20 distributions of strategy choice across treatments. We focus on behavior in Round 1 of the experiment, when subjects have no previous experience with the matching game. We find evidence that subject behavior in Round 1 is purposeful: the empirical distributions in both treatments are statistically different from the distribution predicted by random chance (PDOM: $\chi^2(2) = 46.75$, $p < 0.001$; TDOM: $\chi^2(2) = 59.82$, $p < 0.001$).[31] In particular, subjects are more likely to report their true preferences and less likely to play a truncation strategy compared to random chance.

Result 2. *An equilibrium strategy is played significantly more often when it is risk-dominant.*

By round

We first demonstrate that this treatment effect is not present initially, but rather emerges with subject experience. To do so, we look at subject behavior in the first and last rounds of the experiment separately. In Round 1, truncation is rare (PDOM: 5%, TDOM: 7%) while permutation occurs more regularly and at similar levels (PDOM: 34%, TDOM: 23%) in both treatments. Neither of these treatment differences are statistically significant at conventional levels (truncation: t-test, $p = 0.57$, Mann-Whitney test, $p = 0.57$; permutation: t-test, $p = 0.18$, Mann-Whitney test, $p = 0.18$). In aggregate, we also find that initial behavior does not vary significantly across treatments (Fisher's exact test, $p = 0.39$).

However, there is a significant treatment effect at the end of the session (Fisher's exact test, $p < 0.001$). The observed treatment effect is consistent with subjects using risk-dominance as a selection principle. In Round 20, subjects are more than twice as likely to play a truncation strategy when truncation is risk-dominant (PDOM: 25%, TDOM: 54%). At the same time, in Round 20,

[31] The random distribution assumes that subjects are equally likely to play their four pure strategies.

Treatment	Frequency of Truncation	Rank	Treatment	Frequency of Permutation	Rank
PDOM	0.16	1	TDOM	0.08	1
PDOM	0.19	2	TDOM	0.08	2
PDOM	0.25	3	TDOM	0.12	3
TDOM	0.36	4	PDOM	0.15	4
TDOM	0.41	5	PDOM	0.22	5
TDOM	0.47	6	PDOM	0.25	6

Table 5: Ranking average frequencies of truncation and permutation by session.
PDOM: permutation is risk-dominant. TDOM: truncation is risk-dominant.

no subjects play a permutation strategy when truncation is risk-dominant (TDOM: 0%). Both of these treatment differences are statistically significant at conventional levels (truncation: t-test, p = 0.001, Mann-Whitney test, p = 0.001; permutation: t-test, p = 0.011, Mann-Whitney test, p = 0.011).

By session

An alternative way to measure a treatment effect is to consider experimental sessions (i.e., cohorts) as independent units of observation and rank the sessions by their average frequencies of truncation and permutation. This procedure is shown in Table 5. The three TDOM sessions have higher average truncation rates than the three PDOM sessions, while the three PDOM sessions have higher average permutation rates than the three TDOM sessions. Using a non-parametric rank-sum test, we can reject the null hypothesis of no treatment difference for both cases (p = 0.05). Truncation and permutation strategies are both played significantly more often when they are the risk-dominant equilibrium prediction.

By subject

Each subject submits 20 rank-order lists during the course of the experiment (one in each round). From this data, we can calculate a truncation and permutation rate for each individual subject. The average subject-level truncation rates are 0.20 and 0.42 for the PDOM and TDOM treatments, respectively (t-test, $p < 0.001$; Mann-Whitney test, p = 0.005). The average subject-level permutation rates are 0.21 and 0.09 for the PDOM and TDOM treatments, respectively (t-test, $p < 0.001$; Mann-Whitney test, p = 0.001).

Figure 2 shows the empirical cumulative distribution functions (CDFs) of subject-level truncation and permutation rates. As evident from the pairs of distributions, the first-order stochastic dominance relationships are consistent with the risk-dominance relationships. In other words, the distribution of truncation rates from the TDOM treatment first-order stochastically dominates the distribution of truncation rates from the PDOM treatment (LHS graph) while the distribution of

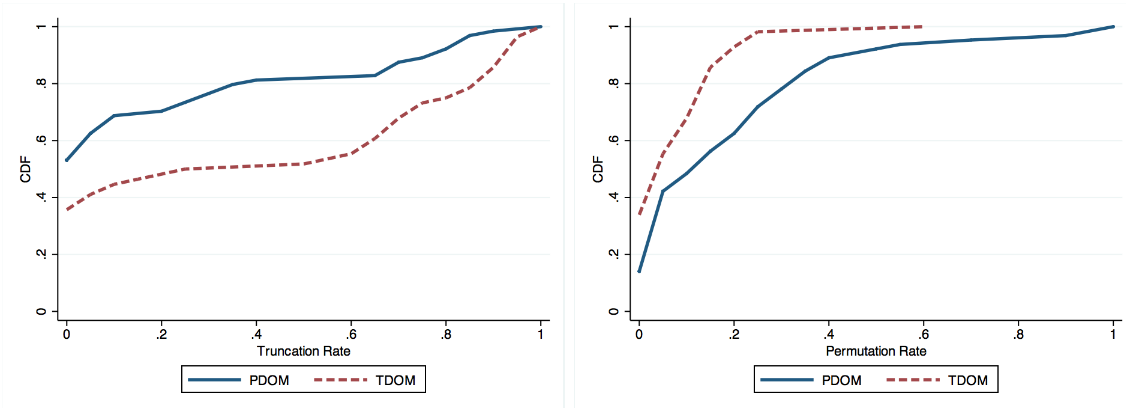

Figure 2: Empirical CDFs of subject-level truncation and permutation rates.
PDOM: permutation is risk-dominant. TDOM: truncation is risk-dominant.

permutation rates from the PDOM treatment first-order stochastically dominates the distribution of permutation rates from the TDOM treatment (RHS graph). We also reject the null hypotheses that subject-level truncation and permutation rates have the same distribution function across treatments (Kolmogorov-Smirnov test, $p < 0.001$).

Finally, we can use the distribution functions to investigate the incidence of "purists" who never play a particular equilibrium strategy. A majority of subjects (53%) never truncate their preference lists in the PDOM treatment, while 36% of subjects never truncate their preference lists in the TDOM treatment. The corresponding numbers for permutation strategies are 14% in the PDOM treatment and 34% in the TDOM treatment. Thus, a higher proportion of subjects abstain from playing a particular strategy when that strategy is risk-dominated.

Result 3. *In both treatments, truncation strategies are played significantly more often and permutation strategies are played significantly less often in later rounds of the experiment.*

We now investigate the effect of learning on subject behavior. Figure 3 shows the average frequencies of different strategies across rounds of the experiment. In the PDOM treatment, only 5% of subjects play a truncation strategy in Round 1 while 25% of subjects play a truncation strategy in Round 20. In the TDOM treatment, the corresponding numbers are 7% in Round 1 and 54% in Round 20. For both treatments, we reject the null hypothesis that truncation rates across the first 10 rounds are the same as truncation rates across the last 10 rounds (PDOM: $\chi^2(1) = 13.8$, $p < 0.01$; TDOM: $\chi^2(1) = 34.5$, $p < 0.01$). Since this truncation time trend is present in both treatments (where risk-dominance yields different predictions), it suggests that payoff-dominance has increasing salience as a selection principle in later rounds of the experiment.[32]

[32]In the context of stag hunt games, Rankin et al. (2000) also find that laboratory subjects focus on payoff-

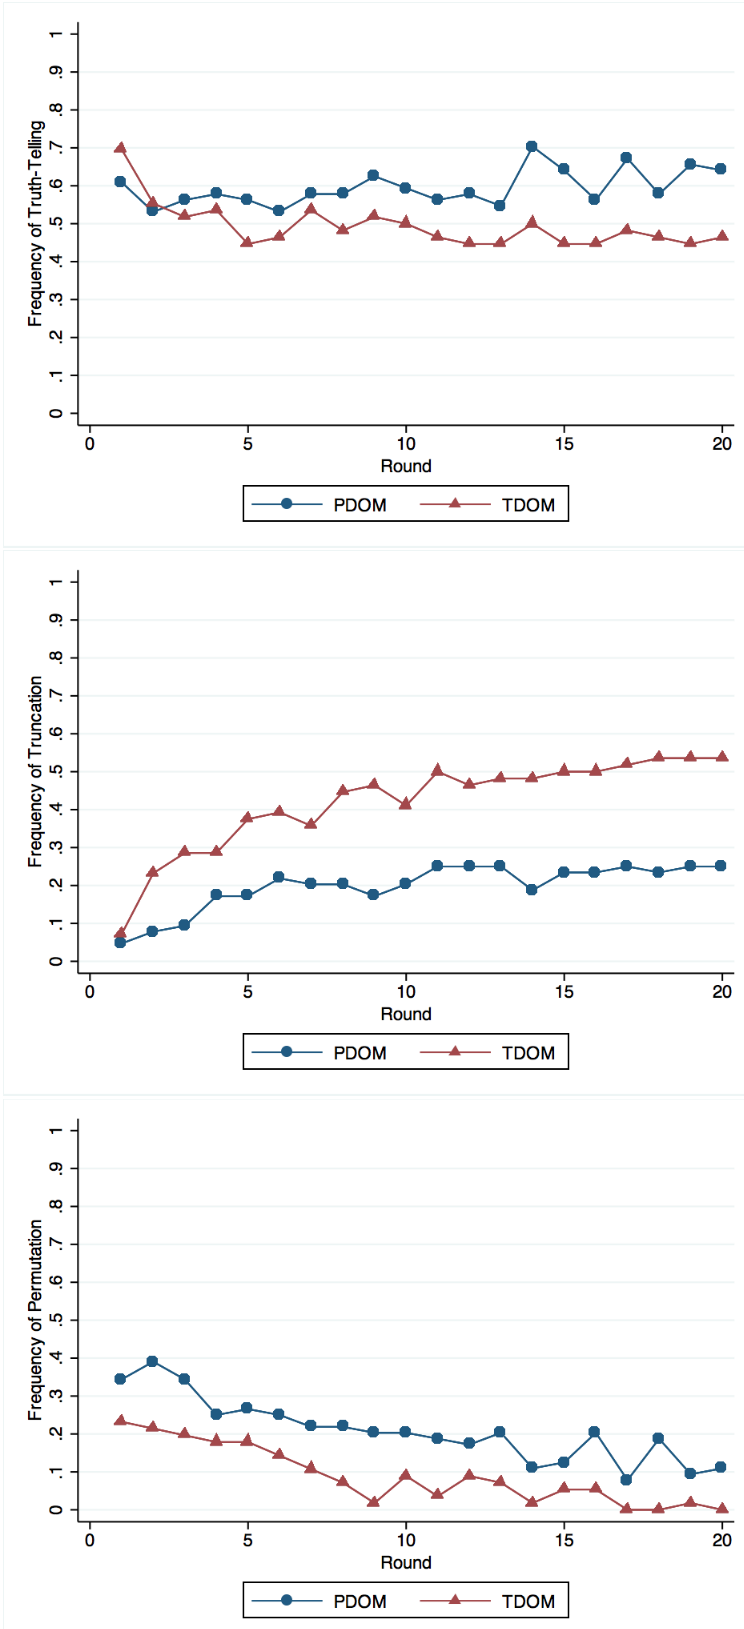

Figure 3: Average frequencies of different strategies across rounds of the experiment. PDOM: permutation is risk-dominant. TDOM: truncation is risk-dominant.

With respect to permutation strategies, we find a decrease in the PDOM treatment from 34% in Round 1 to 11% in Round 20. In the TDOM treatment, although 23% of subjects report a permuted preference list in Round 1, there are no permutations in Round 20. Again, for both treatments, we reject the null hypothesis that permutation rates across the first 10 rounds are the same as permutation rates across the last 10 rounds (PDOM: $\chi^2(1) = 28.9$, $p < 0.01$; TDOM: $\chi^2(1) = 41.2$, $p < 0.01$). As before, the presence of this time trend in both treatments suggests payoff-dominance as the driving force.

We also see that truth-telling rates are remarkably consistent across the experiment. In the PDOM treatment, the average frequency of truthful reporting ranges from a minimum of 53% (in Round 2) to a maximum of 70% (in Round 14). In the TDOM treatment, the minimum and maximum frequencies are 45% (in several rounds) and 70% (in Round 1). Once again, we test for whether truth-telling rates across the first 10 rounds are significantly different from truth-telling rates across the last 10 rounds. For the PDOM treatment, we fail to reject the null hypothesis of no time trend ($\chi^2(1) = 2.03$, $p = 0.155$). For the TDOM treatment, we can reject the null hypothesis of no time trend ($\chi^2(1) = 4.63$, $p = 0.031$). However, the latter result is attributable to the high level of truth-telling in Round 1 and disappears when Round 1 data is removed from the analysis ($\chi^2(1) = 2.17$, $p = 0.140$).

Result 4. *We find limited evidence for protective behavior as a solution concept.*

Overall, truth-telling is the modal strategy in both treatments. A majority (59%) of submitted rank-order lists are truthful in the PDOM treatment and a plurality (49%) of submitted rank-order lists are truthful in the TDOM treatment. The prevalence of truthful behavior is consistent with the hypothesis that subjects use protective strategies. However, since truth-telling is the unique protective strategy in both experimental treatments, any theory of protective behavior would also predict the absence of a treatment effect. As mentioned earlier, we do find a significant treatment effect: subjects play a strategy more often when it is risk-dominant. We also find a statistically significant difference between the average truth-telling rate across treatments (t-test, $p < 0.001$; Mann-Whitney test, $p < 0.001$). The higher rate of truth-telling in the PDOM treatment further suggests that protective behavior alone cannot fully explain the experimental data.

Result 5. *A significantly higher proportion of subject-pairs play a Nash equilibrium when truncation is both payoff-dominant and risk-dominant.*

Table 6 presents data on subject-pairs, with equilibrium strategy profiles in blue. In the TDOM treatment, a majority (62%) of subject-pairs play an equilibrium strategy profile. However, only

dominance rather than other solution concepts.

Strategy Profile	Treatment PDOM	TDOM	Total
(Truth, Truth)	211 (33%)	130 (23%)	341 (28%)
(Truncate, Truncate)	22 (3%)	102 (18%)	124 (10%)
(Permute, Permute)	28 (4%)	8 (1%)	36 (3%)
(Truncate, Truth)	169 (26%)	237 (42%)	406 (34%)
(Permute, Truth)	170 (27%)	55 (10%)	225 (19%)
(Permute, Truncate)	40 (6%)	28 (5%)	68 (6%)
Total	640 (100%)	560 (100%)	1,200 (100%)

Table 6: Empirical distributions of strategy profiles.
PDOM: permutation is risk-dominant. TDOM: truncation is risk-dominant.

34% of subject-pairs play an equilibrium strategy profile in the PDOM treatment. The proportion of markets culminating in a Nash equilibrium is statistically different across treatments (t-test, $p < 0.001$; Mann-Whitney test, $p < 0.001$). This effect is driven almost entirely by the larger incidence of truncation in the TDOM treatment. Indeed, equilibria involving truncation strategies are roughly twice as likely to occur in the TDOM treatment than the PDOM treatment (TDOM: 60% of subject-pairs, PDOM: 29% of subject-pairs). This result has intuitive appeal: efficient coordination is more likely to occur when the recommendations of payoff-dominance and risk-dominance coincide, while coordination failure is more likely to occur when the recommendations of the two selection criteria conflict.

Result 6. *A basic model of reinforcement learning captures the treatment effect found in the experimental data, but not subject learning over time.*

We now investigate whether subject behavior can be described by a basic model of *reinforcement learning*, a low-cognitive theory of decision-making in which agents are more likely to repeat actions that yield more favorable outcomes.[33] In our experiment, at the end of each round, each subject receives feedback about her own payoff but not about the payoff of the other player or the action taken by the other player. Since the theory of reinforcement learning only requires that an agent update her play based on feedback about her own payoffs, it can be directly tested with our experimental data. Alternative theories of learning that require observing or best-responding to other players' past actions (e.g., fictitious play) cannot be explicitly tested in the context of our experiment.

We use the one-parameter reinforcement learning model from Erev and Roth (1998). Each player n has a propensity to play strategy k in round t of the experiment, denoted by $q_{nk}(t)$. For

[33] Although originating from psychology, the theory of reinforcement learning was first introduced into the economics literature by Erev and Roth (1998).

PDOM	Truth	Truncation	Permutation
Simulation, s(1) = 10	38%	25%	37%
Simulation, s(1) = 50	38%	27%	35%
Simulation, s(1) = 100	37%	28%	35%
Experiment	64%	25%	11%
TDOM	Truth	Truncation	Permutation
Simulation, s(1) = 10	37%	43%	20%
Simulation, s(1) = 50	35%	49%	16%
Simulation, s(1) = 100	35%	46%	19%
Experiment	46%	54%	0%

Table 7: Simulated and actual Round 20 behavior.
Round 1: $(p_{Truth} = \frac{1}{3}, p_{Truncate} = \frac{1}{3}, p_{Permute} = \frac{1}{3})$
PDOM: permutation is risk-dominant. TDOM: truncation is risk-dominant.

each player, we assume that the initial propensities are equal across all pure strategies. That is, for each player n and for all strategies k and j, we have that

$$q_{nk}(1) = q_{nj}(1) \tag{1}$$

The reinforcement from receiving a payoff $x \in \mathbb{R}_+$ is represented by the function

$$R(x) = x \tag{2}$$

If player n plays strategy k in round t and receives a payoff of x, propensities are then updated as follows:

$$q_{nj}(t+1) = \begin{cases} q_{nj}(t) + R(x) & j = k \\ q_{nj}(t) & j \neq k \end{cases} \tag{3}$$

The probability that player n chooses strategy k in round t is given by

$$p_{nk}(t) = \frac{q_{nk}(t)}{\sum q_{nj}(t)}, \tag{4}$$

where the sum is taken over all the pure strategies. There are two key features of this model worth highlighting. First, since $R(x)$ is an increasing function, the reinforcement is greater for strategies that yield more profitable outcomes. Second, since $\sum q_{nj}(t)$ is an increasing function of t, the updating of probabilities is more sensitive (i.e., steeper) in earlier rounds and less sensitive (i.e., flatter) in later rounds. The single parameter of this model is the sum of the initial propensities, which we assume is the same for all players and we denote by

PDOM	Truth	Truncation	Permutation
Simulation, s(1) = 10	30%	13%	57%
Simulation, s(1) = 50	29%	15%	56%
Simulation, s(1) = 100	29%	17%	54%
Experiment	64%	25%	11%
TDOM	Truth	Truncation	Permutation
Simulation, s(1) = 10	31%	27%	42%
Simulation, s(1) = 50	31%	35%	34%
Simulation, s(1) = 100	29%	33%	38%
Experiment	46%	54%	0%

Table 8: Simulated and actual Round 20 behavior.
Round 1: ($p_{Truth} = \frac{1}{4}$, $p_{Truncate} = \frac{1}{4}$, $p_{Permute} = \frac{1}{2}$)
PDOM: permutation is risk-dominant. TDOM: truncation is risk-dominant.

$$s(1) = s_n(1) = \sum q_{nj}(1) \tag{5}$$

The parameter $s(1)$ is often called the *strength* of the initial propensities, and it can substantially affect the speed of learning.

To determine whether the one-parameter reinforcement learning model is consistent with our experimental data, we conduct computer simulations for both treatments and for three different values of the free parameter $s(1)$. We use the following three values: $s(1) = 10, 50, 100$.[34] We also use the same payoff structure as the relevant experimental treatment.[35] For each of these six cases, we conduct 1000 simulations in which two simulated players interact for 20 rounds. The simulated players' strategies are randomly determined according to equation (4) and the propensities are updated according to equation (3).

We conduct this exercise for two different sets of initial conditions. First, we depart from the Erev and Roth (1998) framework and assume that initial propensities are equal across the different *types* of strategies (i.e., truth-telling, truncation, permutation). This generates the following mixed strategy for the simulated players in Round 1: ($p_{Truth} = \frac{1}{3}$, $p_{Truncate} = \frac{1}{3}$, $p_{Permute} = \frac{1}{3}$). Second, we maintain the standard assumption that initial propensities are equal across the different pure strategies. This generates a different mixed strategy for the simulated players in Round 1: ($p_{Truth} = \frac{1}{4}$, $p_{Truncate} = \frac{1}{4}$, $p_{Permute} = \frac{1}{2}$).

Tables 7 and 8 report the simulated Round 20 behavior (averaged across both simulated players and all 1000 simulations) alongside the actual Round 20 behavior for both sets of initial conditions. The qualitative results remain unchanged. The model's predictions are also fairly robust to different choices of the free parameter $s(1)$. We observe the same treatment effect that is present in the

[34]For values of $s(1)$ that are sufficiently large relative to the payoffs, the model predicts virtually no learning.
[35]For the PDOM treatment, $x \in \{0, 15, 20\}$. For the TDOM treatment, $x \in \{0, 5, 20\}$.

experimental data: holding the value of $s(1)$ fixed, an equilibrium strategy is played more often in Round 20 when it is risk-dominant. However, we no longer observe the same learning effect from the experiment. Although the TDOM simulations predict that the average truncation rate will increase and the average permutation rate will decrease from Round 1 to Round 20, the PDOM simulations predict the opposite trends.

5 Discussion and Conclusion

In this paper, we apply equilibrium selection concepts to the stable marriage problem. We report three main findings from a laboratory experiment: (1) truth-telling is the most common strategy, (2) truncation and permutation strategies are played significantly more often when they are risk-dominant, and (3) in both treatments, truncation strategies are played significantly more often in later rounds of the experiment. The final point suggests that the salience of payoff-dominance as a selection principle increases with subject experience.

Our experiment can help shed light on several open questions in two-sided matching. In a seminal paper, Roth and Peranson (1999) conduct computational experiments using NRMP submitted rank-order lists from 1987 and 1993-1996. They find that only 0.1% of applicants would have received a different match from the applicant-proposing and hospital-proposing versions of the deferred acceptance algorithm. Assuming that the NRMP rank-order lists accurately reflect participants' true preferences, this exercise suggests that the applicant-optimal and hospital-optimal stable matchings coincide.[36] This result has profound implications for market design: if there is a unique stable matching in markets that are sufficiently large, then there is also no incentive for market participants to behave strategically.

By credibly documenting the use of non-truth-telling strategies, our experimental results are consistent with the empirical regularity of a small core span. In our experimental markets, there are two disjoint stable matchings. In our data, however, 72% of markets have a unique stable matching with respect to the reported preferences (PDOM treatment: 67%, TDOM treatment: 77%). Moreover, our experiment demonstrates that a unique stable matching need not only arise from preference misrepresentation in equilibrium. Rather, any out-of-equilibrium strategy profile in which at least one subject in a pair deviates from truth-telling is sufficient to produce a unique stable matching with respect to the reported preferences.

In addition, previous theoretical work on strategic behavior in matching markets has largely focused on truncation strategies (Coles and Shorrer, 2014; Roth and Rothblum, 1999). There are two common arguments for this emphasis. First, truncation can be profitably implemented

[36]There is now a growing literature on "core convergence" in matching models (e.g. Immorlica and Mahdian, 2005; Kojima and Pathak, 2009; Lee, 2016). Under certain conditions, these papers show that the set of stable matchings shrinks as the size of the market increases.

even with incomplete information about other agents' true preferences or strategic uncertainty about other agents' reported preferences (Coles and Shorrer, 2014; Roth and Rothblum, 1999). Second, it is sufficient to restrict attention to truncation when considering the space of profitable misrepresentation strategies (Roth and Vate, 1991). In other words, any agent who can improve her match partner by deviating from truth-telling can do so by submitting a truncation of her true preferences.

However, the empirical content of truncation strategies remains an open question. Field data is insufficient to address this issue: while centralized matching clearinghouses may provide access to participants' submitted rank-order lists, participants' true preferences are unobserved. Our experiment provides preliminary support for the empirical relevance of truncation strategies and helps highlight the conditions that foster truncation behavior. Our results suggest that truncation should be more likely to arise in settings where participants either have a strong intensity of preference for top-ranked alternatives or have previous experience with the particular mechanism that is being used. Real-world matching protocols often more closely resemble one-shot games for many participants. In the context of the NRMP, however, residency programs usually do have considerable experience based on their participation in the matching process in previous years. The 2016 NRMP Program Director Survey Report indicates that, across all medical specialties, the average number of applicants interviewed was 94 while the average number of applicants ranked was 80. While not conclusive, the NRMP survey data suggests that truncation strategies may play a role in field settings.

References

Atila Abdulkadiroğlu and Tayfun Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.

Itai Ashlagi and Yannai A. Gonczarowski. Stable matching mechanisms are not obviously strategy-proof. *Working Paper*, 2016.

Itai Ashlagi and Flip Klijn. Manipulability in matching markets: conflict and coincidence of interests. *Social Choice and Welfare*, 39(1):23–33, 2012.

Salvador Barberà and Bhaskar Dutta. Protective behavior in matching models. *Games and Economic Behavior*, 8(2):281–296, 1995.

Raymond Battalio, Larry Samuelson, and John Van Huyck. Optimization incentives and coordination failure in laboratory stag hunt games. *Econometrica*, 69(3):749–764, 2001.

Pedro Dal Bó and Guillaume R Fréchette. The evolution of cooperation in infinitely repeated games: Experimental evidence. *The American Economic Review*, 101(1):411–429, 2011.

Colin Camerer. *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, 2003.

Marco Castillo and Ahrash Dianat. Truncation strategies in two-sided matching markets: Theory and experiment. *Games and Economic Behavior*, 2016.

Yan Chen and Tayfun Sönmez. School choice: an experimental study. *Journal of Economic Theory*, 127(1):202–231, 2006.

Peter A Coles and Ran I Shorrer. Optimal truncation in matching markets. *Games and Economic Behavior*, 87:591–615, 2014.

Russell W Cooper, Douglas V DeJong, Robert Forsythe, and Thomas W Ross. Selection criteria in coordination games: Some experimental results. *American Economic Review*, 80(1):218–233, 1990.

Tingting Ding and Andrew Schotter. Matching and chatting: An experimental study of the impact of network communication on school-matching mechanisms. *Games and Economic Behavior*, 2016.

Lester E Dubins and David A Freedman. Machiavelli and the gale-shapley algorithm. *American Mathematical Monthly*, 88(7):485–494, 1981.

Federico Echenique, Alistair J Wilson, and Leeat Yariv. Clearinghouses for two-sided matching: An experimental study. *Quantitative Economics*, 2016.

Ido Erev and Alvin E Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, pages 848–881, 1998.

Clayton R Featherstone and Eric Mayefsky. Why do some clearinghouses yield stable outcomes? experimental evidence on out-of-equilibrium truth-telling. *Working Paper*, 2015.

Clayton R Featherstone and Muriel Niederle. School choice mechanisms under incomplete information: An experimental investigation. *Working Paper*, 2014.

Urs Fischbacher. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, 2007.

Daniel E Fragiadakis and Peter Troyan. Designing mechanisms to make welfare-improving strategies focal. *Working Paper*, 2015.

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

David Gale and Marilda Sotomayor. Ms. machiavelli and the stable matching problem. *American Mathematical Monthly*, 92(4):261–268, 1985.

Glenn W Harrison and Kevin A McCabe. Stability and preference distortion in resource matching: An experimental study of the marriage market. *Research in Experimental Economics*, 8, 1989.

John C Harsanyi and Reinhard Selten. *A general theory of equilibrium selection in games*. MIT Press, 1988.

Avinatan Hassidim, Deborah Marciano-Romm, Assaf Romm, and Ran I Shorrer. "strategic behavior in a strategy-proof environment. *Working Paper*, 2015.

Nicole Immorlica and Mohammad Mahdian. Marriage, honesty, and stability. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 53–62. Society for Industrial and Applied Mathematics, 2005.

Flip Klijn, Joana Pais, and Marc Vorsatz. Preference intensities and risk aversion in school choice: A laboratory experiment. *Experimental Economics*, 16(1):1–22, 2013.

Fuhito Kojima and Parag A Pathak. Incentives and stability in large two-sided matching markets. *American Economic Review*, 99(3):608–627, 2009.

SangMok Lee. Incentive compatibility of large centralized matching markets. *Working Paper*, 2016.

Shengwu Li. Obviously strategy-proof mechanisms. *Working Paper*, 2016.

Stephen Morris, Rafael Rob, and Hyun Song Shin. p-dominance and belief potential. *Econometrica*, pages 145–157, 1995.

Joanna Pais and Ágnes Pintér. School choice and information: an experimental study on matching mechanisms. *Games and Economic Behavior*, 64(1):303–328, 2008.

Frederick W Rankin, John B Van Huyck, and Raymond C Battalio. Strategic similarity and emergent conventions: Evidence from similar stag hunt games. *Games and Economic Behavior*, 32(2):315–337, 2000.

Alex Rees-Jones. Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match. *Working Paper*, 2016.

Alvin E Roth. The economics of matching: Stability and incentives. *Mathematics of operations research*, 7(4):617–628, 1982.

Alvin E Roth. Misrepresentation and stability in the marriage problem. *Journal of Economic Theory*, 34(2):383–387, 1984.

Alvin E Roth. The college admissions problem is not equivalent to the marriage problem. *Journal of Economic Theory*, 36(2):277–288, 1985.

Alvin E Roth. A natural experiment in the organization of entry-level labor markets: regional markets for new physicians and surgeons in the united kingdom. *American Economic Review*, pages 415–440, 1991.

Alvin E Roth and Elliott Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review*, 89(4):748–780, 1999.

Alvin E Roth and Uriel G Rothblum. Truncation strategies in matching marketsin search of advice for participants. *Econometrica*, 67(1):21–43, 1999.

Alvin E Roth and Marilda A Oliveira Sotomayor. *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge University Press, 1992.

Alvin E Roth and John H Vande Vate. Incentives in two-sided matching with random stable mechanisms. *Economic theory*, 1(1):31–44, 1991.

Larry Samuelson. *Evolutionary games and equilibrium selection*. MIT Press, 1998.

David Schmidt, Robert Shupp, James M Walker, and Elinor Ostrom. Playing safe in coordination games: the roles of risk dominance, payoff dominance, and history of play. *Games and Economic Behavior*, 42(2):281–299, 2003.

WELCOME!

Please turn off all electronic devices and place them in your bag or under your desk.

Throughout the experiment, please do not talk to anybody else and please remain silent at all times. If you have any questions, raise your hand and the experimenter will come to personally assist you.

Thank you for participating in this experiment. By showing up on time, you have automatically earned a $7 payment. If you follow the instructions carefully and make good decisions, you can earn additional money. The amount of money that you ultimately earn in this experiment depends on your decisions and the decisions of others. This session will last approximately one hour. At the end of the session, you will be paid privately in cash.

The experiment will be run entirely on the computer and all interactions between yourself and others will take place via the computer terminal. There are a total of 20 rounds in this experiment. At the end of the experiment, one of the 20 rounds will be randomly selected and your monetary payment will be determined based on the outcome of that round. Each of the 20 rounds is equally likely to be selected. Thus, it is in your best interest to take each round seriously. Each round is self-contained: your decisions in one round will not affect your opportunities or earnings in another round.

This experiment is about matching. There are two groups in the experiment: firms and workers. You will be randomly assigned to the role of either WORKER A or WORKER B. Your role will remain the same across all 20 rounds of the experiment. In each round, you will be randomly and anonymously paired with a worker who is assigned to the other role. In other words, if your role is WORKER A then you will be paired with someone in the role of WORKER B in each round (and vice versa).

Your goal in each round is to match with one of the two firms: FIRM A or FIRM B. The roles of the firms are computerized. They are programmed to behave in a certain way, which will be discussed in more detail shortly. You will earn different payments from different matches. The payments corresponding to all possible matches will be shown to you on your screen. There will be four lists on your screen (one for each firm and one for each worker). A firm's list contains two workers and it shows the order in which the firm will be making offers to match with the different workers. A worker's list contains two firms and it shows how much money the worker would earn if matched with the different firms. These lists will remain the same across all 20 rounds of the experiment.

As a worker, if you are matched with the firm in the first position of your list in a given round, you will earn a payment of $20 for that round. If you are matched with the firm in the second position of your list in a given round, you will earn a payment of $15 for that round. Remaining unmatched will result in a payment of $0 for that round.

To determine which firm you are matched with, you will send a message in each round. This message is sent to the computer. It is very important that you understand what a message is since the messages sent by both you and the other worker in your pair determine which firm you are matched with. A message is a ranking of the firms. The message may or may not include all the firms. Thus, although there are two firms you could potentially be matched with, your message can contain either one or two firms.

The way the computer uses the messages to determine which firm you are matched with will be explained below. The computer will go through these steps on its own and you will not observe this process in each round. Instead, you will only see which firm you are matched with and how much money you have earned at the end of each round. However, we will go through these steps so you understand how the matches are calculated.

Before we go through the procedure in detail, we will summarize the main ideas. Essentially, the computer uses the message you submit to decide which firms' offers to accept and reject on your behalf. There are two rules that describe this process.

1. The computer never matches you with a firm that you have not included in your submitted message. This is because, even if that firm makes an offer to match with you, the computer will reject that offer.

2. The computer always matches you with the highest ranked firm (according to your submitted message) that has made you an offer. If you receive an offer from only one firm and that firm is included in your message, then the computer will accept that offer. If you receive offers from both firms and both firms are included in your message, then the computer will accept the offer from the firm that you ranked higher in your message and reject the other offer.

STEP 0: All firms and workers are unmatched. Workers (YOU) send messages to the computer. These messages are a ranking of the firms that the computer will use in the steps below.

NOTE: THE REMAINING STEPS ARE PERFORMED BY THE COMPUTER

STEP 1: <u>Firms propose offers</u>.

Each firm makes an offer to the first worker on its list.

STEP 2: <u>Workers respond to offers</u>.

(a) If a worker receives no offers, then nothing changes. The worker remains unmatched.

(b) If a worker receives one offer, then the computer uses the message of that worker to decide whether or not to accept the offer. For example, suppose that WORKER A receives an offer from FIRM A.

- If WORKER A included FIRM A in its message, then WORKER A is matched with FIRM A.

- If WORKER A did not include FIRM A in its message, then WORKER A is not matched with FIRM A. In this case, WORKER A "rejects" FIRM A.

(c) If a worker receives two offers, then the computer uses the message of that worker to decide which firm to match that worker to. For example, suppose that WORKER A receives offers from both FIRM A and FIRM B.

- If WORKER A included FIRM A in its message but not FIRM B, then WORKER A is matched with FIRM A. WORKER A rejects the offer from FIRM B.

- If WORKER A included FIRM B in its message but not FIRM A, then WORKER A is matched with FIRM B. WORKER A rejects the offer from FIRM A.

- If WORKER A included both FIRM A and FIRM B in its message, then the computer looks at the relative positions of FIRM A and FIRM B in WORKER A's message. If WORKER A's message ranks FIRM A in a higher position than FIRM B, then WORKER A is matched with FIRM A and WORKER A rejects the offer from FIRM B. If WORKER A's message ranks FIRM B in a higher position than FIRM A, then WORKER A is matched with FIRM B and WORKER A rejects the offer from FIRM A.

STEP 3: Unmatched firms propose new offers.

Each firm that is unmatched makes an offer to the second worker on its list.

STEP 4: Workers respond to offers.

(a) If a worker is unmatched, then refer to STEP 2 to determine how the worker decides among offers.

(b) If a worker is currently matched and receives no new offers, then nothing changes. The worker remains matched to whichever firm they were already matched with.

(c) If a worker is currently matched and receives a new offer, then the computer looks at the relative positions of the current match and the new offer in the worker's message. For example, suppose that WORKER A is currently matched with FIRM A and receives a new offer from FIRM B.

- If WORKER A did not include FIRM B in its message, then WORKER A rejects the offer from FIRM B. WORKER A is still matched with FIRM A.

- If WORKER A included FIRM B in its message, then the computer looks at the relative positions of FIRM A and FIRM B in WORKER A's message. If WORKER A's message ranks FIRM A in a higher position than FIRM B, then WORKER A is still matched with FIRM A and WORKER A rejects the offer from FIRM B. If WORKER A's message ranks FIRM B in a higher position than FIRM A, then WORKER A's previous match with FIRM A is broken and WORKER A is now matched with FIRM B.

•

•

•

The procedure continues in this fashion until there are no firms left to make offers. This can happen for two reasons: either both firms are already matched or there is an unmatched firm that has already been rejected by both workers. The final matches for a given round are the matches that are in place when the procedure ends. It is only the final matches that count to determine payments. Matches that are made and then broken do not count for payments.

Note that the computer does not consider your list of payments when deciding which firm to match you with. The computer only uses the rankings from your submitted message (and the submitted message of the other worker you are paired with) to calculate the final matching. Once you are matched to a firm, then your list of payments is used to determine how much money you have earned in that round.

The experimental interface is shown below. The bar at the top of the screen indicates which round the players are currently in. The left hand side of the screen displays the lists for the firms and the workers. Your own list of payments will always be in bold. The right hand side of the screen displays your role in the experiment and asks you to submit a message.

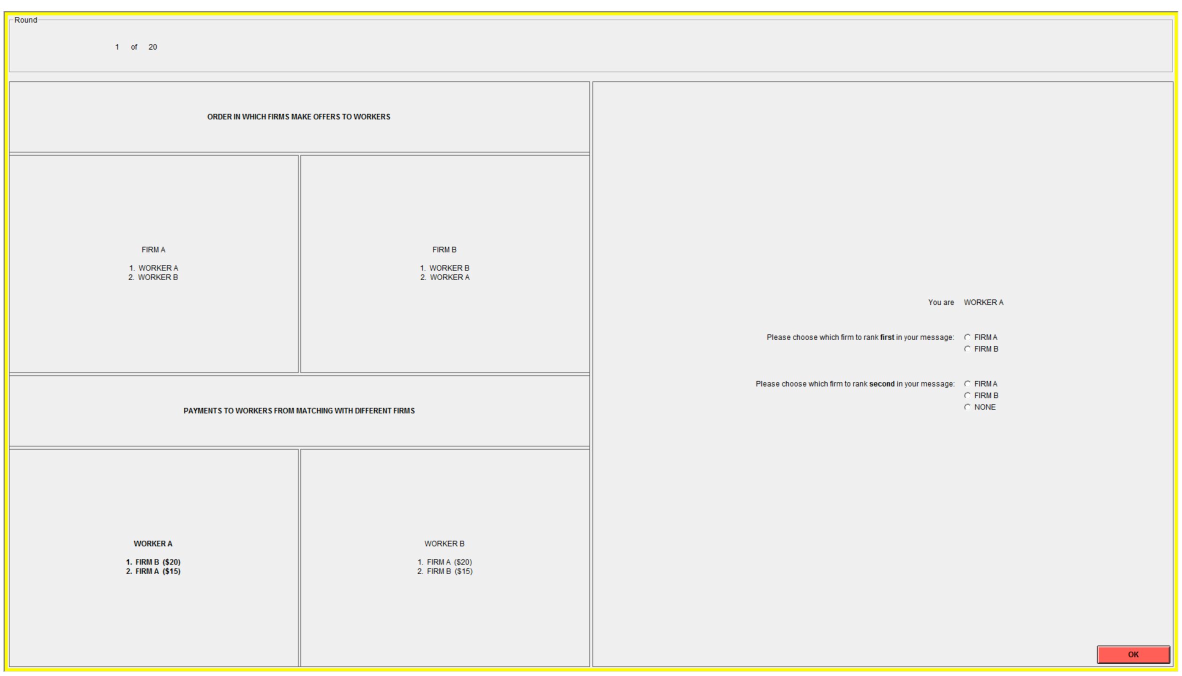

At the beginning of the experiment, there will be a brief demonstration of the procedure that the computer uses to determine the final matchings. You will walk through the steps discussed above to better understand how the messages that are submitted affect which firm you are matched with. Again, keep in mind that you will not have to go through a similar process during the actual experiment. In the experiment, the only action that you will take is to submit a message. The computer will go through these steps on its own to determine which firm you are matched with and it will then report that information to you. The purpose of the example is just to show you in detail the steps the computer is taking to determine the final matchings based on the submitted messages.

To summarize, the order of events in the experiment is as follows:

1. You will go through an example demonstrating the procedure that the computer uses to calculate the final matchings.

2. You are randomly assigned to the role of either WORKER A or WORKER B. Your role will remain the same across all 20 rounds.

3. You learn your payments (as well as the payments of the other worker) for all possible matches. You also learn the order in which the firms will make offers to match with the workers. This information will remain the same across all 20 rounds.

4. You are randomly and anonymously paired with a worker in the other role.

5. You submit a message to the computer which is a ranking of the firms. This ranking can contain either one or two firms.

6. The computer uses the submitted messages to calculate the final matches.

7. The computer reports to you which firm you are matched with and how much money you have earned.

8. You will repeat steps 4-7 a total of 19 times (since there are 20 rounds in the experiment).

9. At the end of the experiment, one of the 20 rounds will be randomly selected and you will be paid your earnings for that round (in addition to the $7 show-up payment). All payments will be made privately and in cash.

If you have any questions at this point, please raise your hand. If not, we will begin the experiment shortly.

Good luck!

WELCOME!

Please turn off all electronic devices and place them in your bag or under your desk.

Throughout the experiment, please do not talk to anybody else and please remain silent at all times. If you have any questions, raise your hand and the experimenter will come to personally assist you.

Thank you for participating in this experiment. By showing up on time, you have automatically earned a $7 payment. If you follow the instructions carefully and make good decisions, you can earn additional money. The amount of money that you ultimately earn in this experiment depends on your decisions and the decisions of others. This session will last approximately one hour. At the end of the session, you will be paid privately in cash.

The experiment will be run entirely on the computer and all interactions between yourself and others will take place via the computer terminal. There are a total of 20 rounds in this experiment. At the end of the experiment, one of the 20 rounds will be randomly selected and your monetary payment will be determined based on the outcome of that round. Each of the 20 rounds is equally likely to be selected. Thus, it is in your best interest to take each round seriously. Each round is self-contained: your decisions in one round will not affect your opportunities or earnings in another round.

This experiment is about matching. There are two groups in the experiment: firms and workers. You will be randomly assigned to the role of either WORKER A or WORKER B. Your role will remain the same across all 20 rounds of the experiment. In each round, you will be randomly and anonymously paired with a worker who is assigned to the other role. In other words, if your role is WORKER A then you will be paired with someone in the role of WORKER B in each round (and vice versa).

Your goal in each round is to match with one of the two firms: FIRM A or FIRM B. The roles of the firms are computerized. They are programmed to behave in a certain way, which will be discussed in more detail shortly. You will earn different payments from different matches. The payments corresponding to all possible matches will be shown to you on your screen. There will be four lists on your screen (one for each firm and one for each worker). A firm's list contains two workers and it shows the order in which the firm will be making offers to match with the different workers. A worker's list contains two firms and it shows how much money the worker would earn if matched with the different firms. These lists will remain the same across all 20 rounds of the experiment.

As a worker, if you are matched with the firm in the first position of your list in a given round, you will earn a payment of $20 for that round. If you are matched with the firm in the second position of your list in a given round, you will earn a payment of $5 for that round. Remaining unmatched will result in a payment of $0 for that round.

To determine which firm you are matched with, you will send a message in each round. This message is sent to the computer. It is very important that you understand what a message is since the messages sent by both you and the other worker in your pair determine which firm you are matched with. A message is a ranking of the firms. The message may or may not include all the firms. Thus, although there are two firms you could potentially be matched with, your message can contain either one or two firms.

The way the computer uses the messages to determine which firm you are matched with will be explained below. The computer will go through these steps on its own and you will not observe this process in each round. Instead, you will only see which firm you are matched with and how much money you have earned at the end of each round. However, we will go through these steps so you understand how the matches are calculated.

Before we go through the procedure in detail, we will summarize the main ideas. Essentially, the computer uses the message you submit to decide which firms' offers to accept and reject on your behalf. There are two rules that describe this process.

1. The computer never matches you with a firm that you have not included in your submitted message. This is because, even if that firm makes an offer to match with you, the computer will reject that offer.

2. The computer always matches you with the highest ranked firm (according to your submitted message) that has made you an offer. If you receive an offer from only one firm and that firm is included in your message, then the computer will accept that offer. If you receive offers from both firms and both firms are included in your message, then the computer will accept the offer from the firm that you ranked higher in your message and reject the other offer.

STEP 0: All firms and workers are unmatched. Workers (YOU) send messages to the computer. These messages are a ranking of the firms that the computer will use in the steps below.

NOTE: THE REMAINING STEPS ARE PERFORMED BY THE COMPUTER

STEP 1: <u>Firms propose offers</u>.

Each firm makes an offer to the first worker on its list.

STEP 2: <u>Workers respond to offers</u>.

(a) If a worker receives no offers, then nothing changes. The worker remains unmatched.

(b) If a worker receives one offer, then the computer uses the message of that worker to decide whether or not to accept the offer. For example, suppose that WORKER A receives an offer from FIRM A.

- If WORKER A included FIRM A in its message, then WORKER A is matched with FIRM A.

- If WORKER A did not include FIRM A in its message, then WORKER A is not matched with FIRM A. In this case, WORKER A "rejects" FIRM A.

(c) If a worker receives two offers, then the computer uses the message of that worker to decide which firm to match that worker to. For example, suppose that WORKER A receives offers from both FIRM A and FIRM B.

- If WORKER A included FIRM A in its message but not FIRM B, then WORKER A is matched with FIRM A. WORKER A rejects the offer from FIRM B.

- If WORKER A included FIRM B in its message but not FIRM A, then WORKER A is matched with FIRM B. WORKER A rejects the offer from FIRM A.

- If WORKER A included both FIRM A and FIRM B in its message, then the computer looks at the relative positions of FIRM A and FIRM B in WORKER A's message. If WORKER A's message ranks FIRM A in a higher position than FIRM B, then WORKER A is matched with FIRM A and WORKER A rejects the offer from FIRM B. If WORKER A's message ranks FIRM B in a higher position than FIRM A, then WORKER A is matched with FIRM B and WORKER A rejects the offer from FIRM A.

STEP 3: Unmatched firms propose new offers.

Each firm that is unmatched makes an offer to the second worker on its list.

STEP 4: Workers respond to offers.

(a) If a worker is unmatched, then refer to STEP 2 to determine how the worker decides among offers.

(b) If a worker is currently matched and receives no new offers, then nothing changes. The worker remains matched to whichever firm they were already matched with.

(c) If a worker is currently matched and receives a new offer, then the computer looks at the relative positions of the current match and the new offer in the worker's message. For example, suppose that WORKER A is currently matched with FIRM A and receives a new offer from FIRM B.

- If WORKER A did not include FIRM B in its message, then WORKER A rejects the offer from FIRM B. WORKER A is still matched with FIRM A.

- If WORKER A included FIRM B in its message, then the computer looks at the relative positions of FIRM A and FIRM B in WORKER A's message. If WORKER A's message ranks FIRM A in a higher position than FIRM B, then WORKER A is still matched with FIRM A and WORKER A rejects the offer from FIRM B. If WORKER A's message ranks FIRM B in a higher position than FIRM A, then WORKER A's previous match with FIRM A is broken and WORKER A is now matched with FIRM B.

•

•

•

The procedure continues in this fashion until there are no firms left to make offers. This can happen for two reasons: either both firms are already matched or there is an unmatched firm that has already been rejected by both workers. The final matches for a given round are the matches that are in place when the procedure ends. It is only the final matches that count to determine payments. Matches that are made and then broken do not count for payments.

Note that the computer does not consider your list of payments when deciding which firm to match you with. The computer only uses the rankings from your submitted message (and the submitted message of the other worker you are paired with) to calculate the final matching. Once you are matched to a firm, then your list of payments is used to determine how much money you have earned in that round.

The experimental interface is shown below. The bar at the top of the screen indicates which round the players are currently in. The left hand side of the screen displays the lists for the firms and the workers. Your own list of payments will always be in bold. The right hand side of the screen displays your role in the experiment and asks you to submit a message.

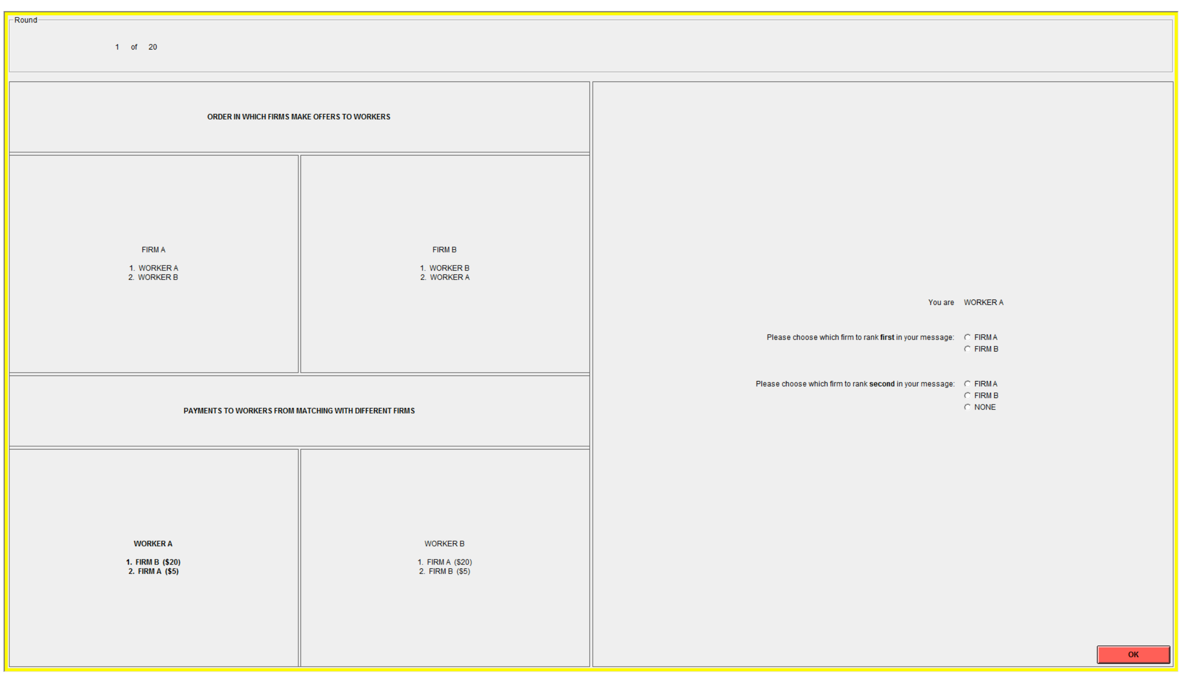

At the beginning of the experiment, there will be a brief demonstration of the procedure that the computer uses to determine the final matchings. You will walk through the steps discussed above to better understand how the messages that are submitted affect which firm you are matched with. Again, keep in mind that you will not have to go through a similar process during the actual experiment. In the experiment, the only action that you will take is to submit a message. The computer will go through these steps on its own to determine which firm you are matched with and it will then report that information to you. The purpose of the example is just to show you in detail the steps the computer is taking to determine the final matchings based on the submitted messages.

To summarize, the order of events in the experiment is as follows:

1. You will go through an example demonstrating the procedure that the computer uses to calculate the final matchings.

2. You are randomly assigned to the role of either WORKER A or WORKER B. Your role will remain the same across all 20 rounds.

3. You learn your payments (as well as the payments of the other worker) for all possible matches. You also learn the order in which the firms will make offers to match with the workers. This information will remain the same across all 20 rounds.

4. You are randomly and anonymously paired with a worker in the other role.

5. You submit a message to the computer which is a ranking of the firms. This ranking can contain either one or two firms.

6. The computer uses the submitted messages to calculate the final matches.

7. The computer reports to you which firm you are matched with and how much money you have earned.

8. You will repeat steps 4-7 a total of 19 times (since there are 20 rounds in the experiment).

9. At the end of the experiment, one of the 20 rounds will be randomly selected and you will be paid your earnings for that round (in addition to the $7 show-up payment). All payments will be made privately and in cash.

If you have any questions at this point, please raise your hand. If not, we will begin the experiment shortly.

Good luck!