LEARNING TO MANIPULATE: OUT-OF-EQUILIBRIUM TRUTH-TELLING IN MATCHING MARKETS

Clayton R. Featherstone, Eric Mayefsky and Colin D. Sullivan

In two-sided settings, market designers tend to advocate for deferred acceptance (DA) over priority mechanisms, even though theory tells us that both types of mechanisms can yield unstable matches in incomplete information equilibrium. However, if match participants on the proposed-to side deviate from equilibrium by truth-telling, then DA yields stable outcomes. In a novel experimental setting, we find out-of-equilibrium truth-telling under DA but not under a priority mechanism, which could help to explain the success of DA in preventing unraveling in the field. We then attempt to explain the difference in behavior across mechanisms by estimating an experience-weighted learning model adapted to this complex strategic environment. We find that initial beliefs drive the difference in agents' ability to find strategic equilibria, rather than alternative explanations such as differences in the learning process.

Keywords: Learning, Experiment, Matching, Deferred acceptance, Priority mechanisms, Stability.

1. INTRODUCTION

Why do some two-sided matching mechanisms continue to be used from year to year while others are abandoned? Although the usual distinction concerns whether a mechanism is stable with respect to the reported preferences, such an explanation is incomplete without also considering whether preferences are truthfully revealed.[1] Previous theoretical literature has looked at large markets to do this; however, we take a different tack by observing strategic preference revelation in the lab. Our evidence suggests that out-of-equilibrium truth-telling under the deferred acceptance mechanism can lead to matches that are more stable than theory predicts.

Two-sided matching mechanisms are widely used in the field. The most well-known example is the National Resident Matching Program (NRMP) which every year makes about 25,000 matches between newly-minted doctors and residency programs in the United States (NRMP, 2009). Once participants have formed their preferences, they submit rank-order lists of acceptable match partners to the NRMP clearinghouse, which then runs those lists through an algorithm, outputting a match. Other examples of two-sided matching include the Association of Psychology Post-doctoral and Internship Centers (APPIC) match (about 2,800 clinical psychologists matched to internship programs per year (APPIC, 2009)), and the New York City Department of Education public high school

[1] One might also bypass truthful preference revelation entirely and simply look at whether a mechanism yields a stable allocation in equilibrium. See, for instance, Roth (1984b), Ergin and Sönmez (2006) and Pathak and Sönmez (2008). Roth (1982) shows that any mechanism that is stable with respect to reported preferences cannot admit truth-telling as a dominant strategy for all players.

match (about 90,000 high school students per year (NYC-DOE, 2009)).[2]

When deciding which mechanism to use in a matching market, the literature has consistently come back to the idea of stability. A *stable* match has no agents who would prefer to remain unmatched (individual rationality) and no *blocking pairs* (pairwise stability), where a blocking pair is two agents who prefer each other to their assigned matches. If agents are free to recontract *ex post*, it is not too hard to see how instabilities might render the match moot, but even if agents must abide by the match, they can sidestep it by anticipating blocking pairs and either formally contracting early or informally prearranging a match.[3] This has been shown both theoretically (Sönmez, 1999; Roth, 1991) and in the lab (Kagel and Roth, 2000). If too many agents leave the match or prearrange, then the clearinghouse will likely be abandoned. Of course, a stable matching mechanism does not necessarily prevent unraveling, but in many real world markets, whether or not a stable mechanism is used seems to make the difference.[4]

Most matching schemes in the field can be classified as either *deferred acceptance (DA)* mechanisms or *priority* mechanisms.[5] DA mechanisms are based on the Gale-Shapley algorithm. One such mechanism, *M-Proposing DA*, is implemented in the following way, denoting the members of the two sides of the market Ms and Ws (Gale and Shapley, 1962):

M-Proposing DA

Step 1: All Ms make an offer to their first-choice W; Ws hold their favorite acceptable offer, rejecting all others.

Step t: Rejected Ms make an offer to their favorite acceptable W that hasn't rejected them yet; Ws hold their favorite acceptable offer from this round and previous rounds, rejecting all others.

STOP: The algorithm stops in the first round where no new offers are made. All held offers become finalized matches.

Priority mechanisms instead use the preferences submitted by the participants to order the set of all possible match pairs. They then try to implement those

[2]For papers on these matches, see Roth (1984a, 1996, 2003); Roth and Peranson (1999); Roth and Xing (1997); Abdulkadiroğlu, Pathak, and Roth (2005); Abdulkadiroglu, Pathak, and Roth (2009).

[3]Usually a pair can do this by agreeing to rank each other first to the clearinghouse. Most mechanisms guarantee that two partners who rank each other first will be matched.

[4]Other causes of early contracting include: insuring over states of the world before payoff relevant information is revealed (Roth and Xing, 1994; Li and Rosen, 1998; Li and Suen, 2000; Suen, 2000), the presence of market power (Roth and Xing, 1994), similar preferences (Hałaburda, 2010), arrival of new agents (Du and Livne, 2010), excess supply of workers combined with insufficient supply of high quality workers (Niederle, Roth, and Ünver, 2009), cultural norms concerning exploding offers (Niederle and Roth, 2009), information transmission through a social network (Fainmesser, 2013), and costs of participation (Damiano, Li, and Suen, 2005).

[5]Another important class of mechanisms, based on linear programming optimization, is not considered here. See Ünver (2001) and Ünver (2005).

match pairs in that order, skipping those that are not feasible due to previously implemented matches (Roth and Sotomayor, 1990). For concreteness, consider the M-*Proposing Priority* mechanism implemented by the following algorithm:[6]

M-Proposing Priority

Step 1: All Ms make an offer to their first-choice W; Ws are permanently matched to their favorite acceptable M who made an offer, rejecting all other offers.

Step t: Rejected Ms make an offer to their favorite acceptable W that has not yet rejected them; matched Ws reject all offers; and unmatched Ws are permanently matched to their favorite acceptable M who made an offer.

STOP: The algorithm stops in the first round where no new offers are made.

A key difference between the M-Proposing DA and M-Proposing Priority algorithms is that DA mechanisms yield matches that are stable with respect to the reported preferences, while Priority mechanisms generally do not. Empirically, markets using DA tend to outlast markets implementing Priority algorithms. For example, Roth (1991) exploited regional variation in medical residency matches in the United Kingdom to demonstrate that regions that adopted DA mechanisms tended to keep using them, while regions that adopted priority mechanisms tended to abandon them after a few years.

However, the simple fact that DA is stable relative to the reported preferences cannot explain why it outlasts priority mechanisms. Under DA, only participants on the proposing side have incentive to report truthfully. The receiving side often fails to reveal truthfully in Bayes-Nash equilibrium (Roth and Rothblum, 1999; Coles and Shorrer, 2014).[7] Furthermore, in equilibrium, neither DA nor priority mechanisms should yield matches that are stable relative to true preferences under incomplete information. Why then does DA persist where Priority fails? Several contributing causes have been considered, but there are still some markets where these explanations are not fully satisfactory.

It could be that preferences are near perfectly correlated on one or both sides of the market. This would push the market toward a unique stable match, thereby removing the incentive to deviate from truth-telling under DA.[8] Although it is intuitive to expect some correlation in preferences, we might also expect a lack of correlation in preferences across matches that are commonly perceived to be of similar quality.

Another possibility is that agents find being unmatched extremely distasteful. Potentially profitable manipulations take a gamble at being unmatched in ex-

[6]The priority ordering for this mechanism ranks potential match pairs in the order of Ms' preferences, with ties broken by Ws' preferences.

[7]Similar results holds for priority mechanisms (Ehlers, 2008).

[8]The simplest way to see this is in the one-to-one case, where it is a straightforward application of the Blocking Lemma and the fact that no individually rational matching can make all the members of one side of the market strictly better off than the unique stable match (Roth and Sotomayor 1990, Lemma 3.5 and Theorem 2.27).

change for a higher probability of matching to a more preferred partner (Roth and Rothblum, 1999). If being unmatched is bad enough, no agent will take this gamble. Even so, in many situations, it is unclear how bad being unmatched is. For instance, in the NRMP match, where hospitals are on the receiving side of the market, unmatched positions can still be filled in the centrally organized aftermarket, known as the "Scramble".

A third option is that the number of stable matches gets small as the market gets large, as established theoretically in Immorlica and Mahdian (2005) and significantly extended in Kojima and Pathak (2009). Although these papers lay out an intuitive mechanism by which core convergence might occur, they do so in the context of a very slowly converging asymptotic (Kadam, 2011); for example, if agents are allowed to list five acceptable members on the other side of the market, as is the case in our experiment, then the Kojima and Pathak (2009) bound on the fraction of agents who could profitably deviate from a truth-telling equilibrium does not go below 1 until the market has in excess of 10^{34} agents.[9] Because of the extreme looseness of this upper bound at more reasonable market sizes, we must instead rely on computational work to give us an idea of how big a market must be for large market results to kick in.

Fortunately, Roth and Peranson (1999) provides just such a benchmark. They show that there is little leeway for manipulation relative to submitted preferences in the NRMP match, although, as they mention, this could be because the submitted preferences had already been manipulated to an equilibrium. To evaluate this possibility, they then look at large simulated markets, finding that markets the size of the NRMP have little room for manipulation, while smaller ones do.[10] Unfortunately, such computational work merely tells us that there is likely a much better bound than the one derived in Kojima and Pathak (2009). How much better remains an open question.

Hence, previous research leaves us reasonably confident that very large markets, such as the NRMP (around 20,000 agents), have very small cores, but leaves us less certain about smaller markets. And there are many such markets; in addition to the small regional matches in the UK (about 150 agents) there are many smaller fellowship matches run by the NRMP where DA also seems to halt unraveling, most of which have fewer than 100 fellowship programs represented, some with multiple positions for each program (Roth, 1991; NRMP, 2009).

A new explanation for the empirical success of DA, which we pursue in this paper, is that match participants on the receiving side of a DA mechanism might truth-tell in an out-of-equilibrium manner, leading to truly stable matches. To

[9]Specifically, the asymptotic states that the upper bound equals $\frac{16\cdot\overline{q}\cdot k}{\log(\overline{q}\cdot n)}$, where $\overline{q}$ is the maximum capacity of any hospital, k is the number of hospitals that each doctor is allowed to list, and n is the number of hospitals. We set $\overline{q} = 1$ and $k = 5$, and solve for the n that makes the bound equal to 1.

[10]See Figure 2 in Roth and Peranson (1999). Further, note that its simulations involve are for one-to-one markets. The asymptotic mentioned in Footnote 9 implies that there is more leeway for manipulation in many-to-one markets, as $\overline{q}$ and k must increase.

confirm this intuition, we will look at strategies used by experimental participants on the receiving side of DA and M-Proposing Priority both in an environment where they should truth-tell and in an environment where they should deviate from truth-telling. We find that truth-telling rates are similarly high in both environments under DA, but that truth-telling rates are both economically and statistically different under Priority. The first result supports our story of out-of-equilibrium truth-telling, while the second demonstrates that the truth-telling is unlikely to be a mere artifact of the lab.

To understand what drives the differences in strategic play, we estimate a structural learning model, derived from the Experience-Weighted Attraction (EWA) model introduced by Camerer and Ho (1999). Our reparametrized EWA decomposes strategic play into initial cognition and learning dynamics. We find that under DA, players fail to identify profitable deviations from truth-telling in their initial assessment of the strategic environment. In addition, they rely more heavily on their faulty assessment, and they are reluctant to explore new strategies.

We would like to emphasize that we think of the out-of-equilibrium truth-telling explanation put forward by this paper as a complement of, rather than a replacement for, the other explanations we have mentioned. The persistence of DA even in small markets implies that there might be something else going on besides the core convergence explanations which have previously been put forward, and we primarily seek to address this gap in understanding.

Before proceeding, we briefly mention how the current paper fits into the previous experimental matching literature. The first two-sided matching experiments date to the early 1990s (Sondak and Bazerman, 1991; Harrison and McCabe, 1996). An experiment that explicitly compares priority and DA mechanisms is described in Kagel and Roth (2000), although their paper focuses more on unraveling behavior than on strategic preference revelation. They do, however, provide a nice demonstration of the intuitive link between stability and persistence. Ünver (2005) runs a similar experiment that also includes linear programming mechanisms. Related experiments include Haruvy and Unver (2007) and Echenique and Yariv (2010), which look at repeated decentralized markets, and Nalbantian and Schotter (1995), which looks at several mechanisms that involve matching with money. Our experiment is perhaps most closely related to Echenique, Wilson, and Yariv (2016), which also looks at strategies in a two-sided matching market. Their design allows agents to go through the DA algorithm as an extensive form game, and their main finding is that agents on the proposing side tend to skip over proposals sub-optimally. Our design treats the DA algorithm as a normal form game, and we focus on the strategies of the receiving side of the algorithm, finding some sub-optimal truth-telling. To our knowledge, we are the first paper to focus on the strategies of the receiving side explicitly. A recent body of literature has developed comparing predicted and actual play under different allocation mechanisms (e.g., Li (2017); Rees-Jones (2017); Duflo (2017); Zhang and Levin (2017)). Finally, we mention several other experiments that focus on strategies used by the proposing side, mainly in the context of school choice, such as Chen and Sönmez (2006); Pais and Pintér (2008); Calsamiglia, Haeringer, and Klijn (2009), and Featherstone and Niederle (2016).

Our findings also contribute to large body of work on preferences for truth-telling. In a meta-analysis of previous experiments, Abeler, Nosenzo, and Raymond (2019) argue that subjects display behavior consistent with a preference for truth-telling for being seen as honest, and suggest that mechanisms that rely on truthful revelation may be effective even when they are not incentive compatible. Our results indicate that the preference for truth-telling may be outweighed when the mechanism is easy to game.

2. TWO MARKETS

In our experiment, we will use M-Proposing DA and M-Proposing Priority in conjunction with two different market structures.[11] Under one structure, theory predicts that the receiving side will deviate from truth-telling in a particular way under both mechanisms, while under the other structure, theory predicts truth-telling. Note that our experimental design will constrain the Ms to truth-tell, focusing on the behavior of the Ws. Because of this design feature, our equilibrium characterizations concern how the Ws respond to the truth-telling from the Ms and whether truth-telling can be sustained in equilibrium for the Ms.

Throughout this section, we will only present theoretical results specific to our experimental markets, but in the Appendix, we show that there is a broad class of symmetric environments in which we expect similar results.[12] Symmetric environments can be thought of as representing realistic situations where match participants have little information about others' preferences. In such settings, the kinds of manipulations that we expect to see in the lab (truncations) are, in the sense of Roth and Rothblum (1999) and Ehlers (2008), fundamental.[13]

2.1. *The uncorrelated market*

Consider a small matching market with 5 Ms and 5 Ws. The true ordinal preferences of each participant are drawn independently from the uniform distribution over rank-order lists that rank $\emptyset$ (the outcome of being unmatched) last. Cardinal payoffs are a decreasing function of ordinal rank only. We call this the *uncorrelated market*.

[11]See the Introduction for definitions of these mechanisms.

[12]See Appendix A for rigorous model definitions, and Appendix B for all proofs. The results in the Appendix are also of some independent interest because they extend the results of Roth and Rothblum (1999) and Ehlers (2008) to show how truncation strategies are not just best-responses to symmetric beliefs, but are also the strategies used in equilibria in which agents use anonymous strategies.

[13]Also see Day and Milgrom (2008) on how such strategies also appear in core selecting auctions.

Before proceeding to characterize equilibrium, we must first introduce a few definitions. A *revelation strategy* is a mapping from true preferences to reported preferences. Now, due to the symmetry of the problem, any equilibrium in which some agent used a strategy that depended only on a match partner's label would seem unnatural. Therefore, think of an agent's true preferences as a six element vector with the outcome of being unmatched, $\emptyset$, as its last entry, and define an *anonymous strategy* to be a revelation strategy that always reports the same permutation of the true preference vector.[14] Further, define a *truncation strategy* to be an anonymous strategy where the permutation simply switches the sixth element and some other element of the true preference. We will also consider it a truncation if the permutation is the identity, that is, truth-telling is also a truncation strategy.

Under M-Proposing DA, the characterization of equilibrium is quite simple, extending the main result of Roth and Rothblum (1999).[15]

PROPOSITION 1 *In the uncorrelated market, under M-Proposing DA, any equilibrium in anonymous, weakly undominated strategies involves truth-telling for each $m \in M$ and truncation for each $w \in W$.*

Under M-Proposing Priority, the best-response of the Ws when the Ms are constrained to truth-tell is similar, extending the main result from Ehlers (2008).

PROPOSITION 2 *In the uncorrelated market, under M-Proposing Priority, if all agents play anonymous, weakly undominated strategies, and all $m \in M$ truth-tell, then all $w \in W$ best-respond to the other agents by playing truncations.*

In the uncorrelated market, then, the unifying principle is that, under both mechanisms, we expect to see the members of W playing truncation strategies.[16]

2.2. *The correlated market*

Now, instead of drawing preferences independently for the members of M, draw only one preference and give it to all members of M. Continue to draw preferences independently for each member of W. We call this the *correlated market*. Propositions 3 and 4 demonstrate that we expect truth-telling for the members of W under both mechanisms.

[14]Note that there is some redundancy in this definition, as the ordering of agents ranked as unacceptable does not matter in any of the mechanisms we consider.

[15]Roth and Rothblum (1999) concerns best response to a certain class of beliefs; our theorem concerns strategies used in a certain class of equilibria.

[16]We might be worried that an experiment that constrains the Ms to truth-tell doesn't have much external validity if such behavior cannot be supported in equilibrium. To this critique, we can provide two statements which are proven in the Appendix. The first is that, at any symmetric equilibrium, the Ms must truth-tell. The second is that the strategic problem of the Ws is the same, regardless of what anonymous, weakly undominated strategies the Ms play, since filtering a uniform distribution through a permutation yields a uniform distribution.

	Truncation (uncorrelated market)	Truth-telling (correlated market)
Priority	9 groups	8 groups
DA	9 groups	8 groups

TABLE I

Experimental treatments

Proposition 3 *In the correlated market, under M-Proposing DA, the unique equilibrium in anonymous, weakly undominated strategies entails truth-telling by all agents.*

Proposition 4 *In the correlated market, under M-Proposing Priority, if all members of M have the same anonymous, weakly undominated strategy, then all members of W best respond by truthfully revealing.*

Proposition 3 follows from realizing that if the members of M must truth-tell, then there is a unique stable match relative to the reported preferences. With a unique stable match, there is no reason to deviate from truth-telling.[17] Proposition 4 follows from realizing that if all members of M play the same revelation strategy, then they will all submit the same reported preferences, which means that a member of W receives all offers in the same round of the M-Proposing Priority algorithm.

To conclude, we might worry that it is unrealistic that all members of M should use the same revelation strategy. The next proposition addresses this concern.

Proposition 5 *In the correlated environment, there exist cardinal payoffs that rationalize an equilibrium where all Ms and Ws truthfully reveal their preferences.*

Intuitively, we know this is so by thinking of a case where the payoff for getting a first-ranked W is more than 5 times the payment for getting a second-ranked W, which in turn is more than 4 times the payment for getting a third-ranked W, etc.

3. EXPERIMENTAL SETUP

Table I shows the four treatments which comprise the experiment's 2×2 design. We switch the profitability of truncation on and off by switching between the correlated and uncorrelated markets. If our hypothesis holds, we would see no significant difference across these markets under M-Proposing DA. It could then be, however, that experimental participants always tell the truth in the lab.

[17]See Footnote 8 for the sketch of the proof.

To control for this, we also observe participant behavior under M-Proposing Priority, where the rationale for deviating from truth-telling seems more straightforward. If we observe a difference in truth-telling across markets under Priority, but not under DA, then we will have shown a real effect.

In the experiment, only Ws will be played by human participants; the Ms will be played by the computer and constrained to truthfully reveal their preferences. Obviously, in real life two-sided matching markets, the proposing side's report to the matching mechanism is not automatic. Under Priority, proposers do not necessarily have dominant strategy incentives to report their preferences truthfully (although as discussed in Section 2, this behavior can occur in equilibrium), and under DA, truthful reporting is a dominant strategy, but there is some experimental evidence that proposing side agents may not propose to all agents in order in an extensive form matching market without frictions (Echenique, Wilson, and Yariv, 2016). We nevertheless use automated proposers playing fixed strategies so that we can focus on the previously unexamined behavior of the receiving side under DA. Using automated Ms reduces the complexity and noise in the decision the participants face. If, as we anticipate, subjects have difficulty learning to successfully manipulate the mechanism in this simplified environment, we are confident they will also have trouble in the more complicated real world markets of interest.

In the lab, each participant plays the same market for 40 rounds with the same group of five players. In every repetition, each W privately learns their new preferences and submits a ranking of some, all or none of the Ms. The computer then generates a match outcome according to the rules of the appropriate mechanism to the treatment. Ws then learn their match outcome, as well as the outcomes of all other Ws. They gain points based on where their match partner appeared in their true preference list for that round, according to payoffs given in Table II. When designing these payoffs, our goal was to find a payoff scheme which provided behavioral incentives that were as comparable as possible between treatments. In Figure 1, we show that we succeeded, relative to the actual behavior observed in the lab.[18] Under the payoff scheme in Table II, the pure strategy equilibria in mixed strategies differ between DA and Priority: it is an equilibrium for all agents to truncate their final two positions (truthfully reporting positions one, two, and three) under *DA-Truncation*, and for agents to truncate three positions under *Priority-Truncation*.[19] In Figure 2, we show average realized payoffs to different types of strategies in each setting. As expected, truth-telling is most profitable in the truth-telling environment; non-truthful truncation is most lucrative in the truncation environment; and non-truncation

[18]Note that a simple reinforcement learning model would predict that the slopes of the curves are much more important than the levels.

[19]We calculate the best response functions by simulation, as in Coles and Shorrer (2014). When all other Ws report their preferences truthfully, the best response is for the final agent to truncate the final three preferences under DA, and to truncate the final two positions under Priority.

Match	1st choice	2nd	3rd	4th	5th	No match
Payoff	32 points	16	8	4	2	0

TABLE II

Payoff table

permutations are least profitable in all environments.

Finally, we address the design choice to allow for repetition, even though most individuals participate in a matching process in the field only once (or perhaps a handful of times in some applications). In the lab, we can adequately mimic neither the stakes faced by participants in real matching markets nor can we realistically allow experimental participants as much time to consider their prospects as they would have in the field. Instead, by having them participate in repeated trials, we allow for participants to learn about the environment and possibly alter their strategy as they progress. One could argue that this makes it unrealistically easy for participants to behave strategically; however, if subjects nonetheless fail to manipulate effectively, we can be confident that manipulation is even more difficult in the field. Moreover, repeated play allows us to model the learning process empirically, identifying characteristics of learning that make predicted equilibria more or less likely.

Briefly, we mention the symmetry of our experimental environments. Non-truncation strategies are not profitable in our setup, but in the field, they might be. Even so, such strategies require much information to implement. Also, though preferences in real-world markets might not look much like those in our experiment, preferences are often tiered. One set (tier) of match partners is clearly preferred to another set, which is preferred to yet another set, but over each tier, preferences are idiosyncratic. In this context, the setup of our experiment can be interpreted as an approximation of at least a sector of the matching market.[20]

All treatments were run at Stanford University during the Spring of 2009. Each session consisted of one or two groups of 5 participants. In sessions with two groups, groups were not mixed during the session, and participants were not informed which other participants were in their group. At the start of each session, participants were read detailed instructions and had to successfully work through the steps of the appropriate mechanism for an example set of reported preferences.[21] Actual play commenced only after all participants completed the exercise and indicated they understood the mechanism rules. Nothing was done to overtly suggest what the treatment variables were, i.e., there was no mention of matching mechanisms or preference distributions other than the ones in use

[20]Additionally, since interview constraints often prevent match participants from evaluating all potential match partners, we might think that pre-match sorting would lead to market segmentation, to similar effect. For more on modeling the interview process, see Lee and Schwarz (2007), Lee and Schwarz (2009), and Coles, Kushnir, and Niederle (2013).

[21]In the lab, we provide a specific context in the hopes of making understanding easier for participants. Proposing side agents (referred to here as Ms) are referred to as "Schools" and the agents receiving offers (here, Ws) are referred to as "Students."

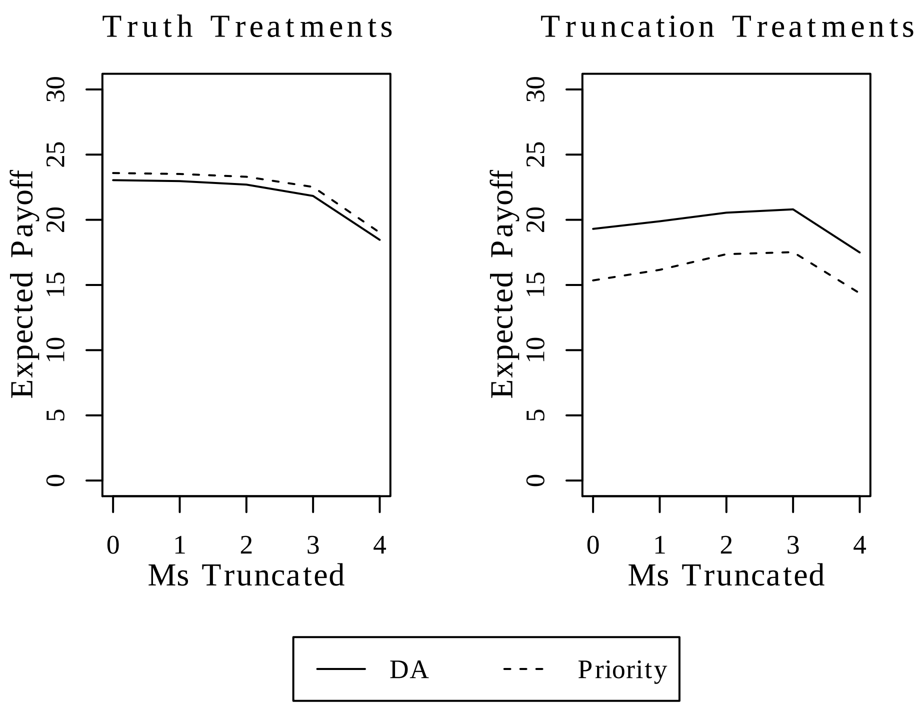

Figure 1: Expected payoff versus number of Ms truncated (empirical)

Expected empirical round payoffs to different strategies across treatments. Expected empirical payoffs are the average what each player would have received in every round under the indicated strategy, holding all other players' strategies constant. *Ms Truncated* indicates the length of the truncation, with 0 representing truth telling {12345}, and 4 representing the truncation strategy {1∅∅∅∅} where only the most preferred match is listed as acceptable. Truncation strategies do not include non-truthful inversions of true preferences.

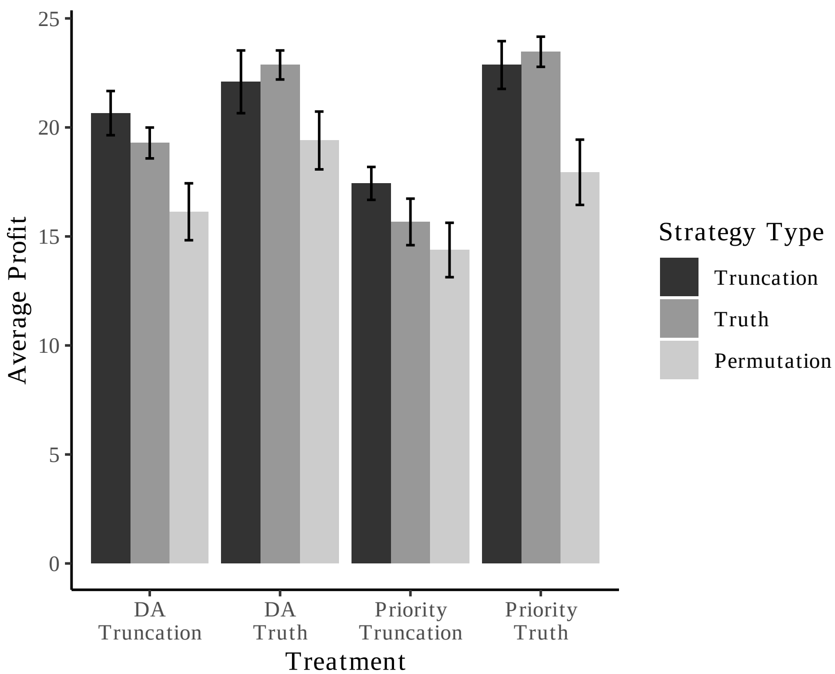

Figure 2: Average Payoff by Strategy Type

Average empirical round payoffs to different strategy types across treatments. Bars represent 95% confidence intervals. Color represents actual subject behavior; treatment indicates mechanism and expected behavior based on the correlation of preferences in the market. For strategy types, *Truncation* indicates a non-truth truncation strategy; *Truth* indicates truth-telling; *Permutation* indicates all other strategies. Average profit is the average of payoffs across all rounds in the treatment environment, including all players who play the indicated strategy in that period.

	DA		Priority
Truth-Telling	66.0%	↔(0.372)	58.4%
	↕(0.200)		↕(0.002)**
Truncation	56.6%	↔(0.001)**	25.3%

TABLE III

Truth-telling rates (all periods)

Numbers in parentheses are p-values from two-tailed Mann-Whitney tests with session-level averages as the units of observation.

	DA		Priority
Truth-Telling	70.2%	↔(0.340)	60.8%
	↕(0.046)**		↕(0.002)**
Truncation	54.7%	↔(0.003)**	19.3%

TABLE IV

Truth-telling rates (last 10 periods)

Numbers in parentheses are p-values from two-tailed Mann-Whitney tests with session-level averages as the units of observation.

in that particular treatment.

During the experimental session, participants could see their preferences for a given round on their computer screen and were reminded of payments for all possible match outcomes. They were then directed to click on radio buttons to rank each of the Ms.[22] After all participants submitted rankings, a results screen showing the participant's match for that round, their point accrual for that round and their total cumulative points would be displayed. At all times, a participant had the ability to see, for all prior rounds, the match outcomes for all participants, her own true preferences, and the rank list she submitted in that round.

4. EXPERIMENTAL RESULTS

4.1. *Overall Truth-telling Rates*

We are most interested in the rate of truth-telling over all periods across the four primary treatments. This value is significantly higher in the DA truncation treatment than in the Priority truncation treatment; however, for the two truth-telling treatments, the differences between the DA and Priority treatments are not statistically significant. Furthermore, the rate difference between the two DA treatments is not statistically significant, while the difference between the two Priority treatments is highly significant.

When we restrict attention to the last ten periods, focusing on the behavior of subjects when they are more experienced, we find qualitatively similar effects.

[22]We did this so that participants would have to click the same number of times regardless of what preference they wished to report. If declaring all Ms unacceptable were too easy, some participants might choose to do this in order to save time and effort.

	DA		Priority
Truth-Telling	16.3%	↔(0.226)	11.1%
	↕(0.673)		↕(0.210)
Truncation	14.3%	↔(0.508)	17.9%

TABLE V

Non-truncation rates

Numbers in parentheses are p-values from two-tailed Mann-Whitney tests with session-level averages as the units of observation.

Statistically, there is a mildly significant difference between the two DA treatments, as well as the high significance between the Priority treatments and the truncation treatments seen in the data for all 40 periods.

Note that for DA, truth-telling rates are slightly lower in the last 10 periods (2% lower) in the truncation treatment, but also 4% higher in the truth-telling treatment. Thus, the significance of the difference in truth-telling rates between the two groups is in some sense as much due to participants in the truth-telling treatment learning to tell the truth as it is those in the truncation treatment learning to truncate. In sum, we only see a significant deviation from the benchmark truth-telling rate under the Priority truncation treatment. Under DA, participants do not respond to the truncation treatment by deviating from truth-telling.

Of course, failure to tell the truth is not synonymous with truncation, and although truncation weakly dominates other non-truth-telling strategies, we do observe some portion of suspects employing "switching" or "dropping" strategies in some rounds. Frequency of this behavior, however, is not significantly different between any of the treatments.

4.2. *Blocking Pairs and Overall Match Stability*

For practical market design, we may be primarily concerned not with the rate at which participants tell the truth, but rather with how successfully a mechanism generates desirable (i.e., stable) match outcomes. One measure of this is the number of blocking pairs present in any given assignment. Since the outcome is never 100% stable in any treatment at any time, the number of blocking pairs is one measure of the degree of stability of a match outcome: a mechanism which generates an outcome that is stable for most participants may still work well enough to be persistent.

Blocking pairs were found to occur significantly more often in the Priority truncation treatment than in the DA truncation treatment or the Priority truth-telling treatment. The two DA treatments were not significantly different in blocking pair frequency; nor were the two truth-telling treatments.

Note that the same M or W can be involved in multiple blocking pairs if there is more than one attainable match partner that they prefer to their actual match partner. However, we do not observe any interesting asymmetries in terms of

	DA		Priority
Truth-Telling	0.47	↔(0.574)	0.59
	↕(0.809)		↕(0.001)**
Truncation	0.49	↔(0.000)**	1.87

TABLE VI

Number of Blocking Pairs per Period

Numbers in parentheses are p-values from two-tailed Mann-Whitney tests with session-level averages as the units of observation.

	DA		Priority
Truth-Telling	2.7%	↔(0.065)*	4.9%
	↕(0.311)		↕(0.030)**
Truncation	3.7%	↔(0.010)**	11.1%

TABLE VII

Percentage of Ms and Ws Unmatched

Numbers in parentheses are p-values from two-tailed Mann-Whitney tests with session-level averages as the units of observation.

which unique agents are involved in multiple blocking pairs: the number of unique Ms involved in blocking pairs is not significantly different than the number of unique Ws for any treatment, and the between-treatment differences are similar qualitatively and in terms of statistical significance when the number of unique Ms and Ws in blocking pairs are considered separately. The total probability of an M or W being unmatched thus follows a similar pattern across treatments.

4.3. *Best Response Frequencies*

Truth-telling rates establish how apt participants are to manipulate, and low non-truth, non-truncation rates establish that these manipulations are, for the most part, some sort of truncation.[23] However, participants who truncate are not automatically maximizing their expected payoff: they may be truncating too much or too little. For the set of payoffs used in the experiment, we can find an equilibrium where all agents truncate symmetrically; however, as out-of-equilibrium strategies may be a best response to other out of equilibrium strategies, we would not necessarily expect sophisticated participants to truncate as if in equilibrium. We instead look at the ability of participants to find the strategy which is a best response to the environment in which they find themselves. If a significant proportion of subjects are able to achieve this in a significant portion of sessions for a certain mechanism, we might reach different conclusions as to their sophistication than we would looking strictly at truth-telling rates (or looking at the frequency of play consistent with theoretical equilibrium, for

[23]The characterization of this other behavior as "non-truthful, non-truncation" is redundant, as truth-telling is one extreme of the set of truncation strategies for participants. We nevertheless use the terminology to ensure clarity.

that matter). Also, we might wonder if there is a great deal of heterogeneity in participant sophistication, or if all participants reported optimal truncations about the same fraction of the time.

However, simply comparing subjects' behavior in an individual round to the optimal behavior possible in that period *ex post* fails to capture the uncertainty which is inherent in truncation strategies—it can be optimal *ex ante* to truncate in each period, even though it may be suboptimal *ex post*. Thus, we consider the participant to be playing optimally in their "environment" if they play the truncation strategy which generates the highest expected utility across some set of rounds they played, given the actual behavior of other participants and generated proposer preferences.

Figure 3a indicates the proportion of participants playing an overall best response at most the indicated proportion of the time for the truncation treatments. For example, approximately 36% of Priority participants never played a best response (compared with about 52% for DA), and 50% of participants played a best response no more than 20% of the time (compared with around 75% for DA). Note that the Priority treatment first order stochastically dominates the DA treatment: for any level of frequency of best response play we consider, more participants best respond at least that frequently in the Priority treatment than in the DA treatment. However, this gap closes when only the last 20 periods are considered, as seen in Figure 3b. Note that this closing of the gap simply implies that under both mechanisms, participants have converged to similarly bad distributions of sub-optimal play.

In the truth-telling treatments (Figures 3c and 3d), truthful reporting is always the unique best response, and much as there was no significant difference in the overall truth-telling rates between DA and Priority in these treatments, there is no noticeable difference in the frequency with which individual subjects play this best response, either in the whole sample or restricting attention to the last 20 periods.

5. LEARNING MODEL

We have shown that subjects learn to manipulate reported preferences advantageously under the Priority mechanism but not under DA, despite theoretical predictions. Understanding how actual behavior departs from theoretical predictions under different allocation mechanisms has become an important area of market design research (e.g., Li (2017); Rees-Jones (2017); Duflo (2017); Zhang and Levin (2017)). However, the process of learning strategic play over time in a market design setting remains poorly understood. The learning process can inform predictions about which equilibria are likely to arise and how long agents will take to find them. In this section, we estimate a structural model and describe the dynamics of learning in a repeated game. We then simulate play under counterfactual conditions to understand what features of the learning process drive the gap in strategic play under different mechanisms. Finally, we incorporate

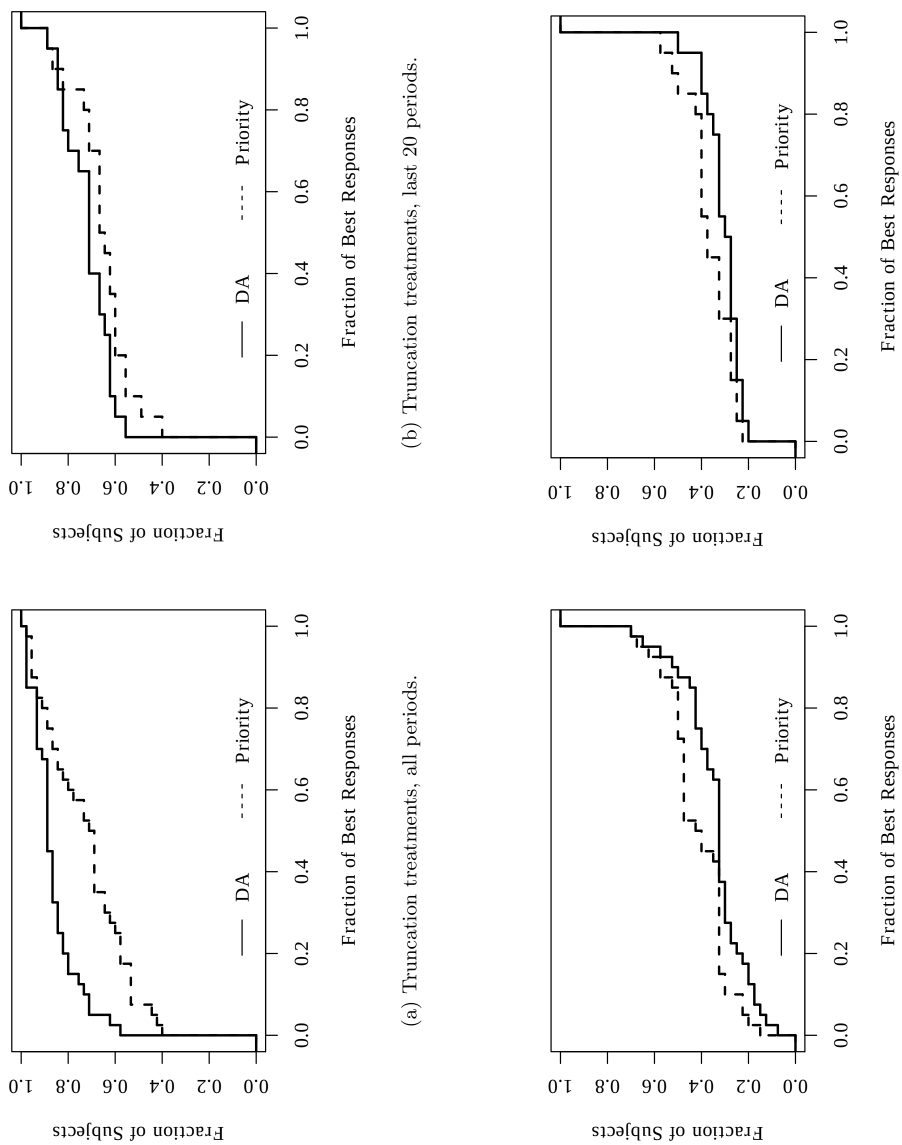

(a) Truncation treatments, all periods.

(b) Truncation treatments, last 20 periods.

(c) Truth-telling treatments, all periods.

(d) Truth-telling treatments, last 20 periods.

Figure 3: Best response frequency CDFs

player heterogeneity to analyze how unsophisticated players may suffer relative to players who learn to play strategically.

5.1. *Parameters and Model Dynamics*

To understand how subjects determine strategies under the different mechanisms and conditions, we estimate a reparametrized Experience-Weighted Attraction (EWA) learning model introduced by Camerer and Ho (1999). EWA is a flexible model incorporating elements of belief-based and choice reinforcement models.[24] We estimate a reparametrized EWA that separately identifies initial cognition and interactive learning.

In the original EWA, the key objects in the model are attractions to strategies. Each agent i begins the game with an initial attraction A to each strategy j, denoted $A_i^j(0)$, derived from pre-game analysis or prior experience. Let s_i^j represent strategy j for agent i, and $s(t)$ represent the set of strategies played in period t. Additionally, define $\pi_i(s_i^j(t), s_{-i}(t))$ as the round t payoffs for player i, which depend on player i's strategy $(s_i^j(t))$ and all other players' strategies $(s_{-i}(t))$.

After each round of play, each agent updates the previous round's attractions using a weighted combination of their prior attraction and the payoff from playing the strategy, according to the recursive formula:

$$A_i^j(t) = \varphi \cdot A_i^j(t-1) + \left[\delta + (1-\delta) \cdot \mathbb{1}_{\{s_i^j\}}(s_i(t))\right] \cdot \pi_i(s_i^j, s_{-i}(t)). \tag{1}$$

The parameter φ represents a discount factor, and determines how quickly previous attractions decay; the parameter δ is an introspection factor, dictating how much the new attractions depend on realized payoffs from the previous round relative to counterfactual payoffs from unplayed strategies. The indicator function $\mathbb{1}_{\{s_i^j\}(s_i(t))}$ is equal to one if agent i played strategy j in round t, and zero otherwise. Thus, the payoff weight associated with realized strategies is one, while the payoff weight of counterfactual strategies is δ.

In this model, attractions map to probabilities of play in each round according to a power form:[25]

$$P_i^j(t+1) = \frac{\left(A_i^j(t)\right)^\lambda}{\sum\limits_{k=1} \left(A_i^k(t)\right)^\lambda}. \tag{2}$$

[24] EWA nests belief-based models, where players form expectations about other players' strategies and choose a best response, and choice reinforcement models, in which past payoffs reinforce successful strategies.

[25] In addition to the power probability form, Camerer and Ho (1999) also describe a logit probability form, which introduces additional parameters. In the interest of parsimony, we use the power probability form for our estimation.

In this equation, the "exploitation factor" λ determines how often a player chooses her more attractive strategies, relative to the probability of exploring less attractive strategies. This dictates the amount of randomness in a player's sequence of strategies: when $\lambda = 0$, the player plays all strategies with equal probability, and as λ increases, the probability of playing the most attractive strategy increases.[26]

Thus, learning dynamics are determined by initial attractions, the weight of previous attractions relative to updating from recent payoffs, and the relative weight of actual and counterfactual payoffs. However, this parametrization fails to distinguish fully between initial cognition and interactive learning.

We define initial cognition to be the process of thinking through a game before any opportunity to learn by actually playing it. This process may depend on mental analyses, the description of game play, any hints and suggestions provided to the players, and beliefs about what other players may do. We distinguish initial cognition from interactive learning, which describes the dynamics of learning over the course of the game based on feedback (in the form of payoffs and information about allocations). The culmination of initial cognition is the set of play probabilities for each possible action j, for each individual or type i, in the first round of play, $\{P_i^j(1)\}_{i,j}$. Although these probabilities are encoded by the initial attractions, $\{A_i^j(0)\}_{i,j}$, and the exploitation factor, λ, the mapping from these parameters to the initial play probabilities is not one-to-one, since the power-form probability function is invariant to multiplying all initial attractions by a common factor.

To pin down the mapping between probabilities and attractions, we define a new parameter that captures the additional information codified in the initial attractions. Let $\|A_i(0)\|$ denote the λ-norm of the vector of initial attractions. That is,

$$\|A_i(0)\| \equiv \left(\sum_{k=1}^{m_i} \left(A_i^k(0) \right)^{\lambda} \right)^{1/\lambda} .$$

This parameter is a measure of the size the initial attractions. In this reparametrization, the initial attractions are no longer free parameters; instead, they are determined by

$$A_i^j(0) = \|A_i(0)\| \cdot \left(P_i^j(1) \right)^{1/\lambda} .$$

The free parameters of this reparametrized learning model are now $\{P_i^j(1)\}_{i,j}$, $\|A_i(0)\|$, λ, φ, and δ. This fully specifies the model, and separates the initial cognition from interactive learning. For simple interpretations of these parameters,

[26]Camerer and Ho (1999) refer to λ as the "exploration" factor. We have changed the name to match the intuition behind the model: it is more likely that the player "exploits" the most attractive strategies (rather than "exploring" new strategies) as λ increases.

we must first discuss the intuition behind the interactive learning component of the model.

Essentially, each attraction is the net present value of the stream of payoffs associated with a strategy. The parameter φ represents the discount rate, while the parameter δ represents how much counterfactual payoffs are weighted relative to realized payoffs. Agents choose an action randomly according to the power-form probability function discussed above and the exploitation factor λ. All of this is sensible, but we have yet to discuss where these discounted sums should start in the first round of play. The initial play probabilities constrain these initial attractions, but don't completely pin them down.

This is the role of $\|A_i(0)\|$. Intuitively, it is the natural way to sum up all of the payoff streams that have been aggregated across the different actions.[27] It tells us how initial cognition will be weighted relative to interactive learning in terms of payoffs from the game. In other words, if the average payoff in a game is \$1 per round, then (very roughly), $\|A_i(0)\|$ tells us how initial cognition is weighted in terms of discounted rounds of interactive play.

5.2. *Estimation*

In order to estimate the parameters of the model through maximum likelihood estimation (MLE), we first need to simplify the parameter space. Many previous papers estimating the EWA model have done so in games with a small strategy space (see, for instance, the sample applications in Camerer and Ho (1999)). With more available strategies, it becomes computationally challenging to estimate the initial attraction to each strategy.[28] In our setting, players select among 325 strategies in each round, creating an intractable estimation problem.[29]

[27]Mathematically, the λ-norm is the NPV required to yield the same probability weight while concentrating the NPVs from all the actions into just one. As such, it is, in some sense, the norm that weights entries in a way that corresponds to probability of play. For instance, note that as the exploitation factor λ grows large, $\|A_i(0)\|$ approaches $\max_j A_i^j(0)$, which makes sense as the maximum attraction is the only one that matters as $\lambda \to \infty$.

[28]In more complex strategy spaces, some papers estimating an EWA model use other approaches to reduce the number of strategies estimated. For example, in the beauty contest game with 101 strategies, authors often lump groups of 10 strategies together such that one initial attraction is estimated for all integers in a range, such as 1–10 (Camerer and Ho, 1999). In our setting, the strategies do not have a clear ordering (as in a 0–100 beauty contest) that would make strategy lumping straightforward. Another popular approach is to assign initial attractions to all pure strategies, rather than estimating them at all (Chen and Khoroshilov, 2003; Cotla, 2015; Wu and Bayer, 2015; Ho, Camerer, and Chong, 2007). Ansari, Montoya, and Netzer (2012) defines the vector of initial attractions as the relative frequency of play in the first round.

[29]The strategy space for each player during each round of play includes any permutation of preferences over all 5 outcomes, and permutations of any set of truncated preferences (as long as at least one preference is listed, and the acceptable matches are listed before any unacceptable ones). The number of possible strategies in a round is $5! + 4 \times \binom{5}{1} + 3! \times \binom{5}{3} + 2! \times \binom{5}{2} + 1! \times \binom{5}{1} = 325$.

However, most (225 of 325) strategies are never played in any round of play, and only 20 strategies are played in the first round of any session. Moreoever, only 11 strategies are played more than once in an initial round, suggesting that initial probabilities of play are concentrated across a small number of strategies. Rather than estimate initial probabilities for each strategy, we estimate initial probabilities for these 11 strategies, and a single initial probability shared uniformly across all other strategies. This drastically reduces the parameter space, while maintaining flexibility to explain a wide range of observed behaviors.[30]

With this setup, we can now estimate 15 parameters for each treatment condition: 11 initial probabilities $P_i^j(1)$ describing the initial cognition process, three scalar parameters (φ, δ, and λ) to describe the learning process, and $\|A_i(0)\|$ identifying the relative weight of initial cognition and learning.[31]

5.3. *Learning Model Parameter Estimates*

Differences in both initial cognition and learning dynamics help explain subjects' failure to manipulate reported preferences under DA. In Table VIII, we present results of the structual estimation. Estimates of φ, λ, and δ characterize learning during play. For ease of interpretation, we present a transformation of λ, $2^{\frac{1}{\lambda}}$, which we describe in detail below. Initial cognition is described in the table with initial probabilities of play into three types of strategies: truth-telling, non-truthful truncation, and permutation strategies.[32]

Initial probabilities of truth telling are similar across the *DA Truth* (55.4%), *DA Truncation* (51.9%), and *Priority Truth* (50.8%) treatments, but much lower under *Priority Truncation* (27.4%). This suggests that before play begins, players in the *Priority Truncation* believe there are profitable deviations from truth-telling. The estimates for initial probabilities of playing non-truthful truncation strategies bear out this finding: subjects under *Priority Truncation* are much more likely to trunctate (46.7%) than under any other treatment. Under all treatment treatments, permutation strategy probabilities range between 21.8% and 28.9% and do not vary enough to drive differences in the truth-telling rate.

In addition, the weight of initial cognition $\|A_i(0)\|$ is higher under DA treatments, indicating that subjects rely more heavily on pre-game analysis when determining their strategies under DA. This reliance on analysis compounds the errors that subjects make in determining their initial probabilities of play in the *DA Truncation* treatment. Subjects under *DA Truncation* play as if they had about 30% more pre-game experience than their counterparts under *Priority Truncation*.

[30]The 11 estimated strategies include all truncation strategies and six permutation strategies that are not predicted by theory. For a list of all estimated probabilities, see Appendix Table XIII.

[31]Note that we only need to estimate 11 probabilities, since the sum of initial probabilities must be one. The probability of playing one of the non-estimated strategies is pinned down by the other estimates. For more details on model estimation, see Appendix C.

[32]Estimation results for individual strategies are shown in Appendix C, Table XIII.

TABLE VIII
Parameter Estimates by Treatment

Parameter	Interpretation	DA Truth	DA Trunc	Priority Truth	Priority Trunc
φ	Discount Factor	0.921 (0.012)	0.881 (0.010)	0.850 (0.017)	0.891 (0.009)
$2^{\frac{1}{\lambda}}$	Probability Doubling Factor	1.547 (0.045)	1.547 (0.050)	1.978 (0.102)	1.716 (0.057)
δ	Introspection Factor	0.004 (0.002)	0.005 (0.002)	0.000 (0.000)	0.002 (0.001)
$\|A(0)\|$	Payoff-Weight of Initial Cognition	64.880 (10.143)	112.319 (17.164)	35.123 (8.001)	85.887 (11.660)
$P^{Truth\text{-}telling}(1)$	Initial Probability of Truth-Telling	0.554 (0.045)	0.519 (0.036)	0.508 (0.053)	0.274 (0.032)
$P^{Truncation}(1)$	Initial Probability Truncation	0.143 (0.031)	0.255 (0.031)	0.288 (0.045)	0.467 (0.061)
$P^{Permutation}(1)$	Initial Probability of Permutation	0.289 (0.034)	0.219 (0.025)	0.218 (0.038)	0.244 (0.026)

Maximum likelihood estimates of reparametrized EWA model, estimated separately by treatment group. Standard errors shown in parenthesis.

As described above, three parameters in our model—φ, λ, and δ—determine the dynamics of the interactive learning process. We find that λ is significantly smaller under Priority than under DA, suggesting that Priority players are more inclined to explore new strategies, while under DA players prefer to exploit their most preferred strategy. This difference may explain why gaps in truncation rates persist after many rounds of play. In Table VIII, we transform λ to make its interpretation more explicit. The transformation $2^{\frac{1}{\lambda}}$ describes the attraction ratio that leads to doubling the probability ratio. Under DA, doubling the probability of play means the attraction to one strategy is about 1.5 times the attraction to another strategy. Under Priority, doubling the probability of play requires a much higher attraction ratio of 1.7-2.0.[33]

Differences between some treatments of the parameters φ and δ are also statistically significant, but the differences are economically less significant and unlikely to explain differences in truncation rates. The introspection factor δ describes how much players are able to learn from unplayed strategies. We find that δ is precisely estimated to be between 0.0002 and 0.005 for all treatments, suggesting that more than 99.5% of learning from any round is from realized payoffs rather than counterfactual learning.

Estimates for discount factor φ range between 0.850 for the *Priority Truth* treatment and 0.921 for the *DA Truth* treatment. The parameter φ dictates how the influence of previous attractions persist over time. To interpret these values,

[33]The ratio of probabilities $\frac{P^i}{P^j}$ of any two strategies i and j is given by $\left(\frac{A_i}{A_j}\right)^{\lambda}$.

we calculate the half-life of the attraction—the number of periods required to halve the influence of the attraction. The half-life of attractions is about 8.4 periods under *DA Truth*, 5.5 periods under *DA Truncation*, 4.3 periods under *Priority Truth*, and 6.0 under *Priority Truncation*.[34] These figures provide some insight into the learning process, but they do not explain the systematic differences in learning to manipulate reported preferences.

5.4. *Simulation*

Having estimated the structural learning model, we can simulate play under counterfactual conditions. We take advantage of this benefit to examine long-term dynamics of learning that are impossible to observe in the lab, and to separately identify the effects of initial cognition and interactive learning. We simulate play under the DA mechanism in the uncorrelated environment, where theory predicts players will truncate their reported preferences. In this setting, we simulate 500 rounds of repeated play, where agents learn according to one of the following learning processes:

1. All *Priority Truncation* treatment estimates
2. *Priority Truncation* initial probabilities, *DA Truncation* scalars, *Priority Truncation* weight
3. *Priority Truncation* initial probabilities, *DA Truncation* weight, *Priority Truncation* scalars
4. *DA Truncation* initial probabilities, *Priority Truncation* weight, *Priority Truncation* scalars
5. All DA Truncation treatment estimates

In Figures 18, 5, and 6, each line represents the average of 500 simulated groups of players under each learning process. Figures 18 and 5 show simulated truth-telling and truncation rates, respectively. Truth-telling rates start off much higher under *DA Truncation* initial probabilities, and slowly decline as players engage with the mechanism over 500 rounds. Even after 500 rounds, however, many players fail to learn successful manipulation strategies, and truth-telling remains higher among players with *DA Truncation* initial probabilities than among the other groups. In Figure 5,

5.5. *Heterogeneity in Learning*

One major assumption in the structural model estimated above is that all players learn according to the same process. In this section, we relax this assumption to explore the role of heterogeneity in the learning process. Indeed, strategyproofness is seen as desirable in part because it protects unsophisticated agents.

[34]The half-life is given by $t_{\frac{1}{2}} = -\frac{\log(2)}{\log(\varphi)}$.

Figure 4: Simulated Truth-Telling Under DA with Hypothetical Learning Parameters

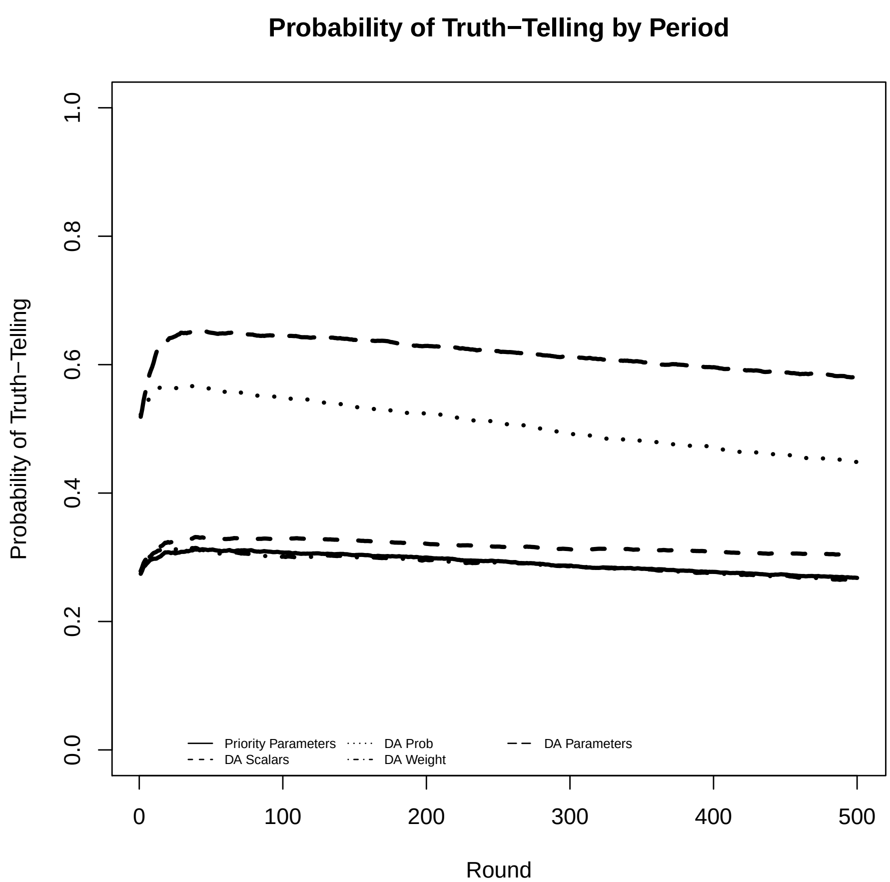

Figure 5: Simulated Truncation Under DA with Hypothetical Learning Parameters

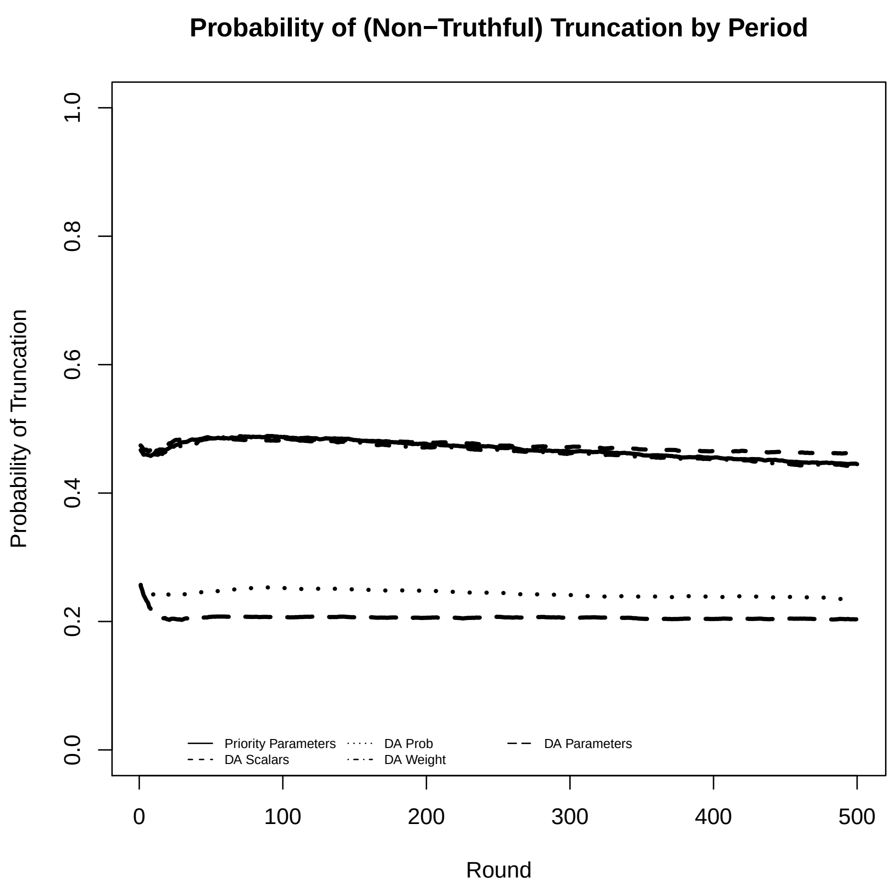

Figure 6: Simulated Blocking Pairs Under DA with Hypothetical Learning Parameters

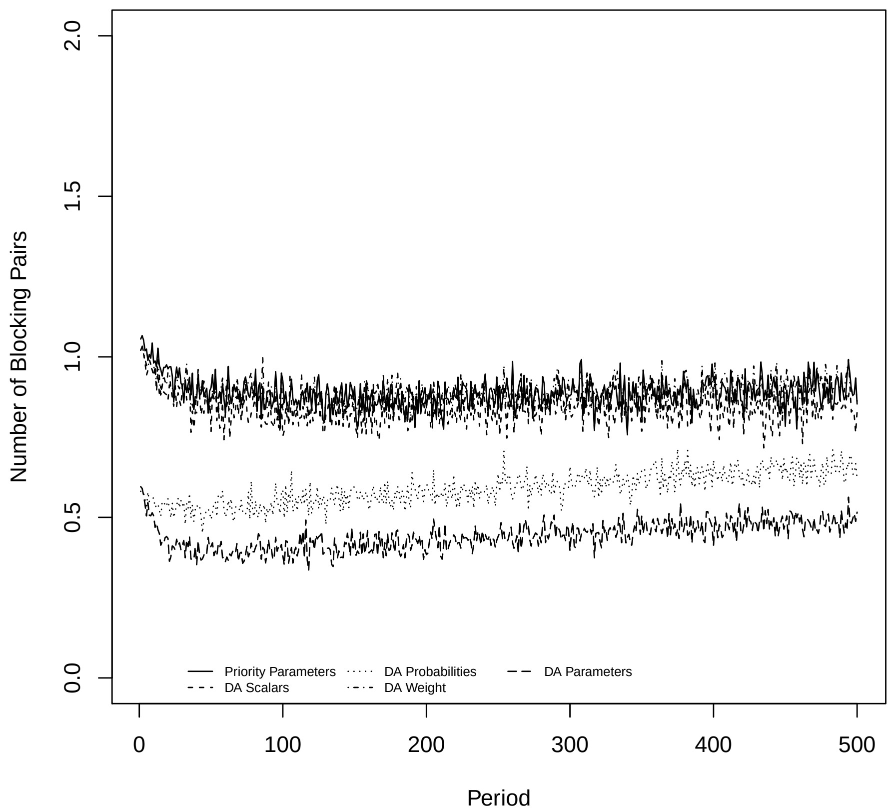

To address this issue, we propose a heterogeneous learning model with two types of agents: Learners and Naifs. Learners choose a strategy in each round according to the EWA learning process described above. They have the opportunity to explore strategies and adjust their play according to the payoffs. Naifs, on the other hand, simply report their true preferences in every (or almost every) round. This model is appealing in its parsimony — it adds only one parameter to the estimation, and allows the resulting estimates to be compared directly. The model described in 5.1 is equivalent to the heterogeneous model, but with the share of Learners constrained to be 100%.

This model is only useful for describing behavior in the uncorrelated market. In the correlated market environment we cannot distinguish between truth-telling due to sophistication and that due to naivete. We estimate the learning model separately by treatment, using the subset of players who play non-truthful strategies at least twice in the forty rounds of play. Players who report truthfully 39 or 40 times are considered naive truth-tellers, and comprise 26.7% of players in the *DA Truncation* treatment and 6.7% of players in the *Priority Truncation* treatment.

Results are shown in Table IX. Table X compares predictions of the homogeneous and heterogeneous models against the data observed in our experiment. To generate the model predictions, we simulate as described above in Section 5.4. In the heterogeneous model, simulated players are randomly assigned a type as a Naif or as a Learner: Naifs report truthfully in every period, while Learners update their probabilities of play according to the EWA model. We find that introducing heterogeneity does not significantly improve model fit.

6. CONCLUSION

Participants in matching markets might not truncate under DA, even when doing so would be significantly profitable. We show this in a simple experimental environment where participants were trained on the mechanism, given ample opportunity to learn through feedback, and were not subject to any randomness that might come from non-straightforward play on the proposing side. Even in this simple setting, players use very little counterfactual analysis, and learning dynamics vary only in players' initial assessment of the game. In the field, where things are more complicated and information is more sparse, we have little reason to think that match participants would be more likely to learn to truncate. These results also suggest that the persistence of DA clearinghouses may rely on participant misoptimization, and that interventions designed to improve understanding could lead to unravelling.

In addition to understanding the persistence of DA in the field, we also think an experiment such as ours feeds into the broader concerns of market design. Whenever a matching mechanism is strategy-proof, it is straightforward for designers to predict agent behavior in the field, since both focality and optimality push towards truth-telling. Sometimes though, strategy-proofness is either not desired

TABLE IX
PARAMETER ESTIMATES BY TREATMENT FOR LEARNERS

Parameter	Interpretation	DA Trunc	Priority Trunc
φ	Discount Factor	0.877 (0.010)	0.891 (0.009)
$2^{\frac{1}{\lambda}}$	Probability Doubling Factor	1.593 (0.054)	1.726 (0.058)
δ	Introspection Factor	0.004 (0.002)	0.001 (0.001)
$\|A(0)\|$	Payoff-Weight of Initial Cognition	142.338 (22.852)	95.566 (13.122)
$P^{Truth\text{-}telling}(1)$	Initial Probability of Truth-Telling	0.394 (0.038)	0.244 (0.031)
$P^{Truncation}(1)$	Initial Probability Truncation	0.322 (0.036)	0.487 (0.061)
$P^{Permutation}(1)$	Initial Probability of Permutation	0.284 (0.029)	0.280 (0.026)

Maximum likelihood estimates of reparametrized EWA model for learners only, estimated separately by treatment group. Sample includes all learners, defined as experimental subjects who played non-truthful strategies at least twice in a session. Learners comprised 73.3% of subjects in the *DA Truncation* treatment and 93.3% of subjects in the *Priority Truncation* treatment. Standard errors shown in parenthesis.

TABLE X
MEASURES OF MODEL FIT

	DA Truncation			*Priority Truncation*		
	Data	Hom. Model	Het. Model	Data	Hom. Model	Het. Model
$P^{\{Truth\}}$ ...						
... in periods 1–10	0.622	0.566	0.600	0.362	0.272	0.289
... in periods 11–20	0.547	0.628	0.620	0.260	0.261	0.272
... in periods 21–30	0.549	0.645	0.627	0.198	0.250	0.262
... in periods 31–40	0.547	0.650	0.631	0.193	0.243	0.253
$P^{\{Truncation\}}$ ...						
... in periods 1–10	0.207	0.232	0.232	0.422	0.461	0.461
... in periods 11–20	0.300	0.206	0.228	0.569	0.468	0.480
... in periods 21–30	0.318	0.203	0.227	0.624	0.477	0.491
... in periods 31–40	0.340	0.204	0.230	0.656	0.478	0.499
Share best responding ...						
... in periods 1–10	0.711	0.737	0.732	0.709	0.660	0.669
... in periods 11–20	0.762	0.740	0.740	0.687	0.671	0.663
... in periods 21–30	0.760	0.741	0.736	0.729	0.667	0.668
... in periods 31–40	0.762	0.745	0.748	0.676	0.669	0.671
Number of blocking pairs ...						
... in periods 1–10	0.071	0.538	0.529	0.349	2.074	2.058
... in periods 11–20	0.104	0.443	0.471	0.382	2.046	2.076
... in periods 21–30	0.118	0.404	0.440	0.360	2.059	2.082
... in periods 31–40	0.100	0.395	0.439	0.404	2.041	2.055

Comparison of statistics in experimental data, simulations from a homogeneous model of learning, and simulations from a heterogeneous model of learning. In the heterogeneous model, a fraction of players play truthfully in all periods; the remaining players learn according to the reparametrized EWA model described in Section 5.1, with learning parameters estimated from experimental data and shown in Table IX. Learners comprised 73.3% of subjects in the *DA Truncation* treatment and 93.3% of subjects in the *Priority Truncation* treatment.

or cannot be achieved due to other design goals. Consider the job of a market designer who has been tasked with creating a two-sided matching mechanism that persists. We can view the current paper as an experiment that would help inform our theoretical designer. Persistence can be intuitively linked to ex post stability, so DA is a natural candidate. Unfortunately, under DA, truth-telling is generally not an equilibrium. Theory provides a set of strategies which could outperform truthful preference revelation: the question is then whether our designer should expect market participants to use these deviations from truth-telling, which is a clear candidate for a focal strategy. If agents use these profitable deviations from truth-telling, then DA will not yield an ex post stable outcome, but if they don't, then it will. To determine which is the more likely outcome, the present lab experiment becomes very informative.

In demonstrating that agents learn to play some deviations from truth-telling, but not others, we bring up the idea that not all equilibria are equal in their predictive power. Depending on the mechanism and environment, agents are sometimes very close to equilibrium play and sometimes not. Some intuitive factors that seem like they should be important for whether a theoretical equilibrium will be realized in the field are focality of truth-telling, obviousness that some deviation from truth-telling will be profitable, difficulty of finding the optimal such deviation, and the profitability of that deviation. Unfortunately, although these factors may guide us intuitively, there is no formal theory for how they might trade off in determining the accuracy of an equilibrium prediction; in fact, most of them are difficult even to define. This is where lab experiments can prove most useful for design. The current paper, for instance, implies that truth-telling is more strongly focal for the receiving side under DA than under Priority. It also shows that under both mechanisms, equilibrium predictions might not hold: under DA, participants truth-tell when they shouldn't, while under Priority, they deviate from truth-telling, but in a sub-optimal way. In short, although the main contribution of this experiment is to show how out-of-equilibrium truth-telling could lead to ex post stability of DA in the field, we also feel that the experiment is the sort of inquiry that should be used in practical market design.

REFERENCES

Abdulkadiroğlu, A., P. Pathak, and A. Roth (2005): “The New York City High School Match,” *American Economic Review*, pp. 364–367.

Abdulkadiroglu, A., P. Pathak, and A. Roth (2009): “Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match,” *The American Economic Review*, 99(5), 1954–1978.

Abeler, J., D. Nosenzo, and C. Raymond (2019): “Preferences for Truth-Telling,” *Econometrica*, 87(4), 1115–1153.

Ansari, A., R. Montoya, and O. Netzer (2012): “Dynamic learning in behavioral games: A hidden Markov mixture of experts approach,” *Quantitative Marketing and Economics*, 10.

APPIC (2009): “2009 APPIC Match Statistics,” http://www.appic.org/match/5_2_2_1_11_match_about_statistics_general_2009.html.

Calsamiglia, C., G. Haeringer, and F. Klijn (2009): “Constrained school choice: An experimental study,” *American Economic Review*.

Camerer, C., and T.-H. Ho (1999): “Experience-Weighted Attraction Learning in Normal Form Games,” *Econometrica*, 67(4), 827–874.

Chen, Y., and Y. Khoroshilov (2003): “Learning under limited information,” *Games and Economic Behavior*, 44(1), 1–25.

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

Coles, P., A. Kushnir, and M. Niederle (2013): “Preference Signaling in Matching Markets,” *American Economic Journal: Microeconomics*, 5(2), 99–134.

Coles, P., and R. Shorrer (2014): “Optimal truncation in matching markets,” *Games and Economic Behavior*, 87, 591–615.

Cotla, C. R. (2015): “Learning in Repeated Public Goods Games — A Meta Analysis,” *Unpublished Manuscript*.

Damiano, E., H. Li, and W. Suen (2005): “Unravelling of dynamic sorting,” *Review of Economic Studies*, 72(4), 1057–1076.

Day, R., and P. Milgrom (2008): “Core-selecting package auctions,” *International Journal of Game Theory*, 36(3), 393–407.

Du, S., and Y. Livne (2010): “Chaos and Unraveling in Matching Markets,” *Arxiv preprint arXiv:1009.0769*.

Dubins, L., and D. Freedman (1981): “Machiavelli and the Gale-Shapley algorithm,” *American Mathematical Monthly*, pp. 485–494.

Duflo, E. (2017): “Richard T. Ely Lecture: The Economist as Plumber,” *American Economic Review*, 107(5), 1–26.

Echenique, F., A. J. Wilson, and L. Yariv (2016): “Clearinghouses for two-sided matching: An experimental study,” *Quantitative Economics*, 7(2), 449–482.

Echenique, F., and L. Yariv (2010): “An experimental study of decentralized matching,” Discussion paper, Discussion paper, Working paper, Caltech.

Ehlers, L. (2008): “Truncation Strategies in Matching Markets,” *Mathematics of Operations Research*, 33(2), 327.

Ergin, H., and T. Sönmez (2006): “Games of school choice under the Boston mechanism,” *Journal of Public Economics*, 90(1-2), 215–237.

Fainmesser, I. (2013): “Social Networks and Unraveling in Labor Markets,” *Journal of Economic Theory*, 148(1), 64–103.

Featherstone, C. R., and M. Niederle (2016): “Boston versus deferred acceptance in an interim setting: An experimental investigation,” *Games and Economic Behavior*, 100, 353 – 375.

Gale, D., and L. Shapley (1962): “College admissions and the stability of marriage,” *American Mathematical Monthly*, pp. 9–15.

Hałaburda, H. (2010): “Unravelling in two-sided matching markets and similarity of preferences,” *Games and Economic Behavior*, 69(2), 365–393.

Hansen, N. (2016): “The CMA Evolution Strategy: A Tutorial,” *CoRR*, abs/1604.00772.

Harrison, G., and K. McCabe (1996): "Stability and preference distortion in resource matching: an experimental study of the marriage problem," *Research in Experimental Economics*, 6, 53–129.

Haruvy, E., and M. U. Unver (2007): "Equilibrium selection and the role of information in repeated matching markets," *Economics Letters*, 94(2), 284–289.

Ho, T. H., C. F. Camerer, and J.-K. Chong (2007): "Self-tuning experience weighted attraction learning in games," *Journal of Economic Theory*, 133(1), 177–198.

Immorlica, N., and M. Mahdian (2005): "Marriage, Honesty, and Stability," in *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, p. 53. Society for Industrial Mathematics.

Kadam, S. V. (2011): "Correlation of Preferences and Honesty in Matching Markets:," *working paper*.

Kagel, J., and A. Roth (2000): "The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment," *Quarterly Journal of Economics*, 115(1), 201–235.

Kojima, F., and P. Pathak (2009): "Incentives and stability in large two-sided matching markets," *The American Economic Review*, 99(3), 608–627.

Lee, R. S., and M. Schwarz (2007): "Signalling preferences in interviewing markets," in *Computational Social Systems and the Internet*, ed. by P. Cramton, R. Müller, E. Tardos, and M. Tennenholtz, no. 07271 in Dagstuhl Seminar Proceedings, Dagstuhl, German.

——— (2009): "Interviewing in two-sided matching markets," .

Li, H., and S. Rosen (1998): "Unraveling in matching markets," *American Economic Review*, 88(3), 371–387.

Li, H., and W. Suen (2000): "Risk sharing, sorting, and early contracting," *Journal of Political Economy*, 108(5), 1058–1091.

Li, S. (2017): "Obviously Strategy-Proof Mechanisms," *American Economic Review*, 107(11), 3257–87.

Mas-Colell, A., M. Whinston, and J. Green (1991): *Microeconomic theory*. Oxford University Press.

Nalbantian, H., and A. Schotter (1995): "Matching and efficiency in the baseball free-agent system: an experimental examination," *Journal of Labor Economics*, 13(1), 1–31.

Niederle, M., and A. Roth (2009): "Market Culture: How Rules Governing Exploding Offers Affect Market Performance," *American Economic Journal: Microeconomics*, 1(2), 199–219.

Niederle, M., A. Roth, and M. Ünver (2009): "Unraveling Results from Comparable Demand and Supply: An Experimental Investigation," *NBER Working Paper*.

NRMP (2009): "Results and Data: 2009 Main Residency Match," `http://www.nrmp.org/data/resultsanddata2009.pdf`.

NYC-DOE (2009): "Statistical summaires: register by grade," `http://schools.nyc.gov/AboutUs/data/stats/Register/CurrentRegisterByGrade/default.htm`.

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

Pathak, P., and T. Sönmez (2008): "Leveling the playing field: Sincere and strategic players in the Boston mechanism," *American Economic Review*, 98, 1636–1652.

Rees-Jones, A. (2017): "Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match," *Games and Economic Behavior*.

Roth, A. (1982): "The economics of matching: Stability and incentives," *Mathematics of Operations Research*, pp. 617–628.

——— (1984a): "The evolution of the labor market for medical interns and residents: a case study in game theory," *The Journal of Political Economy*, 92(6), 991–1016.

——— (1984b): "Misrepresentation and stability in the marriage problem," *Journal of Economic Theory*, 34(2), 383–387.

——— (1989): "Two-sided matching with incomplete information about others' preferences," *Games and Economic Behavior*, 1(2), 191–209.

——— (1991): "A natural experiment in the organization of entry-level labor markets: Regional markets for new physicians and surgeons in the United Kingdom," *The American*

Economic Review, pp. 415–440.

——— (1996): "The NRMP as a Labor Market," *Journal of the American Medical Association*, 275, 1054–1056.

——— (2003): "The origins, history, and design of the resident match," *JAMA: The Journal of the American Medical Association*, 289(7), 909.

Roth, A., and E. Peranson (1999): "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *The American Economic Review*, 89(4), 748–780.

Roth, A., and U. Rothblum (1999): "Truncation strategies in matching markets-in search of advice for participants," *Econometrica*, 67(1), 21–43.

Roth, A., and M. Sotomayor (1990): *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge University Press.

Roth, A., and X. Xing (1994): "Jumping the gun: imperfections and institutions related to the timing of market transactions," *The American Economic Review*, 84(4), 992–1044.

——— (1997): "Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists," *Journal of Political Economy*, pp. 284–329.

Sondak, H., and M. Bazerman (1991): "Power balance and the rationality of outcomes in matching markets," *Organizational Behavior and Human Decision Processes*, 50(1), 1–23.

Sönmez, T. (1999): "Can Pre-Arranged Matches be Avoided in Two-Sided Matching Markets?," *Journal of Economic Theory*, 86(1), 148–156.

Suen, W. (2000): "A competitive theory of equilibrium and disequilibrium unravelling in two-sided matching," *The Rand journal of economics*, 31(1), 101–120.

Ünver, M. U. (2001): "Backward unraveling over time: The evolution of strategic behavior in the entry level British medical labor markets," *Journal of Economic dynamics and control*, 25(6-7), 1039–1080.

——— (2005): "On the survival of some unstable two-sided matching mechanisms," *International Journal of Game Theory*, 33(2), 239–254.

Wu, H., and R.-C. Bayer (2015): "Learning from inferred foregone payoffs," *Journal of Economic Dynamics and Control*, 51(C), 445–458.

Zhang, L., and D. Levin (2017): "Bounded Rationality and Robust Mechanism Design: An Axiomatic Approach," *American Economic Review*, 107(5), 235–39.

APPENDIX A: MODEL AND DEFINITIONS

A ***marriage market under incomplete information*** is a quadruple $(M, W, \mathcal{P}, \lambda)$, where M and W are sets of agents on the two sides of the market, $\mathcal{P}$ is the set of all possible preference profiles for the agents, and λ is a measure over $\mathcal{P}$. We require that all agents in the market find at least one match partner acceptable. An element of $\mathcal{P}$ is a vector $(P_i)_{i\in M\cup W}$ of individual ***preference profiles***. P_m for some $m \in M$ is an ordering over $W \cup \{\emptyset\}$, where $\emptyset$ represents the outcome of being unmatched; P_w for some $w \in W$ is defined similarly. Hence, we can think of some W's preference ordering as an $(|M| + 1)$-vector whose elements are $\emptyset$ and the members of M.[35] A ***matching*** is a function $\mu : M \cup W \mapsto M \cup W \cup \{\emptyset\}$ such that for any $m \in M$ and $w \in W$, we have $\mu(m) \in W \cup \{\emptyset\}$, $\mu(w) \in M \cup \{\emptyset\}$, and $\mu(m) = w \Leftrightarrow \mu(w) = m$. A ***strategy*** for an agent i is a function $\sigma_i : \mathcal{P}_i \mapsto \mathcal{P}_i$ where $\mathcal{P}_i$ denotes the projection of $\mathcal{P}$ onto only agent i's preference profile.

Next, we define an important concept, introduced in Roth and Rothblum (1999), which we use to analyze the information structure of the matching markets used in our experiment. Let the $m \leftrightarrow m'$ operation switch the places of m and m' in the preference of each W and assigns the preferences of m to m' (and vice versa). Let $w \leftrightarrow w'$ be defined analogously. The following definition codifies the idea of a low information environment.

Definition 1 For some $w \in W$, a marriage market (or a distribution over M and W preferences) is ***M-symmetric with respect to w*** if and only if, for any two $m, m' \in M$, $\lambda(P_{-w}|P_w) = \lambda(P_{-w}^{m\leftrightarrow m'}|P_w)$. If this holds for all $w \in W$, we simply call the market ***M-symmetric***. ***W-symmetry*** is analogously defined. If a marriage market is both W-symmetric and M-symmetric, then we call it ***MW-symmetric***.

In such symmetric environments, we want to be able to rule out equilibria where strategies depend on label, as these seem artificial. Formally,

Definition 2 A strategy σ_i is ***anonymous*** if and only if, for any two preferences, P_i and P_i', that list the same number of acceptable match partners, there exists some permutation π such that $\sigma_i(P_i) = \pi(P_i)$ and $\sigma_i(P_i') = \pi(P_i')$.

Note that this definition allows for different permutations to be used when a different number of match partners are acceptable. Of the set of anonymous strategies, in the low information environments we look at in the lab, we will find that we expect a certain type of strategy in equilibrium.

Definition 3 A ***truncation*** is an anonymous strategy where the permutation for a given number of acceptable match partners, $k - 1$, is a composition of permutations that first exchanges the k^{th} position (i.e. $\emptyset$) with the j^{th} position, where $j \leq k$, and then permutes all positions besides k and j in a way that if a position started ranked (above j/between j and k/below k), its permuted position is ranked (above j/between j and k/below k).[36]

Finally, we introduce a technical condition needed for uniqueness (but not existence) of the types of equilibria we will be looking for.

Definition 4 A distribution over preferences is called ***W-thick*** if, for any $w \in W$, $m, m' \in M$, and $m'' \in M \setminus \{m, m'\}$ there is a positive probability that m and m' rank w first, while m'' ranks $w'' \neq w$ first and w'' ranks m'' first. ***M-thick*** is defined analogously. A distribution over preferences is called ***MW-thick*** if it is both M and W thick.

[35]In our context, thinking of preferences as vectors introduces a bit of redundancy since the mechanisms we consider are all individually rational; for example, $(m_1, m_2, m_3, \emptyset, m_4, m_5)$ and $(m_1, m_2, m_3, \emptyset, m_5, m_4)$ are functionally equivalent.

[36]Note that under this definition, a truthful strategy is a truncations.

Thickness is a sufficient condition that prevents an agent from ruling out the possibility that any two potential match partners are her only two stable match partners. Weaker conditions are possible, but thickness itself is quite weak: for instance, it is met when all possible profiles of first choices are drawn with positive probability.

APPENDIX B: PROOFS

LEMMA 1 *Under M-Proposing DA, truth-telling is the only weakly undominated strategy for all $m \in M$.*

PROOF: Assume than the strategy of some $m \in M$ submits a preference $\widetilde{P_m}$ that is not the true preference, P_m. Dubins and Freedman (1981) show that truth-telling cannot yield a worse outcome than a lie. Let k be the first position in the submitted rank-order list that $\widetilde{P_m}$ differs from the true preference, P_m. Let $w = P_m(k)$ and $w' = \widetilde{P_m}(k)$. If all Ws except for w and w' rank m as uncceptable, and w and w' only rank m as acceptable, then m gets w if he submits P_m and w' (which he likes less) if he submits $\widetilde{P_m}$. Hence, we have shown that truth-telling is never worse than a lie and is strictly better given some profile of strategies for the other agents. *Q.E.D.*

LEMMA 2 (Roth, 1989) *Under M-Proposing DA, it is weakly dominated for any $w \in W$ to not list her true first choice first.*

LEMMA 3 *In a marriage market that is M-symmetric with respect to w, if all agents besides w play anonymous strategies, and all $m \in M$ play the same strategy, then the distribution over submitted preferences, $\widetilde{\lambda}(\cdot)$, is also M-symmetric with respect to w.*

PROOF: To prove this, we show why the following equation must hold:

$$\begin{aligned}\widetilde{\lambda}(\sigma_{-w}(P_{-w})|P_w) &= \lambda(P_{-w}|P_w) = \lambda(P_{-w}^{m\leftrightarrow m'}|P_w)\\ &= \widetilde{\lambda}(\sigma_{-w}(P_{-w}^{m\leftrightarrow m'})|P_w) = \widetilde{\lambda}((\sigma_{-w}(P_{-w}))^{m\leftrightarrow m'}|P_w)\end{aligned}$$

The first equality comes from the definition of $\widetilde{\lambda}$, the second from the fact that the true preferences are M-symmetric with respect to w, and the third, again from the definition of $\widetilde{\lambda}$. For the last equality, we must note two things. First, since the $m \leftrightarrow m'$ does not change the rank of $\emptyset$ for the W's, the σ_{-w} operator applies the same permutation to $P_{w'}^{m\leftrightarrow m'}$ as it does to $P_{w'}$. Second, since the M's are all playing the same anonymous strategy, it makes no difference whether we switch the preferences of m and m' before we apply the σ_{-w} operator or after. Hence, σ_{-w} commutes with $m \leftrightarrow m'$.[37] *Q.E.D.*

PROPOSITION 6 *In an M-symmetric marriage market, under M-Proposing DA, there exists an equilibrium in anonymous, weakly undominated strategies that involves truth-telling for each $m \in M$ and truncation for each $w \in W$. Furthermore, if the market is also W-thick, all equilibria in anonymous, weakly undominated strategies are like this.*

PROOF: By Lemma 1, any equilibrium in weakly undominated strategies involves truth-telling by all M's. By Lemma 3, we then know that, at an equilibrium in weakly undominated, anonymous strategies, the distribution of reported preferences, $\widetilde{\lambda}$, is M-symmetric. Then, by the main proposition of Roth and Rothblum (1999), we know that truncation is a best response for all $w \in W$. Furthermore, by Lemma 2, every W must be truthfully ranking her first choice M. Then, by the W-thickness assumption, it is with positive probability that for any $m, m' \in M$, w can only potentially match to m or m'. In these states of the world, we are in Case D of the proof from Roth and Rothblum (1999), which means that truncation strictly dominates non-truncation. *Q.E.D.*

[37]Note that we are not claiming that permutations commute: our interchange operator references school names and not positions in a rank-order list.

Since the uncorrelated market is M-symmetric and W-thick, **Proposition 1 in the main text** is an immediate corollary.

LEMMA 4 *Under M-Proposing Priority, it is weakly dominated for any $w \in W$ to not truthfully rank her first choice M.*

PROOF: In the first round w gets proposals, she will be permanently matched. Ranking her first choice, $m \in M$ first can not hurt her, but failing to do so can hurt her if she also receives a proposal in that round from an $m' \in M$ that she ranked higher than m, but actually likes less. Let m and her declared first choice, $\widetilde{P_w}(1)$, both rank w first, and let all other $m'' \in M$ declare w unacceptable. Ranking m first instead of $\widetilde{P_w}(1)$ is an improvement. *Q.E.D.*

PROPOSITION 7 *In an M-symmetric marriage market, under M-Proposing Priority, if all agents play anonymous, weakly undominated strategies, and in addition, all $m \in M$ truth-tell, then all $w \in W$ can best-respond to the other agents by truncating. If the market is also W-thick, then all of their best responses are truncations.*

PROOF: By Lemma 3, we know that the distribution of reported preferences, $\widetilde{\lambda}$, is M-symmetric with respect to w. Then, by Proposition 3.2 and Remark 3.2 of Ehlers (2008), we know that truncation is a best response for all w. Furthermore, by Lemma 4, every W must be truthfully ranking her first choice M. Then, by the W-thickness assumption, it is true with positive probability that for any $m, m' \in M$, w can only potentially match to m or m'; hence, Equation A2 from Ehlers (2008) must hold strictly, which means that truncation strictly dominates non-truncation. *Q.E.D.*

Since the uncorrelated market is M-symmetric and W-thick, **Proposition 2 in the main text** is an immediate corollary.

LEMMA 5 *Under M-Proposing Priority, any report for any $m \in M$ that does not list all and only all truly acceptable $w \in W$ as acceptable is weakly dominated by one that does.*

PROOF: Consider an arbitrary $m \in M$ submitting a list L with n acceptable match partners which excludes at least one acceptable $w' \in W$. Now consider L', a list identical to L for the first n entries with w' listed in the $(n+1)^{\text{st}}$ position and no acceptable entries thereafter. Under M-Proposing Priority, any set of submissions for other agents resulting in m being matched to a given W when m submits L will also result in M being matched to that W when m submits L'. So L' never generates a worse outcome for m than L. However, consider a set of submissions such that no member of W listed in L ranks m as acceptable, and the submitted preference list of w' lists only m as acceptable. In this case, M-Proposing Priority will match m and w' when L' is submitted and will match m to no one when L is submitted. Since w' is acceptable to m by construction, m achieves a better result in this case by submitting L'.

Now consider some $m \in M$ who lists a truly unacceptable $w \in W$ as acceptable. Removing this w from his list cannot hurt m, since M-Proposing Priority makes permanent matches after each round. Now, let all $w' \in W \setminus w$ declare m unacceptable, let all $m' \in M \setminus m$ declare w unacceptable and let w declare m acceptable. With this strategy profile, m will match to w which he could have avoided by declaring her unacceptable. *Q.E.D.*

LEMMA 6 *Under M-Proposing Priority, if the distribution of reported preferences for all agents besides $m \in M$ are W-symmetric with respect to m, then truth-telling is a best-response for m.*

PROOF: This proof borrows heavily from Roth and Rothblum (1999). First, we lay out a few of the properties of M-Proposing Priority. Consider, P, $w', w \in W$, $m \in M$, and let

		Lie: MPP$[P_m^{w\leftrightarrow w'}, P_{-m}](m)$		
		$= v \notin \{w, w'\}$	$= w$	$= w'$
Truth: MPP$[P_m, P_{-m}](m)$	$= v \notin \{w, w'\}$	Case A	Case B	Impossible
	$= w$	Impossible	Case C	Impossible
	$= w'$	Case D	Case E	Case F

TABLE XI

Table of cases

$v \in (W \setminus \{w, w'\}) \cup \{\emptyset\}$. Denote the match of m when the submitted preferences are P under M-Proposing Priority as MPP $[P]$ (m). Then,

$$\begin{aligned} \text{MPP}[P](m) = v \quad &\Leftrightarrow \text{MPP}[P^{w\leftrightarrow w'}](m) = v \\ \text{MPP}[P](m) = w \quad &\Leftrightarrow \text{MPP}[P^{w\leftrightarrow w'}](m) = w' \end{aligned}$$

Moreover,

$$\begin{aligned} \text{MPP}[P_m^{w\leftrightarrow w'}, P_{-m}](m) = v \quad &\Leftrightarrow MPP[P_m, P_{-m}^{w\leftrightarrow w'}](m) = v \\ \text{MPP}[P_m^{w\leftrightarrow w'}, P_{-m}](m) = w \quad &\Leftrightarrow MPP[P_m, P_{-m}^{w\leftrightarrow w'}](m) = w' \end{aligned}$$

The first set of logical statements follows immediately from the fact that MPP does not give special treatment to any given label. The fact that applying the $w \leftrightarrow w'$ interchange operator to $(P_m^{w\leftrightarrow w'}, P_{-m})$ yields $(P_m, P_{-m}^{w\leftrightarrow w'})$, implies the second set.

Now, let $w \prec_m w'$. Then,

$$(\text{MPP}[P](m) = w) \Rightarrow (\text{MPP}[P_m^{w\leftrightarrow w'}, P_{-m}](m) = w)$$

Moreover,

$$\left(\text{MPP}[P_m^{w\leftrightarrow w'}, P_{-m}](m) = w'\right) \Rightarrow \left(\text{MPP}[P](m) = w'\right)$$

Switching w' and w in a submitted ordering means that w is proposed to in an earlier round. If it was available in the later round, it will still be available in the earlier round, and no one else will be proposing to it in that round. This yields the first logical statement. The second follows from a similar line of reasoning.

Now, consider the outcome for some $m \in M$ for whom $w \prec_m w'$ when he submits a preference that truthfully ranks w and w', MPP$[P](m)$, and when he submits a preference that switches w and w', MPP$[P_m^{w\leftrightarrow w'}, P_{-m}](m)$. Using the formulas we just derived, we summarize what can potentially happen in Table XI, while Table XII tells us what lottery over outcomes m can expect when he truthfully orders w and w' and when he switches their ordering, given that everyone else's preferences are either P_{-m} or $P_{-m}^{w\leftrightarrow w'}$ with equal probability.

Clearly, under every case, if we take symmetry into account, truthfully ordering w and w' either yields an outcome that is equivalent to the outcome achieved with the lie, or weakly stochastically dominates the outcome from the lie. Now, Lemma 5 shows us that an M cannot be hurt by listing all acceptable Ws, so we know that truth-telling is a best response for Ms to W-symmetry.

Note that if we can show that the probability of being in Cases B, D, or E is strictly positive, then we also show that truthfully ordering the Ws strictly stochastically dominates any lie, although we would need a further restriction to weakly undominated strategies to get truth-telling as a unique best response. *Q.E.D.*

Lemma 7 *In an M-thick, W-symmetric marriage market, under M-Proposing Priority, if all $w \in W$ are playing the same weakly undominated, anonymous strategy, and all $m' \in M \setminus \{m\}$ are playing anonymous strategies, then all best responses for $m \in M$ must truthfully rank his true first choice partner.*

	Truth		Lie	
	MPP$[P_m, P_{-m}]$	MPP$[P_m, P_{-m}^{w\leftrightarrow w'}]$	MPP$[P_m^{w\leftrightarrow w'}, P_{-m}]$	MPP$[P_m^{w\leftrightarrow w'}, P_{-m}^{w\leftrightarrow w'}]$
Case A	v	v	v	v
Case B	v	w'	w	v
Case C	w	w'	w	w'
Case D	w'	v	v	w
Case E	w'	w'	w	w
Case F	w'	w	w'	w

TABLE XII

Payoffs for the cases

Proof: By similar logic to Lemma 3, the submitted preferences are W-symmetric with respect to m. Consider the argument of Lemma 6 with regard to the true first choice and some other reported first choice. By the M-thickness assumption and Lemma 4, there is some probability that those two Ws rank m first, meaning that we are in Case E of Lemma 6, meaning that m does strictly better to truthfully rank his first choice. *Q.E.D.*

Lemma 8 *In an M-thick, W-symmetric marriage market, if each $w \in W$ plays the same anonymous, weakly undominated strategy, and each $m' \in M \setminus \{m\}$ truthfully reveals his first choice partner, then under M-Proposing Priority, the only best-response for m' is to truth-tell.*

Proof: By Lemma 2, weakly undominated means that all Ws must truthfully rank their first choice partner. Since, by an argument analogous to Lemma 3, reported preferences must be W-symmetric with respect to m, we conclude through Lemma 6 that m cannot do worse than to truthfully reveal. Further, by Lemma 7, m must also best respond by truthfully ranking his first choice partner at equilibrium. From here, the M-thickness assumption allows us to go the rest of the way in showing that, for any two W's, the probability of being in Case E of Lemma 6 is strictly positive, and that the only best response for m is to truthfully reveal. *Q.E.D.*

Formally, a ***symmetric equilibium*** is one in which any two Ms are playing the same strategy, and any two Ws are playing the same strategy.

Proposition 8 *In an MW-symmetric marriage market, under M-Proposing Priority, there exists a symmetric equilibrium in anonymous strategies that involves truth-telling by the Ms and truncation by the Ws. Furthermore, if the market is MW-thick, then all symmetric equilibria in anonymous, weakly undominated strategies are of this form.*

Proof: If every M is playing the same anonymous strategy, and every W is playing an anonymous strategy, then by Lemma 3,the reported preferences are M-symmetric, and by Ehlers (2008), all Ws can best-respond with a truncation.

Now, consider the problem of finding the best-response of some $w \in W$ to the symmetric M strategies, σ_M, and a profile σ_{-w} in which all members of $W \setminus \{w\}$ are playing the same mixed strategy over truncations. Call this best response $\sigma_w^* (\sigma_{-w}|\sigma_M)$. Solving for the best response is an optimization problem in which w must choose her mix over truncation levels for each possible number of acceptable Ms her preference could hold. The objective is linear in the mixing probabilities,[38] and the set of possible mixing probabilities is closed and convex. Hence, we know that the solution exists, it is convex, and by the Theorem of the Maximum (Mas-Colell, Whinston, and Green 1991, Theorem M.K.6), it is upper hemicontinuous. Hence,

[38]For a given pure strategy profile, w gets an expected payoff. Her expected payoff from a mixed strategy is just a probability-weighted sum of these expected payoffs from pure strategies.

by Kakutani's Fixed Point Theorem (Mas-Colell, Whinston, and Green 1991, Theorem M.I.2), $\sigma_w^*(\sigma_{-w}|\sigma_M)$ has a fixed point. Hence, for any symmetric σ_M, there is a symmetric σ_W where each W is best responding to the other players.

Now, in any such setup, the Ms will not necessarily be best-responding. Since the market is W-symmetric, we know that the reported preferences are W-symmetric, which means that, by Lemma 6, the Ms can best-respond by truth-telling. Hence, we have found a symmetric equilibrium of the sort we were looking for.

Now, if strategies are anonymous and weakly undominated, then M-thickness coupled with Lemmas 7 and 8 requires that all such symmetric equilibria involve Ms truth-telling. Similarly, W-thickness couples with Lemma 2 requires that all such symmetric equilibria involve Ws truncating. *Q.E.D.*

This proposition has an immediate corollary, which is referenced in **Footnote 17 of the main text**.

Corollary (to Proposition 8) *In the uncorrelated market, under M-Proposing Priority, there exists a symmetric equilibrium that involves all Ws playing the same truncation strategy and all Ms truth-telling. Furthermore, all symmetric equilibria in anonymous, weakly undominated strategies are of this form.Also, we can note that so long as the Ms use anonymous, weakly undominated strategies, the Ws still best-respond with truncation. So long as the M strategies don't key in on a label, the Ws view them strategically in the same way as they view truth-telling Ms.*

The big implication here is that if an M believes that the equilibrium played will be a symmetric truncation equilibrium, then truth-telling is the best response. This proposition extends work done in Roth and Rothblum (1999) and Ehlers (2008) to conditions that lead to truth-telling for the proposing side under a priority mechanism.[39] In a broader sense, though, it turns out not to matter whether the Ms truthfully reveal.

Proposition 9 *In a M-symmetric market, under M-Proposing Priority, for any $w \in W$, if for any distinct $m, m' \in M$, P_m and $P_{m'}$ are conditionally independent given P_w and for any $m' \in M$ and $w' \in W$, $P_{m'}$ and $P_{w'}$ are conditionally independent given P_w, and all agents play anonymous, weakly undominated strategies, then w can best-respond with a truncation. Furthermore, if the market is also W-thick, then any best response must be a truncation.*

Proof: Since the preferences of the Ms are all conditionally independent, it must be that for any given number of truly acceptable match partners, all lists with that number of acceptable partners are equally likely. By Lemma 5, the weakly undominated requirement means that the Ms must list all acceptable Ws. The anonymous requirement then means that these lists must be permutations. Running a uniform distribution through a permutation yields a uniform distribution. Hence, the reported preferences of the Ms must be uniformly distributed for each number of acceptable partners, meaning that the reported preferences of the Ms are independent of the strategies they use. Looking back to the proof of Lemma 3, the fact that the Ms' reported preferences are conditionally independent and uniform for each list length, and that M preferences are conditionally independent of W preferences means that we no longer need that all Ms play the same strategy to get the same result. This means, that through a proof very similar to that of Proposition 7, w must best-respond with a truncation. *Q.E.D.*

[39] Roth and Rothblum (1999) and Ehlers (2008) focus on incentives for the receiving side. These papers also assume that reported preferences are M-symmetric instead of assuming that the true preferences are M-symmetric and backing out sufficient conditions to ensure that the reported preferences inherit M-symmetry as well.

This proposition has an immediate corollary, which is references in **Footnote 17 of the main text**.

COROLLARY (to Proposition 9) *In the uncorrelated market, under M-Proposing Priority, if all agents play anonymous, weakly undominated strategies, then all Ws must best-respond with a truncation.*

PROPOSITION (**Proposition 3 in the main text**) *In the correlated market, under M-Proposing DA, the unique equilibrium in anonymous, weakly undominated strategies entails truth-telling by all agents.*

PROOF: Under M-Proposing DA, weakly undominated strategies require that the Ms truthfully reveal (Lemma 1). Hence, under the assumptions, a given W will receive all offers she is going to receive in one round of the algorithm. To see this, first note that the top-ranked W, w_1, will receive all offers in the first round of the algorithm. She will be matched to her declared favorite M, and since this is a declared top-top match, the algorithm will never break it up. In the next round, the second ranked W, w_2, will receive offers from all other Ms. She will accept her declared favorite M who proposes, and the algorithm will never break this match (since the only potential M that w_2 might defect to is matched to w_1, who he prefers, and w_1 was given her declared top M). And so on. So at some point in the algorithm, a W's preference is used to choose a favorite M from a set of Ms that higher ranked Ws have not yet taken. There is no gain to not truthfully revealing, as our member of W is facing a static decision problem. Since every W has a one-in-five chance of being the last ranked W by all Ms, there is always a positive loss to dropping. *Q.E.D.*

PROPOSITION (**Proposition 4 in the main text**) *In the correlated market, under M-Proposing Priority, if all members of M have the same anonymous, weakly undominated strategy, then all members of W best respond by truthfully revealing.*

PROOF: Under M-Proposing Priority, weakly undominated for the Ms means that all women are listed as acceptable (Lemma 5). Under the assumptions, a member of W will receive all offers in one round of the algorithm. There is no gain to not truthfully revealing then, as our member of W is facing a static decision problem. Since every W has a positive probability of being the last ranked W by all Ms, there is always a positive loss to dropping any M. *Q.E.D.*

PROPOSITION (**Proposition 5 in the main text**) *In the correlated environment, there exist cardinal payoffs that rationalize an equilibrium where all Ms and Ws truthfully reveal their preferences.*

PROOF: For each M, consider a payoff vector $\pi = (p_1, p_2, p_3, p_4, p_5)$ which is constructed as $p_5 = 1$, $p_4 = p_5 \cdot |M| + 1$, $p_3 = p_4 \cdot |M| + 1$, etc. In the correlated M-Proposing Priority environment, each M has a $1/|M|$ chance of being the first choice of any W. Thus, from the perspective of an M with payoffs described by π, even in the worst case when all other Ms also rank M's first choice as first, the M would still prefer the $1/|M|$ chance of getting its first choice than a certainty of getting its second choice. Similarly, an M failing to get its first choice would prefer the $1/|M$— chance of getting its second choice to a certainty of getting its third choice, and so on. Hence all Ms truthfully reveal, and by the previous Proposition, the Ws must as well. *Q.E.D.*

APPENDIX C: MODEL ESTIMATION

The reparameterized EWA model suggests the need to estimate parameters for the learning process δ, ϕ, and λ; an initial probability of play for each strategy (with initial probabilities either shared across individuals or estimated separately for each subject); and $\|A_0\|$ (again, either shared across individuals or estimated separately by subject or group), representing the weight of initial cognition in units of payoff amounts. In a Bayesian model, this initial cognition would be akin to pseudo-observations of play from previous rounds.

This suggests a parameter space of at least 329 dimensions, with still higher dimensionality if we allow probabilities of play and the weight of initial cognition to vary across individuals.[40] This is computationally intractable due to the large number of initial probabilities, even when we assume assume all initial probabilities are shared by all players in a treatment.

However, most (225 of 325) strategies are never played in any round of any treatment. Moreover, only 20 strategies are ever played in the first round of any session, and only 11 strategies are played more than once in any first round. This suggests that estimating initial probabilities for all 325 strategies is not only computationally infeasible, but also not necessary for us to understand the dynamics of play. Instead, for each of the four treatment groups, we estimate the initial probabilities of play for all strategies played more than once in any initial round, and a single joint attraction toward playing all other strategies. This reduces the search space to 15 dimensions (three learning parameters, 11 probabilities, and the initial cognition weight).[41]

Let us denote strategies by five digits, denoting the true preference ranks of the player's submitted preferences by the digits themselves, and the submitted preferences by the order the digits. Let the symbol ∅ represent a match listed as unacceptable in the submitted preference list. For instance, the strategy {12345} represents complete truth-telling, while {12354} represents a permutation strategy with the least preferred options listed in reverse order. A truncation strategy such as {123∅∅} consists of listing only the most preferred three preferences. These 11 strategies played more than once in an initial round include both complete truth-telling {12345} and the four possible truncation strategies: {123∅∅}, {12∅∅∅}, {1234∅}, {1∅∅∅∅}. The six remaining strategies are permutations, or combinations of permutation and truncation: {21345}, {213∅∅}, {13245}, {21435}, {12354}, {23145}. Thus, this estimation strategy allows us to measure differences in initial probability to truth telling and various truncation strategies, and to use these parameters to follow the trajectory of attractions over the course of the game. For the remaining strategies—those played either once or not at all in an initial round—we estimate a single initial attraction in each treatment. While limiting our estimation of individual attractions to repeated initial strategies requires a *post hoc* justification, we believe this is necessary to make the model tractable, and allows the use of this model in a much more complex space than usual. This approach allows us to capture the differences between truth-telling and truncation that we care about, while significantly simplifying the strategy space. Estimating a joint attraction for all unplayed strategies, also allows the model to scale with payoff values. Thus, this approach is flexible to applications with different payoffs.

C.1. *Technical Details*

To maximize over the rugged likelihood terrain, we implement the stochastic, derivative-free Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) optimizer. CMA-ES is designed to be robust to local optima, ridges, and discontinuities in ill-conditioned and non-separable problems (Hansen, 2016). We estimate standard errors using a numerical approximation of the Hessian, and transform to our reparametrized EWA via the delta method. We executed all maximum likelihood estimation in Java.

[40]The number of possible strategies in a round is $5! + 4 \times \binom{5}{1} + 3! \times \binom{5}{3} + 2! \times \binom{5}{2} + 1! \times \binom{5}{1} = 325$

[41]Note that we actually want to estimate 12 probabilities that are mutually exclusive and comprehensively exhaustive, and therefore must sum to 1. By estimating 11 of the probabilities directly, we get the 12^{th} for free).

For the treatment-level estimation, we directly estimate initial probabilities of 10 of the 11 strategies played initially more than once, and an additional probability shared among all other strategies. We estimate one strategy (truth-telling, $\{12345\}$) indirectly, by taking one minus the sum of the other probabilities. This decision was merely practical: our optimizer accepts simple boundaries (a minimum and a maximum) for each of the estimated parameters, so we run the risk at each iteration of the optimizer to have the sum of the directly-estimated probabilities sum to more than one. By leaving out the most commonly played strategy, we reduce the frequency of this event. When the sum of the randomly-drawn probability proposal points is greater than one, we instruct the log likelihood function to return an arbitrarily large negative value, encouraging the optimizer to seek elsewhere.

Table XIII: Parameter Estimates by Treatment

Parameter	Interpretation	DA Truth	DA Trunc	Priority Truth	Priority Trunc
φ	Discount Factor	0.9207	0.8808	0.8496	0.8913
		(0.0118)	(0.0097)	(0.0173)	(0.009)
λ	Exploitation Factor	1.5885	1.5886	1.016	1.2836
		(0.1069)	(0.1169)	(0.0764)	(0.0784)
δ	Introspection Factor	0.004	0.005	0.0002	0.0015
		(0.0016)	(0.002)	(0.0001)	(0.0006)
$\|A(0)\|$	Payoff-Weight of Initial Cognition	64.8797	112.3189	35.1234	85.8873
		(10.1432)	(17.1641)	(8.0008)	(11.6596)
$P^{\{12345\}}(1)$	Initial Probability of Truth-Telling	0.5541	0.5188	0.5077	0.2742
		(0.0451)	(0.0363)	(0.0529)	(0.0319)
$P^{\{1234\varnothing\}}(1)$	Initial Probability of One-Point Truncation	0.0338	0.0552	0.0357	0.0556
		(0.0154)	(0.0158)	(0.0169)	(0.0155)
$P^{\{123\varnothing\varnothing\}}(1)$	Initial Probability of Two-Point Truncation	0.0597	0.1276	0.1356	0.1905
		(0.0206)	(0.0233)	(0.0332)	(0.0273)
$P^{\{12\varnothing\varnothing\varnothing\}}(1)$	Initial Probability of Three-Point Truncation	0.0412	0.0572	0.0873	0.1868
		(0.0169)	(0.0164)	(0.0262)	(0.0267)
$P^{\{1\varnothing\varnothing\varnothing\varnothing\}}(1)$	Initial Probability of Four-Point Truncation	0.0088	0.0154	0.0298	0.034
		(0.0087)	(0.0085)	(0.0163)	(0.0128)
$P^{\{21345\}}(1)$	Initial Probability of {21345}	0.1077	0.077	0.0899	0.086
		(0.0265)	(0.0184)	(0.0274)	(0.0187)
$P^{\{213\varnothing\varnothing\}}(1)$	Initial Probability of {213∅∅}	0.0064	0.0158	0.0085	0.0098
		(0.0064)	(0.0085)	(0.0084)	(0.0061)
$P^{\{21435\}}(1)$	Initial Probability of {21435}	0.0136	0.017	0.0146	0.0138
		(0.0094)	(0.0086)	(0.0106)	(0.007)
$P^{\{12354\}}(1)$	Initial Probability of {12354}	0.0331	0.0124	0.046	0.0205
		(0.0149)	(0.0077)	(0.0199)	(0.0091)
$P^{\{23145\}}(1)$	Initial Probability of {23145}	0.0078	0.0018	0.0082	0.017
		(0.0072)	(0.0032)	(0.0082)	(0.0081)
$P^{\{13245\}}(1)$	Initial Probability of {13245}	0.0265	0.0326	0.0199	0.0339
		(0.0138)	(0.0121)	(0.0129)	(0.0118)
$P^{\{\text{other}\}}(1)$	Initial Probability of Other Strategies	0.0003	0.0002	0.0001	0.0002
		(0.0001)	(0.0001)	(0.0001)	(0.0001)

Maximum likelihood estimates of reparametrized EWA model, estimated separately by treatment group. Standard errors shown in parenthesis.