An Experimental Study of Decentralized Matching

Federico Echenique* Alejandro Robinson-Cortés† Leeat Yariv‡§

October 25, 2022

Abstract

We present an experimental study of decentralized matching markets, such as labor or marriage markets. Experimental participants are informed of everyone's preferences and can make arbitrary non-binding match offers that get finalized when a period of market inactivity has elapsed. Several insights emerge. First, stable outcomes are prevalent. Second, while centralized clearinghouses commonly aim at implementing extremal stable matchings, our decentralized markets most frequently culminate in the *median* stable matching. Third, preferences' cardinal representations impact the stable partners participants match with. The dynamics underlying our results exhibit successive blocking pairs, with agents accounting for the likelihood a match sticks.

Keywords: Decentralized Matching, Experiments, Market Design.

*Department of Economics, UC Berkeley, fede@econ.berkeley.edu

†Department of Economics, University of Exeter Business School, a.robinson-cortes@exeter.ac.uk

‡Department of Economics, Princeton University, lyariv@princeton.edu

§We thank Drew Fudenberg, Jacob Goeree, Matthew Jackson, Gabriel Katz, Erik Madsen, Sergio Montero, Muriel Niederle, Charles Plott, Andrew Schotter, and Alistair Wilson for useful conversations and suggestions. Financial support from the National Science Foundation (SES 0963583), and the Gordon and Betty Moore Foundation (grant 1158) is gratefully acknowledged.

1 Introduction

1.1 Overview

This paper presents an experimental investigation of decentralized matching markets. We consider two-sided markets, such as labor markets consisting of workers and firms, marriage markets comprising women and men, and so on. We study the outcomes emerging from free interactions between market participants. Would unconstrained decentralized interaction produce stable outcomes?[1] When there are multiple stable matchings, would some have stronger drawing power? What dynamics would decentralized interactions follow? Our experiments are designed to answer these questions.

The motivation for our study is straightforward. Decentralized matching markets are everywhere: many labor markets and private school systems, are by and large decentralized. Likewise, marriage markets in many western countries operate in a decentralized fashion. Even centralized matching systems are often preceded or followed by decentralized interactions.[2] Furthermore, given the importance of stability in the design of *centralized* matching markets—e.g., for the design of clearinghouses for certain entry-level labor markets and public school systems—we wish to understand when a market would, on its own, result in a stable outcome and, when it does, which stable outcome would emerge. Understanding market parameters that generate different outcome features would indicate when market intervention in the form of centralization would be particularly beneficial.

In our baseline experiments, each side of the market is composed of 8 participants. Each participant is fully informed of all participants' cardinal preferences. Payoffs are designed

[1]A stable outcome corresponds to a pairing of agents from both market sides such that no agent prefers to sever their partnership and no pair of agents prefer to pair with one another over remaining with their assigned partners.

[2]See Echenique, Gonzalez, Wilson, and Yariv (2022) for a description of how pre-match interviews might influence the medical residents' centralized match outcomes in the US.

so that each market participant has either one, two, or three stable match partners. Furthermore, for each ordinal preference profile in the market, there are several cardinal utility representations, differing in the utilitarian welfare of each side of the market as well as the marginal returns from matching with one partner as opposed to a more preferred one. In each round, agents on each side are free to make match offers to any agent on the other side of the market, up to one offer at a time, as well as accept any offer that arrived, up to one at a time. Markets end after 30 seconds of inactivity. Payoffs are determined according to the matches created, where unmatched agents receive a payoff of 0.

The cooperative theory of one-to-one matching offers predictions on the set of plausible outcomes, namely the set of stable matchings, or the core, under two basic premises on the underlying markets: that all agents are completely informed of other participants' preferences (as well as their own) and that agents can freely match with one another. Our design aims at mimicking as much as possible these premises. In particular, agents have complete information on everyone's preferences and agents are free to make offers to one another in a rather unconstrained manner.[3] Our design allows us to inspect the organic selection of stable outcomes, when they are reached, and the endogenous path that generates them.

Three main insights come out of our experiments. First, *stable outcomes appear in a predominance of cases*. Stable matchings occur in 88% of our markets. When there are few stable matchings—one or two—the fraction of stable matching is even higher, exceeding 90%. Furthermore, emergent unstable matchings are very close to stable in terms of payoffs and number of blocking pairs. Our markets are complex enough that eyeballing a stable matching is extremely challenging. From a computational perspective, finding the set of stable matchings is a hard problem in general, see Gusfield and Irving (1989). It is therefore interesting that

[3]This is reminiscent of some of the original general equilibrium experiments, which examined whether markets reach an equilibrium without imposing constraints on the sequential actions that lead them there, see our discussion below.

market forces, when left to their own devices, do find a stable matching, and do so relatively quickly: in time, rarely exceeding 5 minutes, as well as in offer volume, which averages 46 offers per market.

Second, in markets in which each agent has three stable match partners, *the median stable matching emerges as the modal outcome*:[4] 77% of pairings are between median stable match partners, and 80% of markets converge to the median stable matching. This is particularly interesting when contrasted with the leading clearinghouse used in the field, the Deferred Acceptance (DA) algorithm of Gale and Shapley (1962), which implements one of the extremal stable matchings that is most preferred by one market side.

Our last insight pertains to the selection and the cardinal representation of preferences. Stability is an ordinal concept. Thus, the set of stable matchings in a market does not depend on the participants' cardinal assessments of partners. In the lab, these *cardinal representations have a strong effect on the selection of stable matchings*. In particular, the side of the market that has "more to lose" by forgoing their most preferred matching—say, because their marginal loss from shifting from their most preferred stable matching to a less preferred one is greater than the other side's—tends to establish its most preferred matching more frequently.

We report results from two sets of additional treatments that investigate our main results' robustness, in terms of both market size and the bargaining power market participants have. The first set of treatments involves larger markets, with 15 participants, instead of 8, on each side. The main findings from our baseline markets continue to hold. Stable matchings, and in particular median stable matchings, are very frequent. The second set of treatments allows agents only on one market side to make offers. Pairings between stable partners still occur habitually at a rate of 87%. However, the fraction of markets that converge to full stability

[4] Stable matchings are ordered so that there are two stable matchings, each most preferred by one market side and least preferred by the other. The *median* stable matching is ranked by all the agents in between the other two.

is substantially lower and stands at 26%. Perhaps surprisingly, despite the absence of offers from one side of the market, median stable outcomes prevail, occurring in 83% of the final stable matchings. Furthermore, 65% of stable pairings are between median stable partners.

Our design tracks the dynamic path of offers and counteroffers by which our markets reach stability. As already mentioned, convergence to stability is rather quick. We use discrete choice models to explain the making of, and responses to, offers. We find that participants are strategically sophisticated: when making an offer, they appear to put themselves in the place of the recipient of the offer, and gauge whether the offer is profitable to the receiver. That is, they weigh both payoff and yield when evaluating offers.[5]

Taken together, our results indicate that decentralized interactions yield stability for many market structures. Furthermore, they do so efficiently, converging quickly and with relatively few offers. Intervention via centralized clearinghouses may still be beneficial, however. First, decentralized interactions tend to produce the median stable matchings, while common matching clearinghouses tend to implement extremal stable matchings. To the extent that a market designer favors one market side—e.g., young residents in the medical match, or students in a school choice setting—centralized intervention may be beneficial. Second, while the number of offers in our experimental markets is relatively small, of the order of the number of possible pairs, scaling this number to large markets suggests many offers, which may entail substantial costs. Centralized clearinghouses are helpful in eliminating the need for targeted individual offers and speeding up market-wide outcomes.

[5]We also use simulations to examine whether dynamic models that have been offered in the theoretical literature fit our data, namely the Roth and Vate (1990), and Ackermann, Goldberg, Mirrokni, Röglin, and Vöcking (2011) dynamics, as well as an alternative version of Gale and Shapley's 1962 DA in which proposers are randomly selected from both sides of the market. While they are useful benchmarks, none of the three models explains the features of the empirical dynamics we observe in our experiment satisfactorily.

1.2 Literature Review

Theoretically, our experimental design corresponds to the cooperative model of matching markets, see Roth and Sotomayor (1990). The model provides clear predictions: outcomes coincide with the core of the market, the set of stable matchings. In that respect, our experimental results provide a strong experimental validation of the theory underlying the stability notion.

Several propose particular dynamic decentralized processes by which one-to-one matchings are created; see, e.g., Haeringer and Wooders (2011), Ferdowsian, Niederle, and Yariv (2022), and Pais (2008). These papers usually impose some structure on the process by which offers are made and accepted. The main focus of this literature is on the identification of conditions under which stability is likely to arise through equilibrium. Complete information of the prevailing preferences, as in our experiments, allows for stability to emerge in equilibrium, while more stringent demands on preferences and equilibrium selection are required for stability to be the unique prediction. To the extent of our knowledge, the literature is silent on the selection of stable matchings when multiple ones exist in the market.

While the experimental literature on matching markets has grown rapidly in recent years (Hakimov and Kübler, 2021), there are only a few studies of decentralized markets. Kagel and Roth (2000) analyze the transition from decentralized matching to centralized clearing-houses, when market features lead to inefficient matching through unraveling. Nalbantian and Schotter (1995) analyze several procedures for matching with transferable utility, decentralized matching among them, where agents have private information about payoffs. Nalbantian and Schotter (1995) include private negotiations between potential match partners. Offers in our treatment are private as well; only accepted offers become public, but they are non-binding. Agranov, Dianat, Samuelson, and Yariv (2022) allow for transfers in small decentralized markets that follow protocols similar to ours, where information about preferences is either complete or incomplete. Transfers and incomplete information make stability elusive, particularly when preferences are submodular. Pais, Pintér, and Veszteg (2020) study the effects of information and costly-offer frictions on outcomes and find that enough frictions may make stability difficult to achieve. Finally, Niederle and Roth (2009) also look at an incomplete information setting in which one side of the market (the firms) makes offers to the other side (the workers) over three experimental periods. They study the effects of offer structure on the information that gets used in the final matching and consequent market efficiency.[6]

There is also a methodological link between the current paper and some of the experimental work studying financial markets and general equilibrium predictions in the lab, see for instance Smith (1962), Plott and Smith (1978), or the survey in Chapter 6 of Kagel and Roth (1995). As in our paper, the underlying predictions of general equilibrium theory pertain to outcomes, and by and large shine through in experiments; this despite the precise dynamics leading to these outcomes not having been imposed by the experimenters.

2 Theoretical preliminaries

We start by reviewing the underlying cooperative matching model and the theoretical results that are pertinent to our paper.

Let F and C be disjoint, finite sets. We call the elements of F "foods" and the elements of C "colors." We use the language of foods and colors in our experimental design, but these sets can stand for firms and workers in labor markets, men and women in heterosexual marriage markets, etc. A *matching* is a function $\mu : F \cup C \to F \cup C$ such that for all $f \in F$ and $c \in C$,

1. $\mu(c) \in F \cup \{c\}$,

2. $\mu(f) \in C \cup \{f\}$,

[6] There is a growing experimental literature studying centralized matching systems, e.g., Bergstrom, Bergstrom, and Garratt (2013), Harrison and McCabe (1996), Chen and Sönmez (2006), Haruvy and Ünver (2007), Pais and Pintér (2008), Echenique, Wilson, and Yariv (2016), Featherstone, Mayefsky, and Sullivan (2022), and Featherstone and Niederle (2016).

3. $f = \mu(c)$ if and only if $c = \mu(f)$.

Whenever a is unmatched under μ, we write $\mu(a) = a$; when f and c are matched under μ, then $c = \mu(f)$ (and $f = \mu(c)$). Let $\mathcal{M}$ denote the set of all matchings.

A *preference relation* is a linear order (a complete, transitive, and antisymmetric binary relation). In particular, we assume preferences are *strict*. A preference relation for a food $f \in F$, denoted $P(f)$, is understood to be over the set $C \cup \{f\}$, with f representing the possibility of being unmatched. Similarly, for $c \in C$, $P(c)$ denotes a preference relation over $F \cup \{c\}$. For simplicity and consistency with our experimental design, we assume that each food (color) prefers any color (food) over remaining unmatched. A *preference profile* is a list P of preference relations for foods and colors, i.e.,

$$P = \left((P(f))_{f \in F}, (P(c))_{c \in C} \right).$$

Denote by $R(f)$ the weak version of $P(f)$. That is, $c' \, R(f) \, c$ if either $c' = c$ or $c' P(f) c$. The definition of $R(c)$, the weak version of $P(c)$, is analogous.

Fix a preference profile P. We say that a pair (c, f) *blocks* μ if $c \neq \mu(f)$, $c \, P(f) \, \mu(f)$, and $f \, P(c) \, \mu(c)$. In words, (c, f) is a blocking pair if c and f prefer to be matched to one another over their assigned matches under μ. A matching is *stable* if there is no pair that blocks it.[7] Denote by $S(P)$ the set of all stable matchings.

Gale-Shapley Theorem (Gale and Shapley, 1962) *$S(P)$ is non empty, and there are two matchings μ_F and μ_C in $S(P)$ such that, for all $f \in F$, $c \in C$, and $\mu \in S(P)$,*

$$\mu_F(f) \, R(f) \, \mu(f) \, R(f) \, \mu_C(f),$$

$$\mu_C(c) \, R(c) \, \mu(c) \, R(c) \, \mu_F(c).$$

[7] We ignore individual rationality since we restrict attention to preferences under which all agents prefer to be matched to another agent rather than remaining unmatched.

The matchings μ_F and μ_C coincide when the market has a unique stable matching. The matching μ_F is called *food optimal*, while μ_C is called *color optimal*. The matching μ_F is preferred by all foods to any other stable matching, and all colors prefer any stable matching to μ_F. Analogously for μ_C. The proof of the Gale-Shapley Theorem is constructive, and uses what is often referred to as the Deferred Acceptance (DA) algorithm to identify one of the extreme matchings, μ_F or μ_C. Beyond its theoretical role in establishing existence, DA is the algorithm often used in centralized markets. For instance, the National Resident Matching Program uses a variation of it (Roth and Peranson, 1999).

The set of *stable partners* of an agent a is the set of agents that are matched to a under some stable matching, i.e., $\{\mu'(a) \mid \mu' \in S(P)\}$. Likewise, a pair $(f, c) \in F \times C$ is said to be a *stable pair*, or a *stable match*, if f and c are matched under some stable matching.[8] Therefore, in a stable matching, all matches are stable, and all agents are matched to a stable partner. However, in unstable matchings, it might be reasonable for some agents who are not stable partners to be matched: from an individual perspective, it makes sense to form matches in which no agent has a blocking partner, regardless of whether the match is stable.

Consider a preference profile P for which $S(P)$ has an odd number K of matchings, and denote the partners of agent a in each of these matchings by $a_1, \ldots, a_K$ (which may not all be distinct). A *median stable matching* is a matching $\mu \in S(P)$ such that, for all agents $a \in F \cup C$, $\mu(a)$ is a's median partner among a's stable partners under $P(a)$. That is, $\mu(a)$ occupies the $\frac{K+1}{2}$-th place in a's preference among $a_1, \ldots, a_K$. We refer to $\mu(a)$ as the *median stable partner* of a.

In general, median stable matchings are guaranteed to exist, see Teo, Sethuraman, and Tan (2001) (in fact, they also exist when K is even). Median stable matchings present a compromise between the two sides of the market. Interestingly, there are no known simple

[8] Note the distinction between a *match* and a *matching*. A match refers to a pair (f, c) who are matched. A matching refers to the function describing all the matches in a market.

algorithms that generate median stable matchings. Certainly, one can search for all stable matchings of a market and then identify a median one. From a computational perspective, however, this can potentially be quite demanding as the problem of finding all stable matchings is computationally hard (see Gusfield and Irving, 1989, for general references; Irving and Leather, 1986 show that determining the number of all stable matchings is generally #P-complete; while Cheng, 2008 shows that finding the median stable matching is hard). These results contrast with the problem of finding a color- or food-optimal stable matching, which can be done in polynomial time by using DA.

The notion of stability, as well as the ranking of the different stable matchings, are ordinal in nature. In particular, the theory does not allow for refined predictions on the basis of *how much* agents prefer certain partners to others.

3 Experimental Design

Our experimental design corresponds to a decentralized one-to-one, two-sided market.[9] The two sides are termed colors and foods, and contain the same number of participants each. In each round, each participant is randomly assigned a role: "red," "blue," etc. if a color; "apple," "banana," etc. if a food. A participant can match with one and only one participant from the other side of the market, each match resulting in a potentially different monetary payoff. All participants observed all potential payoffs from a numerical matrix on the experimental interface. If a participant is unmatched, they earn a payoff of 0.

We implemented a *baseline* design with 8 participants on each side of the market. We also implemented two variations intended to check for robustness with respect to market size and bargaining power. In the baseline design, over the course of the experiment, participants

[9]The instructions and the set of payoff matrices we used are available at https://sites.google.com/site/decentralizedmatching/.

are free to propose a match to anyone on the other side of the market. At any point in time, participants observe all current matches through a panel of the experimental interface. Importantly, participants can make an offer while (tentatively) matched, and offers can be made to any member of the opposing side of the market, including participants who are already matched. If a matched agent accepts a new offer, their existing match is undone. When receiving an offer, a participant has 10 seconds to respond. Each market ends after 30 seconds of inactivity.

We ran five baseline experimental sessions, each consisting of 2 practice rounds and 10 real rounds.[10] Markets were designed with two objectives in mind. First, in order to see whether cardinal representations of preferences matter, we implemented multiple cardinal representations of the same ordinal markets. Specifically, for any ordinal preference—say, red prefers apple to banana—there are many ways by which these preferences can be presented cardinally. For example, red receiving \$50 and \$10, or \$5 and \$4, from matching with apple and banana, respectively, would both correspond to the same ordinal ranking. Our design entailed 6 underlying ordinal descriptions of markets with 17 different cardinal representations. Second, in order to study the endogenous selection of stable matchings, we designed our experimental markets so that all participants had either one, two, or three possible stable partners. When all agents have only one stable partner, the market as a whole has a unique stable matching.

For each fixed number of stable partners (ranging from 1 to 3), we use several markets differing in market participants' ordinal and cardinal preferences. The following is a general description of the match payoffs used. Table 1 summarizes all our experimental treatments.

[10]Four of the sessions included 32 participants, so we ran two markets at the same time, each consisting of 16 participants (8 foods and 8 colors), for a total of 20 non-practice experimental rounds per session. In the remaining session we only had 16 participants for a total of 10 non-practice experimental rounds. Hence, in total, we ran 90 experimental rounds across five baseline sessions. Due to software malfunctions, 5 markets were not presented as intended to participants; we drop these from the data. Thus, in total, we have 85 rounds of experimental market data. Including data from the dropped rounds does not alter results qualitatively.

Unique stable match partner. We used 4 different ordinal markets: assortative preference markets, where participants on each side of the market agree on the ranking of participants on the other market side; markets with assortative preferences on one side, where only the members of one side of the market are in agreement; a market including a fully egalitarian matching, providing all agents the same payoff, which is unstable; and a "generic" market with a unique stable matching without agents agreeing on the ranking of others on either side. Each of the first three markets was implemented via one cardinal representation. The last, generic market was implemented via four cardinal treatments. In two of these we varied how aligned the interests were across the market: that is, if μ is the stable matching, we computed the correlation of the vectors $(u_i(\mu))_{i \in F}$ and $(u_{\mu(i)}(\mu))_{i \in F}$, where $u_i(\mu)$ denotes the payoff of agent i in the matching μ. We created one market in which the correlation was -0.9 and one where it was 0.9. We used two additional cardinal representations: one in which, for each agent, the difference in utilities between matching with the agent's k'th and $k+1$'th choices was 20¢ and one in which these marginal differences were 70¢. Altogether, we used 7 different cardinal markets with a unique stable matching.

Two stable match partners. We used one ordinal market in which each agent had two possible stable match partners. These were constructed so that there were two 4×4 embedded markets, where any agent within a submarket preferred to match with anyone from that submarket over anyone from the other. We varied the overall utilitarian efficiency of each matching, the utilitarian efficiency of foods relative to colors from each matching, the distribution within each matching,[11] as well as the marginal loss for either side of the market from switching from their more preferred stable matchings to their less preferred ones (higher for foods or for colors). Overall, we used 6 cardinal markets of this sort.

[11] Since egalitarian motives appear frequently in experiments, we were concerned that some form of altruism would be driving our results. We therefore designed payoffs so that in some treatments, fully egalitarian matchings were unstable (see the description of our markets with a unique stable matching above). Furthermore, we included cardinal representations in which certain *stable* matchings were more egalitarian than others.

Table 1: Description of treatments

Agents per side	*Proposing sides*	*#Exp. Mkts.*	*#Ordinal Mkts.*	*#Cardinal Mkts.*	*#Sessions*	*#Particp's*	*#Stable matchings*	*Avg. #stable partners*
Baseline								
*Unique stable matching**								
8	2	30	4	7	5	144	1	1.00
Two embedded 4-by-4 markets†								
		35	1	6	5	144	2×2	1.75
5 stable matchings & 3 stable partners								
		20	1	4	3	80	5	3.00
Unilateral offers								
8	1	15	1	2	4	112	1	1.00
		25	1	4	4	112	2×2	1.75
		23	1	3	5	144	5	3.00
Large markets‡								
15	2	8	3	3	3	90	1	1.00
		4	2	2	2	60	3	2.93
Total		160	11	22	13	378		

Notes: the table reports all experimental markets (round within a session). Markets are grouped into three broad treatments: *baseline* (8-by-8 with two-sided proposals), *unilateral offers* (8-by-8 with one-sided proposals), and *large markets* (15-by-15 with two-sided proposals). For each treatment, the table reports: the number of experimental markets, and the ordinal and cardinal representations used, the number of sessions in which at least one market in the corresponding treatment was played, the number of participants who played in at least one market, the number of stable matchings, and the average number of stable partners. The last row reports the totals across all experimental markets.

*One of the ordinal markets (out of 4) had a salient egalitarian unstable matching, which was used in 4 experimental markets (out of 30). Eight experimental markets (out of 30) were run with two alternative ordinal payoff matrices: four of them had one-sided aligned preferences, and the other four had two-sided aligned preferences.

†In each of the two embedded markets, there were eight agents, four on each side. Three agents on each side had two stable partners, and one agent had one, for an average of 1.75 stable partners per agent.

‡In the markets with a unique stable matching, five (out of 8) had one-sided aligned preferences. In the markets with multiple stable matchings, in one of the markets every agent had three stable partners, while in the other one some agents had two stable partners and most had three (average = 2.87).

Three stable match partners. We used one ordinal market represented cardinally in 4 ways: one in which the marginal differences between utilities derived from matching with one's k'th and $k+1$'th most preferred partners was 20¢, one in which it was 70¢, one in which for foods it was 20¢ and for colors 70¢, and one in which these differences were 20¢ for both market sides, but colors' payoffs were all shifted up by $1. In these markets, while each individual has precisely three distinct stable partners, there are five different market-wide stable matchings.

In addition to the baseline treatments, with 8 participants on each side who could all make match offers, we implemented two additional treatments that differed from the baseline in either market size or the relative bargaining power each market side had.

Large markets. We ran several treatments with larger markets, containing 15 participants on each side. We concentrated on markets with either a unique matching or 3 stable matchings. We used 3 distinct markets with a unique stable matching: assortative preferences on one side, assortative preferences with a fully egalitarian unstable matching, and preferences that were not assortative on either side. We used two markets with 3 stable matchings (in one market, all agents had three stable partners; in the other one, most agents had three stable partners, while some had two).

Unilateral offers. We also ran several sessions with 8×8 markets and payoffs as in our baseline treatment in which only foods could make offers. Otherwise, the market operated as in our baseline treatments.

All sessions took place at the California Social Science Experimental Laboratory (CASSEL), using a modification of the multi-stage software. All participants were UCLA undergraduates and each participant participated in only one session. The average payment per participant was $40 in our baseline treatments, $60 in the large 15×15 market treatments, and $39 in the treatments in which only foods were able to make offers.[12] All of these were combined with a $5 show-up fee.

4 Market Outcomes

Three main findings emerge from our experiments. First, across our treatments, *most market outcomes are stable*. Furthermore, market outcomes that are unstable are close to stable; they are close in the sense that both the number of blocking pairs, as well as the unrealized payoff gains from not forming them, are small. Second, in our treatments with three stable partners, *most agents are matched to their median stable partner*. Surprisingly, even when we "handicap" one side of the market, so that it cannot make any offers, we continue to see the median as the modal outcome. Last, *cardinal representations of preferences affect the particular stable matchings that get selected*. Specifically, higher cardinal incentives to colors make the color optimal matching more likely to be selected; similarly for foods.

In all our treatments, learning across rounds did not appear to have significant effects on neither outcomes nor behavior. All of our results are therefore presented from an aggregation across all 10 rounds.

4.1 Stability in Experimental Markets

Virtually all agents match through our markets' operations, and a large fraction of markets culminates in a stable matching, as shown in Table 2. The table summarizes the overall outcomes in our baseline treatments. Over 99% of agents are matched when markets terminate. Furthermore, 88.24% of markets are fully stable: no agent in the whole market has a blocking partner. Markets with three stable partners exhibit slightly fewer market-wide stable out-

[12]Standard deviations were $3, $10, and $4 respectively.

Table 2: Outcomes in baseline treatment

	Unique stable matching	*Two embedded 4-by-4 markets*	*5 stable matchings & 3 stable partners*	*All baseline*
#Mkts. with stable matching / #Mkts	27 / 30	33 / 35	15 / 20	75 / 85
% Mkts. with stable matching	90.00	94.29	75.00	88.24
Avg. % pairs w/o blocking partners (BPs)	95.32	98.65	90.00	95.24
Avg. % pairs w/o BPs | unstable matching	53.24	76.39	60.00	61.25
Avg. % agents with ≥ 1 *BP*	2.50	0.71	5.00	2.35
Avg. % agents with ≥ 1 *BP | unstable matching*	25.00	12.50	20.00	20.00
Avg. # of BPs per agent | ≥ 1 *BP*	1.35	1.00	1.23	1.22
Avg. % unmatched agents	0.42	0.36	0.00	0.29
Avg. % unmatched agents | ≥ 1 *BP*	4.76	25.00	0.00	6.43

Notes: the table reports the following final outcomes for each treatment in the *baseline* treatment: (i) number and (ii) percent of final matchings that are stable; avg. number of final pairs (matches) in which no agent has a blocking partner across (iii) all markets and (iv) markets with an unstable final matching; average percent of agents with at least one blocking partner across (v) all markets and (vi) markets with an unstable final matching; (vii) average number of blocking partners per agent across agents with at least one blocking partner; average number of unmatched agents (viii) among all agents, and (ix) among agents who have at least one blocking partner.

comes, 75%. As a measure of stability at the match or pair level, we check the proportion of matched pairs in which no member was part of a blocking pair. In over 95% of matches, no member has a blocking partner, with little variation across markets. The treatment with the lowest proportion of blocking partners is the one in which every agent has three stable partners. Even in this case, in 90% of matches, no agent can form a blocking pair.[13]

Not all our markets reached full stability, but markets that were not fully stable were fairly close to stable, as shown in Figure 1. Even in unstable markets, Table 2 indicates that in the majority of ultimately matched pairs (61.25%), no agent has a blocking partner. Indeed, only 20% of agents has a blocking partner when markets finalize (compared to 2.35% of agents across all markets, culminating in stable or unstable matchings). Another way to measure the proximity to a stable matching is by looking at the maximal number of disjoint blocking pairs in a matching. The left panel of Figure 1 presents the empirical cumulative distribution functions (CDFs) of the maximal number of disjoint blocking pairs across our different treatments.[14] The figure contains the distribution of blocking pairs pertaining to *markets in which outcomes were not fully stable*.[15] In our baseline treatment, most outcomes are close to stable: most markets culminating in unstable outcomes have only one or two disjoint blocking pairs.

We can also use payoffs to measure the distance to stability. For every agent, we compare the payoff they received from their final match, with the one they would have received had they matched with their most preferred blocking partner. That is, we compute the maximum loss due to instability per agent. Agents in markets that did not culminate in a stable matching

[13]The use of color and food labels in our markets did not seem to have any effect. For example, if one considers banana and mango to be associated with yellow, apple and cherry with red, and kiwi and pear with green, there is no significant increase in the corresponding matches relative to any other classification.

[14]In our baseline treatments, when no agents are matched, the maximal number of disjoint blocking pairs is 8, the number of potential disjoint pairs. When a stable matching is in place, there are 0 blocking pairs.

[15]Including all markets in the figure would not allow us to visualize the fine-grained distribution of blocking pairs, as the prevalence of stable matchings implies a large spike at zero blocking pairs.

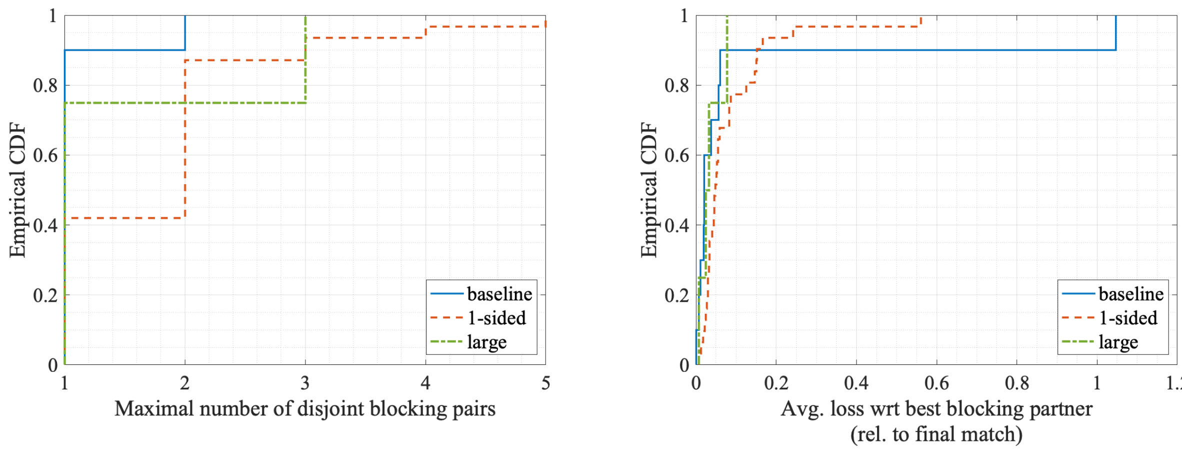

Figure 1: Distance to stability in unstable markets

lost on average 12.82% of their final payoff, equivalent to around 50¢. By contrast, the same average across all markets is 1.51%, equivalent to around 6¢. The right panel of Figure 1 reports the CDFs of this measure across all unstable markets. Even in unstable markets, most agents had very small payoff losses due to instability. Indeed, the median loss is equivalent to 2.08% of the final payoff, which comes down to around 9¢.[16]

In terms of the duration of market interactions, convergence to the stable matching was rapid. On average, markets terminated after 2.78 minutes (sd = 1.35), including the final 30 seconds of inactivity; 89% of markets did so before the five minute mark. We return to the dynamics underlying these observations in Section 7.

4.2 The Emergence of Median Stable Matches

The median stable matching has strong drawing power and cardinal representations of preferences affect outcomes, as shown in Table 3. The top panel of the table reports the distribution

[16]Figure 4 in the Appendix shows CDFs for alternative measures of distance to stability across the distinct treatments: total number of blocking pairs, and the average loss with respect to the best blocking partner in absolute terms, relative to the average market payoff, and relative to the average of the same measure across random matchings (which with great likelihood are unstable). Roughly speaking, all measures point to the same conclusion: the majority of agents in markets that did not reach full stability had few blocking partners and did not incur in great payoff losses.

of final matchings and matched pairs across the baseline markets in which each agent has three stable match partners. A key difference between the markets is the level, and marginal, differences of payoffs within each side of the market. We use the notation of "x-y marginals" to denote a market in which the marginal difference in utilities between one partner and the next-best partner was x cents for foods and y cents for colors. In the 20-20 marginals market with 100 color shift (labeled as $20\text{-}20_{+100}$ in the table), a 100¢ (or \$1) was added to the color payoffs of the 20-20 market.

The median stable matching is the modal outcome: 80% of stable matchings correspond to the median stable matching, entailing *all* participants being matched to their median stable partner. Moreover, *none* of these markets converged to an extremal stable matching, food-optimal or color-optimal. Recall that the baseline markets in which each agent has three stable partners have five market-wide stable matchings. Hence, there are two stable matchings that are neither the median nor extremal. We emphasize that, in these markets, a median matching had to be "discovered" by participants: it is challenging to read the set of stable matchings from looking at the payoff tables (an 8×8 table with 128 numerical entries).

Shifting attention to individual match-level outcomes, we observe similar patterns. 77% of the final matched pairs that were stable correspond to median stable partners. From the perspective of individual behavior, agents care about stable partners, not market-wide matchings. It is remarkable that the vast majority of stable matches—which comprise 97% of all final matched pairs—correspond to the median. That is, in the vast majority of cases, agents ended up in stable matches in which neither of the two parties were matched to their most preferred stable partner: no side of the market systematically got their way. Of the remaining stable matches, the majority were color-optimal stable matches (20% out of 23%), and the rest were food-optimal matches.[17]

[17] The preferences in these markets, both ordinal and cardinal, were not symmetric across the two sides, which explains the imbalance.

Table 3: Selection of stable matching and cardinal effects

Marginal Utility Differences (foods–colors)	20–20	20–20$_{+100}$	20–70	70–70	All
Baseline					
# Markets	5	5	5	5	20
Market-level outcomes (matchings)					
% stable	60	40	100	100	75
% median | stable	67	100	60	100	80
% non-extremal | stable	100	100	100	100	100
% food-optimal | stable	0	0	0	0	0
% color-optimal | stable	0	0	0	0	0
Individual-level outcomes (matches)					
% stable	95	92	100	100	97
% median | stable	64	63	80	100	77
% food-optimal | stable	11	0	0	0	3
% color-optimal | stable	24	37	20	0	20
Unilateral offers (foods propose)					
# Markets	0	5	9	9	23
Market-level outcomes (matchings)					
% stable	–	20	11	44	26
% median | stable	–	0	100	100	83
% non-extremal | stable	–	100	100	100	100
% food-optimal | stable	–	0	0	0	0
% color-optimal | stable	–	0	0	0	0
Individual-level outcomes (matches)					
% stable	–	93	83	93	89
% median | stable	–	49	61	79	65
% food-optimal | stable	–	14	21	19	19
% color-optimal | stable	–	38	18	2	16

Notes: The table reports (i) the percentage of markets which final matching corresponds to the median, non-extremal, food-optimal, or color-optimal stable matching, and (ii) the percentage of final matches that are food-optimal, color-optimal, or median stable matches (i.e., part of a stable matching). The tables reports results for all experimental markets with five stable matchings, in *baseline* (top panel) and *unilateral offers* (bottom panel). The results are disaggregated into cardinal treatments, according to the marginal utility of each side (payoff difference from less preferred to more preferred partner). The second column corresponds to markets in which utility differences are 20¢ on each side, with the payoffs of colors shifted upwards by $1.

Stability depends on agents' ordinal preferences alone, not on their cardinal representation. Of course, our experimental participants receive different monetary payments depending on who they match with, and one would expect the different monetary magnitudes to play a role. It is interesting, then, that in our data, the prevalence of stable matchings does not seem to depend on any specific cardinal representation of preferences. But *which* stable matching gets selected does depend on the cardinal utility representation, as shown in Table 3. Participants on the side of the market (foods or colors) having "more to lose" tend to partner with their most preferred match at higher frequencies (significant at any conventional levels).

5 Robustness to Bargaining Power and Market Size

We investigate the robustness of our results along two dimensions: the bargaining power afforded to each side of the market, and the market's size. As for bargaining power, in our *unilateral design*, only one side of the market is allowed to make offers. Specifically, only foods are able to make offers to colors. The effect of market size is studied by means of our *large-market design*, which resembles our baseline treatments but with 15 agents on each market side. Table 4 summarizes our findings.

Unilateral Offers. A natural possible explanation for the prevalence of the median stable matches is that, in our experiments, participants on both sides of the market could make offers. In a sense, they had equal bargaining power. In contrast, in Gale and Shapley's DA, only one side makes offers while the other side decides which offers to accept. DA produces the optimal matching for one side, so it is possible that allowing both sides of the market to make offers is at the root of the frequent median matches we observe. This conjecture is interesting, particularly in view of the fact that several real-world matching markets that operate in a decentralized manner allow one side greater, if not sole responsibility for making

Table 4: Outcomes in unilateral offers and large markets

	Unilateral offers	*Large markets*
#Mkts. with stable matching / #Mkts	32 / 63	8 / 12
% Mkts. with stable matching	50.79	66.67
Avg. % pairs w/o blocking partners (BPs)	77.65	93.33
Avg. % pairs w/o BPs $\mid$ unstable matching	54.57	80.00
Avg. % agents with ≥ 1 BP	13.10	3.61
Avg. % agents with ≥ 1 BP $\mid$ unstable matching	26.61	10.83
Avg. # of BPs per agent $\mid \geq 1$ BP	1.23	1.04
Avg. % unmatched agents	1.39	0.00
Avg. % unmatched agents $\mid \geq 1$ BP	3.71	0.00

Notes: the table reports the following final outcomes for markets with *unilateral offers* and in *large markets*: (i) number and (ii) percent of final matchings that are stable; avg. number of final pairs (matches) in which no agent has a blocking partner across (iii) all markets and (iv) markets with an unstable final matching; average percent of agents with at least one blocking partner across (v) all markets and (vi) markets with an unstable final matching; (vii) average number of blocking partners per agent across agents with at least one blocking partner; average number of unmatched agents (viii) among all agents, and (ix) among agents who have at least one blocking partner.

offers: e.g., the job market for academics, the marriage market in certain cultures, etc. To explore the effects of bargaining or proposal power, in our *unilateral offers* treatments, only foods were allowed to propose matches. While stability rates are not as high as in our baseline experiments, we find that outcomes are fairly stable and also correspond to median outcomes.

With unilateral offers, outcomes are qualitatively similar to those in our baseline treatments, as seen in the bottom panel of Table 3 and the left column of Table 4. In these treatments, 99% of participants are matched. In 78% of the final matches, no agent is in a blocking pair, compared with 95% in the baseline markets in which both sides can make offers. At the market level, 51% of the markets culminate in a stable matching, compared with 88% in the baseline treatments.[18] As Figure 1 shows, the observed unstable matchings are close to stable, although not as close as in our baseline treatments. Across the markets that fail to reach stability, on average, around a quarter of agents has at least one blocking partner, compared to a fifth in the baseline treatments. Similar to the baseline treatments, agents who have blocking partners, have 1.2 such partners on average. The convergence in unilateral markets is somewhat faster than in our baseline treatments: markets terminate after 2.21 minutes (sd = 1.05), including the final 30 seconds of inactivity, around 20% earlier than in the baseline treatment.

Regarding selection, two important observations emerge. First, in markets with four stable matchings (two per each embedded submarket), the stable matchings favored by foods, the proposers, are much more frequent. In the baseline treatments, 27% converge to the food-optimal stable matching, 21% to the color-optimal, and the rest to a non-extremal stable matching.[19] By contrast, in markets in which only foods can make offers, 79% of the markets

[18]The corresponding percent of final matches in which no one is in a blocking pair (namely, final matchings that are stable) for markets with a unique stable matching, two embedded 4×4 markets, and with five stable matchings and three stable partners are 80% (47%), 92% (76%), and 61% (26%).

[19]In these markets, non-extremal stable matchings are the ones in which one of the 4×4 submarkets is at the food-optimal stable matching, and the other submarket is at the color-optimal one.

converge to the food-optimal stable matching, and none to the color-optimal stable matching (the rest to one of the two non-extremal ones). The same finding emerges when looking at individual matches across all markets, not only those reaching stability. In the baseline markets, among all the agents who have two stable partners (75%), 68% of final matches are food-optimal, and the rest are color-optimal. With unilateral offers, 88% of the final matches are food-optimal.

For the markets with two embedded 4×4 markets, results are consistent with extremal outcomes, but this is no longer true in markets with three stable partners. Indeed, for markets in which each agent has three stable partners, median outcomes are again the most common. As the bottom panel of Table 3 shows, 65% of final matches that are stable correspond to median stable matches when only foods are able to make offers, compared to 77% when both sides can do so. Similarly, among markets that reach full stability, 83% reach the median stable matching, compared to 80% in the baseline treatments. There is also an increase in the number of matches that are optimal for foods, the proposing side, from 3% to 19%, but this comes mainly at the expense of matches that are optimal for colors, the receiving side. Nonetheless, cardinal incentives matter for the distribution of outcomes as they do in our baseline treatments: when foods have more to lose from forgoing higher-ranked partners, color-optimal matches are less common.

Market Size. The outcomes in our large-market treatment are very similar to those in our baseline experiments, as seen in the right column of Table 4. Our large markets involve 30 participants, which is quite large as experimental markets go. They do not, of course, come near the size of some real-world matching markets—such as, for example, the medical residents market—but it is a comfort that duplicating the size of our baseline markets does not upset our main results.

In our large-market treatments, in 93% of the final matches, no agent was part of a blocking pair. In terms of market-wide outcomes, while 67% of matchings are fully stable, the number of blocking pairs within unstable markets is, again, very small, see Figure 1. Our large experimental markets exhibit at most 3 disjoint blocking pairs, and in over 60% of cases just one. Indeed, as Table 4 reports, even in markets that fail to reach full stability, the vast majority of agents have no blocking partners (89%, as compared to 96% across all markets, unstable and stable). Agents who have blocking partners have, on average, just one; and there are no agents left unmatched. Interestingly, the seemingly more complex markets with three stable matchings always culminate in a stable matching in our large 15×15 markets.

As one might expect, convergence is slower in these large markets. On average, markets terminate after 5.35 minutes (sd = 3.06), including the final 30 seconds of inactivity, and half do so before the five-minute mark.

As for the selected stable match partners, the results are similar to those in our baseline treatments, although arguably more extreme. In the large market treatments, *every* agent who has three stable partners (93%) matches with their median stable partner. Hence, the market-wide stable matchings are exclusively median stable matchings.

6 Cardinal Incentives and Social Preferences

A matching delivers cardinal (monetary) payoffs to market participants, and these payoffs may be more or less fair, they may be more or less equal across participants. The degree of fairness of a matching can reasonably be expected to influence whether it is chosen: a taste for egalitarian outcomes is well-documented in experimental economics (the literature has suggested different types of social preferences, for surveys see, e.g., Chapter 4 in Kagel and Roth, 2020, and Fehr and Gächter, 2000). Fairness considerations may, in particular, be important given our finding that median matches are very common.

Our design entailed several markets with identical ordinal preference profiles, but different cardinal utility representations. Focusing on such variations, we have already seen that the side that faces steeper cardinal incentives is more likely to achieve its optimal stable matching. We now inspect the impact of payoff *distributions*. We show that equality of payoffs across match partners makes a matching more likely, *as long as it is stable.*

We first analyze markets with a salient egalitarian, but unstable, matching (see Section 3). In these markets, one unstable matching entails identical payoffs to all market participants, and comparable utilitarian welfare to that generated by the unique stable matching.[20] Participants consistently fail to select the unstable egalitarian matching. In our baseline treatments, none of these markets end up in the unstable, albeit more equal, matching outcome.

Our larger markets, in which more payoff variation could be introduced, generate similar findings. For example, we ran one 15×15 market in which there was an unstable matching under which all agents received exactly \$4; there was a unique stable matching in which the average payoff was also \$4, which was much more unequal.[21] While the final outcome in these markets is not always stable—in some instances a few blocking pairs remain—participants clearly avoid the egalitarian unstable matching.

Second, we use the markets with two stable matchings to assess the effects of cardinal utilities on the selection of food- and color-optimal stable matchings.[22] Markets are more likely to converge to the extremal stable matching, food- or color-optimal, which has the lowest payoff variation. That is, when the dispersion of payoffs is relatively high in, say, the color-optimal matching, markets tend to achieve the food-optimal matching, and vice-versa.

[20]By specifying identical payoffs to all market participants, we avoid the challenge of identifying the individuals that are relevant for participants' social preferences.

[21]Its Gini index was 26. For a country's income distribution, this is in the Scandinavian range, but it appears starkly unequal compared to perfect egalitarianism with no utilitarian efficiency loss.

[22]Recall that each of these markets was constructed by embedding two smaller 4×4 markets, which allows us to gain more payoff variations, see Section 3. In these markets, 99% of final matches are stable. In addition, the embedding of the "submarkets" was effective—indeed, there are only 6% of cross-offers and 1% of cross-matches (the acceptance probability of cross-offers is 0.05, compared with 0.42 of non-cross offers).

Specifically, we compute the *coefficient of variation* of agents' payoffs in a given matching: the standard deviation of payoffs divided by their mean. The coefficient of variation is a "scale free" measure of the dispersion in payoffs across agents. We then compute the ratio of the coefficient of variation at the food-optimal stable matching over that at the color-optimal stable matching. Markets with a high ratio are those in which the payoff variation in the food-optimal stable matching is high, relative to that in the color-optimal stable matching. When the ratio is above the median ratio in our data, the food-optimal matching obtains 11.7% of the time, and the color-optimal 35.3%. When it is below the median, they obtain 43.8% and 6.3%, respectively. The difference between these values is significant at any conventional level of confidence. The implication is that when the variance in payoffs at the color-optimal matching is relatively high, we tend to get more food-optimal outcomes, and vice-versa.

To summarize, our results suggest that egalitarian, or fairness, considerations play a role in selecting outcomes, but they are not so strong as to trump stability.[23]

7 Market Dynamics

So far, our discussion focused on final market outcomes. In this section, we describe the dynamics generating the outcomes we observe.

7.1 Evolution of Offers, Responses, and Matches

Markets reach stability gradually, with the vast majority of blocking opportunities, especially the most profitable ones, vanishing during the initial stages in a round, as shown in Figure 2. The figure plots the total number of blocking pairs, the maximum number of disjoint blocking

[23]Our results cannot be directly compared with the experimental bargaining literature that has documented a preference for fairness. The nature of bargaining in our experiments is substantially different from other experiments, such as the dictator game, where the compromise between the two sides is obvious. In fact, as noted, identifying the set of stable matchings in our markets is arguably extremely challenging.

pairs, the number of agents in at least one blocking pair, and the average payoff loss with respect to the best blocking partner of each agent. The solid lines correspond to averages across all baseline markets, and the shaded regions to a standard deviation above and below the average at every point in time. On average, after 20% of markets' duration has elapsed, half of all blocking opportunities vanish, and the average gains per agent from forming blocking pairs are cut in half. Nonetheless, over half the agents still have a potential blocking pair. When 90% of a market's duration has elapsed, on average, there are only around four agents who still have a blocking partner, and each has one or two blocking opportunities. Most markets end with no blocking pairs. After a relatively short interval of time, most blocking pairs are traded away.[24]

Offers echo the picture emerging from Figure 2. More offers are made initially: about half are made within the first 40% of markets' duration. Table 5 provides an overall summary of offer features throughout market operations.

Markets with a larger number of stable matchings exhibit more offers, repeated offers, and matches, as shown in the top panels of Table 5. On average, 59 offers are made, and 25 matches are formed in markets with five stable matchings. In markets with a single stable matching, on average, only 45 offers are made, and 16 matches are formed, while in the two four-by-four markets, there are around 40 offers and 16 matches.[25]

The offer volume contrasts with that of Gale and Shapley's DA. Under DA, our markets with multiple stable matchings require a relatively low number of offers, depending on which side is proposing, 13–14 for markets with five stable matchings, and 10–12 for the two four-by-four markets, while the markets with a unique stable matching require 18–21 offers. These

[24] A similar image emerges when differentiating baseline markets according to their number of stable matchings; as well as for our additional treatments entailing larger markets and unilateral offers. See Figure 5 in the Appendix.

[25] These figures may reflect the complexity of our markets, where more complex markets require more activity to converge. Recall that, even though the two four-by-four markets feature four stable matchings, they effectively operate as two smaller markets, each with four participants.

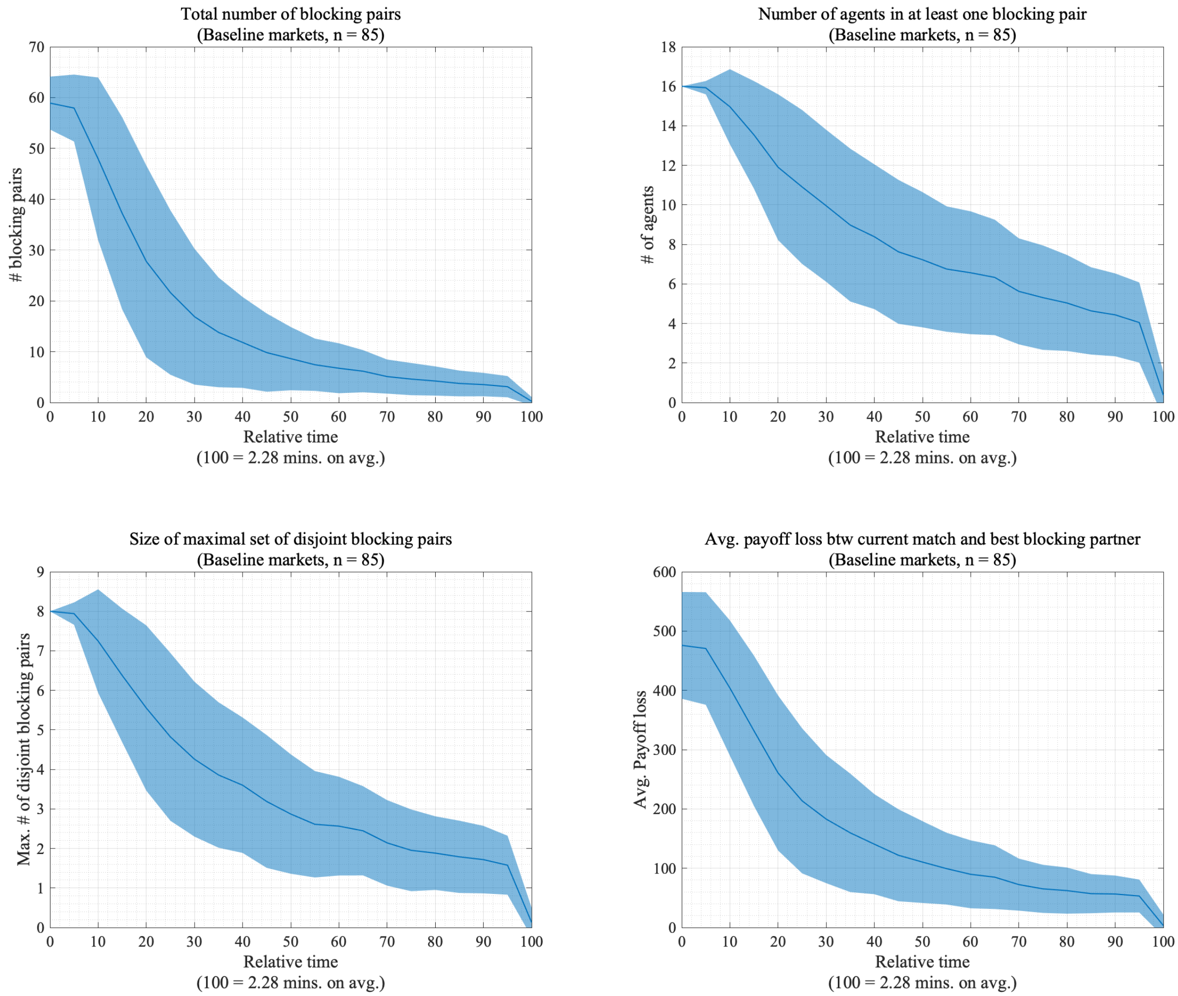

Figure 2: Distance to stable matching over time

numbers are not a good predictor of the volume of offers in our experimental markets, which is consistent with the final outcomes we observe not coinciding with the extremal stable matchings DA produces. Another contrast with the DA algorithm is the volume of repeat offers we observe. Under DA, an offer is never repeated. In our data, around a third of offers are made to agents whom the proposer had already made an offer to previously. While the acceptance rate of offers is around 43%, repeated offers are accepted at lower rates of around 29%. Across our treatments, however, 19% of matches are repeat matches.

Table 5: Summary statistics of dynamics

	Unique stable matching	*Two embedded 4-by-4 markets*	*5 stable matchings & 3 stable partners*	*All baseline*
# Mkts.	30 (35.3%)	35 (41.2%)	20 (23.5%)	85 (100.0%)
# offers	44.8	39.7	59.2	46.1
# matches	15.9	15.9	24.6	18.0
% accepted offers	41.3	46.5	40.4	43.2
% repeated offers	29.0	33.6	36.8	32.7
% accepted $\mid$ *repeated*	28.1	30.3	27.3	28.8
% repeated matches	16.2	17.6	24.6	18.8
% offers to blocking partners	65.5	65.0	63.5	64.8
% offers only-proposer beneficial	31.1	29.8	33.6	31.2
% proposer is matched	28.5	36.5	42.3	35.0
% proposer is matched $\mid$ *#offer* > 1	53.3	59.0	63.2	58.0
% receiver is matched	42.2	43.6	46.4	43.8
% receiver is matched $\mid$ *#offer* > 1	65.5	61.8	63.7	63.5
% proposer is active	88.1	78.7	84.7	83.5
% offer is downward	76.7	77.8	61.1	73.5
% offer is Gale-Shapley	42.8	48.6	31.8	42.6
% offer skips someone	35.6	23.4	41.0	31.8

Notes: The table reports averages of the following variables across markets in the *baseline* treatment: number of offers; number of matches (same as number of accepted offers); % of offers that are accepted; % of offers that are repeated (proposer had already proposed to the receiver previously); % of offers that are accepted among those that are repeated; % of matches that are repeated (proposer and receiver had already been matched previously); % offers made to a blocking partner; % of offers that would only be beneficial to the proposer if accepted; % of offers in which the proposer/receiver is matched (conditioning on offers other than the first of each proposer); % offers in which the proposer is *active* (a proposer is *active* if there exists an agent they prefer to their current match whom they have not made an offer to); % offers that are *downward* (an offer is *downward* if the proposer has not proposed to an agent who they prefer less than the receiver); % offers that are *Gale-Shapley* (an offer is *Gale-Shapley* if the receiver is the proposer's most preferred agent among the ones the proposer has not proposed to); and % offers that *skip someone* (an offer *skips someone* if there exists an agent the proposer has not made an offer to and whom the proposer prefers more than the receiver).

Offers are nearly always beneficial for proposers, frequently made to blocking partners, and seemingly independent of whether either is matched, as shown in the middle panels of 5. Roughly 65% of offers are made to blocking pairs while 31% of offers are beneficial only to the proposer.[26] Furthermore, agents continue making offers while matched. Taking out the first offer made by every agent, the majority of offers, 58% of them, are made by matched agents. Similarly, offers are made regardless of whether the recipient of the offer is matched or not.

The bottom panel of Table 5 provides further insights on *when* participants decide to make offers. In line with our observations so far, participants make offers when there are conceivable benefits to doing so, but their behavior is distinctively different from that prescribed by the DA algorithm. We call an agent *active* if there exists another agent they prefer to their current match whom they have not proposed to previously.[27] Across all our baseline treatments, most proposers (84%) are active. This proportion is higher than that of proposers who are matched, in line with our observation that agents make offers while paired.

In terms of *whom* offers are made to, we classify offers into three overlapping categories: *downward*, *Gale-Shapley*, and whether they *skip someone*. An offer is *downward* if the proposer has not proposed to an agent whom they prefer less than the receiver. Intuitively, offers are downward if a proposer is going down their preference ranking when making an offer. The condition for an offer being downward relates only to offers made before, not proposals that were accepted: even if a proposer makes an offer to a target they prefer to their current match, their offer might still be downward. While the vast majority of the offers we observe are downward (74%), it is not uncommon for agents to go up their ranking. An offer is said to be *Gale-Shapley* if the receiver is the proposer's most preferred agent among the ones the proposer

[26]The remaining 4% can be catalogued as "mistakes," in that these offers would not be beneficial to the proposer if accepted.

[27]In the DA algorithm, by construction, all agents who make offers are active. However, the converse does not hold: matched active agents cannot make offers.

has not proposed to.[28] A minority of offers are Gale-Shapley (43%), implying that agents either skip potential partners or make repeat offers throughout the market. Accordingly, an offer is said to *skip someone* if there exists an agent the proposer has not made an offer to whom they prefer to the receiver. Around 32% of offers skip someone.

In sum, while agents tend to go down their preference lists when making proposals, they generally do not follow the order of proposals proscribed by the DA algorithm: they frequently make repeat offers and skip agents on their preference list. Similarly, while the vast majority of proposals are beneficial to both the proposer and the receiver, a non-negligible fraction of offers ignore the incentives faced by receivers and are beneficial only to the proposers.

Offer features vary over time. Figure 3 displays all offers made, in all of our baseline markets. On the x-axis, we plot the relative time of each proposal, where 0 indicates the start of a round and 100 the time at which the last proposal occurs.[29] The top two panels have the receiver's rank on the proposer's preference list on the y-axis, where 1 stands for the most preferred agent and 8 for the least preferred. The bottom two panels have the proposer's rank on the receiver's preference list on the y-axis. The left two panels indicate which offers are rejected (in red) and which are accepted (in blue). The right two panels indicate whether offers are made to blocking partners (in green), beneficial only to the proposer and not to the receiver (in yellow), beneficial only to the receiver (in red), or beneficial to neither (in gray).[30]

Figure 3 suggests several messages. First, offers are much more likely to be accepted by receivers who rank their proposers highly. Second, proposers internalize receivers' responses. While many of the offers are made to the proposer's top choice, 45% to be precise, a significant fraction is made to the second and third choices, even at early stages of the market (overall,

[28]The label follows from noting that in the DA algorithm, all offers are Gale-Shapley.

[29]On average, the last proposal took place after 2.28 minutes in our baseline markets; 30 seconds before the round terminates.

[30]The percentages displayed in the legend are across all offers made, pooling all the markets together. Hence, they differ slightly from those in Table 5, which average across markets.

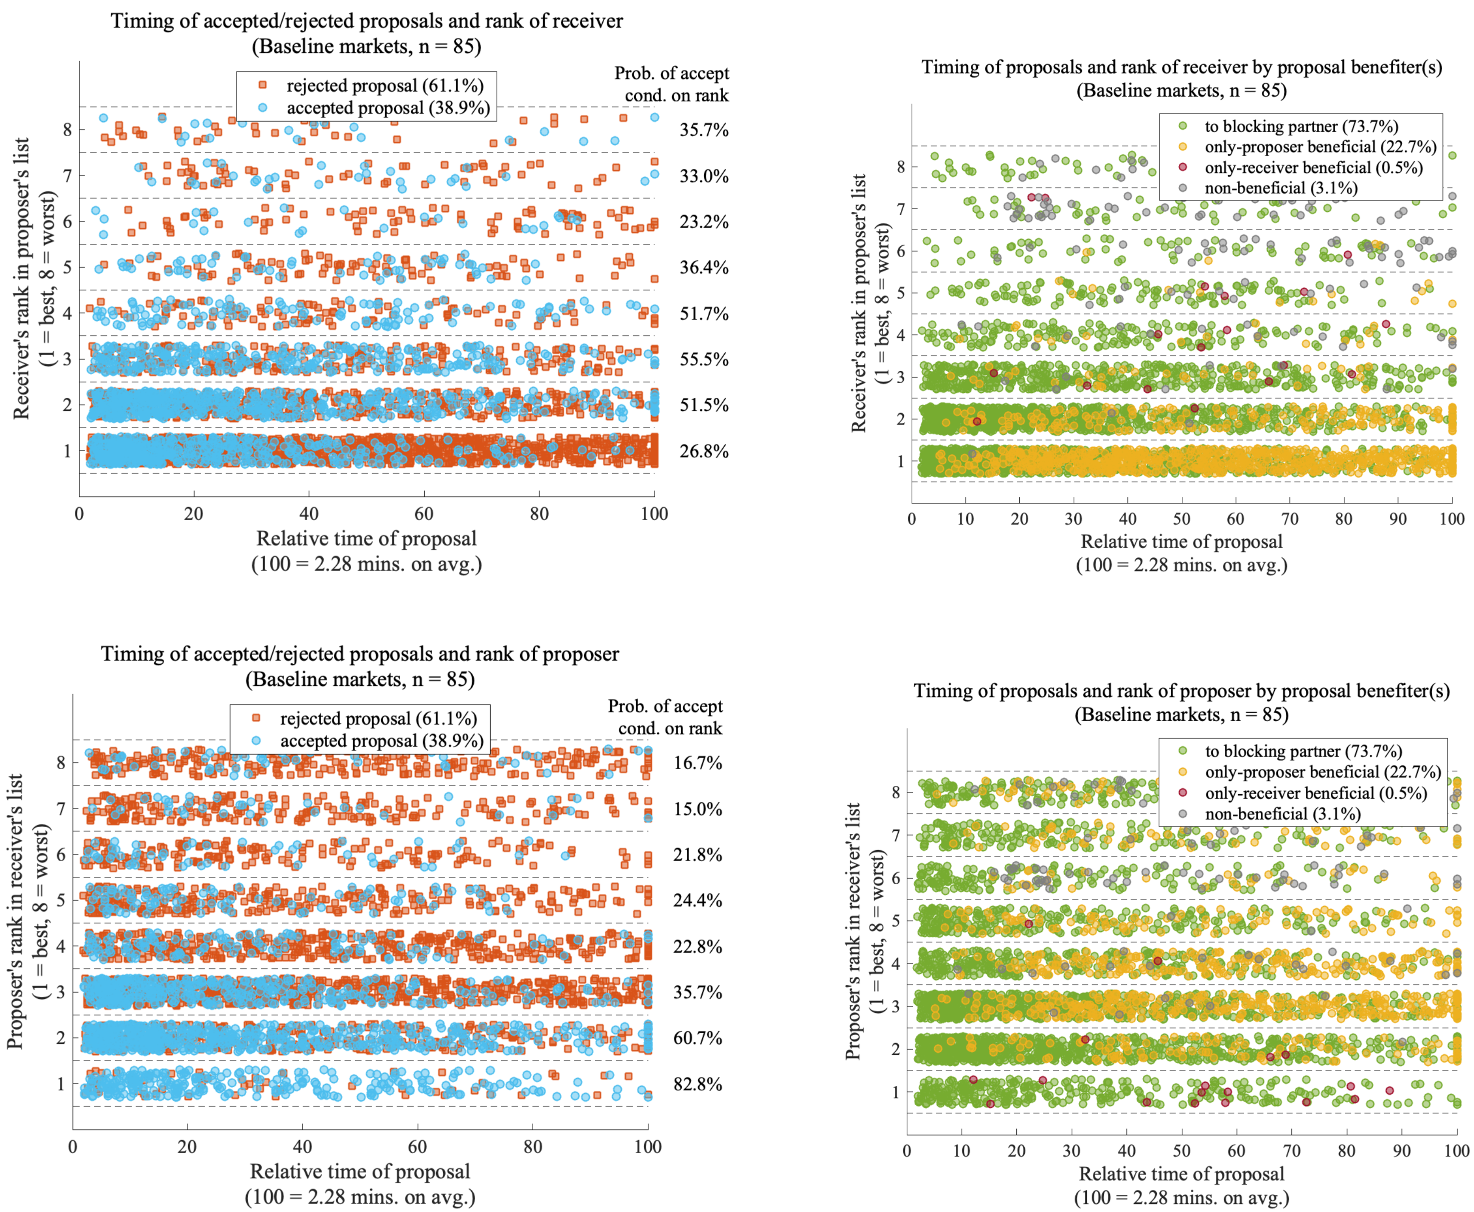

Figure 3: Timing of proposals, rank of receiver (top) and of proposer (bottom)

85% of offers are made to one of the proposer's three top choices). Offers made to the second, third, and fourth top choices of a proposer are much more likely to be accepted than those made to a proposer's most preferred partner: over 50% acceptance rate compared with 27%. Third, proposers do not internalize receivers' responses fully. A substantial fraction of offers are made to receivers for whom proposers are low ranked. Such offers are often rejected. In fact, the majority of offers that are only beneficial to proposers are made to proposers' most preferred partners, disregarding receivers' preferences, at later stages of the market. These offers are frequently rejected.

Next, we inspect how our markets unfold over time and evaluate several theoretical models of market dynamics. We then analyze agents' choices to make or accept offers over time.

7.2 Dynamic Models and Simulations

The theoretical literature on stabilization dynamics in matching markets is limited.[31] Three dynamic models guide our analysis of the evolution of matchings in our experimental markets. First, we consider a version of the Gale and Shapley (1962) DA algorithm in which proposers are chosen randomly from both sides. Second, we consider the Random Paths to Stability (RPS) model of Roth and Vate (1990). Third, we consider the Random Best Reponse (RBR) dynamics proposed by Ackermann et al. (2011), which allows (myopic) optimization in a process resembling RPS.

Two-Sided Random Deferred Acceptance (Two-RDA). In the algorithm we consider, a proposer is chosen uniformly at random from either side among all agents who are active.[32] Proposers make offers to agents they prefer the most among those they have not proposed

[31] There is an active literature on dynamic matching markets, see Baccara and Yariv (2022) and our literature review. However, relatively little attention has been dedicated to how agents reach stable outcomes, analogous to the literature on learning in game theory, or tâtonnement dynamics in competitive markets.

[32] As before, an agent is active if there exists someone they have not proposed to, whom they prefer to their current match.

Table 6: Market Simulations—Two-Sided Random Deferred Acceptance (Two-RDA), Random Paths to Stability (RPS), and Random Best Response (RBR)

	Experiment	*Two-RDA*	*RPS*	*RBR*
Market Activity				
# *offers*	59.2	54.2	37.6	38.2
# *matches*	24.6	30.7	37.6	38.2
% *accepted offers*	40.4	56.4	100.0	100.0
% *repeated offers*	36.8	0.0	17.6	24.1
% *accepted* | *repeated*	27.3	-	100.0	100.0
% *repeated matches*	27.2	12.8	17.6	24.1
# *offers to blocking partners*	37.3	30.7	37.6	38.2
% *offers to blocking partners*	63.5	56.4	100.0	100.0
% *offers only-proposer beneficial*	33.6	43.6	0.0	0.0
% *proposer is active*	84.7	100.0	99.2	98.7
% *offer is downward*	61.1	100.0	69.8	67.8
% *offer is Gale-Shapley*	31.8	100.0	15.1	22.6
% *offer skips someone*	41.0	0.0	82.3	73.0
Market-level outcomes (matchings)				
% *stable*	75.0	46.9	100.0	100.0
% *median* | *stable*	80.0	63.9	59.2	55.3
% *non-extremal* | *stable*	100.0	89.2	83.8	83.6
% *food-optimal* | *stable*	0.0	5.4	0.6	3.7
% *color-optimal* | *stable*	0.0	5.4	15.7	12.7
Individual-level outcomes (matches)				
% *stable*	96.9	93.8	100.0	100.0
% *median* | *stable*	75.6	51.2	71.5	69.4
% *fruit-optimal* | *stable*	2.5	19.8	0.6	3.7
% *color-optimal* | *stable*	18.8	22.6	28.0	26.8

Notes: The table reports average activity and outcome measures across the *Baseline* markets with five stable matchings and three stable partners (see Table 3). The first column reports the experimental data, and the second through fourth report simulations using the Two-Sided Random Deferred Acceptance (Two-RDA), the Random Paths to Stability (RPS) of Roth and Vate (1990), and the Random Best Response (RBR) of Ackermann et al. (2011). For each algorithm and market, we ran 10,000 simulations. In Two-RDA, proposers are chosen uniformly at random. In RPS, blocking pairs are chosen with probability proportional to $\exp(\lambda g_{f,c})$, where $g_{f,c}$ denotes the total net gain of blocking pair (f, c) and $\lambda = 0.0175$. In RBR, agents who have at least one blocking partner are chosen with probability proportional to $\exp(\lambda g_a)$, where g_a is the maximum net gain of agent a across all its blocking partners and $\lambda = 0.00425$. We choose λ in both cases to roughly match the number of offers made to blocking partners in the experimental markets. See Tables 9, 10, and 11 in the Appendix to see results with distinct distributions for each algorithm.

to previously. In turn, agents accept offers when the proposers are ranked higher than their current matches. The algorithm terminates when the set of active agents is empty.[33]

Results from Two-RDA generate similar offer volumes to those observed in the data, but are inconsistent with other data features, as shown in Table 6. The table reports activity and outcome measures for our baseline markets. The first column of table 6 reports the experimental data, and the second column reports the average across 10,000 simulations of Two-RDA. The volume of market activity in Two-RDA is similar to what we observe in our data: the average number of offers in the data is 59, and 54 in Two-RDA. Despite the slightly higher number of offers, in the data, offers are less likely to be accepted, translating into a lower number of matches than in Two-RDA. Our experimental markets converge to stable matchings more often than Two-RDA, which yields stability in less than half of the simulations.[34] When Two-RDA does converge to a stable outcome, the median is the modal outcome both at the market and the individual-level. However, conditional on achieving stability, the frequency of the median stable matching is still far lower than in the data.

In the Appendix, we also report results from a variation of Two-RDA that captures agents' cardinal incentives to form blocking pairs. Specifically, we run simulations in which an active agent a is chosen to be a proposer with probability proportional to g_a or $\exp(\lambda g_a)$, where g_a is the gain a would obtain if their next proposal were accepted, and $\lambda > 0$ is a fixed parameter. The resulting simulations generate a higher frequency of stable matchings, but they fail to replicate other aspects of the data, in particular the volume of offers. Table 9 in the Appendix reports the results for the proportional case, and for different values of λ.

Random Paths to Stability (RPS). This model assumes that blocking pairs are formed at random. Starting from some matching at time t, say μ_t, the set of all blocking pairs is

[33] See the Appendix for a formal description of the algorithm.

[34] In contrast with the regular DA, Two-RDA is not guaranteed to result in a stable matching.

tabulated, and one is formed at random. That is, the corresponding color and food in that blocking pair get matched and their partners in μ_t (if they exist) are unmatched. The resulting matching is μ_{t+1}, and the process continues iteratively. Roth and Vate (1990) proved that these dynamics converge to a stable matching with probability one.[35]

RPS does a poor job at predicting the offer volume and the distribution of ultimate stable matchings that we see in our experimental data. In order to give naïve dynamics such as RPS a chance at explaining our data, we consider versions of RPS in which the probability a blocking pair forms depends on the welfare gain for the agents participating.

In the Appendix, we also report results from versions of RPS in which the probability that a blocking pair forms depends on the welfare gain of blocking partners. In particular, we consider a version in which the probability a blocking pair forms is proportional to the sum of payoff gains of the blocking partners. We also consider a version in which that probability is logistic. Namely, the probability any blocking pair (f, c) forms is proportional to $\exp(\lambda g_{f,c})$, where $g_{f,c}$ is the sum of f and c's payoffs from matching and λ is a sensitivity parameter.

Table 10 in the Appendix reports results from simulations of the original RPS, as well as its two variants, including alternative values of the sensitivity parameter λ of the logistic variant. The logistic model seems to fit the data best, perhaps due to the additional degree of freedom its sensitivity parameter affords. The third column of Table 6 reports the results of the simulated markets with this logistic variant with $\lambda = 0.0175$, chosen to approximately match the number of offers made to blocking partners in our experimental markets.

Results from the logistic variant of RPS are consistent with some features of the data. By design, outcomes are stable. Furthermore, the median stable matching is the modal outcome. Nonetheless, the frequency of median stable matching is substantially lower than what we

[35] Rudov (2022) shows that, in fact, this prediction cannot be refined further: under mild conditions, any unstable matching can reach any stable matching through these dynamics.

observe in the data and the dynamics of offers and responses is different.[36] As Table 6 shows, there are fewer repeated offers and matches, somewhat more downward offers and more offers exhibiting skips.

Random Best Response (RBR). This dynamic model, due to Ackermann et al. (2011), is an alternative to RPS in which, instead of randomly choosing blocking pairs, a random agent is selected at each stage. That agent's most preferred blocking pair is then formed, if one exists. Specifically, given a matching μ_t, we tabluate the set of agents that have at least one blocking partner, and choose one at random. The next matching, μ_{t+1}, is obtained by matching the chosen agent with their most preferred blocking partner. RBR converges with probability one to a stable matching, just as RPS.

We consider versions of RBR in which cardinal payoff information is allowed to play a role. The standard version of RBR does not and, as with RPS, is not in line with the offer volume aspect of our data. We consider two variants analogous to those we consider for RPS. Let g_a denote the net gain agent a would obtain if matched to the most preferred blocking partner. In the two variants of RBR, we choose each agent a who is part of a blocking pair with probability proportional to g_a, or $\exp(\lambda g_a)$, where λ is a sensitivity parameter. Table 11 in the Appendix reports simulation results for these different variants of RBR, allowing for an array of λ values. The fourth column in Table 6 reports results for the logistic variant with $\lambda = 0.00425$, which yields simulated offer volumes to blocking partners resembling those observed in our experimental markets.

As can be seen from Table 6, the logistic variant of RBR yields results similar to those of RPS. Its outcomes are consistent with some global features of our data. As with RPS, it yields stable outcomes by design. Additionally, it generates a high frequency of median stable

[36] Notably, the frequency of median stable matchings does not depend on the sensitivity parameter λ.

matchings, albeit at a substantially lower frequency than that observed in the data. However, like RPS, this version of RBR fails to replicate many features of the dynamics we observe.

Taken together, while neither of the three classes of dynamic models we consider replicate our experimental findings precisely, they shed some light into the dynamics underlying our experimental results. First, both cardinal and ordinal incentives matter in the experiment. While agents frequently go down their preference lists as in the DA, they also focus on exploiting blocking opportunities, especially the most profitable. Second, the amount of repeated offers observed in the experiment cannot be explained by exploiting blocking opportunities alone. Third, while the median stable matching is the modal outcome under each of the models, none predict it at frequencies as high as those appearing in the data.

7.3 Individual Dynamics

As Figure 2 and our prior analysis indicated, offers appear to be driven by the volume of potential blocking pairs in the market. We now analyze the main determinants driving market participants to make and accept offers.

Determinants of Offer Targets. Offers are driven by yield—the chance that the offer will be accepted—and payoff, as well as by past offers and matches, as Table 7 shows. The table reports the results of multiple conditional logits explaining offer targets. We use the following regressors. Proposer's and receiver's *Payoff Advantage (PA)* is the change in payoff to the agent if the offer is accepted. We include it as a dummy, indicating whether it is positive (i.e., the proposer or receiver find the match profitable given the payoff of their current match), and also differentiate gains from losses by splitting the variable into its positive and negative components. In addition, we include the receiver's rank in the proposer's rank-order list (higher rank means less preferred), whether the receiver is currently matched, whether the pair formed by the proposer and the receiver are a blocking pair, whether the pair has been matched previously, the number of proposals the proposer has made to the receiver previously, whether the pair formed by the proposer and the receiver is a stable pair, and whether the proposal is downward, Gale-Shapley, or skips someone (as defined above). We estimate several specifications, and test their fit in-sample and out-of-sample using the mean-squared error, the percentage of choices in which the alternative chosen in the data has the maximum predicted choice probability, and the average probability of correctly predicting the data.

The average marginal effects reported in Table 7 confirm some of the observations already made. First, proposers not only take their own preferences into account, but also those of the receivers. The probability of proposing to a blocking partner is around 0.12 higher than to someone who is not. Likewise, proposers are more likely to target receivers that find it more profitable to match with them. Second, proposers are more likely to propose to receivers who they have already proposed to in the past, and, especially, to receivers whom they have already been matched with. Third, proposers are more likely to make downward or Gale-Shapley proposals. Last, and perhaps surprisingly, proposers are more likely to make proposals to receivers who are stable partners. As already noted, it is extremely challenging to identify stable partners in our experimental markets. The draw of stable partners indicates that stability may be intrinsically attractive.[37]

Determinants of Offer Responses. Receivers are more likely to accept proposals that are more profitable in monetary terms, disregarding proposers' payoffs, and are more selective the more offers they receive over time, as Table 8 illustrates. The table reports the results of binary logits explaining which offers are accepted. We include similar regressors to those in Table 7, with the addition of a dummy variable capturing whether the proposal is the first

[37] In the Appendix, we also replicated our analysis here including cardinal payoff information for both proposers and receivers, see Tables 12 and 13. Results are qualitatively the same.

Table 7: Proposal Conditional Logits

	(1)	(2)	(3)	(4)	(5)	(6)
Proposer's PA > 0	0.257***	0.138***	0.112***	0.047	0.045	0.106***
	(0.036)	(0.034)	(0.033)	(0.030)	(0.030)	(0.035)
Receiver's rank (in prop's list)	-0.075***	-0.076***	-0.074***	-0.071***	-0.064***	-0.060***
	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)
Receiver is matched		-0.010	0.087***	0.102***	0.039*	0.038*
		(0.016)	(0.028)	(0.020)	(0.020)	(0.022)
Receiver's PA > 0		-0.047	-0.151***	-0.167***	-0.153***	-0.181***
		(0.054)	(0.044)	(0.046)	(0.041)	(0.047)
Are blocking pair (BP)		0.271***	0.276***	0.318***	0.286***	0.304***
		(0.044)	(0.042)	(0.044)	(0.040)	(0.050)
$\max\{\textit{Rec's PA}, 0\}$			0.033***	0.035***	0.019***	0.027***
			(0.006)	(0.005)	(0.005)	(0.005)
$\min\{\textit{Rec's PA}, 0\}$			0.050***	0.047***	0.034***	0.034***
			(0.008)	(0.008)	(0.007)	(0.008)
Matched previously				0.054***	0.037***	0.067***
				(0.011)	(0.010)	(0.012)
# previous proposals (prop. to rec.)				0.018***	0.018***	0.026***
				(0.003)	(0.003)	(0.007)
Are stable partners (SP)					0.078***	0.087***
					(0.007)	(0.008)
Proposal is downward						0.085***
						(0.023)
Proposal is Gale-Shapley						0.109***
						(0.035)
Proposal skips someone						-0.006
						(0.041)
Observations	31,506	31,506	31,506	31,506	31,506	31,506
Adj. R^2	0.311	0.374	0.389	0.410	0.420	0.442
MSE *(sample)*	8.699	7.949	7.757	7.586	7.487	7.234
MSE *(2-fold* $\times$ *valid)*	8.701	7.958	7.770	7.608	7.513	7.257
MSE *(future* $\mid$ *present)*	8.207	7.493	7.260	7.260	7.142	6.737
MSE *(present* $\mid$ *future)*	9.341	8.481	8.271	8.131	8.054	7.905
$\%CorrMaxCP$ *(sample)*	44.501	55.523	56.464	57.115	55.885	58.538
$\%CorrMaxCP$ *(2-fold* $\times$ *valid)*	44.551	55.195	56.065	57.265	55.970	58.603
$\%CorrMaxCP$ *(future* $\mid$ *present)*	48.800	59.657	60.400	61.029	58.400	62.686
$\%CorrMaxCP$ *(present* $\mid$ *future)*	41.361	51.294	51.836	53.506	52.129	54.007
$Avg\,\mathbb{P}(OK\,Pred)$ *(sample)*	34.097	39.338	40.402	42.214	43.047	45.277
$Avg\,\mathbb{P}(OK\,Pred)$ *(2-fold* $\times$ *valid)*	34.108	39.340	40.394	42.173	42.991	45.232
$Avg\,\mathbb{P}(OK\,Pred)$ *(future* $\mid$ *present)*	33.736	39.667	41.212	42.747	43.896	46.761
$Avg\,\mathbb{P}(OK\,Pred)$ *(present* $\mid$ *future)*	35.215	39.745	40.547	41.433	41.842	43.694

Notes: Table reports average marginal effects of conditional logits. The response variable indicates the receiver of every proposal in the data. Standard errors are clustered at participant level. *, **, and *** stand for 90%, 95%, and 99% confidence levels, respectively. The table also reports the mean-squared error (MSE) of the predicted choice probability, percentage of choices in which the predicted probability of the alternative chosen in the data is the greatest among all alternatives ($\%CorrMaxCP$), and the average probability of correctly predicting the data ($Avg\,\mathbb{P}(OK\,Pred)$). Each is computed in the estimation sample and out of the sample using: random two-fold cross-validation, predicting the final five rounds with the first five rounds, and the first five rounds using the final five. See the Appendix for more details.

proposal made by the proposer to the receiver, and the total number of proposals the receiver has received so far.

Larger monetary gains to receivers are associated with higher acceptance probabilities, even for proposers who are ranked similarly, echoing the importance of cardinal payoffs in our markets. History matters for receivers: they are more likely to accept offers from proposers with whom they had already matched, but less likely to accept from those whom had already rejected. In general, receivers are less likely to accept offers as the receive more of them. Similar to proposers, receivers seem to be drawn to proposers who are stable partners, accepting their offers at higher rates.[38]

8 Discussion and Conclusions

We provide an empirical benchmark for the performance of decentralized matching markets using an array of lab experiments. We document three main findings. First, decentralized markets very often culminate in stable matchings on their own, absent any centralized intervention. The stabilization process is quick, in terms of both time and market activity. Second, the median stable matching has very strong drawing power and is frequently selected. Third, cardinal incentives impact the distribution of selected matchings that are not the median. Roughly speaking, the side of the market that has "more to lose" from forgoing their favorite stable matching, is more likely to implement it. We also describe the dynamics leading to stability. By and large, participants form successive blocking pairs. However, participants are perhaps more sophisticated than suggested by the naïve dynamics the literature has proposed, with proposers strategically targeting receivers who value them highly and responders taking into account past market activity.

[38]Adding participant fixed effects to these regressions does not alter the results and decreases their predictive power. This suggests that there is no significant unobserved heterogeneity across participants acting as receivers; see Table 14 in the Appendix.

Table 8: Acceptance Binary Logits

	(1)	(2)	(3)	(4)	(5)	(6)
Receiver's PA > 0	0.761***	0.586***	0.568***	0.850***	0.718***	0.698***
	(0.032)	(0.042)	(0.042)	(0.190)	(0.164)	(0.162)
Receiver is matched	0.208***	0.408***	0.240***	0.238***	0.239***	0.213***
	(0.028)	(0.033)	(0.037)	(0.037)	(0.033)	(0.031)
$\max\{$*Rec's PA*$, 0\}$		0.078***	0.038***	0.038***	0.037***	0.034***
		(0.008)	(0.010)	(0.010)	(0.008)	(0.008)
$\min\{$*Rec's PA*$, 0\}$		0.069**	0.032	0.029	0.003	-0.007
		(0.028)	(0.029)	(0.030)	(0.021)	(0.017)
Proposer's rank (in rec's list)			-0.037***	-0.037***	-0.030***	-0.016***
			(0.005)	(0.005)	(0.004)	(0.004)
Proposer's PA > 0				0.230	0.135	0.112
				(0.145)	(0.125)	(0.126)
Are blocking pair (BP)				-0.292	-0.189	-0.196
				(0.187)	(0.163)	(0.163)
Matched previously					0.154***	0.135***
					(0.020)	(0.021)
First proposal (prop. to rec.)					-0.009	-0.003
					(0.029)	(0.029)
# previous proposals (prop. to rec.)					-0.039*	-0.041**
					(0.020)	(0.020)
# previous proposals (total to rec)					-0.018***	-0.015***
					(0.005)	(0.005)
Are stable partners (SP)						0.111***
						(0.013)
Observations	4,147	4,147	4,147	4,147	4,147	4,147
Adj. R^2	0.317	0.351	0.370	0.371	0.397	0.410
MSE *(sample)*	15.407	14.307	13.702	13.690	13.117	12.794
MSE *(2-fold* $\times$ *valid)*	15.409	14.318	13.717	13.706	13.235	12.927
MSE *(future* $\mid$ *present)*	16.039	14.786	14.154	14.136	13.832	13.333
MSE *(present* $\mid$ *future)*	14.987	14.134	13.540	13.557	12.967	12.835
$\%CorrMaxCP$ *(sample)*	74.608	79.551	80.420	80.395	81.215	81.601
$\%CorrMaxCP$ *(2-fold* $\times$ *valid)*	74.605	79.560	80.420	80.394	80.762	81.590
$\%CorrMaxCP$ *(future* $\mid$ *present)*	73.672	80.069	79.269	79.326	79.783	81.154
$\%CorrMaxCP$ *(present* $\mid$ *future)*	75.292	78.923	80.676	80.634	81.845	81.135
$Avg\,\mathbb{P}(OK\,Pred)$ *(sample)*	69.187	71.301	72.473	72.493	73.614	74.303
$Avg\,\mathbb{P}(OK\,Pred)$ *(2-fold* $\times$ *valid)*	69.181	71.296	72.458	72.480	73.550	74.237
$Avg\,\mathbb{P}(OK\,Pred)$ *(future* $\mid$ *present)*	67.867	69.723	70.952	70.989	71.871	72.604
$Avg\,\mathbb{P}(OK\,Pred)$ *(present* $\mid$ *future)*	70.159	72.592	73.695	73.680	74.986	75.625

Notes: Table reports average marginal effects of binary logits. The response variable is an indicator of whether a proposal was accepted. Standard errors are clustered at participant level. *, **, and *** stand for significantly different to zero at a 90%, 95%, and 99% confidence level, respectively. The table also reports the mean-squared error (MSE) of the predicted probability, percentage of choices in which the predicted probability of the alternative chosen in the data is the greatest among all alternatives ($\%CorrMaxCP$), and the average probability of correctly predicting the data ($Avg\,\mathbb{P}(OK\,Pred)$). Each is computed in the estimation sample and out of the sample using: random two-fold cross-validation, predicting the final five rounds with the first five rounds, and the first five rounds using the final five. See the Appendix for more details.

Our findings are important from a market design perspective. While our experimental markets often culminate in a stable matching, it is rarely an extremal stable matching. If there are reasons for a market designer to desire extremal stable matchings—say, favoring young interns in a labor market environment or students in a school-choice setting—a centralized clearinghouse such as the commonly used DA may be beneficial. Indeed, most real-world stable clearinghouses implement an extremal stable matching for submitted preferences. In addition, while our decentralized markets took relatively little time and few offers to converge, scaling those durations, offer volumes, and necessary market turnover to large matching markets may come at a substantial efficiency cost. Centralization clearly allows a market designer to establish a matching rapidly.

Our results also have implications for the theory of dynamic stabilization in matching markets. Existing models that generate stable matchings through the sequential formation of blocking pairs (such as Roth and Vate, 1990, or Ackermann et al., 2011), or which prescribe an order for the proposals made by agents (such as the DA algorithm), do not match basic features of our data. There is therefore room for further theoretical work that provides foundational guidance on the selection of stable matchings in decentralized markets and takes into account cardinal, not only ordinal, assessments of partners.

Our experiments are designed as a benchmark for decentralized interactions in an idealized setting, allowing an examination of the cooperative theoretical predictions. Nonetheless, we hope our design opens the door for further investigations that take into account various frictions frequently present in applications: offer costs, rematching costs, incomplete information, and the like.

References

Ackermann, H., P. W. Goldberg, V. S. Mirrokni, H. Röglin, and B. Vöcking (2011): "Uncoordinated two-sided matching markets," *SIAM Journal on Computing*, 40, 92–106. 4, 33, 34, 37, 43, 53

Agranov, M., A. Dianat, L. Samuelson, and L. Yariv (2022): "Paying to Match: Decentralized Markets with Information Frictions," *mimeo*. 5

Baccara, M. and L. Yariv (2022): "Dynamic matching," *Online and Matching-Based Market Design*. 33

Bergstrom, C. T., T. C. Bergstrom, and R. J. Garratt (2013): "Choosing partners: A classroom experiment," *The Journal of Economic Education*, 44, 47–57. 6

Chen, Y. and T. Sönmez (2006): "School choice: an experimental study," *Journal of Economic Theory*, 127, 202–231. 6

Cheng, C. T. (2008): "The generalized median stable matchings: finding them is not that easy," in *Latin American Symposium on Theoretical Informatics*, Springer, 568–579. 9

Echenique, F., R. Gonzalez, A. J. Wilson, and L. Yariv (2022): "Top of the Batch: Interviews and the Match," *American Economic Review: Insights*, 4, 223–38. 1

Echenique, F., A. J. Wilson, and L. Yariv (2016): "Clearinghouses for two-sided matching: An experimental study," *Quantitative Economics*, 7, 449–482. 6

Featherstone, C. R., E. Mayefsky, and C. D. Sullivan (2022): "Why do some clearinghouses yield stable outcomes? Experimental evidence on out-of-equilibrium truth-telling," *mimeo*. 6

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. 6

Fehr, E. and S. Gächter (2000): "Fairness and retaliation: The economics of reciprocity," *Journal of Economic Perspectives*, 14, 159–181. 24

Ferdowsian, A., M. Niederle, and L. Yariv (2022): "Decentralized matching with aligned preferences," *mimeo*. 5

Gale, D. and L. S. Shapley (1962): "College admissions and the stability of marriage," *The American Mathematical Monthly*, 69, 9–15. 3, 4, 7, 33

Gusfield, D. and R. W. Irving (1989): *The stable marriage problem: structure and algorithms*, MIT Press. 2, 9

Haeringer, G. and M. Wooders (2011): "Decentralized job matching," *International Journal of Game Theory*, 40, 1–28. 5

Hakimov, R. and D. Kübler (2021): "Experiments on centralized school choice and college admissions: a survey," *Experimental Economics*, 24, 434–488. 5

Harrison, G. W. and K. A. McCabe (1996): "Expectations and fairness in a simple bargaining experiment," *International Journal of Game Theory*, 25, 303–327. 6

Haruvy, E. and M. U. Ünver (2007): "Equilibrium selection and the role of information in repeated matching markets," *Economics Letters*, 94, 284–289. 6

Irving, R. W. and P. Leather (1986): "The complexity of counting stable marriages," *SIAM Journal on Computing*, 15, 655–667. 9

Kagel, J. H. and A. E. Roth (1995): *The Handbook of Experimental Economics*, vol. 1, Princeton University Press. 6

——— (2000): "The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment," *The Quarterly Journal of Economics*, 115, 201–235. 5

——— (2020): *The Handbook of Experimental Economics*, vol. 2, Princeton University Press. 24

Nalbantian, H. R. and A. Schotter (1995): "Matching and efficiency in the baseball free-agent system: an experimental examination," *Journal of Labor Economics*, 13, 1–31. 5

Niederle, M. and A. E. Roth (2009): "Market culture: How rules governing exploding offers affect market performance," *American Economic Journal: Microeconomics*, 1, 199–219. 6

Pais, J. (2008): "Incentives in decentralized random matching markets," *Games and Economic Behavior*, 64, 632–649. 5

Pais, J. and Á. Pintér (2008): "School choice and information: An experimental study on matching mechanisms," *Games and Economic Behavior*, 64, 303–328. 6

Pais, J., A. Pintér, and R. F. Veszteg (2020): "Decentralized matching markets with(out) frictions: a laboratory experiment," *Experimental Economics*, 23, 212–239. 6

Plott, C. R. and V. L. Smith (1978): "An experimental examination of two exchange institutions," *The Review of Economic Studies*, 45, 133–153. 6

Roth, A. E. and E. Peranson (1999): "The redesign of the matching market for American physicians: Some engineering aspects of economic design," *American Economic Review*, 89, 748–780. 8

Roth, A. E. and M. Sotomayor (1990): *Two-Sided Matching: A Study in Game-Theoretic Modelling and Analysis*, Econometric Society Monographs. 5

Roth, A. E. and J. H. V. Vate (1990): "Random paths to stability in two-sided matching," *Econometrica*, 1475–1480. 4, 33, 34, 36, 43, 52

Rudov, K. (2022): "Fragile Stable Matchings," *mimeo.* 36

Smith, V. L. (1962): "An experimental study of competitive market behavior," *Journal of Political Economy*, 70, 111–137. 6

Teo, C.-P., J. Sethuraman, and W.-P. Tan (2001): "Gale-shapley stable marriage problem revisited: Strategic issues and applications," *Management Science*, 47, 1252–1267. 8