Paying to Match: Decentralized Markets with Information Frictions*

Marina Agranov[†] Ahrash Dianat[‡] Larry Samuelson[§] Leeat Yariv[¶]

December 2, 2025

Abstract

We experimentally study decentralized one-to-one matching markets with transfers. We vary the information available to participants, complete or incomplete, and the surplus structure, supermodular or submodular. Several insights emerge. First, while markets often culminate in efficient matchings, stability is more elusive, reflecting the difficulty of arranging attendant transfers. Second, incomplete information and submodularity present hurdles to efficiency and especially stability; their combination drastically diminishes stability's likelihood. Third, matchings form "from the top down" in complete-information supermodular markets, but exhibit many more and less-obviously ordered offers otherwise. Last, participants' market positions matter far more than their dynamic bargaining styles for outcomes.

Keywords: Matching Markets, Transfers, Incomplete Information, Experiments.

JEL: C78, C92, D02, D47.

*We thank Jeongbin Kim and Pellumb Reshidi for superb research assistance. We are also grateful to the Editor, Noam Yuchtman, and two anonymous reviewers for their many helpful suggestions. We acknowledge financial support from the International Foundation for Research in Experimental Economics and from the NSF, through grant SES 1629613.

[†]Division of the Humanities and Social Sciences, California Institute of Technology, Email: magranov@hss.caltech.edu

[‡]Department of Economics, University of Essex, Email: a.dianat@essex.ac.uk

[§]Department of Economics, Yale University, Email: larry.samuelson@yale.edu

[¶]Department of Economics, Princeton University, Email: lyariv@princeton.edu

1 Introduction

1.1 Overview

A large and influential literature studies matching markets with complete information. It assumes individuals know not only their own preferences over possible allocations, but also all others' preferences. This assumption is often made for tractability and is reasonable for some applications. In large school choice systems, for example, historical data on admission patterns might be persistent and therefore predictable. Nevertheless, there are numerous applications in which *incomplete* information regarding others' preferences, at times even one's own, might be present. Newly-minted doctors seeking a residency in a specific specialty may not be privy to the idiosyncratic preferences of both their peers and the hospitals they interview with. Similarly, fresh academic Ph.D.'s may not be fully informed of other applicants' preferences or of universities' particular research and teaching needs. Outside the labor market, potential adoptive parents are unlikely to know the preferences over child characteristics of other prospective adopters, or of the birth mothers placing children for adoption. In marriage markets, an entire consulting industry builds on the premise that individuals frequently misread others' preferences, even after lengthy interactions. While incomplete information is prevalent across these applications, there is a paucity of empirical data to elucidate its impacts.

We report results from an array of lab experiments in which participants match through decentralized interactions, as in the examples above. Our treatments vary the surplus structure and the initial information available to market participants about others' preferences. We illustrate the substantial effect limited information has on outcomes and the regularities it produces.

Since many markets operate through transfers—for instance, labor markets often entail targeted wages and housing markets generate exchanges via prices—we consider decentralized matching markets allowing for transfers. While common wisdom suggests that transfers assist in equilibrating economic systems, they also allow for various bargaining frictions to emerge when negotiations are unconstrained. Differences in bargaining styles, failures to follow up with appropriate offers, and the like could, in principle, affect the generated match surpluses and their split between parties. Furthermore, in practice, implementing the "right" transfers might be challenging: agents often interact with a small set of potential partners and have limited information on what others should reasonably expect to receive. The challenge of finding the "right" transfers is compounded by agents' incomplete information about others' preferences. Our design allows us to discern how transfers evolve and get determined.

Experimental Design. While our study speaks to many applications, field data are difficult to gather. Rarely do we know market participants' true preferences, their detailed interactions, or the information available to them. The lab setting allows us to control for different market features and observe how individuals respond.

We consider two-sided one-to-one markets in which agents can make match offers to one another. Agents in the market are characterized by their "types," which can stand for education level or expertise in labor markets, age in child-adoption processes, or pizzazz in the marriage setting. All agents prefer to be matched with higher-type agents: they generate a greater match surplus.

Going back to Becker (1973), the matching literature has often assumed match surpluses exhibit either supermodularity or submodularity; see Chade, Eeckhout, and Smith (2017) and Chiappori (2017) for literature surveys. We accordingly consider two surplus structures: supermodular and submodular. In supermodular markets, higher types experience a greater marginal increase in generated surplus from matching with higher types. Consequently, the utilitarian efficient matching is positively assortative. In submodular markets, the reverse holds: it is the lower-type agents who experience the greatest marginal benefit from increasing their partner's quality. In such markets, utilitarian efficiency implies negatively assortative matchings.

In our experimental markets, interactions are decentralized: individuals can make offers to one another in an unscripted fashion, with the sole restriction that any agent can make or accept only one offer at a time. An offer indicates a split of the match surplus. Accepted offers generate tentative matches—they can be broken by making a different offer that gets accepted, or by accepting an alternative offer. A market terminates after 30 seconds of inactivity. We use this matching protocol for three reasons. First, it echoes the narrative underlying the cooperative underpinnings of stable matchings with complete information: any coalition can be formed throughout a market's duration. Second, it provides participants ample opportunities to learn about others' match qualities, which is important for some of our analysis. Last, it captures features of many real-world matching settings, where pairs negotiate with a variety of partners before settling on a permanent match. For example, employees can negotiate with one firm and switch to negotiations with another, home buyers inspecting a property can decide to forgo that property and inspect others, and individuals seeking a life partner often date multiple individuals sequentially.[1]

In our complete-information treatments, participants know the surplus generated by each pair of participants in the market. In our incomplete-information treatments, participants enter the market knowing the profile of match surpluses they can generate. That is, they know their own type. They do not, however, know the types of other participants. In particular, at the start of the market, they do not know who will help them generate which surplus. Learning occurs through market interactions. Whenever an offer is accepted, both partners learn the match surplus they jointly generate and the resulting payoffs they receive.[2] The full market observes the creation of matches, as well as the surplus that recipients of accepted offers receive, but not the total surplus such matches generate.

Empirical Regularities. Several insights emerge from our analysis. As we describe in our literature review below, prior experimental work supports the impression that, at least without transfers, decentralized interaction yields stable outcomes. Our results show that matching markets with transfers often produce utilitarian efficient matchings, but stable outcomes are less common.[3] Thus, there is a sense in which the common wisdom on the benefits of transfers holds water in our data, albeit at the cost of stability. This effect is particularly salient when the surplus is submodular. Specifically, with complete information, 94% of our supermodular markets culminate in efficient, positively assortative outcomes and 78% of market results are stable. When preferences are submodular, only 73% of our markets yield efficient, negatively assortative outcomes and a mere 14% of ultimate outcomes are stable.

When information is incomplete, the theoretical literature offers little guidance as to the dynamics and outcomes one might expect; see our review of the literature. In particular, there is little work inspecting the information revealed through the operations of a market that starts

[1] While markets' termination depends on all market participants ceasing activity, any particular pair can stop searching for alternatives sooner by maintaining their pairing till the market ends.

[2] Offers still entail the specification of a transfer. Absent information on match surpluses, certain offers can lead to negative payoffs. We allow participants to cancel such offers ex-post. We do not see participants "gaming" the experiments through this option.

[3] Utilitarian efficient matchings maximize the sum of match surpluses. Stability with transfers corresponds to a matching and a profile of transfers such that no pair of agents can find a transfer that, if implemented, would make their match desirable relative to what they already receive. It is well known that stable matchings are utilitarian efficient (e.g., Roth and Sotomayor, 1990).

with incomplete information. In simple incomplete-information markets such as ours, where participants enter with considerable information and every accepted or rejected offer yields more information, interactions could plausibly lead to complete information. In Section 3, we introduce our notion of incomplete-information stability. Even though this notion is arguably permissive, we show that under conditions frequently satisfied in our data, incomplete-information stability indeed coincides with complete-information stability. We consider alternative stability notions in Section 6.

In our data, incomplete information and submodularity constitute substantial obstacles to both efficiency and stability, as shown in Section 4. Indeed, with incomplete information, in markets with a supermodular surplus structure, 84% of outcomes are efficient, while 54% are stable. With submodular surpluses, 39% of outcomes are efficient, and a mere 3% are stable. In other words, the combination of incomplete information with submodularity nearly eliminates the likelihood of achieving stable outcomes, despite participants' exposure to information throughout markets' operation. This observation opens the door for richer theories of matching with incomplete information: ones that allow for market learning, but still yield outcomes that are not complete-information stable.

The analysis of offer dynamics in Section 5 provides a lens into how market outcomes are generated. When information is complete, markets governed by supermodular surpluses "match from the top," in the sense that high-type agents (who generate the greatest surplus) match first, followed by medium-type agents, and finally low-type agents.[4] In contrast, the sequenceing of matches is more haphazard in submodular markets. Many more offers are made throughout those market operations. Early matches that "stick" are often between high-type agents, and generate a barrier to achieving market-wide efficiency. Medium-type agents follow suit, while low-type agents tend to match late, regardless of their partner's type.

When information is incomplete, regardless of the underlying surplus structure, initial matches are often transient and broken later in the market's operations as participants learn. This leads to far more market activity. Matches between incongruent types, low and high, occur earlier. Interestingly, the enhanced market activity we observe implies that, by the end of our markets' operations, most participants learn with whom they can generate which surplus. However, even

[4]This sequence resembles the Marshallian path studied by Plott, Roy, and Tong (2013).

matches formed later in our incomplete-information markets, by agents who are then well informed, do not echo the patterns seen in their complete-information counterparts.

Why do outcomes not mimic those in our complete-information markets, despite effectively complete information? In Section 6, we consider several possible channels and show that our data are consistent with simple satisficing heuristics of participants. Based on their type, participants aim at a certain payoff and actively search for partners only below their satisficing threshold. This observation echoes evidence documented in individual decision-making; see Caplin, Dean, and Martin (2011). In our setting, satisficing emerges through dynamic market interactions.

A large literature, scientific and popular, suggests the importance of bargaining strategies in determining outcomes—for reviews, see Bertrand (2011) and Shell (2006). One may wonder whether our participants' individual tendencies help explain their obtained payoffs. We indeed see large variation in individuals' bargaining styles, in terms of offer volume, number of matches created throughout a market, how demanding offers are, and the timing of interactions. Nonetheless, in Section 7, we show that none of these has any substantial explanatory power for outcomes. Ultimately, market features and one's position in the market—namely, the type—appear to explain the lion's share of the payoffs participants receive. This result may be comforting to some: while individual characteristics and access to opportunities for self-improvement may greatly affect one's consequences, they do so only insofar as they alter one's position, or "quality," in the market. Once in the market, outcomes appear resilient to individual variations that go beyond participants' contributions to the surpluses they can generate.

Implications. Taken together, our results carry several important implications for market design. First, when transfers are available, stability through decentralized interaction is not guaranteed and depends crucially on the underlying preferences in the market. In particular, centralized clearinghouses may be especially beneficial when preferences are submodular. Second, incomplete information presents a substantial hurdle for achieving stability. Interventions directed at alleviating information frictions may therefore prove useful. Such interventions could entail allowing for more communication between participants, introducing public releases of information about market attributes, etc. Alternatively, centralization of the matching process may be particularly important when participants are not perfectly informed. The existing literature rarely considers clearinghouses with imperfectly-informed participants, and the consideration of those may be necessary. Last, educating participants to achieve their maximal potential "quality" as they enter a matching market may be of far greater importance than educating participants on how to negotiate their interactions once in the market.

1.2 Related Literature

A large theoretical literature considers matching markets with complete information, with or without transfers; see Roth and Sotomayor (1990). Incomplete information poses many challenges to the analysis of matching markets. Even the basic stability notion is difficult to translate when agents are not fully informed of everyone's preferences. Pairs that would block a matching with complete information—pairs of agents that prefer one another over their allocated match—may not form if agents cannot discern who would prefer a match with them. Furthermore, any attempt to form a blocking pair allows other participants to partially infer the preferences of agents involved, possibly others' when those are correlated.

Starting from Liu, Mailath, Postlewaite, and Samuelson (2014), several recent papers suggest stability notions for incomplete information markets, mostly restricting attention to incomplete information only on one market side. Bikhchandani (2017) considers matching markets without transfers. Liu et al. (2014), and relatedly Alston (2020) and Liu (2022), consider matching with transfers as we do. At the heart of the stability notions in these papers is the requirement that any agent's "deviation" from a matching, with its attendant transfers, ultimately leaves her (weakly) worse off. The stability notions in these papers are defined with respect to a given matching and a fixed specification of the uncertainty agents have about others' types. They do not consider the possibility that the match formation process may reveal information, rendering the specification of uncertainty endogenous and possibly leading to effectively complete information.[5]

Stability notions are interesting due to the presumption that, in frictionless decentralized markets, agents would ultimately converge to a stable outcome. Nonetheless, only a few papers explicitly model the decentralized process that generates such stable outcomes. Roth and Vate (1990) offer non-strategic dynamics that yield stable outcomes, whereby at each stage, a random blocking pair is implemented. Complete information is implicitly assumed, as all blocking pairs are

[5]Immorlica, Leshno, Lo, and Lucier (2020) suggest a regret-free notion of stability without transfers when one side of the market is subject to incomplete information and can perform costly search. Matching mechanisms are recast as price-discovery processes, illustrating a class that yields approximately regret-free stable outcomes.

considered at every step. Ferdowsian, Niederle, and Yariv (2025) consider a decentralized market game in which firms and workers interact in a dynamic fashion. They illustrate the hurdles incomplete information or time frictions generate in establishing stability as a property of equilibrium outcomes. Chen and Hu (2020) focus on incomplete-information settings and construct an adaptive learning process leading to stable outcomes as they define them.

Hakimov and Kübler (2021) present a general overview of the experimental matching literature. Several papers in that literature share features with our design.[6] Concurrent work by He, Wu, Zhang, and Zhu (2024) considers decentralized matching with complete information, allowing for transfers between participants. He et al. (2024)'s design resembles that of our complete-information treatment, although their experimental markets have a pre-specified time of operation.[7] In line with our results, they observe far more stable outcomes in assortative than in anti-assortative markets. In contrast with our results, they observe a substantial fraction of unmatched individuals, potentially due to the limited-market horizon. Outside the matching context, Agranov and Elliott (2021) identify the important impacts of transfers on bargaining outcomes.

Nalbantian and Schotter (1995) report on experiments emulating the market for professional baseball players in the free-agent years. They analyze several allocation procedures with transfers, including a decentralized procedure in which agents are informed of their own bargaining outcomes, but not of others'. Pais, Pintér, and Veszteg (2020) study decentralized markets with multiple stable matchings. There are no transfers and markets operate over a fixed duration. In some treatments, participants are privately informed of their preferences. Incomplete information in their setting does not affect the stability or the efficiency of the final outcome. However, similar to what we observe, it boosts market activity. Niederle and Roth (2009) study unraveling by considering an incomplete-information setting in which one side of the market makes proposals to the other side over three experimental periods and information is exogenously and incrementally released. They illustrate the impact of proposal structure on matching outcomes and efficiency. Incomplete information has also been considered in centralized markets, mostly focusing on school-choice mechanisms and, therefore, absent transfers. See, for example, Chen and Sönmez (2006), Pais and Pintér (2008), and work that followed.[8]

[6] There are also several papers studying decentralized matching with complete information and no transfers, see Echenique, Robinson-Cortés, and Yariv (2025) and references therein.

[7] In their design, participants whose match is broken in the last 15 seconds can make offers within 15 seconds.

[8] There is also recent empirical work that highlights the consequences of incomplete information in school choice,

Table 1: Match Surpluses (Efficient Allocation in Bold)

	Kiwi	Mango	Plum
Blue	**8**	16	24
Crimson	16	**32**	48
White	24	48	**72**

(a) Supermodular

	Kiwi	Mango	Plum
Blue	8	32	**56**
Crimson	32	**48**	64
White	**56**	64	72

(b) Submodular

2 Experimental Protocol

In our experiments, participants engage in a sequence of decentralized markets varying in two dimensions: the surplus structure, which is either supermodular or submodular, and the information available. In some matching markets, participants are fully informed of everyone's preferences, in others they receive only partial information. We describe each of these dimensions in turn. Sample experimental instructions are presented in the Online Appendix.

Our two-sided one-to-one matching markets are comprised of six participants, three on each side. Each participant has a role: a color or a food. These roles capture the market sides and can be metaphors for firms and workers in labor markets, men and women in the marriage market, etc. There are three payoff-relevant types for each role. Specifically, the three color types are Blue, Crimson, and White. The three food types are Kiwi, Mango, and Plum. A matching specifies which colors and foods are unmatched, as well as a mapping between the remaining colors and foods that is one-to-one: if a color type, say Blue, is assigned to a food type, say Kiwi, then the reverse holds as well: Kiwi is assigned to Blue.[9]

Table 1 depicts the match surpluses governing preferences in our markets. Panel (a) corresponds to our supermodular surpluses, while panel (b) captures our submodular surpluses. In each, the entry corresponding to each color and food pair is the match surplus generated by that pair. Unmatched agents receive a payoff of 0.

We consider two information structures. In our complete-information treatments, participants are fully informed of surpluses generated by all pairs. In particular, they can readily rank their potential partners in terms of the surpluses they would jointly generate. In our incomplete-information treatments, participants know their own type: at the market's outset, they are in-

see Kapor, Neilson, and Zimmerman (2020).

[9]In our experimental results, the incidence of unmatched participants is only 0.5% across all sessions.

formed of the three possible surpluses they can generate with the three potential partners on the other side of the market. However, they do not know which partner generates which surplus, nor which of the remaining two types have been assigned to which other agent on their own side of the market. Formally, we permute the type labels in markets and inform participants only of their own type. The structure of payoffs is transparent, as well as the randomization procedure.

The general structure of each session is as follows. The governing match surpluses, either supermodular or submodular, as well as the information structure, are fixed throughout the session. Participants engage in one practice round, followed by 10 real rounds, each corresponding to a new market (of the same sort in terms of both preferences and information), with a freshly drawn random set of participants within the session. Participants maintain their role, a color or a food, throughout the session. However, in each round, they are randomly assigned one of the three types corresponding to their role.

In our complete-information treatments, each participant observes the full surplus matrix throughout every round.[10] In our incomplete-information treatments, each participant only observes her own possible match surpluses at the start of each round.[11] As we soon explain, interactions in the market can reveal some information on who generates which surplus.

We now describe the rules of the matching protocol. In each round, participants start off unmatched. Each participant is free to make at most one match proposal to any individual of the opposite role at any given time. In the complete-information treatments, a match proposal specifies how the match surplus will be split among the two individuals—i.e., the proposer's payoff and the responder's payoff, summing up to the match surplus. In the incomplete-information treatments, a match proposal only specifies the responder's payoff, with the attendant proposer payoff revealed if the responder accepts. Thus, the proposer bears the payoff risk under incomplete information. This design choice was made to echo many applications in which the proposing side has limited information on the returns to her proposal. For instance, firms offering employment often cannot assess workers' abilities. Nonetheless, we limited the scope of risk by allowing proposers to immediately retract an offer if it turned out to generate a strict loss. For example, suppose Blue

[10]To prevent participants from reacting to cosmetic features of the surplus matrix, we shuffled the rows and columns of the surplus matrix between rounds. In particular, specific colors and foods correspond to different match-surplus profiles in each round. As a result, the efficient matchings do not always coincide with the diagonal (in the positively assortative case) or anti-diagonal (in the negatively assortative case) of the surplus matrix.

[11]For example, Blue in panel (a) of Table 1 would know that her possible match surpluses are 8, 16, and 24. However, she would not know which surplus corresponds to which match.

and Kiwi generate a match surplus of 8. If Blue offers Kiwi a payoff of 16, and Kiwi accepts, then Blue earns a payoff of -8. Since Blue earns a negative payoff, we would allow her to unilaterally cancel the match.[12] This design choice was made both for practical reasons, in order to limit the liability participants face, and also to mimic applications in which catastrophic relationships can be severed promptly. For instance, a worker who does not have their presumed credentials can be fired quickly by the firm that hired them in many labor-market settings. Whenever an offer that gives a positive payoff for both parties is accepted, the match can be dissolved only upon one of its participants making or receiving an alternative offer that is accepted.

A match generating a loss certainly conveys some information to the proposer. In general, whenever a match proposal is made and accepted, both the proposer and the responder observe their realized payoffs. Other participants observe the formation of the match and the responder's payoff. This is the main channel through which learning occurs under incomplete information.

In terms of market activity, participants are free to accept or reject match proposals they receive from other individuals. Any match proposal that is not accepted or rejected within 10 seconds automatically disappears and is effectively interpreted as a rejection. An acceptance forms a match. Participants who are currently matched can still make and receive new match proposals, with their current match persisting until such a new match proposal is accepted, at which point (only) the current match is dissolved. A round ends after 30 seconds of inactivity. In incomplete-information treatments, in order to allow participants to avoid excessive waiting, we introduced a "move on to next round" button. Whenever at least 5 participants pressed the button, the market for that group and round terminated.[13] A participant's payoff in their current match at the end of a round is their final payoff for that round. If a participant is unmatched at the end of a round, their final payoff is 0 for that round.

At the end of each session, participants also complete two risk-elicitation tasks and two altruism tasks. We use a version of the Gneezy and Potters (1997) investment task to measure risk preferences and a dictator game to measure altruism.[14] All payoffs in the experiment are

[12]In our data, 88% of negative offers were canceled by the proposers immediately. Other cases were associated with immediate new offers by the proposers. Either way, we never see negative payoffs for any participant at any round.

[13]We were particularly concerned about the sensitivity of our incomplete-information treatments to specific individuals prolonging rounds after long waits. Whenever 5 of the 6 group members are ready to halt market activity, we suspected the choice of the sixth individual to delay the end time would not affect outcomes. As will be seen, our incomplete-information treatments indeed took much longer in terms of both time and market activity.

[14]Specifically, in each of the two risk-elicitation tasks, participants were provided with 200 tokens that they had to allocate between a safe project, returning token for token, and a risky project. In the first task, each token invested in

expressed in tokens. For the main part of the experiment, 10 tokens correspond to $1. For the elicitation tasks, one token corresponds to $0.01. Participants' total earnings in the experiment include a $10 show-up payment, the sum of their final payoffs across the 10 rounds, and their payoffs from the auxiliary elicitation tasks.

The experimental sessions were run at the Experimental and Behavioral Economics Laboratory (EBEL) at UC Santa Barbara between July 2018 and October 2018. Each treatment, characterized by a combination of underlying preferences and information structure, was run in 4 sessions, each consisting of 12 participants. This generated 16 sessions overall and a total of 192 participants. Each session lasted approximately 90 minutes, and paid an average of $24.75, in addition to a $10 show-up fee. The experiment was programmed and conducted with the oTree software (Chen, Schonger, and Wickens, 2016).

3 Theoretical Preliminaries: Stability with Information Frictions

3.1 Complete Information Benchmark

In both our supermodular and submodular markets, types can be ordered. In Table 1, for example, White generates higher surplus than Crimson, who generates higher surplus than Blue, for any food. Similarly, Plum generates higher surplus than Mango, who generates a higher surplus than Kiwi, for any color. Since labels of different types are shuffled across rounds, and the two sides of our markets are fully symmetric, we refer to the types generating the highest surplus as *high* types, those generating the intermediate surplus as *medium* types, and those generating the lowest surplus as *low* types.

Let C be the set of colors and F the set of foods, with generic elements f and c. Let $\pi : C \times F \to \mathbb{R}$ identify the surplus associated with each pair of agents. A *matching* (with no unmatched agents) is a bijection $m : C \to F$.[15] An *outcome* is a matching m and a pair of transfer schemes $\tau_C : C \to \mathbb{R}$ and $\tau_F : F \to \mathbb{R}$.

the risky project returned 2.5 tokens with 50% probability, and 0 otherwise. In the second task, each token invested returned 3 tokens with 40% probability, and 0 otherwise. Altruism was elicited using two dictator-game elicitations, one with an endowment of 100 tokens and another with an endowment of 200 tokens.

[15] This formulation does not allow unmatched agents. Since in our design any partner is valuable, and agents could interact till satisfied, we view it as unrestrictive. In fact, in our experimental data, across all sessions, 0.005 of the participants are unmatched.

With complete information, an outcome (m, τ_C, τ_F) is *stable* if it satisfies

$$\begin{aligned} \tau_C(c) + \tau_F(m(c)) &= \pi(c, m(c)) \quad \forall c \in C & (1) \\ \tau_C(c), \tau_F(f) &\geq 0 \quad \forall c \in C, f \in F & (2) \\ \tau_C(c) + \tau_F(f) &\geq \pi(c, f) \quad \forall (c, f) \in C \times F. & (3) \end{aligned}$$

The first condition is a feasibility condition ensuring that transfers in a matched pair do not exceed the generated surplus. The second condition is an individual rationality requirement, asserting that no agent prefers to be unmatched. The interpretation is that if this condition were violated, the agent in question would either have not proposed or not accepted the match. The third condition is a no-blocking condition that ensures there is no pair of agents who can increase their payoffs by matching with one another and appropriately dividing the attendant surplus. The interpretation is that if this condition were violated, one of the two agents would propose to match with a transfer profile that would increase both of their payoffs.

A matching m is a *stable matching* if it is consistent with some stable outcome. In each of our markets, there is a unique stable matching. This stable matching maximizes the sum of surpluses, and so we refer to it as the utilitarian efficient, or simply efficient, matching. In our supermodular markets, the stable matching is the positive assortative matching (PAM). In our submodular markets, the stable matching is the negative assortative matching (NAM).[16] In each market, there are many stable outcomes that pair the unique stable matching with various transfers. Appendix B characterizes the transfers consistent with a stable outcome. Positive assortative matching, together with transfers that split each surplus equally, constitutes a stable outcome of the supermodular market. In the submodular case, no stable outcome splits each surplus equally. Instead, in matches between low- and high-type agents (Blue with Plum and White with Kiwi in Table 1), the high types (White and Plum) must each receive more than half the surplus.

Without transfers, if surpluses are split using some fixed ratios between the two sides of the market,[17] there is a unique stable matching in both the supermodular and submodular markets: the PAM that matches individuals of congruent types. In both markets depicted in Table 1, the

[16] This observation is generally true for supermodular and submodular surplus structures, where the (generically) unique stable matching is PAM and NAM, respectively.

[17] That is, we can restrict $\tau(c) = \alpha\pi(c, m(c))$ and $\tau(f) = (1-\alpha)\pi(m(f), f)$ for all c and f, where $\alpha \in (0, 1)$.

unique stable matching would match the highest types, White and Plum, then the medium types, Crimson and Mango, and finally the low types, Blue and Kiwi.

3.2 Accounting for Information Frictions

With incomplete information, our requirements for an outcome (m, τ_C, τ_F) to be stable naturally include conditions 1–2 that

$$\tau_C(c) + \tau_F(m(c)) = \pi(c, m(c)) \qquad \forall c \in C \tag{4}$$

$$\tau_C(c), \tau_F(f) \geq 0 \qquad \forall (c, f) \in C \times F. \tag{5}$$

The first, again, ensures that payoffs are feasible, while the second ensures that payoffs are individually rational.[18]

To formulate a counterpart of the no-blocking condition, we say that an *outcome is blocked* if there is a match and a pair of transfers that would increase the payoffs of both agents, such that one of the agents *knows* that this match/transfer combination will increase her payoff. The idea is that, as in the case of complete information, the informed agent would make such a proposal, knowing the resulting match would yield a gain that the receiver would realize (recall that match proposals reveal to the responder the induced payoff). The incomplete-information stability concepts in the literature—see, for instance Liu et al. (2014) or Liu (2022)—share this conceptual foundation, with the details and differences of various notions revolving around what it means to "know" that a proposed match will increase one's payoff. Our definition of incomplete-information stability accounts for the fact that, as they interact, agents learn information that can help them identify blocks. We adopt a demanding notion of a block, and hence a permissive notion of stability.

Each time a match is formed, both agents learn the surplus available in that match. As a result, an agent who has formed matches with two of the agents on the other side of the market has complete information about the surpluses she can generate with each partner: she has observed two of the surpluses and knows the set of three possible surpluses. An agent who has been matched with two of the agents on the opposite side of the market thus knows of the presence of any block that involves her. Let $\tilde{C} \subseteq C$ be the colors who, over the course of the interaction, have matched

[18]In our incomplete-information treatment, an agent accepting an offer knows her payoff, and hence can reject negative payoffs, while a proposing agent is allowed to nullify the offer if its payoff is revealed to be negative.

with at least two opponents, and hence have complete information about the surpluses they can generate. Let $\tilde{F} \subseteq F$ be defined analogously for foods. The no-blocking requirement that completes our definition of *incomplete-information stability* is then that no blocks are known to the agents:

$$\tau_C(c) + \tau_F(f) \geq \pi(c,f) \qquad \forall c \in \tilde{C},\ f \in F \tag{6}$$

$$\tau_C(c) + \tau_F(f) \geq \pi(c,f) \qquad \forall c \in C,\ f \in \tilde{F}. \tag{7}$$

It follows immediately from the definitions that if an outcome (m, τ_C, τ_F) is complete-information stable, then it is also incomplete-information stable. Intuitively, reducing the amount of information available to the agents cannot make it easier for them to block a matching, and hence cannot shrink the set of stable outcomes. Conversely, if (m, τ_C, τ_F) is incomplete-information stable and either all colors or all foods have matched with at least two agents, i.e.,

$$(C \subseteq \tilde{C}) \ \vee \ (F \subseteq \tilde{F}), \tag{8}$$

then (m, τ_C, τ_F) is also complete-information stable. This follows from noting that if at least one of the participants in each possible match knows the surplus that would be generated by the match, then incomplete-information stability conditions (6)–(7) reduce to the complete-information stability condition (3).

As the interaction proceeds, in principle, agents can draw additional inferences about payoffs beyond the surpluses they observe in the matches they make.[19] We could exploit these additional inferences to formulate a more permissive notion of blocking and, hence, a more restrictive stability concept. However, we suspect that such inferences would be challenging for participants, and so are unwilling to rely on a such a more restrictive stability concept. In fact our results are strengthened by the finding outcomes in our data are inconsistent even with our straightforward, and more permissive, notion of incomplete-information stability.

In what follows, we use conditions (4)–(7) to show that our incomplete-information markets do not routinely produce incomplete-information stable outcomes.

[19] For example, agents learn the payoff of the recipient of any match that is formed. If food i sees that color j earned a payoff of 25 from a match, it knows the surplus from that match is at least 25. If i has at some point matched with food k and learned that the surplus of that match is 32, then i has enough information in the supermodular market to identify the surplus i generates with each partner.

4 Aggregate Outcomes

We start by describing the aggregate outcomes in our experimental markets. We use all 10 rounds played in each session and omit the practice round.

4.1 Efficiency and Stability

Table 2 depicts the fraction of efficient outcomes, those corresponding to a matching that maximizes the sum of match surpluses, i.e. the fraction of PAM for supermodular markets and NAM for submodular markets. The table also depicts the fraction of (complete information) stable outcomes, consisting of both an efficient matching and a stable transfer profile. For the latter, we also display the fraction of outcomes that are "almost stable" in that the matching is efficient, and no pair can generate more than a small additional amount, 2 or 5 tokens, by blocking the observed outcome.[20] We return to incomplete-information stability when discussing the information available to participants through market interactions in the next subsection.

Table 2: Final Outcomes (allocations)

	Complete Information		Incomplete Information	
	Supermodular mean (se)	Submodular mean (se)	Supermodular mean (se)	Submodular mean (se)
Efficiency	0.94 (0.03)	0.73 (0.06)	0.84 (0.03)	0.39 (0.09)
Matching				
PAM	0.94 (0.03)	0.01 (0.01)	0.84 (0.03)	0.08 (0.05)
NAM	0.00 (0.00)	0.73 (0.06)	0.00 (0.00)	0.39 (0.09)
Stability				
exact	0.78 (0.08)	0.14 (0.08)	0.54 (0.06)	0.03 (0.02)
- 2 tokens	0.89 (0.05)	0.33 (0.10)	0.68 (0.05)	0.06 (0.04)
- 5 tokens	0.93 (0.04)	0.58 (0.10)	0.78 (0.05)	0.19 (0.08)

Notes: Bootstrapped robust standard errors are reported in parentheses with clustering at the session level. An efficient outcome is one that maximizes the sum of match surpluses. PAM indicates positive-assortative matching, while NAM indicates negative-assortative matching. An allocation is (complete information) stable (or stable $-x$ tokens) if the matching is utilitarian efficient and no pair can generate any gain (of more than x tokens) by blocking the observed outcome.

Table 2 suggests several insights. First, a substantial fraction of all of our markets culminates

[20]Throughout, we follow Cameron, Gelbach, and Miller (2008) and bootstrap standard errors, clustering at the session level.

in efficient matchings. For the most part, matchings are positively assortative when PAM is efficient and negatively assortative when NAM is efficient.

Second, incomplete information poses an important obstacle for efficiency, particularly for submodular markets. For supermodular markets, moving from complete to incomplete information generates a reduction of 10 percentage points in the fraction of efficient matchings. In contrast, for our submodular markets, the fraction of efficient matchings almost halves, dropping by 34 percentage points.

Third, there are fewer stable than efficient outcomes. Much of the instability we observe is due to unstable transfer profiles, as captured by the difference between the fraction of outcomes that entail an efficient matching and those that are stable. For example, with complete information, 6% of supermodular markets do not culminate in an efficient matching, while 16% of supermodular markets yield an efficient matching but with unstable transfers. This wedge is even more pronounced for submodular markets. With complete information, 27% of markets do not culminate in efficient matchings, while 59% produce an efficient matching with unstable transfers.

Fourth, the surplus structure has important impacts on whether emergent outcomes are stable. Submodular markets generate significantly and substantially fewer stable outcomes than do supermodular markets, especially under incomplete information, with differences that are more pronounced than the comparisons of efficient outcomes. With complete information, 78% of supermodular markets culminate in a stable outcome, while for submodular markets, the corresponding figure is 14%. This is the case even when allowing for minimal "mistakes" in terms of the transfer profile. With incomplete information, the fraction of stable outcomes in supermodular markets falls to 54%, while the fraction of stable outcomes in submodular markets nearly vanishes, standing at 3%.

These comparisons are mirrored by the pattern of blocking pairs. For supermodular markets, unstable outcomes are associated with an average of 1.56 blocking pairs with complete information and an average of 1.65 blocking pairs with incomplete information. The corresponding figures for our submodular markets are 1.67 and 2.86. In particular, with incomplete information, submodular preferences yield significantly more blocking pairs than supermodular preferences (at the 1% significance level). The patterns are more pronounced when considering efficient, but unstable, outcomes. In those, we see an average of one blocking pair with complete information

TABLE 3: Discovering Partners in Incomplete Information Markets

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Supermodular					
all players know partners (frac of markets)	0.14	0.40	0.51	0.58	0.60
all players on one side know partners (frac of markets)	0.30	0.69	0.79	0.83	0.88
av # of players who know their partners	3.61	4.90	5.21	5.31	5.41
Submodular					
all players know partners (frac of markets)	0.21	0.43	0.64	0.68	0.85
all players on one side know partners (frac of markets)	0.40	0.73	0.85	0.86	0.95
av # of players who know their partners	3.68	4.81	5.39	5.44	5.68

Notes: For each type of market, supermodular or submodular, we report the fraction of markets in which all players discover their surplus from matching with each possible partner in the corresponding quintile (first row), all players on at least one side discover their surplus from matching with each possible partner in the corresponding quintile (second row), and the average number of players (out of six) who discover their partners' surpluses in each quintile (third row).

and an average of 1.25 blocking pairs with incomplete information when surpluses are supermodular. The figures for our submodular markets are 1.32 and 1.93, both significantly higher (at the 1% significance levels). A similar pattern emerges when considering the average surplus generated by blocking pairs of unstable outcomes, with incomplete information and submodularity leading to greater average surplus available across blocking pairs.[21] These results suggest that, even with fairly unrestricted and decentralized interactions, market participants facing incomplete information are not internalizing enough information to attain complete-information stability, a point we now turn to.

4.2 Learning through Interactions and Incomplete-Information Stability

As incomplete-information markets evolve, participants have opportunities to learn others' types. Indeed, whenever an offer is accepted, individuals learn the resulting payoffs and can thus identify the type of individual on the other side of the market. Table 3 documents this type of learning throughout quintiles of our sessions. In both supermodular and submodular markets, by the third quintile, all participants are fully informed about the surpluses they can generate with all those on the other side of the market in over 50% of markets, with an average of over 5 out of 6 participants

[21] In supermodular markets that culminated in an unstable outcome, the average surplus available to blocking pairs was 7.06 tokens with complete information and 7.19 tokens with incomplete information. For submodular markets, the corresponding figures were 5.33 and 7.90 tokens. In particular, the greatest number of blocking pairs and available surplus per blocking pair were observed for submodular markets with incomplete information.

being informed. By the termination of each market, these numbers are even higher.[22]

Incomplete-information stability coincides with complete-information stability when agents on at least one side all know the surpluses they can generate with each partner; see Section 3.[23] As shown in Table 3, by markets' termination, at least one side is informed as such in 88% of our supermodular markets and 95% of our submodular markets. Even if we assume that agents forget the first few matches, these figures remain high. For example, if participants forget their first match, 76% of supermodular markets and 86% of submodular markets culminate in at least one side being fully informed. These numbers go down to 68% and 70%, respectively, if participants forget their first three matches.

In principle, participants could hasten their learning about the surpluses available in various matches by designing early offers to reveal information, to be subsequently used in forming stable matches. Participants learn the surplus from a match only when an offer is accepted. Therefore, if participants were making early offers with the specific aim of revealing information, we would expect to see more generous early offers designed to ensure acceptances and hence learning. As we show in Appendix A, that is not the case. In incomplete-information markets, offers are not significantly more generous initially. In fact, they are not significantly more generous than offers observed in our complete-information markets. Furthermore, rejection rates are not uniformly lower in initial stages of our incomplete-information markets. [24] Nearly complete learning thus arises naturally in the course of the search for desirable matches. Despite this, we see many matches that are not stable.

4.3 Payoffs

Table 4 provides aggregate statistics of participants' payoffs and transfers. The top panel of the table depicts the average payoffs by type. Theoretically, all stable outcomes of our supermodular markets entail higher types gaining greater payoffs. In our submodular markets, however, a uniform distribution over stable outcomes would yield only 50% of stable transfer profiles exhibiting

[22]Learning in submodular markets appears slightly faster than in our supermodular markets. However, this difference is significant only in the last quintile. Indeed, regression analysis bootstrapping errors and clustering at the session level suggests a corresponding p-value lower than 0.05 only when looking at the last quintile.

[23]This observation does not rely on participants' higher order beliefs. Inference about information sets available to other participants requires arguably more complex considerations.

[24]For instance, focusing on the first quintile, for complete-information submodular markets, 33% of offers from M-type agents are rejected and 39% of offers from high-type agents are rejected. For incomplete-information submodular markets, the corresponding figures are 31% and 39%, respectively.

Table 4: Players' Final Payoffs

	Complete Information		Incomplete Information	
	Supermodular mean (se)	Submodular mean (se)	Supermodular mean (se)	Submodular mean (se)
Payoffs by Type				
Low Type	4.11 (0.06)	16.07 (0.32)	4.41 (0.13)	14.72 (0.51)
Medium Type	15.83 (0.10)	23.81 (0.21)	15.60 (0.45)	24.19 (0.12)
High Type	35.86 (0.08)	38.87 (0.39)	35.29 (0.24)	36.33 (0.46)
Equal Splits				
All allocations	0.79 (0.04)	0.32 (0.01)	0.67 (0.05)	0.39 (0.05)
First two deals	0.69 (0.04)	0.62 (0.05)	0.24 (0.01)	0.24 (0.03)
Efficient allocations	0.81 (0.03)	0.31 (0.01)	0.69 (0.04)	0.29 (0.02)
Stable allocations	0.87 (0.01)	0.30 (0.02)	0.80 (0.05)	0.33 (0.17)

Notes: Average payoffs by players' types are reported in the first three rows. Bootstrapped robust standard errors (clustered at the session level) and reported in the parentheses. Equal splits correspond to ratios of payoffs between two matched players falling between 45% and 55%. For the first two deals in the incomplete-information markets, we focus on deals in which the proposer receives non-negative payoffs.

this monotonic pattern (see our Appendix B for details on the constraints stability with transfers imposes). Interestingly, when considering average payments, in all of our treatments we see high types garnering significantly greater payoffs than medium types, who end up with significantly greater payoffs than low types. Nonetheless, when we consider these comparisons market by market, we see some differences, echoing our theoretical predictions. When information is complete, all of our supermodular markets, but only 83% of our submodular markets, end up with higher types gaining greater payoffs. When information is incomplete, 93% of our supermodular markets and 61% of our submodular markets generate greater payoffs for higher types.

Equal splits are common in experimental work (e.g. Andreoni and Bernheim, 2009). While equal splits are consistent with stability for supermodular markets, that is not the case for submodular markets, a point we return to in Section 6.1. In fact, were participants simply following an egalitarian norm and always proposing an equal split of match surpluses, we would expect to observe PAM outcomes in all our markets, as described in Section 3. Table 2 suggests that egalitarian norms cannot be mediating all of our results. Indeed, in our submodular markets, we see PAM outcomes quite rarely. In fact, while the transfers required for stable outcomes seem elusive, particularly in submodular markets, the overall efficiency in those markets is far greater than under PAM. While PAM would generate a total surplus of 128 tokens, the top panel of Table 4 indicates an average total surplus of 157.5 tokens under complete information and an average total surplus of 150.5 tokens under incomplete information (recall that there are two participants of each type in each market). In other words, transfers do allow for greater efficiency, even in cases in which stability appears the most challenging to achieve.

The bottom panel of Table 4 displays explicitly the frequency of (roughly) 50 – 50 splits of the match surplus in the final outcomes observed across markets. In order to allow for some margin of error, we consider offers that range from 45% to 55% of the surplus. Since there are many opportunities for participants to learn their match surpluses with others through markets' operations—see Section 4.2 above—we consider the possibility of equal splits for incomplete-information markets as well. Such roughly equal splits are prevalent, albeit not universal, in our supermodular markets: hovering around 80% with complete information and around 70% with incomplete information. They are far less common in our submodular markets, comprising fewer than 40% of the outcomes regardless of the information available.[25]

Figure 1 describes the cumulative distribution functions corresponding to payoffs across our treatments. While complete-information treatments generally yield more pronounced "jumps" the underlying preferences have a clear impact. With supermodular preferences, many of the payoffs take one of three values. Indeed, when superimposing the distribution of payoffs that would be generated by equally splitting the surplus between agents in the efficient matching (the thick solid black line, corresponding to equal masses receiving payoffs of 4, 16, and 36), we see substantial overlap. Submodular preferences, however, yield a nearly uniform distribution of payoffs, particularly when information is incomplete. With complete information, we do see a large fraction of individuals receiving a payoff of 24, an observation we soon elaborate on. Furthermore, the payoff distributions corresponding to submodular preferences first order stochastically dominate those corresponding to supermodular preferences. This, again, suggests that participants are not simply implementing the efficient matching and splitting the surplus in those markets. Indeed, the resulting distribution (depicted in the thick solid grey line, corresponding to one third of the population receiving a payoff of 24 and two thirds receiving a payoff of 28) would second order, but not first order, stochastically dominate the corresponding distribution for our

[25]Interestingly, when looking at initial accepted offers in our markets, we see far more equal shares of match surpluses for matches between congruent types in our submodular markets. In the first two accepted offers, equal sharing occurs in over 89% of deals between participants of the same type. Nonetheless, even in these initial deals, participants of different types do not split the surpluses equally. In particular, low types matching with high types get only 20% of the surplus. We provide more details on behavior at the start of our markets' operations in Appendix A.1.

Figure 1: Cumulative Distributions of Players' Payoffs

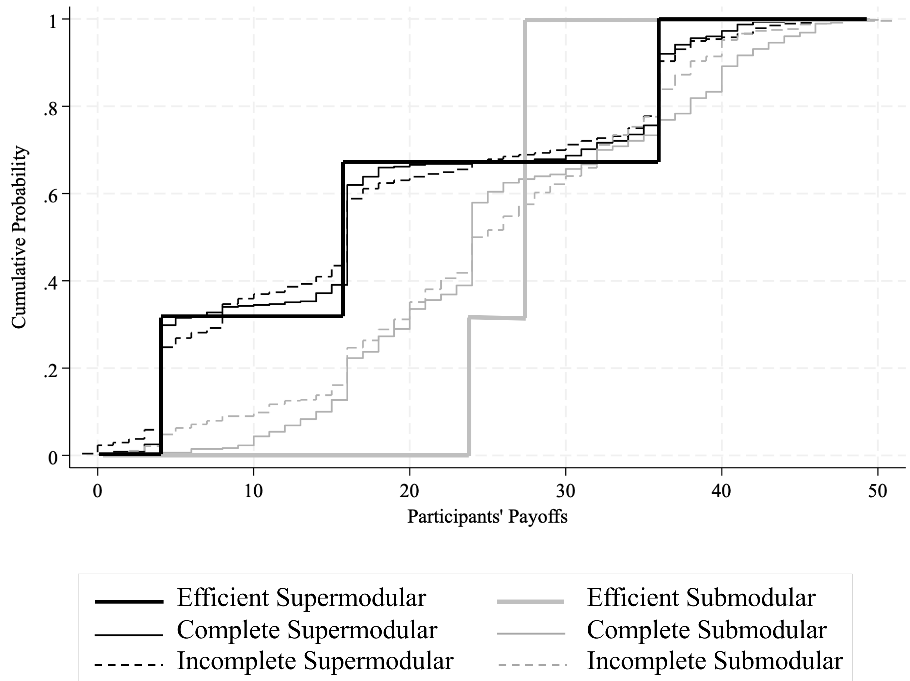

supermodular markets.

Figure 1 presents payoff distributions aggregated across all market participants. At first blush, it indicates similarity in overall payoff distributions between complete- and incomplete-information markets of the same surplus variety. Figure 2 presents the distribution of payoffs by types of agents in markets that culminated in "nearly stable" outcomes, ones in which the matching is efficient and there is a stable allocation of the surpluses such that each participant's payoff is at most 2 tokens away from what they would receive in that allocation.[26] The figure depicts how observed payoffs relate to the range of payoffs corresponding to stable outcomes, captured by the red solid line in each panel (see Appendix B for computational details). In line with our discussion above, the distribution of payoffs for all agents and across our markets exhibits a spike at the midpoint of stable allocations. However, the figure conveys several additional messages. Considering our complete-information treatments, payoffs corresponding to the midpoint of the interval of stable transfers are far more prevalent in our supermodular markets. In fact, in our submodular markets, while over 70% of the medium-type agents receive the mid-point of their stable payoffs, low-

[26] Allowing for a small margin of error ensures that the samples we discuss are of sufficient size.

Figure 2: Distribution of Payoffs by Types in Stable (-$2) Allocations

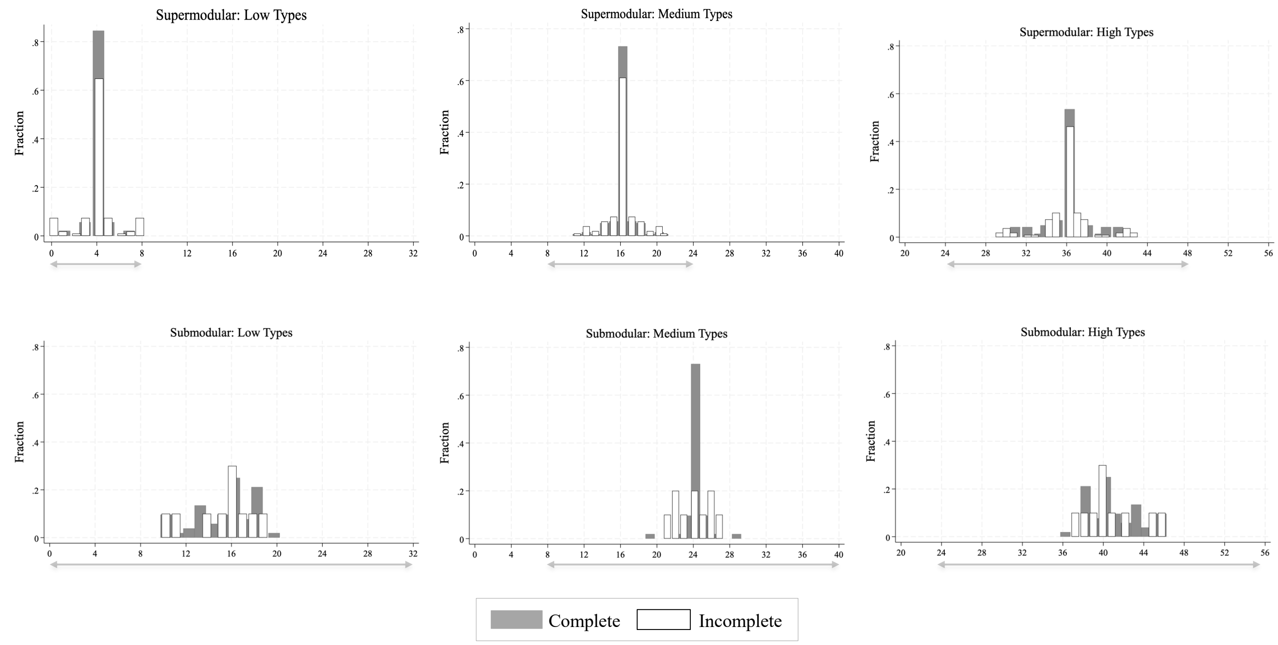

Notes: The arrows below the horizontal axis of each graph depict the range of payoffs corresponding to stable outcomes predicted by theory.

and high-type agents have far more disperse payoff distributions. With incomplete information, payoffs are noticeably more diffuse, particularly when preferences are submodular. In fact, for submodular preferences, for all types, no allocation carries more than 30% of our data.

4.4 Summary of Aggregate Outcomes

Overall, participants often establish efficient matchings. Stability is more elusive, indicating that it is easier to settle on an efficient matching than to arrange the attendant transfers. Despite substantial information flows through markets' operations, incomplete information presents hurdles to efficiency and especially stability. The combination of incomplete information and submodularity reduces stability to a tiny percentage of outcomes.

Table 5: Market Characteristics

	Complete Information		Incomplete Information	
	Supermodular mean (se)	Submodular mean (se)	Supermodular mean (se)	Submodular mean (se)
Market duration (in sec)	158.6 (12.1)	171.5 (13.9)	191.6 (14.1)	292.9 (9.2)
Nb of offers extended				
overall per market	24.8 (1.3)	24.1 (2.6)	42.4 (3.7)	45.9 (4.6)
by Low type	9.6 (1.2)	10.7 (1.7)	19.9 (1.3)	21.0 (2.3)
by Medium type	8.0 (0.7)	8.4 (0.6)	14.2 (1.7)	15.8 (1.9)
by High type	7.2 (0.7)	5.2 (0.3)	8.3 (0.8)	9.1 (1.2)
Nb of offers accepted				
overall per market	5.6 (0.2)	8.8 (1.1)	15.9 (0.9)	17.3 (1.0)
by Low type	1.7 (0.0*)	2.1 (0.2)	5.9 (0.3)	5.7 (0.3)
by Medium type	2.1 (0.1)	3.1 (0.4)	5.9 (0.4)	5.7 (0.3)
by High type	1.8 (0.1)	3.5 (0.5)	4.2 (0.4)	6.0 (0.3)

Notes: Bootstrapped robust standard errors (clustered at the session level) are reported in parentheses. The value of robust st error in * is 0.04.

5 Market Activity and Match Evolution

Table 5 summarizes the aggregate features of activity in our markets. As can be seen, market activity took, on average, between 3 and 5 minutes. Our incomplete-information markets correspond to significantly longer durations and a greater volume of offers made and accepted. Our submodular markets correspond to somewhat longer times to reach an outcome. However, the volume of offers made or accepted does not seem to hinge on the underlying preference structure. At the aggregate level, there are no great differences across agent types in terms of overall offers made or accepted. One exception corresponds to the low-type agents, who make significantly more offers than the high-type agents, particularly when information is incomplete.

We start by considering the progression of matches over the course of a market's operation. In order to depict our markets' evolution, we inspect quintiles of activity in each market's operation. Namely, for each round in our sessions, we consider the time between the first and last moments of activity. We split that time span into five quintiles, thereby effectively normalizing the length of activity in each of our experimental markets.

Figure 3 reports the mean number of transient and permanent matches formed at each quintile of market duration across treatments, for each combination of matched types—LL denotes two low-type agents, LM denotes one low- and one medium-type agent, and so on. Several obser-

Figure 3: Transient and Permanent Matches, by Quintiles

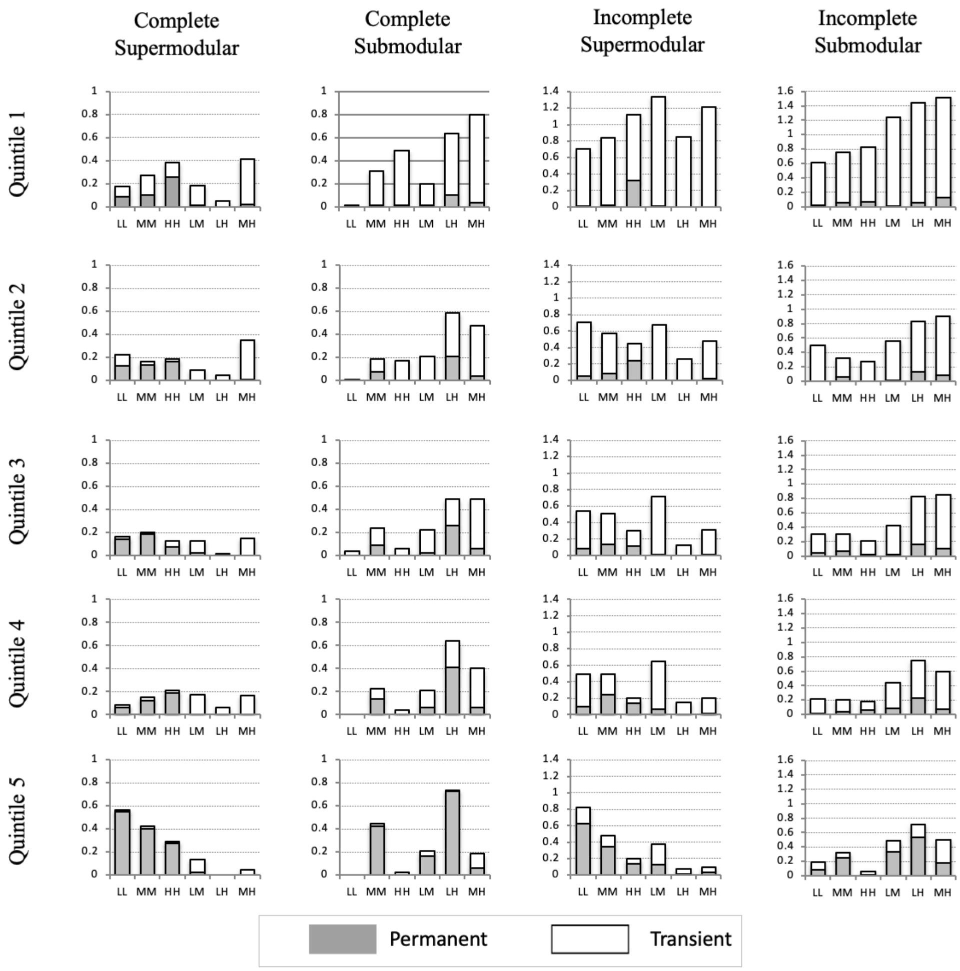

Notes: Quintiles are defined relative to the total time the average market takes to reach a final outcome, where average market duration is calculated as the mean time difference between the last and first activity in that market. The white parts of the bars indicate the transient matches, i.e. those that are broken later on. Grey bars depict the final matches that remain intact until the market's termination.

vations emerge. First, there are substantial differences in match dynamics across supermodular and submodular markets. The complete-information treatments are particularly telling. With supermodular preferences, both transient and permanent matches are first formed between the high-type agents. They are followed by medium-type agents, and last by low-type agents. In contrast, with submodular preferences, extremal types form matches first: we see the most early transient matches between medium- and high-type agents and the most early permanent matches between low- and high-type agents. These are followed by matches between congruent medium-type agents. The early permanent matches between low- and high-type agents occur at substantially lower rates than the corresponding congruent and permanent early matches in our supermodular markets, particularly when accounting for the fact that each market contains two incongruent low- and high-type pairs. Last, we see a substantial fraction of matches between low- and medium-type agents at later market phases. These late matches are not part of stable outcomes.

The second observation from Figure 3 is that incomplete-information treatments are associated with far more offers, in every quintile and particularly early on. Some of these offers act to reveal information about the types of agents on the other side of the market—indeed, early-quintile offers are frequently transient. This presumably explains why some features of the complete-information markets remain. Specifically, we see more transient and permanent early matches involving high-type agents. In contrast, the lower-type agents take more time to find partners.[27]

6 Potential Mechanisms for Instability

Our analysis suggests ample scope for learning in our markets, even if participants do not perfectly record all of their experiences. The dearth of stable outcomes in our incomplete-information treatments is thus puzzling. The current section suggests potential explanations.

Matching theory provides limited guidance on the equilibrating dynamics one might expect in decentralized markets with transfers. We first show that the common tendency of lab participants to split surpluses equally does not account for the patterns we observe. Next, we illustrate that sufficiently pessimistic beliefs might be consistent with our markets' outcomes. Finally, we demonstrate behaviors in our data consistent with participants following a satisficing heuristic, reducing substantially their efforts once hitting a (type-dependent) payoff target.

6.1 Equal-splits hypothesis

As noted in Section 4.3, equal splits are a common feature in experimental studies (e.g. Andreoni and Bernheim, 2009). Table 4 shows that while equal splits align with stability in supermodular

[27] Repeat offers occur frequently, even in later parts of markets' operations. Consider the average fraction of repeat offers in the second half of each round. In complete-information supermodular markets, these range from 0.27 for low-type agents to 0.34 for medium-type agents to 0.13 for high-type agents. In incomplete-information supermodular markets, these range from 0.34 for low-type agents to 0.29 for medium-type agents to 0.25 for high-type agents. Our submodular markets yield nearly identical frequencies of repeat offers.

markets, they do not in submodular markets. If participants were merely adhering to an egalitarian norm by always proposing equal splits of match surpluses, we would expect to observe PAM outcomes across all markets. However, Table 2 indicates that egalitarian norms alone cannot fully explain our results.

Still, the draw to equal splits may affect how markets evolve. As Figure 4 shows, participants in certain market roles are more inclined to offer equal splits, particularly in early stages of markets' operation. The figure describes the distribution of offers—as a percentage of the surplus—across quintiles in our complete information treatments. We include an analogous figure pertaining to the distribution of offers in our incomplete information treatments in the Appendix A.2. The same patterns emerge.

As can be seen, in supermodular markets, there is a fraction of low-type participants who make 50-50 offers throughout a market's duration. The prominence of 50-50 offers is much lower for other types. For sub-modular markets, it is the medium-type participants who exhibit a pronounced draw to equal-split offers, but this tendency is far less pronounced for other types. These patterns are fairly stable throughout a round's duration, although 50-50 offers are somewhat lower for medium-type participants in sub-modular markets during the second to fourth quintiles. Thus, an egalitarian norm cannot explain the patterns in our data even in early stages of markets' operation.

6.2 Stability with Pessimistic Beliefs

Certainly, if participants ignore all the information they receive, any individually-rational outcome can be rationalized as incomplete-information stable. As already discussed, simply forgetting the first few matches participants form cannot explain the outcomes we observe. As we now show, a form of biased beliefs may.

Suppose that instead of using all the information participants learn from their market interactions, they assemble only the minimal information consistent with the observed final matching. Suppose, further, that participants form pessimistic expectations (from the point of view of blocking), recognizing the possible match surpluses in the market. Namely, when observing a match in which the receiver realizes a payoff of x, they assume the matched partner receives the maximal feasible payoff of $72-x$. Then, in those incomplete-information outcomes that fail to be complete-

Figure 4: Cumulative Distributions of Proposers' Shares across Offers Made

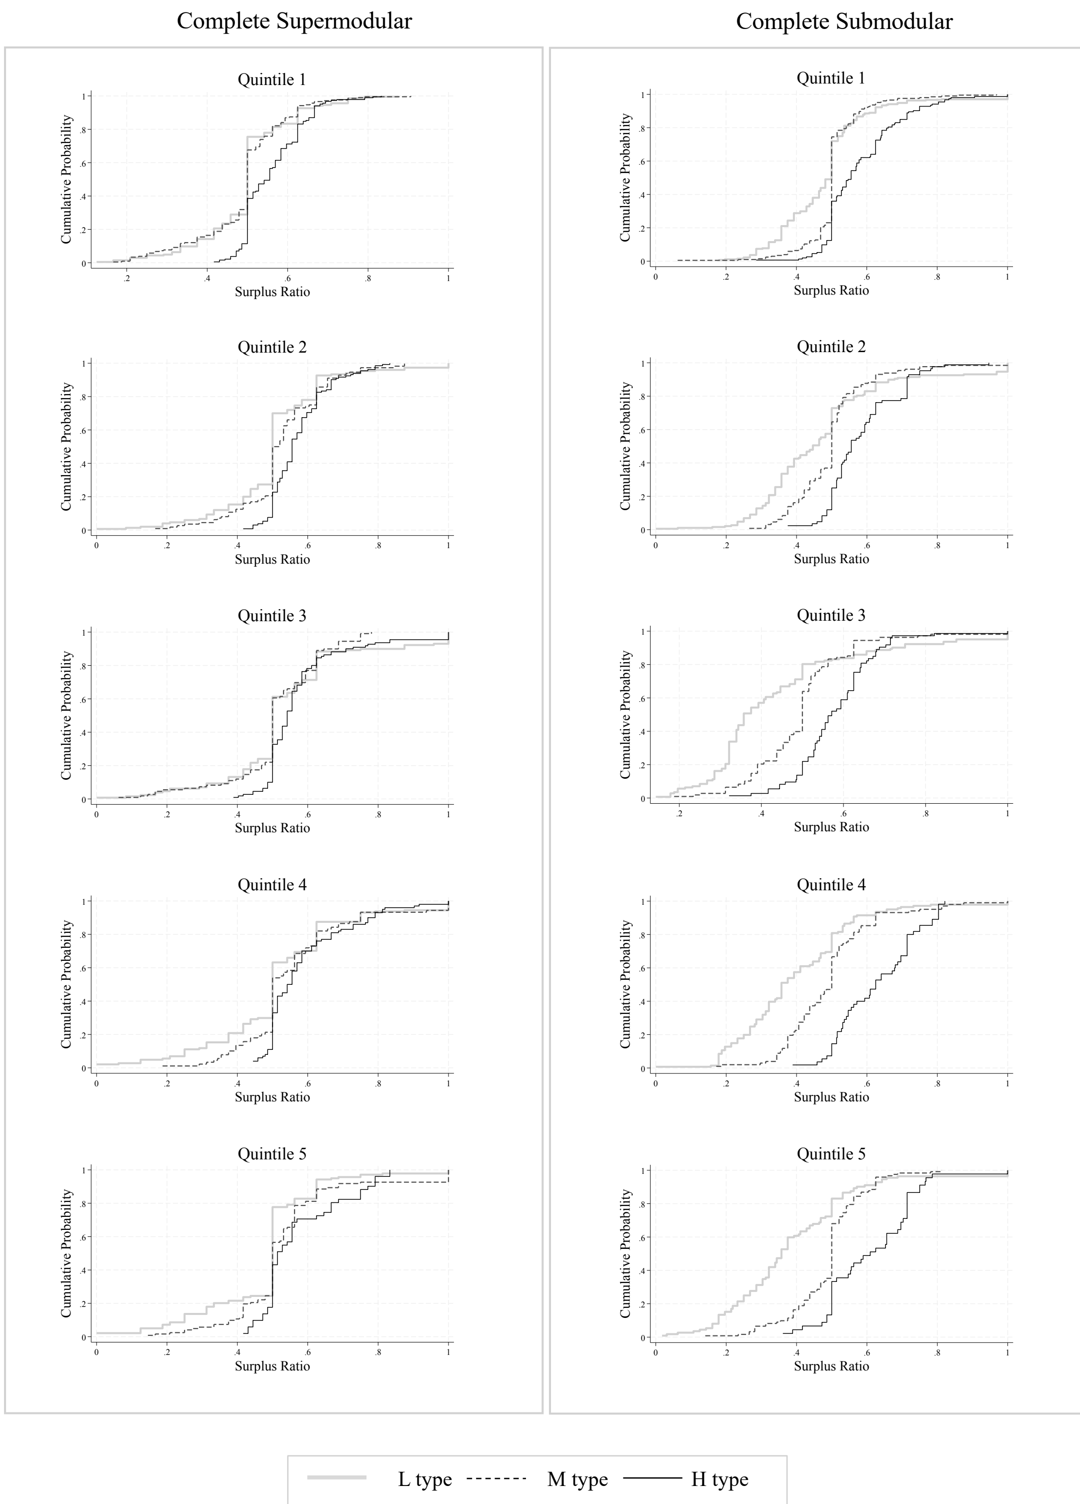

information stable, there is an average of only 0.59 blocking pairs in the supermodular markets and 1.04 blocking pairs in the submodular markets. These numbers are considerably smaller than those reported above for our complete-information markets (1.56 and 1.67, respectively). This indicates that there are beliefs consistent with the information revealed under which participants come as close to (incomplete-information) stable outcomes as they do under complete information. In practice, we might expect participants to be less pessimistic and to glean yet more information from the matching process (though perhaps not as much as do the hyper-rational agents in theoretical treatments of incomplete-information stability such as Liu et al., 2014). We leave as an important topic for future work the more refined investigation of whether the behavior we observe reflects stability on the part of timid, under-inferring agents, or failures of less timid and essentially fully informed agents to achieve stability.

While pessimistic beliefs provide an effective way to organize the market outcomes we observe, they do not, on their own, explain the dynamics that markets follow.[28] In what follows, we illustrate a simple heuristic consistent with the market activity we observe.

6.3 Satisficing Patterns

We now consider a satisficing heuristic, whereby participants target a particular payoff, possibly dependent on their market type. Caplin et al. (2011) document the use and prevalence of satisficing heuristics in individual choices. Our design allows for satisficing behavior to emerge in a market setting, where the satisficing threshold depends on global market features. Satisficing would suggest that participants make offers before hitting their target payoff, and cease making offers once that target is achieved.

For each market, we compute the number of offers made by players of each type—low, medium, or high—conditional on their current payoff, excluding unmatched individuals. We then group offers in bins corresponding to participants' current payoff. We normalize the number of offers by the volume of offers made by the corresponding type across all rounds. Table 6 presents the resulting fraction of offers by each type in each of our treatments, conditioning on the current pay bin. Horizontal lines in the table correspond to a satisficing rule, allowing for a 21% mistake rate.

[28]For instance, pessimistic beliefs at the start of each market's activity could correspond to each participant suspecting the minimal match surplus available to them. However, initial offers in our incomplete-information markets significantly exceed half that surplus for each of the three types. Furthermore, initial offers do not reflect a persistent fraction other than 50% of that "pessimistic" surplus.

That is, fewer than 21% of offers are observed among participants whose payoffs exceed those indicated by the line. As can be seen, similar qualitative conclusions result from using lower mistake thresholds.[29] The offer rates documented in Table 6 naturally reflect a combination of how active people are conditional on their current pay and how frequently they find themselves with that pay.

Table 6: Frequency of Offers Made Conditional on Current Pay

	Supermodular					
	Low Type		Medium Type		High Type	
Current Pay	Complete	Incomplete	Complete	Incomplete	Complete	Incomplete
1 to 4	0.87	0.62	0.03	0.07	0.00	0.03
4 to 8	0.13	0.27	0.26	0.22	0.00	0.04
8 to 12	0.00	0.06	0.29	0.23	0.03	0.15
12 to 16	0.00	0.01	0.34	0.32	0.05	0.08
16 to 24	0.00	0.02	0.07	0.14	0.34	0.25
24 to 32	0.00	0.01	0.00	0.01	0.39	0.32
32 to 40	0.00	0.00	0.00	0.00	0.19	0.12
40 and above	0.00	0.00	0.00	0.01	0.00	0.01

	Submodular					
	Low Type		Medium Type		High Type	
Current Pay	Complete	Incomplete	Complete	Incomplete	Complete	Incomplete
1 to 4	0.30	0.24	0.00	0.05	0.00	0.00
4 to 8	0.16	0.22	0.00	0.11	0.00	0.00
8 to 12	0.14	0.13	0.04	0.09	0.00	0.01
12 to 16	0.28	0.25	0.24	0.17	0.01	0.02
16 to 24	0.12	0.13	0.57	0.37	0.05	0.15
24 to 32	0.00	0.04	0.15	0.19	0.35	0.52
32 to 40	0.00	0.00	0.00	0.02	0.54	0.27
40 and above	0.00	0.00	0.00	0.00	0.05	0.02

Notes: Fractions represent the frequency of offers of a particular type in a particular market conditional on current-payoff bin. The horizontal line corresponds to the satisficing thresholds with a 21% mistake threshold: payoffs beyond the threshold lead to fewer than 21% of offers.

Estimated satisficing thresholds are remarkably consistent across complete- and incomplete-information markets across types, with only one exception corresponding to low types in supermodular markets. This consistency is reassuring and suggests that market positions are among the strongest forces shaping our participants' behavior, an idea we explore further in the following section. In fact, the estimated satisficing thresholds closely approximate the jumps in participants' payoff distributions depicted in Figure 1.[30] While the jumps resemble those implied by stability

[29]We use a 21% cutoff rather than 20% since for medium types in submodular markets with incomplete information, the natural dividing line lies slightly above 20%. Except for these types and for high types in supermodular markets under complete information, a 15% cutoff would have yielded the same results. Similar levels are derived when considering the first and second half of each market's operations separately. In addition, we do not see substantial differences across early and later rounds.

[30]In Appendix A.3, we perform a coarser exercise. We take the average payoff corresponding to each information

Table 7: Variance Decomposition of Players' Payoffs

	Complete Info		Incomplete Info	
	Supermodular	Submodular	Supermodular	Submodular
Indicators for L, M, and H types	0.94	0.86	0.74	0.80
Own Offers (frac)	0.01	0.06	0.12	0.17
Greed	0.01	0.02	0.01	0.01
First Deal (quintile)	0.03	0.03	0.00	0.00
Last Deal (quintile)	0.02	0.00	0.13	0.01
Risk measures	0.00	0.00	0.00	0.00
Altruism measures	0.00	0.02	0.00	0.00
overall R^2	0.96	0.84	0.91	0.70
nb obs	$n = 478$	$n = 480$	$n = 480$	$n = 480$

Notes: Shapley R^2 decomposition is reported for indicators for the market position of players as well as bargaining styles. Own Offers (frac) is the average number of offers made by each participant type as a fraction of all offers made in that specific market. Greed is the share of surpluses demanded in offers averaged across all rounds. First Deal and Last Deal represent the quintile in which the first and last matches were created, respectively. Risk and altruism measures include responses to the auxiliary investment tasks and Dictator games performed at the end of each session.

in supermodular markets, they diverge sharply in submodular markets. This divergence offers one channel for the greater instability we observe when preferences are submodular. A promising direction for future work is to examine how behavioral heuristics interact with broader classes of preference and information structures.

7 Bargaining Styles

We now turn to the various bargaining styles participants used and how they affect outcomes in our markets. The main message is that, while individuals vary widely in how they approach our markets, in terms of the number of offers they make, their insistence on higher shares, etc., their payoffs are not strongly affected by these differing interaction styles. By and large, payoffs are determined by the governing match surpluses: supermodular or submodular, and participants' position in the market in terms of their types.

Table 7 describes results from a variance decomposition of players' payoffs into different features of participants for each market type: the share demanded in offers, how active participants are in the market (the prevalence of offers they generate), the time at which they are active during the markets' operations, as well as their risk and altruism attitudes inferred from our auxiliary

and preference structure, as well as an agent's type, as that agent's threshold. Agents with payoffs below the threshold make offers at substantially higher rates than agents with payoffs above that threshold.

elicitations in each session.[31] As the table illustrates, the bulk of the variation in players' payoffs is explained by the variation in participants' types.[32]

These results are interesting in view of the vast literature suggesting that bargaining strategies may be important in various contexts, most notably in labor markets; see Bertrand (2011) and Shell (2006). While our settings are certainly simplistic in many ways, they suggest a very different story. Market position appears to have a dominant impact on outcomes. Who one matches with, and with what transfer, depends importantly on the role one occupies in the market, while being less insensitive to how one bargains, particularly in complete information markets. Thus, to the extent that a person's type is shaped by pre-market choices—education, skills, and similar investments—these pre-market decisions, rather than individual bargaining styles, are what ultimately drive outcomes. This is not to suggest that bargaining styles are irrelevant. Especially in settings with incomplete information, persistence in making offers (captured by Own Offers in Table 7) still accounts for 12%-17% of the variation in payoffs.

8 Conclusions

We report an array of results examining outcomes and the dynamics of interactions in decentralized matching markets. We allow for transfers and consider markets with complete and incomplete information, and with supermodular or submodular surplus structures.

Without transfers, prior literature has suggested that decentralized interactions often culminate in stable, and hence efficient, outcomes. Our first main message is that, in the presence of transfers, a substantial fraction of outcomes is still efficient, while fewer outcomes are stable. Participants often establish the matches required for efficiency, but find it more difficult to settle on transfers that induce stable outcomes. Our second message is that limited information and the shift from a supermodular to a submodular surplus structure have important impacts. Limited information constitutes an important hurdle for stability, even when participants have ample opportunities to learn about others through market interactions. Stability is similarly more chal-

[31]Numbers in each column do not always sum up to 1 due to rounding. We have only 478, rather than 480, observations for supermodular markets with complete information since one pair in this treatment did not agree on their surplus split and remained unmatched. In Appendix A.4, we also include similar analysis in which we account for the prevalence of *accepted* offers, or deals, instead of arbitrary offers.

[32]In Appendix A.4, we include related regression results similarly illustrating that participants' market positions are the only feature that organizes our payoff data coherently.

lenging to achieve with submodular surpluses. Together, incomplete information and submodularity are deadly for stability—in those markets, less than 5% of ultimate outcomes are stable. Our third message pertains to the dynamics of interactions. When surpluses are supermodular, high types tend to match first, followed by medium types, and then by low types. When surpluses are submodular, there is substantial variation in the order of match formation and resulting payoffs. Indeed, the ultimate payoff distribution is nearly uniform. Participants' behavior is consistent with a satisficing heuristic in which payoff targets are type dependent. The final message emerging from our data regards bargaining styles. While we see variation in bargaining patterns across individuals, those seem to have little impact on final payoffs. The main determinants of participants' payoffs in our markets are the governing surplus structure and their positions, or types.

We believe our study opens the door for several further investigations, both theoretical and experimental. As described in our literature review, several recent papers have suggested some natural stability notions allowing for incomplete information on one side. The results from those papers, as well as the notion we suggest, cannot be easily reconciled with our data. In part, this may be due to the fact that the suggested notions are agnostic about the decentralized process that takes place in the market. In particular, existing solution concepts typically take as fixed a specification of uncertainty and a candidate stable outcome, ignoring the fact that the residual uncertainty is likely to depend on the process by which the candidate outcome is attained. We hope our experimental results might be useful in refining those stability notions to richer settings. Our results point to the need for a model allowing agents to glean complete information about the other side of the market through the matching process, while settling on an outcome that is not complete-information stable. Our experiments also suggest that, when information is incomplete, there might be great value for centralization if (complete-information) stability is an objective. Nonetheless, the current literature offers little insights on the design of matching markets in the presence of incomplete information, for any objective. We believe this would be a fruitful direction for future research.

Our experiments examine bargaining frictions due to transfers and incomplete information. There are certainly other frictions that would be interesting to inspect. In applications, offers are often costly—in labor markets, employers frequently have limited interview slots or limited offers to make, regardless of their yield; in marriage markets, dating is costly in terms of both time and money. Naturally, frictions pertaining to offer costs may dampen market activity. That, in turn, would limit learning even further, as participants "try out" fewer matches. The inefficiencies we identify may then offer lower bounds on the inefficiencies one might expect in practice. There are other details that our basic design does not incorporate: whether offers are exploding, communication channels between market participants, various information structures available, etc. We hope our study sets the stage for further investigations of such additional frictions.

Substantively, our findings highlight that the combination of incomplete information and submodular preferences can pose significant barriers to achieving both efficiency and equity. In contexts where matching quality is indeed submodular—such as when the most severe patients gain disproportionately from treatment by top doctors, or when children with greater learning challenges benefit more from highly skilled teachers—informational frictions can be particularly detrimental. In such environments, centralized intervention may play an especially valuable role.

References

Agranov, M. and M. Elliott (2021): "Commitment and (in)Efficiency: A Bargaining Experiment," *Journal of the European Economic Association*, 19, 790–838.

Alston, M. (2020): "On the Non-existence of Stable Matches with Incomplete Information," *Games and Economic Behavior*, 120, 336–344.

Andreoni, J. and B. D. Bernheim (2009): "Social Image and the 50–50 Norm: A Theoretical and Experimental Analysis of Audience Effects," *Econometrica*, 77, 1607–1636.

Becker, G. S. (1973): "A Theory of Marriage: Part I," *Journal of Political Economy*, 81, 813–846.

Bertrand, M. (2011): "New Perspectives on Gender," in *Handbook of Labor Economics, vol 4b*, ed. by O. Ashenfelter and D. Card, Amsterdam: Elsevier, 1543–1590.

Bikhchandani, S. (2017): "Stability with One-sided Incomplete Information," *Journal of Economic Theory*, 168, 372–399.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008): "Bootstrap-based improvements for inference with clustered errors," *The Review of Economics and Statistics*, 90, 414–427.

Caplin, A., M. Dean, and D. Martin (2011): "Search and Satisficing," *American Economic Review*, 101, 2899–2922.

Chade, H., J. Eeckhout, and L. Smith (2017): "Sorting through Search and Matching Models in Economics," *Journal of Economic Literature*, 55, 493–544.

Chen, D. L., M. Schonger, and C. Wickens (2016): "oTree—An Open-source Platform for Laboratory, Online, and Field Experiments," *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Chen, Y. and T. Sönmez (2006): "School Choice: An Experimental Study," *Journal of Economic Theory*, 127, 202–231.

Chen, Y.-C. and G. Hu (2020): "Learning by Matching," *Theoretical Economics*, 15, 29–56.

Chiappori, P.-A. (2017): *Matching with Transfers: The Economics of Love and Marriage*, Princeton University Press.

Echenique, F., A. Robinson-Cortés, and L. Yariv (2025): "An Experimental Study of Decentralized Matching," *Quantitative Economics*.

Ferdowsian, A., M. Niederle, and L. Yariv (2025): "Strategic Decentralized Markets: The Effects of Information Frictions," *Mimeo*.

Gneezy, U. and J. Potters (1997): "An Experiment on Risk Taking and Evaluation Periods," *The Quarterly Journal of Economics*, 112, 631–645.

Hakimov, R. and D. Kübler (2021): "Experiments on Centralized School Choice and College Admissions: A Survey," *Experimental Economics*, 24, 434–488.

He, S., J. Wu, H. Zhang, and X. Zhu (2024): "Decentralized Matching with Transfers: Experimental and Noncooperative Analyses," *American Economic Journal: Microeconomics*, 16, 406–439.

Immorlica, N., J. Leshno, I. Lo, and B. Lucier (2020): "Information Acquisition in Matching Markets: The Role of Price Discovery," *SSRN 3705049*.

Kapor, A. J., C. A. Neilson, and S. D. Zimmerman (2020): "Heterogeneous Beliefs and School Choice Mechanisms," *American Economic Review*, 110, 1274–1315.

Liu, Q. (2022): "Stability and Bayesian Consistency in Two-Sided Markets," *American Economic Review*, 110, 2625–2666.

Liu, Q., G. J. Mailath, A. Postlewaite, and L. Samuelson (2014): "Stable Matching with Incomplete Information," *Econometrica*, 82, 541–587.

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.

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

Pais, J. and Á. Pintér (2008): "School Choice and Information: An Experimental Study on Matching Mechanisms," *Games and Economic Behavior*, 64, 303–328.

Pais, J., Á. Pintér, and R. F. Veszteg (2020): "Decentralized Matching Markets with(out) Frictions: A Laboratory Experiment," *Experimental Economics*, 23, 212–239.

Plott, C., N. Roy, and B. Tong (2013): "Marshall and Walras, Disequilibrium Trades and the Dynamics of Equilibration in the Continuous Double Auction Market," *Journal of Economic Behavior and Organization*, 94, 109–205.

Roth, A. E. and M. Sotomayor (1990): *Two-Sided Matching: A Study in Game-theoretic Modeling and Analysis*, Econometric Society Monographs Vol 18, Cambridge University Press.

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

Shell, G. R. (2006): *Bargaining for Advantage: Negotiation Strategies for Reasonable People*, Penguin.

A Additional Analysis

A.1 Do Markets Start the Same Way?

Since most of our results consider the entirety of participants' market interactions, one may wonder whether the impact of the surplus variety occurs through market dynamics alone. In this section, we show that even initial activity varies by the sort of surplus that governs preferences. We focus on complete-information markets, since initial interactions in incomplete-information markets are naturally noisy due to the limited knowledge participants have regarding their own and others' preferences.

Figure A.1 illustrates the frequency of matches formed and offers made at the start of our markets. As can be seen, different patterns emerge across our markets. For example, in our supermodular markets, initial matches are mostly formed between stable partners, with a predominance of matches and offers between high types. In contrast, in our submodular markets, while early matches between high types are still prevalent, most offers occur between low and high types, or between medium and high types. Furthermore, there is roughly an equal share of matches between low and high, medium and high, or two medium types.

Figure A.1: First Two Offers and Deals with Complete Information

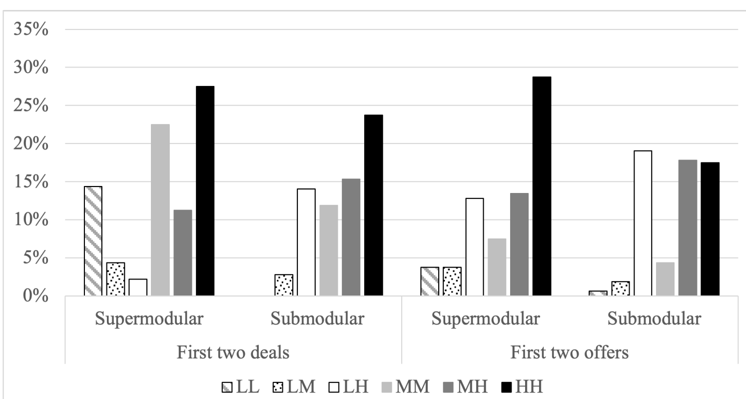

Notes: The frequencies for non-congruent matches (LM, LH, MH) are divided by two, since there are twice as many opportunities to make such offers/deals than congruent ones (LL, MM, and HH).

Table A.1 suggests that the profile of transfers also differs across market classes from the start. Interestingly, submodular markets are characterized by more equal splits at the start of interactions, even though, as we showed in the paper, this pattern reverses as markets evolve. We consider more aspects of equal-splits dynamics in the next subsection.

Table A.1: Frequency of Equal Splits in First Deals/Offers

	First two deals		First two offers	
	Supermodular	Submodular	Supermodular	Submodular
LL	0.78		1.00	0.00
LM	0.71	1.00	0.67	0.50
LH	0.29	0.18	0.54	0.43
MM	0.83	0.95	0.83	0.86
MH	0.47	0.61	0.58	0.70
HH	0.75	0.89	0.65	0.79

Notes: Equal split corresponds to between 45% and 55% of the total surplus.

A.2 Equal Splits in Incomplete Information Markets

Figure A.2 provides the analogue of Figure 4 for our incomplete information markets. The patterns are similar.

The figure speaks to the learning dynamics that take place in our markets. Recall that participants cannot learn the surplus from a rejected offer. Therefore, if participants were making offers with the simple aim of learning early on, we would expect to see more generous early offers, with substantially lower rejection rates. As Figure A.2 illustrates, in incomplete-information markets, offers are not significantly more generous initially. In fact, they are not significantly more generous than offers observed in our complete-information markets. Furthermore, rejection rates are not uniformly lower in initial stages of our incomplete-information markets.

Figure A.2: Cumulative Distributions of Proposers' Shares across Offers Made in Incomplete Information Markets

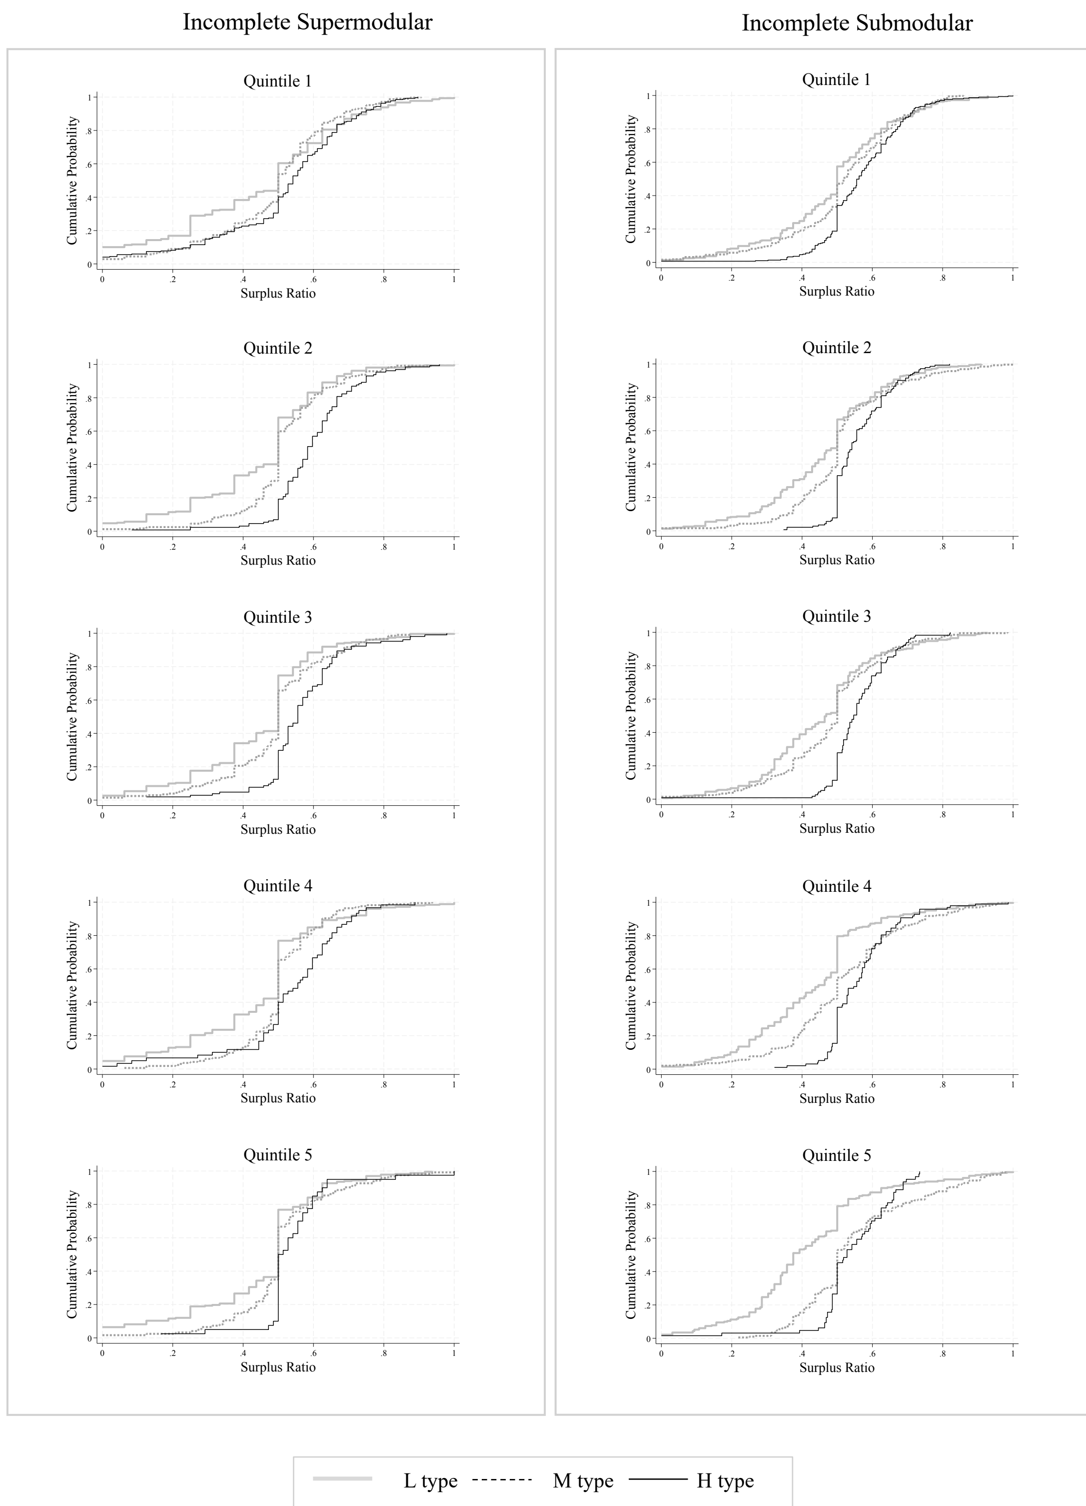

Notes: We focus on offers in which the surplus split is between 0 and 1.

A.3 Satisficing

As a coarser alternative to our analysis in the text, consider using the average payoff for each information and preference structure, along with an agent's type, as that agent's threshold. Table A.2 shows that agents with payoffs below their threshold make offers at significantly higher rates than those with payoffs above it.

Table A.2: Frequency of Offers Made Conditional on the Current Pay Relative to the Satisficing Level

	Complete Information					
	Supermodular			Submodular		
	L type	M type	H type	L type	M type	H type
below cutoff	0.87	0.92	0.98	0.88	0.85	0.91
above cutoff	0.13	0.08	0.02	0.12	0.15	0.09
	Incomplete Information					
	Supermodular			Submodular		
	L type	M type	H type	L type	M type	H type
below cutoff	0.62	0.84	0.93	0.76	0.80	0.92
above cutoff	0.38	0.16	0.07	0.24	0.20	0.08

Notes: Satisficing cutoffs are defined as the average payoff of each type in a market (see Table 3). We exclude cases in which market participants' current payoff is zero.

A.4 Bargaining Styles

Table A.3 reports estimated coefficients from a linear regression model explaining payoffs in different markets and for different types by different features of participants' bargaining behaviors. As can be seen, no features coherently organizes the data beyond participants' positions in the market.

Table A.3: The Effect of Bargaining Styles on Players' Payoffs

Complete Supermodular	Frac Offers by	Nb Matches	Greed	Timing of First Deal	Timing of Last Deal
L type	-0.17	0.82	0.00	-0.04	0.01
M type	-1.77	-2.52**	0.00	0.14*	-0.04
H type	-13.33**	20.10	0.01	-0.02	0.32**
Complete Submodular	Frac Offers by	Nb Matches	Greed	Timing of First Deal	Timing of Last Deal
L type	1.05	-3.67	-0.02	-0.28*	-0.33
M type	1.18	**-6.04***	-0.04**	0.22	-0.19
H type	-2.05	2.70*	-0.01	0.43	0.04
Incomplete Supermodular	Frac Offers by	Nb Matches	Greed	Timing of First Deal	Timing of Last Deal
L type	-1.44	2.20	**0.03***	-0.57	0.15
M type	**-10.64***	**-9.13***	**0.13***	2.34**	-0.21
H type	-5.72	**-9.30***	0.04	-2.35	**-0.49***
Incomplete Submodular	Frac Offers by	Nb Matches	Greed	Timing of First Deal	Timing of Last Deal
L type	-17.72	-11.63**	0.06	0.09	-0.68
M type	**-14.3***	**-8.66***	-0.04	2.06	**-0.82***
H type	**-9.96***	-1.44	0.14**	**-2.43***	0.39

Notes: Each column reports the estimated coefficients from linear regressions in which the dependent variable is participants' payoffs and the independent variables include the constant, the dummies for medium and high types (leaving the low types as the base group), the round and the round squared, two measures of risk attitudes per participant, two measures of altruism per participant, and the variables indicated at the top of the column. Frac Offers captures the fraction of offers extended by each participant type in the specific market. Nb Matches is the total number of matches made throughout the market by each participant type. Greed is the share of surpluses demanded in offers by each participant type. Timing of First and Last Deals is the quintile in which the first and the last deal was made by each participant type. Bootstrapped robust standard errors (clustered at the session level) are reported using the following color-codes: *** indicates significance at 1% level, ** at 5%, * at 10%, and xx indicates non-significant coefficients.

Table A.4: Variance Decomposition of Players' Payoffs, Robustness Check 1

	Complete Info		Incomplete Info	
	Supermodular	Submodular	Supermodular	Submodular
Indicators for L, M, and H types	0.90	0.74	0.74	0.72
Own Offers (frac)	0.01	0.05	0.12	0.16
Greed	0.06	0.17	0.03	0.10
First Deal (quintile)	0.02	0.03	0.00	0.00
Last Deal (quintile)	0.01	0.00	0.11	0.01
Risk measures	0.00	0.00	0.00	0.00
Altruism measures	0.00	0.02	0.00	0.00
overall R^2	0.96	0.83	0.91	0.70
nb obs	$n = 418$	$n = 443$	$n = 390$	$n = 428$

Notes: Shapley R^2 decomposition is reported for indicators for the market position of players as well as bargaining styles. Own Offers (frac) is the average number of offers made by each participant type as a fraction of all offers made in that specific market. Greed is the share of surpluses demanded in offers by a participant in a particular round. First Deal and Last Deal represent the quintile in which the first and last matches were created, respectively. Risk and altruism measures include responses to the auxiliary investment tasks and Dictator games performed at the end of each session.

Tables A.4 and A.5 present two robustness checks of the variance decomposition exercise reported in the main text. Table A.4 relies on a subject–round–specific greed measure rather than a subject–specific one, resulting in a smaller sample size since rounds in which participants made no offers are excluded. Table A.5, instead, uses the fraction of accepted offers (deals) made by each participant, rather than the fraction of offers. As can be seen, results are qualitatively identical and participants' market position still appear as the main factor explaining payoff outcomes.

Table A.5: Variance Decomposition of Players' Payoffs, Robustness Check 2

	Complete Info		Incomplete Info	
	Supermodular	Submodular	Supermodular	Submodular
Indicators for L, M, and H types	0.94	0.83	0.78	0.97
Own Deals (frac)	0.00	0.10	0.08	0.00
Greed	0.01	0.02	0.02	0.00
First Deal (quintile)	0.03	0.03	0.00	0.00
Last Deal (quintile)	0.02	0.01	0.13	0.02
Risk measures	0.00	0.00	0.00	0.00
Altruism measures	0.00	0.01	0.00	0.00
overall R^2	0.96	0.84	0.91	0.68
nb obs	$n = 478$	$n = 480$	$n = 480$	$n = 480$

Notes: Shapley R^2 decomposition is reported for indicators for the market position of players as well as bargaining styles. Own Deals (frac) is the average number of deals made by each participant type as a fraction of all offers made in that specific market. Greed is the share of the surplus demanded in offers by a participant averaged across all rounds. First Deal and Last Deal represent the quintile in which the first and last matches were created, respectively. Risk and altruism measures include responses to the auxiliary investment tasks and Dictator games performed at the end of each session.

B Predictions from Stability with Transfers

As described in Section 3, a stable outcome is an individually rational outcome that has no blocking pairs. When transfers are available, there exist many transfer profiles that sustain stable outcomes in both our supermodular and submodular markets. The associated stable matching, however, is unique for both market types. The table below shows the system of inequalities for each market type that identifies the set of stable transfer profiles. We use x_i to denote the payoff of market participant i, using the labels in Table 1 in the main text. The first column corresponds to our supermodular market, the second to our submodular market.

Supermodular market	Submodular market
$x_{\text{Kiwi}} \geq 0$	$x_{\text{Kiwi}} \geq 0$
$x_{\text{Mango}} \geq 0$	$x_{\text{Mango}} \geq 0$
$x_{\text{Plum}} \geq 0$	$x_{\text{Plum}} \geq 0$
$x_{\text{Blue}} \geq 0$	$x_{\text{Blue}} \geq 0$
$x_{\text{Crimson}} \geq 0$	$x_{\text{Crimson}} \geq 0$
$x_{\text{White}} \geq 0$	$x_{\text{White}} \geq 0$
$x_{\text{Kiwi}} + x_{\text{Blue}} \geq 8$	$x_{\text{Kiwi}} + x_{\text{Blue}} \geq 8$
$x_{\text{Kiwi}} + x_{\text{Crimson}} \geq 16$	$x_{\text{Kiwi}} + x_{\text{Crimson}} \geq 32$
$x_{\text{Kiwi}} + x_{\text{White}} \geq 24$	$x_{\text{Kiwi}} + x_{\text{White}} \geq 56$
$x_{\text{Mango}} + x_{\text{Blue}} \geq 16$	$x_{\text{Mango}} + x_{\text{Blue}} \geq 32$
$x_{\text{Mango}} + x_{\text{Crimson}} \geq 32$	$x_{\text{Mango}} + x_{\text{Crimson}} \geq 48$
$x_{\text{Mango}} + x_{\text{White}} \geq 48$	$x_{\text{Mango}} + x_{\text{White}} \geq 64$
$x_{\text{Plum}} + x_{\text{Blue}} \geq 24$	$x_{\text{Plum}} + x_{\text{Blue}} \geq 56$
$x_{\text{Plum}} + x_{\text{Crimson}} \geq 48$	$x_{\text{Plum}} + x_{\text{Crimson}} \geq 64$
$x_{\text{Plum}} + x_{\text{White}} \geq 72$	$x_{\text{Plum}} + x_{\text{White}} \geq 72$
$x_{\text{Kiwi}} + x_{\text{Mango}} + x_{\text{Plum}} + x_{\text{Blue}} + x_{\text{Crimson}} + x_{\text{White}} \geq 112$	$x_{\text{Kiwi}} + x_{\text{Mango}} + x_{\text{Plum}} + x_{\text{Blue}} + x_{\text{Crimson}} + x_{\text{White}} \geq 160$

C ONLINE APPENDIX: Sample Instructions

Welcome. Welcome to EBEL and thank you for participating in today's experiment. Please place all of your personal belongings away so that we can have your complete attention. Please use the laptops as instructed. In particular, please do not attempt to browse the web or use programs unrelated to the experiment.

Guidelines. You will be paid in private and in cash at the end of the experiment. The amount that you ultimately earn in the experiment depends on your decisions, the decisions of others, and random chance. You have each earned a $10 payment for showing up on time. You will be using laptops for the entire experiment, and all interactions between yourself and others will take place via the laptop?s terminal. Please DO NOT socialize or talk.

Overview. Today's experiment is about matching. The main part of the experiment consists of 11 rounds: one practice round and 10 actual rounds. We will also ask you to complete several simple tasks at the end.

At the beginning of the experiment, you will be randomly assigned a role: either a color or a food. Your role will remain fixed across all rounds of the experiment. At the beginning of each round, you will be randomly assigned a type:

- If you are a color, you can be one of three types: Blue, Crimson, or White.
- If you are a food, you can be one of three types: Kiwi, Mango, or Plum.

Your type (but not your role) may change across each round of the experiment.

In each round, you will be randomly assigned to a group of six: three colors (one Blue, one Crimson, and one White) and three foods (one Kiwi, one Mango, and one Plum). Your goal in each round is to match with an individual from the opposite role. If you are a color, your goal is to match with a food. If you are a food, your goal is to match with a color.

Match Payoff.

	KIWI	MANGO	PLUM
BLUE	8	16	24
CRIMSON	16	32	48
WHITE	24	48	72

A color and food who match together earn a match payoff (in tokens). **You will earn a different payoff depending on whom you match with and how you choose to split the match payoff.**

Example 1:

	KIWI	**MANGO**	PLUM
BLUE	8	**16**	24
CRIMSON	16	32	48
WHITE	24	48	72

Blue and Mango can match together, with the **sum of payoffs** for Blue and Mango yielding **16 tokens**. For example, Blue could receive a payoff of 10 tokens and Mango would then receive a payoff of 6 tokens. Another example: Blue could receive a payoff of 7 tokens and Mango would then receive a payoff of 9 tokens.

Example 2:

	KIWI	MANGO	**PLUM**
BLUE	8	16	24
CRIMSON	16	32	**48**
WHITE	24	48	72

Crimson and Plum can match together, with the **sum of payoffs** for Crimson and Plum yielding **48 tokens**. For example, Crimson could receive a payoff of 30 tokens and Plum would then receive a payoff of 18 tokens. Another example: Crimson could receive a payoff of 23 tokens and Plum would then receive a payoff of 25 tokens.

Example 3:

	KIWI	MANGO	PLUM
BLUE	8	16	24
CRIMSON	16	32	48
WHITE	**24**	48	72

White and Kiwi can match together, with the **sum of payoffs** for White and Kiwi yielding **24 tokens**. For example, White could receive a payoff of 18 tokens and Kiwi would then receive a payoff of 6 tokens. Another example: White could receive a payoff of 9 tokens and Kiwi would then receive a payoff of 15 tokens.

Information about the payoff matrix. At the beginning of each round, we will reshuffle the rows and columns of the matrix.

	KIWI	MANGO	PLUM
BLUE	8	16	24
CRIMSON	16	32	48
WHITE	24	48	72

Examples:

	KIWI	MANGO	PLUM
BLUE	8	16	24
CRIMSON	16	32	48
WHITE	24	48	72

	KIWI	MANGO	PLUM
BLUE	16	32	48
CRIMSON	24	48	72
WHITE	8	16	24

	KIWI	MANGO	PLUM
BLUE	24	16	8
CRIMSON	48	32	16
WHITE	72	48	24

	KIWI	MANGO	PLUM
BLUE	48	32	16
CRIMSON	72	48	24
WHITE	24	16	8

All possible reshufflings are equally likely. At the beginning of each round, you will not know how the matrix was reshuffled, but you will know which payoffs you might receive.

Examples. Suppose you are Blue, and you learn that your payoffs might be 16, 32, 48. It might be that a match with Kiwi generates 16, a match with Mango generates 32, and a match with Plum generates 48.

	KIWI	MANGO	PLUM
BLUE	16	32	48
CRIMSON	24	48	72
WHITE	8	16	24

It might be that a match with Kiwi generates 48, a match with Mango generates 32, and a match with Plum generates 16.

	KIWI	MANGO	PLUM
BLUE	48	32	16
CRIMSON	72	48	24
WHITE	24	16	8

All possible reshufflings are equally likely. While you will know your possible payoffs, you will not know which payoff corresponds to which matching. Your payoffs will appear on the top part of your interface. Note that no other subject in your group observes your possible match payoffs. Each person only sees his or her own possible match payoffs.

Rules of the Experiment. In each round, you will start off unmatched. You are free to make at most one match proposal to any individual of the opposite role at any given time.

A match proposal must specify what your partner will get. If your proposal is accepted, then you will get the **match payoff minus what your partner will get**. You are also free to accept or reject match proposals that you receive from other individuals. Any match proposal that is not accepted or rejected **within 10 seconds** will automatically disappear. When a match is accepted, both the proposer and the receiver see the payoffs they get from the match.

Example 1: Suppose you are Blue, and you know your possible match payoffs are 16, 32, and 48. If you make an offer of 10 to Kiwi, which gets accepted, you might see that you get 6 and Kiwi gets 10. This implies that you and Kiwi have a match payoff of 6+10 = 16.

Example 2: Suppose you are Blue, and you know your possible match payoffs are 16, 32, and 48. If you make an offer of 10 to Kiwi, which gets accepted, you might see that you get 38 and Kiwi gets 10. This implies that you and Kiwi have a match payoff of 38+10 = 48.

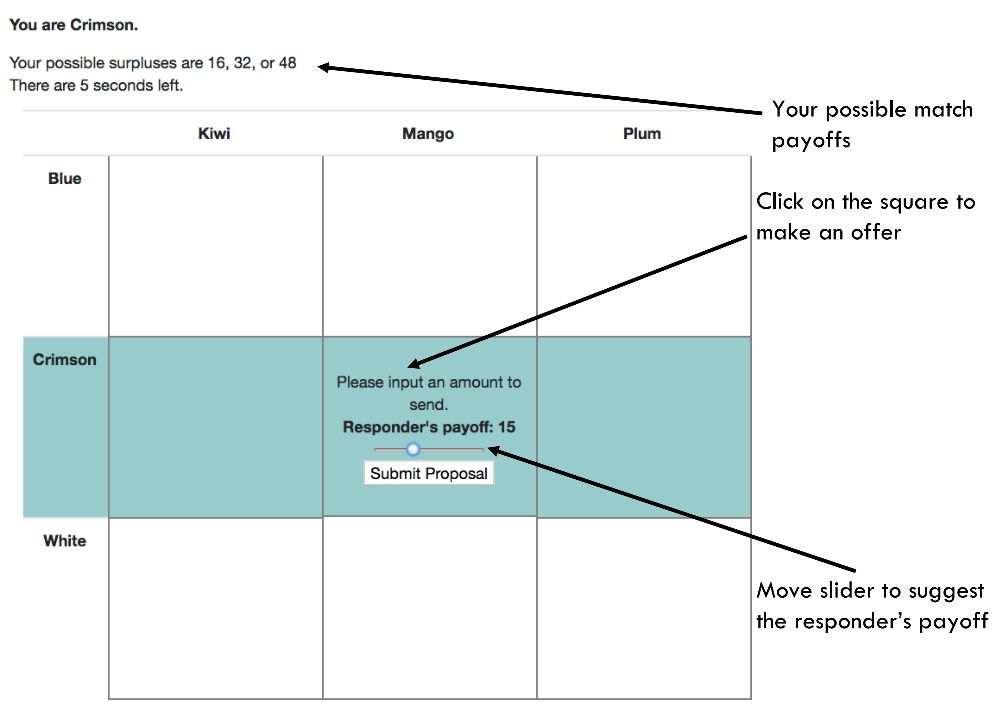

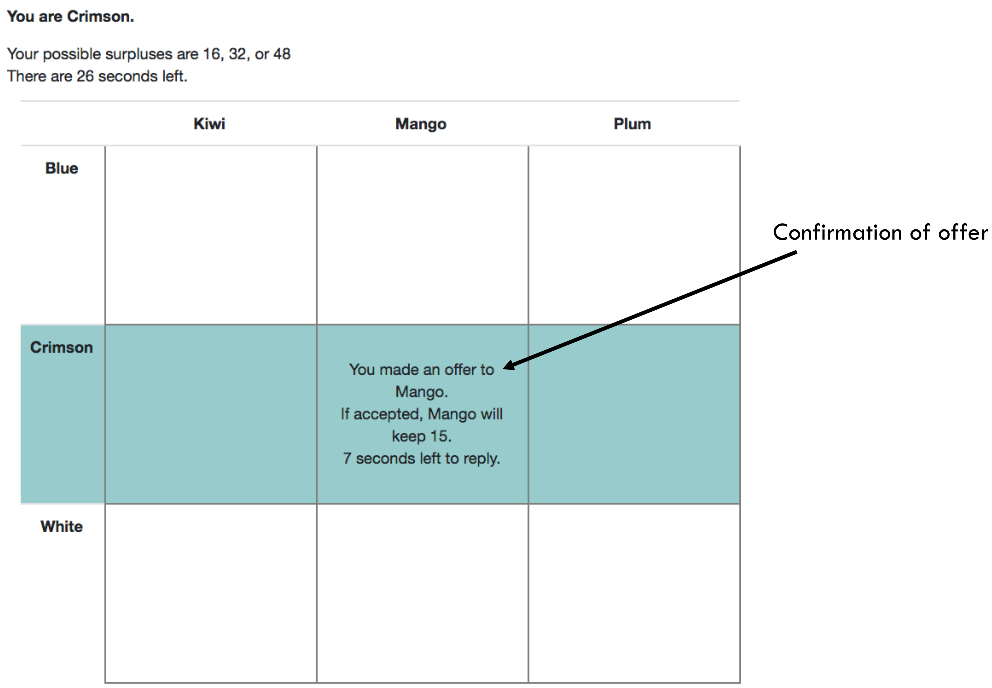
You are Crimson.
Your possible surpluses are 16, 32, or 48
There are 26 seconds left.
Kiwi
Mango
Plum
Blue
Crimson
White
You made an offer to Mango.
If accepted, Mango will keep 15.
7 seconds left to reply.
Confirmation of offer

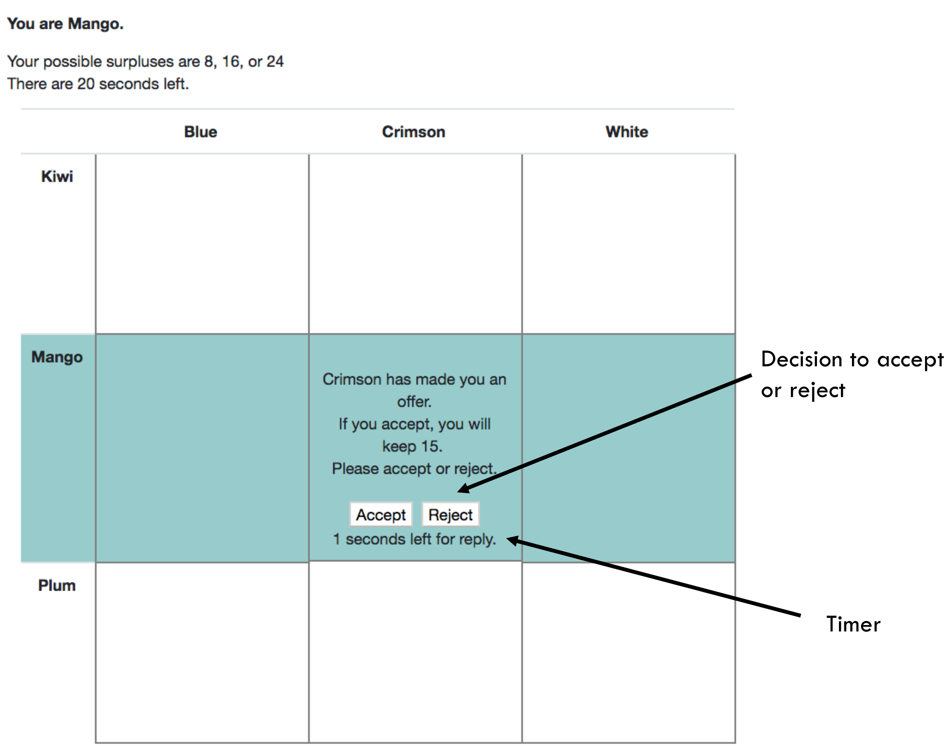
You are Mango.
Your possible surpluses are 8, 16, or 24
There are 20 seconds left.
Blue
Crimson
White
Kiwi
Mango
Plum
Crimson has made you an offer.
If you accept, you will keep 15.
Please accept or reject.
Accept
Reject
1 seconds left for reply.
Decision to accept or reject
Timer

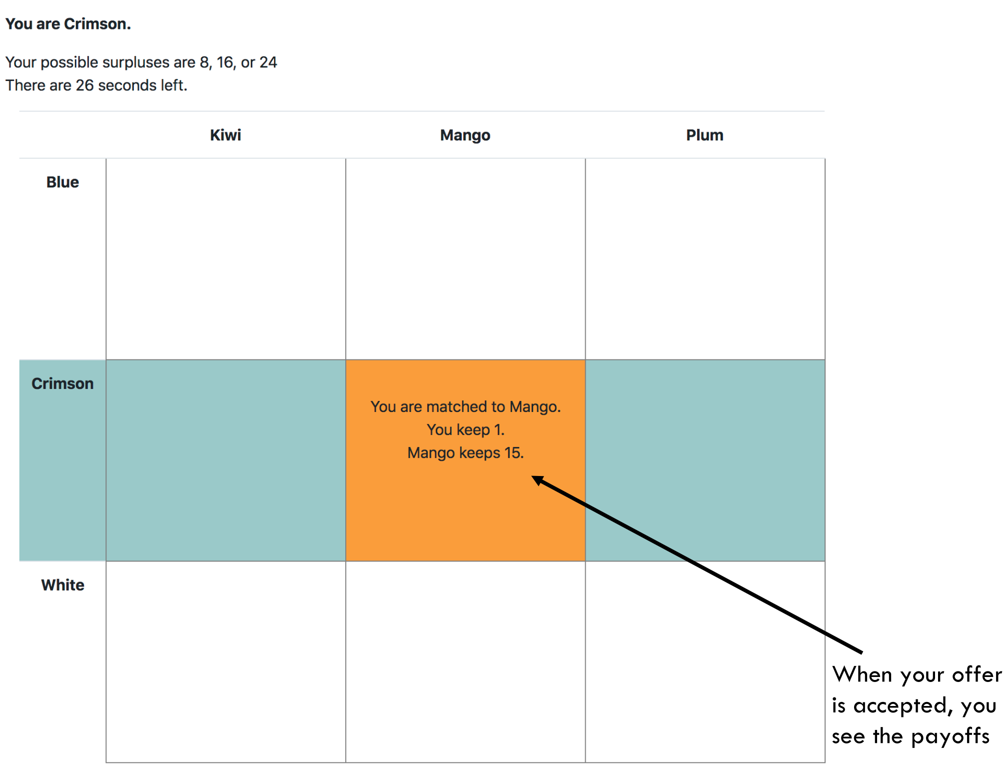
You are Crimson.
Your possible surpluses are 8, 16, or 24
There are 26 seconds left.
Kiwi
Mango
Plum
Blue
Crimson
You are matched to Mango.
You keep 1.
Mango keeps 15.
White
When your offer is accepted, you see the payoffs

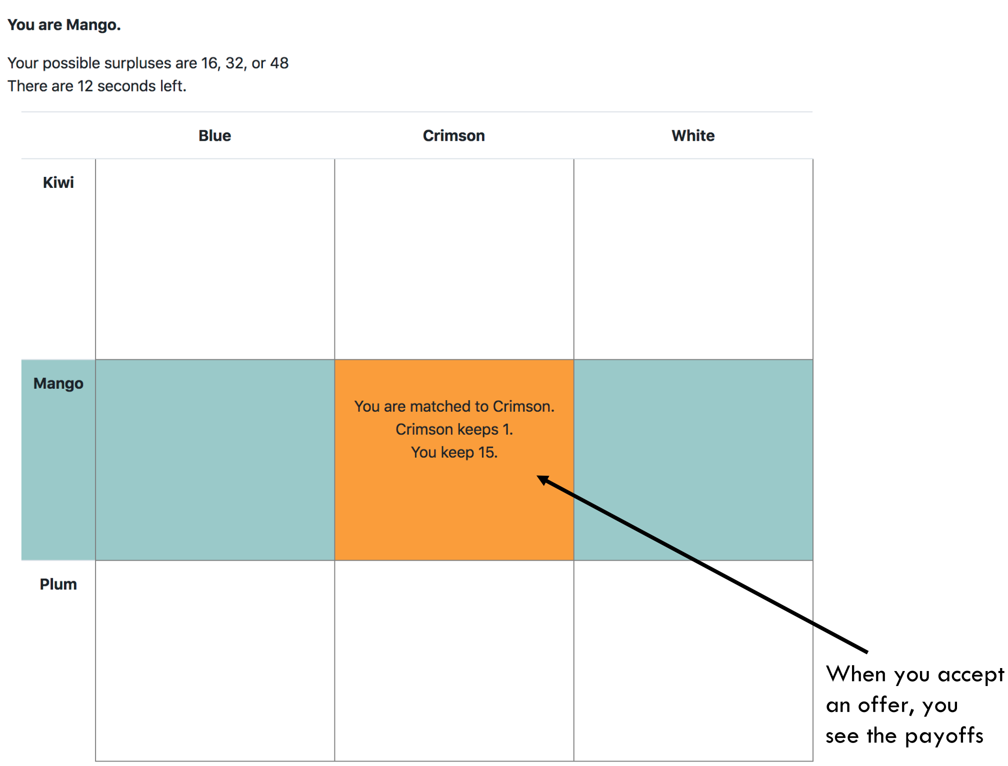
You are Mango.
Your possible surpluses are 16, 32, or 48
There are 12 seconds left.
Blue
Crimson
White
Kiwi
Mango
You are matched to Crimson.
Crimson keeps 1.
You keep 15.
Plum
When you accept an offer, you see the payoffs

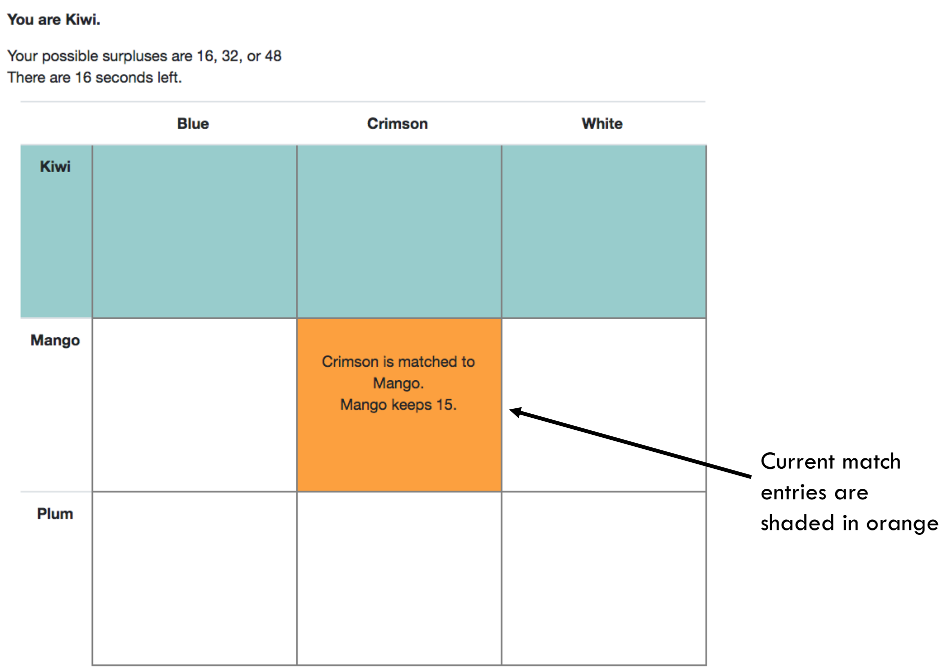
You are Kiwi.
Your possible surpluses are 16, 32, or 48
There are 16 seconds left.
Blue
Crimson
White
Kiwi
Mango
Plum
Crimson is matched to Mango.
Mango keeps 15.
Current match entries are shaded in orange

Rules of the Experiment, cont... If you proposed a match which was accepted and in which your own earnings are negative, then you will have an opportunity to unilaterally break this match by clicking a "cancel" button.

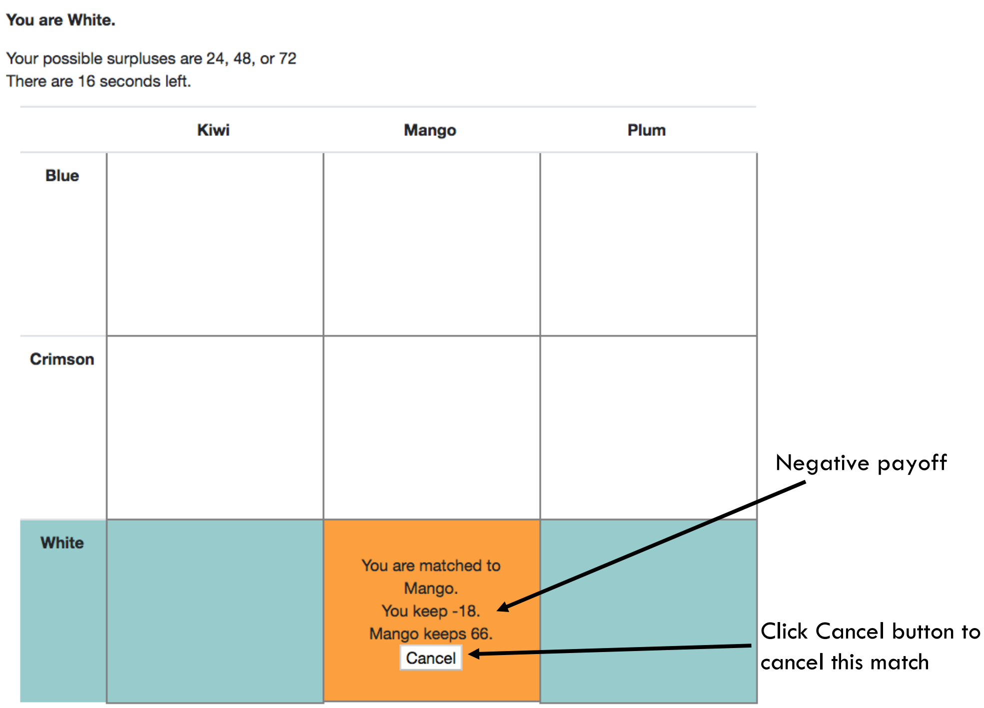

If you are currently matched, you can still make a match proposal to another individual. If the match proposal is accepted, then your current match will be broken. If you are currently matched, you can still receive match proposals from other individuals. If you accept a match proposal, then your current match will be broken. When you accept a match, or someone else accepts a match with you, any proposal you had made or received before will be automatically rejected.

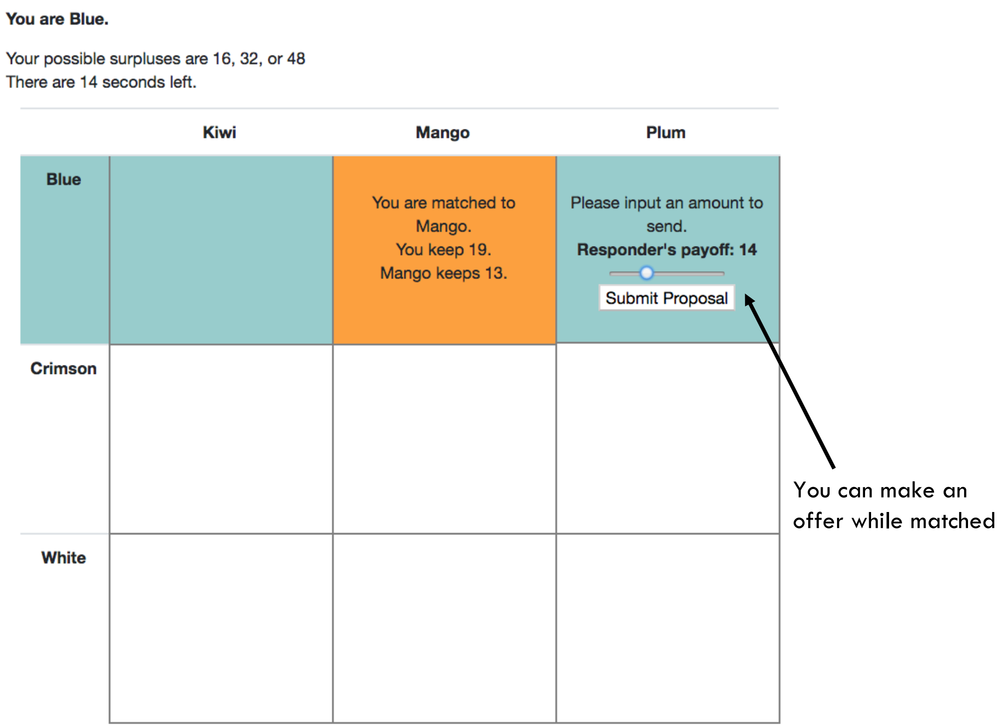
You are Blue.
Your possible surpluses are 16, 32, or 48
There are 14 seconds left.
Kiwi
Mango
Plum
Blue
You are matched to Mango.
You keep 19.
Mango keeps 13.
Please input an amount to send.
Responder's payoff: 14
Submit Proposal
Crimson
White
You can make an offer while matched

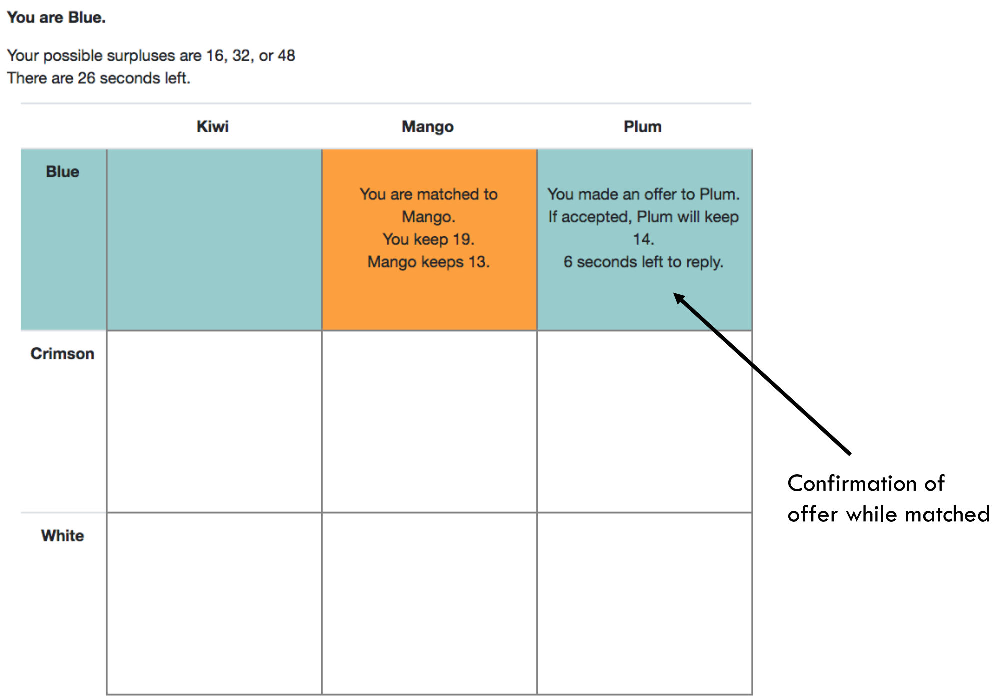
You are Blue.
Your possible surpluses are 16, 32, or 48
There are 26 seconds left.
Kiwi
Mango
Plum
Blue
You are matched to Mango. You keep 19. Mango keeps 13.
You made an offer to Plum. If accepted, Plum will keep 14. 6 seconds left to reply.
Crimson
White
Confirmation of offer while matched

A round ends after 30 seconds of inactivity (i.e., after 30 seconds in which no new match proposals have been made). Whenever you are ready to proceed, you can also press the "Move on to Next Round" button. When at least 5 in your group press the button, you will automatically move to the next round (without waiting the 30 seconds). If any activity occurs after you press the button, the button will get released.

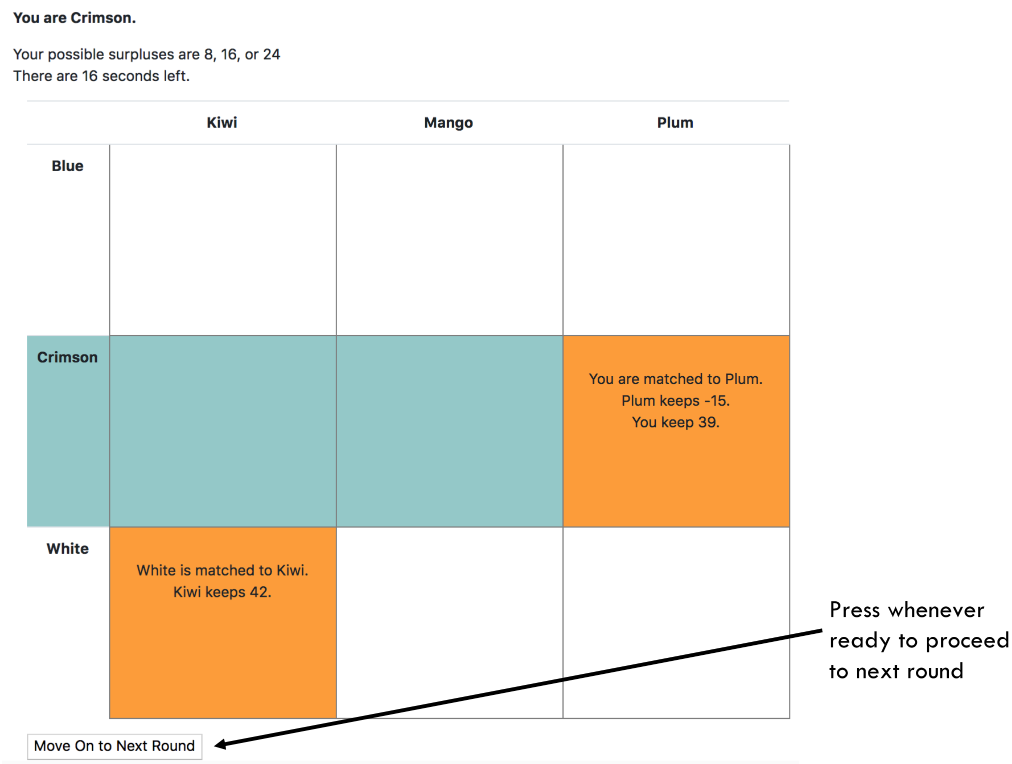

Payoffs for a Round. Your payoff when the round ends is your final payoff for that round. If you are not matched at the end of the round, your final payoff for that round is 0. Your payoff from temporary matches that are made and then broken during a round do not count for your final payoff.

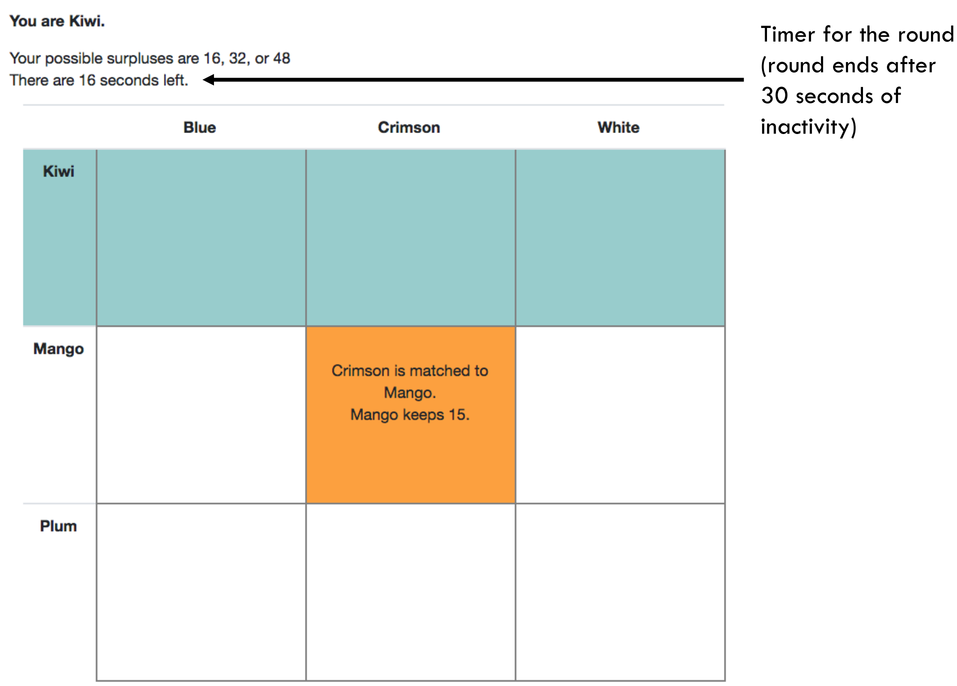

Post-Experiment. At the end of the experiment, you will be paid the sum of your payoffs across rounds (excluding the practice round). You will also be asked to complete several simple tasks at the end. You can earn additional money based on your decisions in these tasks.

Your Earnings. Your total earnings in the experiment are the sum of the following amounts:

- $10 show-up payment
- payoff from 10 main rounds: 10 tokens = 1 dollar
- payoff from the simple tasks:

You need not tell any other participant how much you earned. If there are no questions, we will now begin the actual experiment.