Learning in Matching*

Shan Gui Simin He Gaoji Hu†

October 28, 2026

Abstract

We experimentally study how learning shapes match formation under incomplete information. In decentralized matching markets, an uninformed agent can learn from her own prospective match, from other pairs' matches, or from the absence of a match. We design experiments that isolate these channels and uncover the behavioral challenges each poses. Relative to a no-learning benchmark, the first two channels add little difficulty, whereas learning from the absence of a match substantially impedes correct matching. Within each channel, we also vary how much information learning reveals, and find that partial revelation typically outperforms full revelation in producing correct matches.

JEL Classification: C78, C92, D47, D82, D83.

Keywords: matching, incomplete information, learning, experiment.

*We thank Emiliano Catonini, Yi-Chun Chen, Wooyoung Lim, Qingmin Liu, Vikram Manjunath, Juanjuan Meng, John K.-H. Quah, Tayfun Sönmez, and audiences at various seminars and conferences for their helpful comments and suggestions. **Gui** acknowledges the financial support from the National Natural Science Foundation of China (No. 72503130). **He** acknowledges the financial support from the National Natural Science Foundation of China (No. 72473089). **Hu** acknowledges the financial support from the National Natural Science Foundation of China (No. 72622016, No. 72394391). The experiment is preregistered (AEARCTR-0016825). All remaining errors are our own.

†School of Economics, Shanghai University of Finance and Economics, and the Key Laboratory of Mathematical Economics (SUFE), Ministry of Education, China. **Gui:** guishan@sufe.edu.cn. **He:** he.simin@mail.sufe.edu.cn. **Hu:** hu.gaoji@gmail.com.

1 Introduction

Matching markets determine who gets which jobs, who fills which school places, and who marries whom. These outcomes shape individuals' careers and lives, making it important to understand how matches form. A large literature studies matching markets under complete information, assuming that agents know the characteristics and preferences relevant to matching. In many real-world markets, however, information is incomplete. For example, firms may not fully observe applicants' productivity when hiring, and supervisors may not fully observe prospective students' research ability. Agents must therefore make matching decisions under uncertainty based on what they know, and that knowledge can itself evolve as they observe others' behavior. This creates an interplay between learning and matching.

Agents can learn from others' behavior through different channels. In labor markets, for example, a firm may infer a candidate's attractiveness from the candidate's willingness to accept its offer. Beyond its own match, the firm may also learn from observing where other candidates take jobs, as these appointments can reveal market conditions that the firm itself faces. More demanding still, even the absence of a match can be informative: an unfilled position may lead market participants to update their beliefs if they recognize that the position would have been filled under certain market conditions.

None of these inference problems arises under complete information. A coherent theory of matching under incomplete information must accommodate them because they determine which matches form and which do not. Recent theoretical work, pioneered by Liu et al. (2014) and Liu (2020), has begun to incorporate these inference considerations into stability concepts for matching under incomplete information. Specifically, a market settles at a stable outcome if no one wants to move given what they know, as in the complete-information setting, and additionally, if no one can learn anything further from the fact that no one is moving.

This stream of theories posits full sophistication, as most classical theory does. However, whether real-world agents draw inferences sophisticatedly as above in decentralized markets is an open question. The answer matters both for explaining how matching markets operate and for designing them under incomplete information. In this paper, we use controlled experiments to isolate different learning channels and uncover the behavioral challenges associated with each.

We first identify three canonical learning channels that can arise from publicly observable matching outcomes: learning from a participant's own match, learning from another pair's match, and learning from the absence of another pair's match. We then examine whether human participants can successfully perform the inferences required by these theoretical channels in decentralized matching markets. When they fail, we further ask which agents make mistakes, what types of mistakes they make, and how these mistakes evolve with experience.

Existing experiments have examined decentralized matching, but almost entirely under complete information. The few studies that consider incomplete information do not address the inference problems described above. The most closely related study, Agranov et al. (forthcoming), investigates how information frictions and surplus structures affect efficiency and stability in markets with transfers. Learning in their setting is evaluated by comparing matching outcomes against complete-information concepts. Our experiment advances this literature by making the learning channels themselves the primary treatment variable, and by evaluating matching outcomes against the inference each market requires rather than against a complete-information benchmark.

Experimental design. Studying learning in matching markets is challenging because beliefs are typically unobservable, and observed matching outcomes may reflect multiple underlying mechanisms. Controlled experiments allow us to isolate distinct learning channels by varying the payoff structure while holding the market environment fixed.

We study decentralized one-to-one matching markets with four participants, two

on each side. Three equally likely states determine matching payoffs. Participants on one side observe the realized state and those on the other do not. Participants may freely make, accept, and reject matching proposals and dissolve existing matches until a market settles. Matches are publicly observable as they form and dissolve, whereas proposals and responses behind them remain private. Uninformed participants must therefore infer the state from the matches they observe or from those they do not.

We construct payoff configurations that provide a clean behavioral measure of each inference required. This measure depends only on which matches ultimately form and not on the sequence of proposals and responses that produced them. Hence, our identification strategy requires no commitment to a particular stability concept or matching protocol.

The primary treatment dimension in our design is the learning channel, with different markets isolating different inference requirements. We construct three market types that require learning from a participant's own match (*Own*, for short), from another pair's match (*Other*), or from the absence of a match (*No*). We compare these markets with a benchmark in which no learning is required (*Plain*) and also examine environments in which multiple learning channels operate jointly (*Compound*).

As a second treatment dimension, we vary the extent of information revelation, distinguishing between *Full* revelation, where learning uniquely pins down the realized state, and *Partial* revelation, where learning only narrows down the set of possible states. We aim to identify two opposing mechanisms: *Full* revelation may improve market behavior by increasing certainty and confidence, while *Partial* may perform better because it can support optimal matching even under imperfect belief updating.

A summary of all market types and subtypes is provided in Table 1. In a within-subject design, 192 subjects each played 20 markets, one per round, with different orders to control for order effects.

Experimental results. Our results show that, in the absence of learning (*Plain*), markets nearly always generate the predicted matches. Introducing learning require-

Table 1: Market types based on learning channel and extent of information revelation

	Full revelation	Partial revelation
Plain	–	–
Learning from own match	Own-Full	Own-Partial
Learning from other's match	Other-Full	Other-Partial
Learning from no match	No-Full	No-Partial
Compound learning	Own+Other	–
	Own+No	–

ments has little effect under *Own* and *Other* but substantially impedes predicted match formation under *No* and *Compound*. Most mistakes occur in markets involving *No* (that is, *No* or *Own+No*), indicating that learning from the absence of a match is particularly challenging. A plausible explanation is that *No* requires participants to infer the state from non-occurring matches, which is cognitively more demanding than the inferences required in *Own* or *Other*. The analyses of market dynamics support this interpretation.

Furthermore, the extent of information revelation matters. Within both *Other* and *No*, markets reach predicted outcomes more often in *Partial* than in *Full*. This suggests that the greater scope for optimal matching despite imperfect belief updating dominates the potential benefit of greater certainty and confidence from uniquely identifying the realized state. Thus, more precise information does not necessarily improve matching outcomes when the remaining uncertainty still allows optimal decisions. Together, these findings show that both the learning channels and the extent of information revelation are important determinants of decentralized match formation.

Implications. Our findings have three implications for the study and design of matching markets under incomplete information. First, participants' matching behavior is largely consistent with the predictions, providing behavioral support for theoretical approaches that incorporate learning into stability concepts and market dynamics.

Second, learning is substantially more challenging when participants must draw inferences from the absence of matches. This highlights the importance of accounting not only for the information available to agents but also for their ability to process it, particularly when inference requires counterfactual reasoning. Finally, our results suggest that the design of matching markets should consider these cognitive limitations. Making relevant information more explicit, rather than requiring participants to infer it from the absence of potential matches, may facilitate more effective matching in decentralized markets.

Related Literature. Our paper relates to two strands of literature: matching under incomplete information and experimental studies of decentralized matching markets.

The first strand develops the theoretical foundation for matching under incomplete information. Since the seminal work of Liu et al. (2014), a growing literature has incorporated incomplete information into classical matching theory, substantially extending the framework of Gale and Shapley (1962) and Shapley and Shubik (1971). Subsequent research has studied adaptive matching processes (Chen and Hu, 2020, 2024), Bayesian stability (Liu, 2020; Bikhchandani, 2017; Alston, 2020), epistemic foundations (Pomatto, 2022; Wang, 2023), two-sided incomplete information (Chen and Hu, 2023; Park, 2025), structural properties (Hu, 2026; Peralta, 2025), welfare and efficiency (Chen and Ho, forthcoming), and broader cooperative analyses of incomplete information (Liu, 2025); see Liu (2024) for a recent survey. Our paper complements this literature by identifying the fundamental learning channels that arise as markets evolve and providing, to our knowledge, the first experimental test of these channels. Because our identification strategy does not rely on a particular stability concept, it applies to a broad class of theories of matching under incomplete information.[1]

The second strand of literature studies decentralized matching markets experimentally. Existing experiments examine whether decentralized markets converge to stable outcomes across a variety of market institutions (Nalbantian and Schotter, 1995;

[1]We provide one such theory reflecting our experimental design in Supplemental Appendix D.

Otto and Bolle, 2011; Pais, Pintér and Veszteg, 2020; Dolgopolov et al., 2024; He et al., 2024; Agranov et al., forthcoming; Echenique, Robinson-Cortés and Yariv, 2025). Most of these studies assume complete information when evaluating market outcomes, even when incomplete information is introduced experimentally. For example, Nalbantian and Schotter (1995) and Agranov et al. (forthcoming) study transferable-utility markets with incomplete information, while Pais, Pintér and Veszteg (2020) consider a non-transferable environment in which participants know only their own preferences. Nevertheless, all of these papers evaluate market outcomes using complete-information stability or efficiency concepts. By contrast, our experiment is designed around theories of matching under incomplete information by making learning channels the primary treatment variable.

2 Experimental Protocol

In our experiments, participants take part in a sequence of decentralized matching markets under incomplete information. In these markets, some participants face payoff uncertainty, as they do not know the payoffs from potential matches. Therefore, it is nontrivial for them to decide whether and with whom to match, and they must draw inferences based on the observables in each market. The markets vary along two dimensions: the required learning channels and the extent of information revelation. This section describes the market environment, canonical learning channels, treatment variations, and the experimental procedure. The experimental instructions are provided in Supplemental Appendix B.1.

2.1 Matching environment

We study two-sided, one-to-one matching markets. Each market consists of four participants, two on each side, represented by two fruits and two colors. These sides can represent workers and firms in a labor market, as well as women and men in a marriage market. Each participant can either remain unmatched or match with a

single participant from the opposite side.

Each market features one of three equally likely states represented by three types of weather: Sunny, Cloudy, or Rainy. State information is one-sided incomplete: fruits observe the weather, whereas colors do not. In labor markets, for instance, workers may have private information about their effort costs or productivity that firms cannot observe. Matching payoffs depend on the partnership and the realized state. For example, Table 2 reports payoffs for each possible match under each weather, with the fruit's payoff listed first. If Lemon matches with Green and the weather is Sunny, Lemon earns 1 and Green earns 2. By contrast, a Lemon–Yellow match yields a payoff of -10 to both participants in every state. Remaining unmatched yields a payoff of zero.

Table 2: Matching market: Example 1 (realized weather: Sunny)

	Yellow		Green		
Lemon	-10	-10	1	2	Sunny
	-10	-10	1	4	Cloudy
	-10	-10	-1	-8	Rainy
Mango	-10	-10	-10	-10	Sunny
	-10	-10	-10	-10	Cloudy
	-10	-10	-10	-10	Rainy

The decentralized matching process works as follows. At the beginning of each market, a payoff table is displayed publicly, a weather is randomly drawn and shown only to fruits, and all participants start unmatched. Participants may freely make, accept, and reject proposals to and from the opposite side, as well as dissolve existing matches. When a participant makes a proposal, the receiver has 15 seconds to accept or reject it. The proposer cannot make another proposal until the current one is resolved. An accepted proposal creates a temporary match. A proposal expires if it is rejected or receives no response within 15 seconds. Matched participants can

unilaterally dissolve the partnership; a newly formed match supersedes any existing matches involving either participant. The roles of proposer and receiver are not fixed: A participant awaiting a response to her own proposal may still accept or reject incoming proposals. Similarly, a participant who has received a proposal may make a proposal to another participant before responding.

Within each market, participants observe matches as they form and dissolve, and a full history of temporary matches is also publicly observable. The market ends when temporary matches remain unchanged for 60 seconds.[2] All temporary matches become final when the market ends. In addition to this public information, each participant privately observes all proposals she sends and receives. An experimental screenshot is provided in Supplemental Appendix B.2.

2.2 Three canonical learning channels

Because participants on the color side do not observe the state and thus face payoff uncertainty, a good matching decision requires inference. This section introduces three canonical channels through which participants draw inferences: from a participant's own match, from another pair's match and, most demanding of all, from the absence of a match.

Let us first examine the matching market example in Table 2. Lemon observes the realized state Sunny and may propose to Green because matching yields a positive payoff. Green, who is uninformed, is unwilling to match with Lemon due to the possibility of receiving -8 under Rainy.[3] But she observes Lemon's proposal and infers that the weather cannot be Rainy, since Lemon would not propose in that weather. Consequently, Green accepts Lemon's proposal because matching yields a positive payoff for her under both Sunny and Cloudy. The same match and inference arise if Green first proposes to Lemon and learns from Lemon's acceptance. Either way,

[2]Alternatively, after the first 60 seconds, the market can end if all participants press the "Agree to Proceed" button.

[3]To be more precise, the expected payoff under the prior is negative.

Green learns from her own match formation.

Learning can also occur through a more demanding channel. In the example of Table 3, Mango is willing to match with Green in every weather. Green therefore cannot learn about the weather from Mango's proposal or from his acceptance of her probing proposal, and the possibility of the negative payoff of -8 in Rainy prevents the Mango–Green match from forming. In contrast, Yellow needs no inference at all: although she is uninformed, she is willing to match with Lemon in every weather, so the Lemon–Yellow match forms. Once Green observes it, she infers that the weather must be Sunny or Cloudy, because Lemon would not form this match under Rainy. Consequently, Green rules out the negative payoff and becomes willing to match with Mango. Compared with the example in Table 2, here Green learns nothing from Mango's willingness to match but instead learns from the Lemon–Yellow match. This inference is more demanding because it requires Green to extract information from the behavior of another pair that is not directly relevant to Green's own matching decision.

Table 3: Matching market: Example 2 (realized weather: Sunny)

	Yellow		Green		
Lemon	1	2	-10	-10	Sunny
	1	4	-10	-10	Cloudy
	-1	8	-10	-10	Rainy
Mango	-10	-10	1	2	Sunny
	-10	-10	1	4	Cloudy
	-10	-10	1	-8	Rainy

Learning can be even more demanding. In the example of Table 4, Mango is willing to match with Green in every weather, so Green learns nothing from his willingness. Moreover, because Lemon is unwilling to match with Yellow under Sunny, they do not form a match, ruling out Green's learning from witnessing another pair's

match. However, Green can recognize that Lemon and Yellow would have matched under Cloudy or Rainy. Observing the absence of the Lemon–Yellow match therefore allows Green to infer that the weather is Sunny. Consequently, Green rules out the possibility of receiving negative payoffs and becomes willing to match with Mango. Compared with the markets in Tables 2 and 3, here Green must recognize that a match would have formed in some states and draw an inference from its absence. This inference is even more demanding because it requires Green to extract information from the absence of matching activity, rather than from an observed action or match.

Table 4: Matching market: Example 3 (realized weather: Sunny)

	Yellow		Green		
Lemon	−1	2	−10	−10	Sunny
	1	4	−10	−10	Cloudy
	1	8	−10	−10	Rainy
Mango	−10	−10	1	2	Sunny
	−10	−10	1	−4	Cloudy
	−10	−10	1	−8	Rainy

2.3 Treatment variation

In our experiment, each subject experiences a sequence of decentralized matching markets, varying along the learning channel and the extent of information revelation. The payoff table and realized state determine which treatment a market belongs to.

Our primary treatment variation is across the three learning channels introduced in Section 2.2: from a participant's own match (*Own*, for short), from another pair's match (*Other*) and from the absence of a match (*No*). As a benchmark, we also construct markets in which learning is unnecessary for decision making because all states imply the same optimal decision (*Plain*). Finally, we construct markets in which an uninformed participant must draw inferences through multiple channels to evaluate

a potential match (*Compound*). We study Compound markets for two reasons. First, they allow us to examine more complex learning environments than those with a single channel. Second, interactions between learning channels may generate additional market behavior that cannot be observed when only one channel operates. Overall, we have five market types.

While the learning channel constitutes our primary treatment dimension, we also vary the extent of information revelation. In some markets, learning allows uninformed participants to pin down the realized state, which we refer to as *full revelation (Full)*. In others, learning only narrows down the set of possible states without identifying the exact state, which we refer to as *partial revelation (Partial)*. This distinction may affect matching decisions through two opposing mechanisms. On the one hand, residual uncertainty under partial revelation may reduce participants' confidence and impair decision quality (*certainty-and-confidence* mechanism). On the other hand, partial revelation may facilitate matching when multiple possible states favor the same decision, allowing participants to make an optimal matching decision even with imperfect belief updating (*favorable-state* mechanism). We therefore investigate whether the extent of information revelation systematically affects matching behavior.[4]

For our experimental treatments, we vary the extent of information revelation within each of the three canonical learning channels, yielding six treatments. For Compound, our $2 \times 2 \times 3$ matching markets allow us to construct payoff configurations in which two learning channels operate jointly: Own+Other and Own+No.[5] The two Compound treatments feature full revelation only. Therefore, we have two subtypes for each market type, except for Plain. Table 5 illustrates the market types and subtypes. In all markets except Plain, Green must learn through one or two channels.

[4]By construction, under partial revelation, all possible states remaining after learning imply the same optimal decision for the uninformed participant. Thus, partial revelation preserves decision certainty despite leaving some uncertainty about the underlying state.

[5]Jointly identifying Other and No requires at least three distinct pairs and is therefore not feasible in our $2 \times 2 \times 3$ markets.

Table 5: Illustrative payoff tables for market subtypes (realized weather: Sunny in all)

	Yellow		Green		
Lemon	−10	−10	−10	−10	Sunny
	−10	−10	−10	−10	Cloudy
	−10	−10	−10	−10	Rainy
Mango	−10	−10	1	2	Sunny
	−10	−10	1	4	Cloudy
	−10	−10	1	8	Rainy

(a) Plain

	Yellow		Green		
Lemon	1	2	−10	−10	Sunny
	1	4	−10	−10	Cloudy
	1	8	−10	−10	Rainy
Mango	−10	−10	1	2	Sunny
	−10	−10	1	4	Cloudy
	−10	−10	1	8	Rainy

(b) Plain

	Yellow		Green		
Lemon	−10	−10	1	2	Sunny
	−10	−10	−1	−4	Cloudy
	−10	−10	−1	−8	Rainy
Mango	−10	−10	−10	−10	Sunny
	−10	−10	−10	−10	Cloudy
	−10	−10	−10	−10	Rainy

(c) Own-Full

	Yellow		Green		
Lemon	−10	−10	1	2	Sunny
	−10	−10	1	4	Cloudy
	−10	−10	−1	−8	Rainy
Mango	−10	−10	−10	−10	Sunny
	−10	−10	−10	−10	Cloudy
	−10	−10	−10	−10	Rainy

(d) Own-Partial

	Yellow		Green		
Lemon	1	2	−10	−10	Sunny
	−1	4	−10	−10	Cloudy
	−1	8	−10	−10	Rainy
Mango	−10	−10	1	2	Sunny
	−10	−10	1	−4	Cloudy
	−10	−10	1	−8	Rainy

(e) Other-Full

	Yellow		Green		
Lemon	1	2	−10	−10	Sunny
	1	4	−10	−10	Cloudy
	−1	8	−10	−10	Rainy
Mango	−10	−10	1	2	Sunny
	−10	−10	1	4	Cloudy
	−10	−10	1	−8	Rainy

(f) Other-Partial

	Yellow		Green		
Lemon	−1	2	−10	−10	Sunny
	1	4	−10	−10	Cloudy
	1	8	−10	−10	Rainy
Mango	−10	−10	1	2	Sunny
	−10	−10	1	−4	Cloudy
	−10	−10	1	−8	Rainy

(g) No-Full

	Yellow		Green		
Lemon	−1	2	−10	−10	Sunny
	−1	4	−10	−10	Cloudy
	1	8	−10	−10	Rainy
Mango	−10	−10	1	2	Sunny
	−10	−10	1	4	Cloudy
	−10	−10	1	−8	Rainy

(h) No-Partial

	Yellow		Green		
Lemon	1	2	−10	−10	Sunny
	−1	4	−10	−10	Cloudy
	1	8	−10	−10	Rainy
Mango	−10	−10	1	2	Sunny
	−10	−10	1	−4	Cloudy
	−10	−10	−1	−8	Rainy

(i) Own+Other

	Yellow		Green		
Lemon	−1	2	−10	−10	Sunny
	1	4	−10	−10	Cloudy
	−1	8	−10	−10	Rainy
Mango	−10	−10	1	2	Sunny
	−10	−10	1	−4	Cloudy
	−10	−10	−1	−8	Rainy

(j) Own+No

We use a within-subject design in which each participant experiences one practice round and 20 payoff-relevant markets, one per round. The 20 markets are generated from two sets of parameters, each with a structure similar to that shown in Table 5. To control for order effects, we use a Latin-square design that prescribes four treatment orders. Appendix A reports the parameters and the design of treatment orders.

2.4 Identification strategy

Our payoff configurations allow us to cleanly identify the learning channels, provided that participants are payoff maximizers, that is, they pursue matches that yield positive expected payoffs whenever such matches exist. To illustrate, consider the payoff configurations in Table 5. Three features are common to all 20 markets. First, at most two matches are feasible, Lemon–Yellow and Mango–Green; the remaining two pairs, Lemon–Green and Mango–Yellow, are infeasible because their payoffs are uniformly negative. Second, whenever Lemon–Yellow yields positive payoffs to both players under Sunny, its formation requires no inference by Yellow because her optimal decision is the same in every weather state. Third, in every market except Plain, the Mango–Green match requires Green to draw the relevant inference correctly through the learning channel(s) we designed.

Identification therefore centers on whether Mango–Green forms as predicted. However, because the formation of this match may depend on whether Lemon–Yellow forms in Other and No markets, identification also requires examining the full market outcome. Importantly, we do not require Lemon–Yellow to form before Mango–Green in order to identify learning in the Other markets, even though this ordering provides the most intuitive illustration. If Mango–Green forms first, Green can still update after observing the subsequent formation of Lemon–Yellow and choose to maintain the match. Thus, our identification strategy depends on which matches ultimately form, rather than on a particular sequence of proposals or match formation.

2.5 Procedures

The experiment was conducted at the Shanghai University of Finance and Economics in September and October 2025. We recruited 192 subjects from the Economics Lab's subject pool through Ancademy, a platform for social science experiments. We ran eight sessions of 24 subjects each. In each session, subjects were randomly divided into two independent matching groups of 12. Each subject participated in only one session. Treatment orders were randomly assigned at the matching group level. Appendix A.2 reports all randomization details.

The experiment was computerized using oTree (Chen, Schonger and Wickens, 2016) and conducted in Chinese. English translations of the instructions and screenshots appear in the Supplemental Appendix B. Participants received computerized instructions and completed comprehension questions before the experiment. The same experimenter oversaw all sessions.

At the start of the experiment, six subjects in each matching group were randomly assigned to the fruit role and six to the color role, with roles fixed throughout. At the start of each round, the 12 subjects were randomly partitioned into three markets of four and assigned the labels Lemon, Mango, Yellow, and Green. Each subject received an initial endowment of 10 points before the decentralized matching process started. The total round payoff equaled the endowment plus the matching payoff and was always non-negative.

At the end of the experiment, subjects completed a short demographic survey. Four of the 20 rounds were randomly selected for payment, with points denominated directly in Chinese yuan (CNY). Average earnings were 70.51 CNY, including a 20 CNY participation bonus (approximately 10.53 USD). Sessions lasted 50–60 minutes.

3 Theoretical Foundation: A Discussion

This section clarifies the theoretical considerations underlying our experimental design. Classical matching models, such as those of Gale and Shapley (1962) and Shapley and Shubik (1971), assume complete information, under which agents know the desirability of potential partners. In many matching markets, however, this knowledge is unavailable. Participants must therefore infer the value of a match from what others do, as illustrated in Section 2.2. These inference problems do not arise under complete information, but under incomplete information they shape which matches form. A coherent theory of stable matchings under incomplete information must therefore accommodate inferences at different levels of sophistication—a program begun in recent work pioneered by Liu et al. (2014) and Liu (2020).

Our experiment is not designed to test any particular stability theory. Instead, we isolate canonical learning channels that arise across different theoretical approaches. Supplemental Appendix D sketches a theoretical framework that accommodates Own, Other, and No. In terms of stability concepts, broadly speaking, Own enters through the definition of a blocking pair. That is, when an uninformed agent evaluates a potential match with an informed one, she need only consider states under which the informed agent is willing to match. No is then used to characterize market stability. That is, a market can be viewed as stable only when no blocking pair forms (and no existing match dissolves), and additionally, no agent can further update beliefs from the absence of activity. In terms of the dynamic matching process, agents may update information through all three channels, separately or jointly, in decentralized matching markets.

What participants can learn depends crucially on what they can observe. In our experiment, participants publicly observe the formation and dissolution of matches, but not the detailed proposals and responses that generate them. This information structure reflects observability in many real-world markets. For example, in mar-

riage and labor markets, realized matches are often publicly observable, whereas the underlying proposals, acceptances, and negotiations are not. Moreover, our design follows the cooperative (or reduced-form) perspective in matching theory dating back to Gale and Shapley (1962), which abstracts from the details of the matching process and focuses on its outcomes. In the market illustrated in Table 2, for instance, the Lemon–Green match yields the same matching outcome and information update for Green regardless of who initiates the proposal. As Liu (2025) argues, reduced-form models isolate the economically relevant implications of strategic interaction without imposing unnecessary assumptions on the adjustment process.

Once attention is restricted to publicly observable match evolution, the three canonical learning channels illustrated in Section 2.2 are exhaustive. Any information conveyed by match evolution must arise from one of three sources: an agent's own match, others' matches, or the absence of a match. More complex forms of learning may combine these channels, but they can still be decomposed into these three building blocks.[6]

Finally, we discuss two key choices in our experimental design motivated by theoretical considerations. First, our decentralized matching markets satisfy a convergence property. Existing learning-and-blocking dynamics (Roth and Vande Vate, 1990; Chen and Hu, 2020, 2024) show that adaptive matching processes in which agents continually pursue profitable deviations converge to outcomes satisfying the corresponding stability concepts (Gale and Shapley, 1962; Liu et al., 2014; Liu, 2020). Although we do not impose any stability notion a priori, each of our experimental markets should converge to a unique outcome, which underpins our identification strategy in Section 2.4.[7]

[6]To facilitate the identification of learning channels, we focus on match formation rather than combining formation with dissolution; accordingly, all markets begin with all agents unmatched. Moreover, because the absence of one's own match provides no additional information in our experimental configurations, the No markets focus on learning from the absence of another pair's potential match.

[7]As with the stability concepts, we avoid committing to a particular theory of dynamics. The

Second, our experimental design is robust to a broad range of risk preferences. Although some theories assume risk neutrality, our design does not rely on this assumption. In all treatments in Section 2.3, summarized in Table 5, even moderately risk-seeking participants may avoid matching unless they can rule out states with negative payoffs, because the potential losses are larger in magnitude than the corresponding gains. At the other extreme, highly risk-averse participants who match only when a positive payoff is guaranteed are also accommodated, because uninformed participants in our design must rule out all negative-payoff states before matching. In either case, the experimental design and analysis remain unchanged.

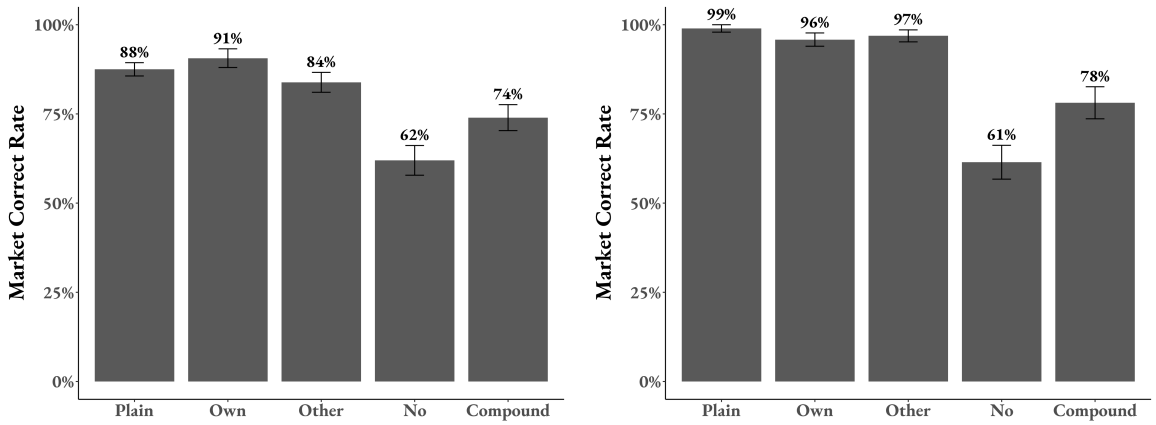
4 Results

Under our identification strategy, whether learning occurs as predicted depends only on which matches ultimately form. In each market, we call a market outcome *correct* if every match in it yields positive payoffs to both participants in the realized state, and any match that fails to form would yield a negative payoff to at least one participant. An individual participant's outcome is *correct* if it coincides with her assignment under the correct market outcome. We use this benchmark to examine the experimental results along three dimensions: (i) at the market level, whether each market reaches the correct matching outcome; (ii) at the individual level, who makes mistakes and what types of mistakes they make; and (iii) from the dynamic perspective, the timing and order of correct match formation. Together, these analyses show whether and how participants learn through each channel and help identify the sources of their failures.

path-to-stability literature since Roth and Vande Vate (1990) typically assumes naive or myopic agents who form or revise a match whenever a profitable opportunity arises. By contrast, decentralized matching models such as Lauermann and Nöldeke (2014) and Ferdowsian, Niederle and Yariv (2025) allow for strategic behavior. This distinction is largely muted in our markets because each agent has at most one acceptable partner; the relevant decision is whether to match rather than whom to match with.

4.1 Market-level performance

We first examine whether market correct rates vary across the primary market types. Figure 1 reports market correct rates for *Plain*, *Own*, *Other*, *No*, and *Compound*, using all 20 rounds in panel (a) and the last 10 rounds in panel (b). Bars report means and error bars indicate ± one standard error, both computed at the 12-player matching-group level.



(a) All 20 rounds (b) Last 10 rounds

Figure 1: Market correct rates by market type

Our results show that reaching the correct market outcome is hardest under the market type *No* and is also relatively difficult under *Compound*, whereas *Own* and *Other* show little additional difficulty relative to *Plain*. Across all 20 rounds, no market type reaches a 100% correct rate. The market correct rate is the lowest under *No* (62%), and is significantly lower than under *Compound* (74%), which is in turn lower than under *Plain* (88%). By contrast, the correct rates under *Own* and *Other* do not differ significantly from *Plain*. This general pattern persists in the last 10 rounds. Market correct rates under *Plain*, *Own*, and *Other* are all close to 100% and indistinguishable from one another, suggesting that any initial difficulty associated with these channels largely disappears with experience. In contrast, the correct rate remains at 61% under *No*, significantly lower than the 78% under *Compound*; both

remain significantly below *Plain*, *Own*, and *Other*. We report the corresponding statistical tests in Table 6.

Table 6: Market correct rates: pairwise tests across market types

	Own	Other	No	Compound
Plain	0.071	0.293	< 0.001	< 0.001
Own		0.015	< 0.001	0.002
Other			0.001	0.016
No				0.013

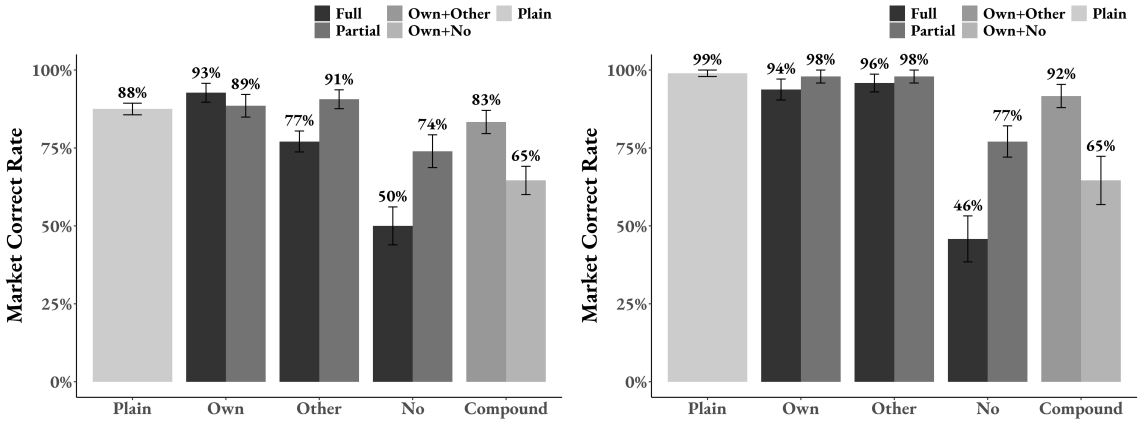
(a) All 20 rounds

	Own	Other	No	Compound
Plain	0.083	0.317	< 0.001	0.002
Own		0.655	< 0.001	0.006
Other			< 0.001	0.001
No				0.021

(b) Last 10 rounds

Notes: Each cell reports the p-value from a two-sided Wilcoxon signed-rank test at the matching-group level ($n = 16$).

We next decompose the correct rate of each market type into its subtypes. For *Own*, *Other*, and *No*, the subtypes differ in the extent of information revelation, and could be either *Full*, in which learning uniquely identifies the realized state, or *Partial*, in which some uncertainty remains. For *Compound*, the two subtypes differ in their composition and could be either *Own+Other* or *Own+No*. Plain has no subtype. Figure 2 reports market correct rates by subtype for all 20 rounds in panel (a) and the last 10 rounds in panel (b). As above, bars report means and error bars indicate $\pm$ one standard error at the matching-group level.



(a) All rounds

(b) Last 10 rounds

Figure 2: Market correct rates by subtypes

We find that *Partial* markets perform better than *Full* markets under *Other*

and *No*, but not under *Own*. And *Own+No* is more difficult than *Own+Other*. *Own* shows no significant difference in market correct rates between Full and Partial, either across all 20 rounds or in the last 10 rounds. Under *Other*, Partial has a significantly higher market correct rate than Full across all 20 rounds, by 14 percentage points (pp), but this difference disappears in the last 10 rounds as participants gain experience. Under *No*, Partial has a significantly higher market correct rate than Full in both samples, by 24 pp across all 20 rounds and by 31 pp in the last 10 rounds. For the two Compound subtypes, the market correct rate is 18 pp higher under *Own+Other* than under *Own+No* across all 20 rounds, and this difference remains significant in the last 10 rounds, suggesting that the difficulty associated with *No* persists in Compound learning.[8] Table 7 reports the corresponding statistical tests.

Table 7: Comparisons across subtypes

		Within Type	Own+Other	Own+No
Own	Full	0.294	0.078	< 0.001
	Partial		0.199	0.001
Other	Full	0.002	0.115	
	Partial		0.121	
No	Full	0.011		0.011
	Partial			0.228
Compound		0.001		

(a) All 20 rounds

		Within Type	Own+Other	Own+No
Own	Full	0.317	0.706	0.006
	Partial		0.180	0.002
Other	Full	0.564	0.317	
	Partial		0.180	
No	Full	0.005		0.082
	Partial			0.302
Compound		0.009		

(b) Last 10 rounds

Notes: Each cell reports the p-value from a two-sided Wilcoxon signed-rank test at the matching-group level ($n = 16$).

Result 1 (market-level correct rates)**.** *For market types, the ranking of market correct rates is Plain = Own > Other > Compound > No across all 20 rounds, and Plain = Own = Other > Compound > No in the last 10 rounds.*

For market subtypes, Partial and Full do not differ under Own. Under Other, Partial outperforms Full across all 20 rounds, but the difference disappears in the last

[8]The correct rates of the Compound subtypes generally lie between those of their corresponding non-Compound channels. In particular, *Own+No* performs significantly worse than *Own*, but does not differ significantly from *No*. Across all 20 rounds, *Own+No* differs significantly only from *No-Full* among the two *No* subtypes. *Own+Other* does not differ significantly from either *Own* or *Other*.

10 rounds. Under No, Partial outperforms Full in both samples. Within Compound, Own+Other outperforms Own+No in both samples.

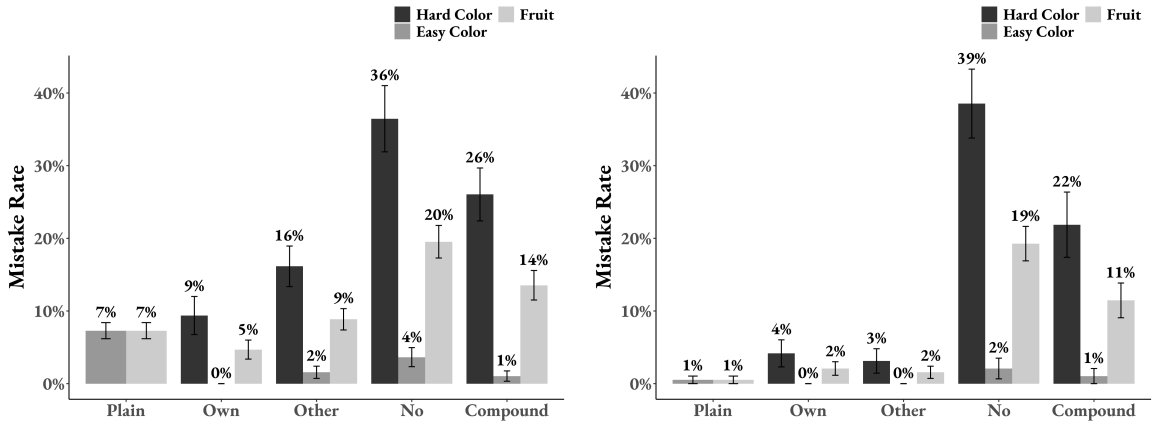
4.2 Individual mistakes

The market-level results show that incorrect outcomes are concentrated in *No* and *Compound* and persist with experience. To understand where these failures originate, we turn to individual mistakes. We say that an individual makes a *mistake* if her final match is not *correct*. We proceed in three steps. First, across market types, we analyze mistake rates by participant role, distinguishing fruits from colors and further classifying colors as *Hard* or *Easy* depending on whether learning is required to achieve the correct match.[9] These results motivate a focus on *Hard* colors, and underscore the importance of learning for match formation. Second, we examine how their performance differs across market subtypes. Finally, we study the nature and evolution of mistakes by asking what types of mistakes participants make and when these mistakes occur over the experimental rounds.

4.2.1 Mistake rates by participant role

In a four-player matching market, two players are fruits who know the state and hence the matching payoffs, while the other two are colors who face payoff uncertainty. A color is classified as *Hard* if she must draw an inference to identify the correct matching decision, and *Easy* otherwise. Figure 3 displays individual mistake rates by market type and participant role, for all 20 rounds in panel (a) and the last 10 rounds in panel (b). Note that in *Plain*, neither color needs to learn, so we pool the two colors and use their mistake rate as a benchmark for comparison with Hard colors in the other market types.

[9]The overall individual mistake rate, pooled across roles, closely approximates one minus the market-level correct rate. As a result, the corresponding analysis largely mirrors the market-level results reported in Section 4.1. We therefore relegate the pooled individual mistake results to Supplemental Appendix C.



(a) All 20 rounds

(b) Last 10 rounds

Figure 3: Individual mistake rates by market type and by participant role

We find that hard colors account for most individual mistakes, and their performance across market types largely mirrors the market-level results.[10] Across all 20 rounds, mistake rates for Easy colors are below 4% under every market type, whereas those for Hard colors range from 9% under *Own* to 36% under *No*. Compared with the 7% mistake rate under *Plain*, Hard-color mistake rates are significantly higher under *Other*, *No*, and *Compound*, while *Own* does not differ significantly from *Plain*. In the last 10 rounds, Easy-color mistake rates remain below 3% under every market type, and those for Hard colors fall to 4% under *Own* and to 3% under *Other*. In contrast, Hard-color mistake rates remain high at 39% under *No* and 22% under *Compound*, with *No* significantly higher than *Compound*. These results suggest that individuals' learning from the absence of another pair's match, as well as learning through two channels jointly, remains challenging even with experience. We report the corresponding statistical tests in Table 8.

[10]We do not analyze fruits' mistake behavior separately, since it is mechanically reflected in colors' mistake rates. Specifically, in each market type, the fruit mistake rate equals the average of the two colors' rates.

Table 8: Mistake rates for Hard colors: pairwise tests across market types

	Own	Other	No	Compound
Plain	0.852	0.009	< 0.001	< 0.001
Own		0.020	< 0.001	0.002
Other			0.002	0.010
No				0.069

(a) All 20 rounds

	Own	Other	No	Compound
Plain	0.046	0.271	< 0.001	0.001
Own		0.655	< 0.001	0.005
Other			< 0.001	0.001
No				0.023

(b) Last 10 rounds

Notes: Each cell reports the p-value from a two-sided Wilcoxon signed-rank test at the matching-group level ($n = 16$).

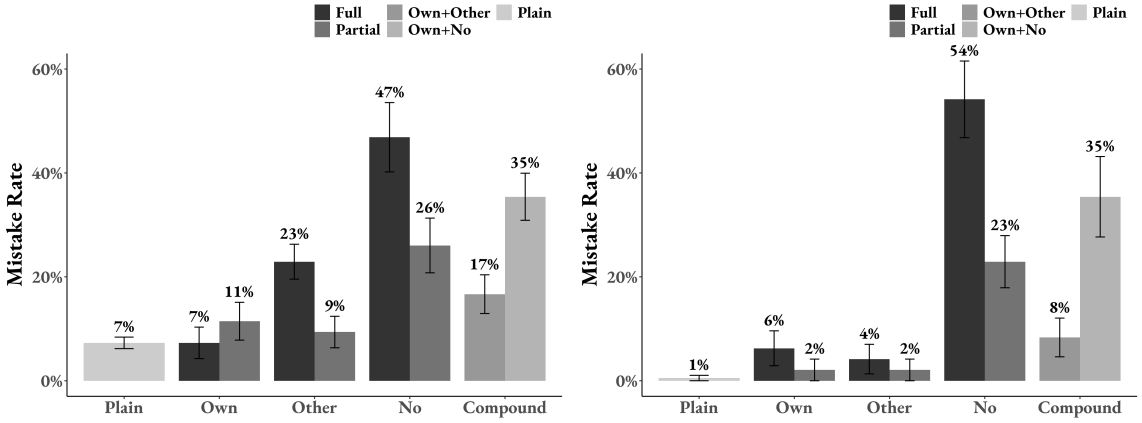
Result 2 (Mistake rates by participant role and market types)**.** *For all market types, Hard colors account for most mistakes. The ranking of mistake rates for Hard colors is Plain = Own < Other < Compound = No across all 20 rounds, and Plain = Own = Other < Compound < No in the last 10 rounds.*

4.2.2 Mistake rates for Hard colors by market subtype

The previous section shows that the mistake-rate patterns for Hard colors closely mirror the market-level results, which is expected given that Hard colors account for most mistakes. We now investigate whether mistake rates for Hard colors also differ by market subtype, which allows us to assess how the extent of information revealed and the complexity of the learning environment affect learning performance. Figure 4 plots mistake rates by subtype, with panel (a) showing all 20 rounds and panel (b) showing the last 10 rounds.

Subtype differences for Hard colors again closely parallel the market-level subtype differences discussed in Section 4.1. Under *Own*, Hard-color mistake rates do not differ significantly between *Full* and *Partial* in either sample. Under *Other*, *Partial* has a significantly lower mistake rate than *Full* by 14 percentage points (pp) across all 20 rounds, but the gap falls and is no longer significant in the last 10 rounds as participants gain experience. Under *No*, *Partial* has a significantly lower mistake rate than *Full* in both samples, by 21 pp across all 20 rounds and 31 pp in the last 10 rounds.

For the two *Compound* subtypes, Hard colors make significantly more mistakes under *Own+No* than under *Own+Other* in both samples, by 18 pp across all 20 rounds and 27 pp in the last 10 rounds. Relative to their non-Compound counterparts, *Own+Other* does not differ significantly from either *Own* or *Other*, whereas *Own+No* has significantly higher mistake rates than *Own* but does not differ significantly from *No-Partial* and is only marginally lower than *No-Full*. Corresponding statistical tests are reported in Table 9.



(a) All rounds (b) Last 10 rounds

Figure 4: Mistake rates for Hard colors by subtype

Table 9: Mistake rates for Hard colors: comparisons across market subtypes

		Within Type	Own+Other	Own+No
Own	Full	0.401	0.081	0.001
	Partial		0.110	< 0.001
Other	Full	0.003	0.220	
	Partial		0.107	
No	Full	0.024		0.054
	Partial			0.150
Compound		0.001		

(a) All 20 rounds

		Within Type	Own+Other	Own+No
Own	Full	0.317	0.706	0.006
	Partial		0.180	0.002
Other	Full	0.564	0.317	
	Partial		0.180	
No	Full	0.005		0.061
	Partial			0.342
Compound		0.009		

(b) Last 10 rounds

Notes: Each cell reports the p-value from a two-sided Wilcoxon signed-rank test at the matching-group level ($n = 16$).

Result 3 (Hard-color mistake rates by market subtypes)**.** *For market subtypes, Hard-color mistake rates do not differ between Partial and Full under Own. Under Other,*

the mistake rate is lower in Partial than in Full across all 20 rounds, but the difference disappears in the last 10 rounds. Under No, the mistake rate is lower in Partial than in Full in both samples. Within Compound, the mistake rate is lower under Own+Other than under Own+No in both samples.

4.2.3 Mistake types and evolution across rounds

To further understand matching failures and the underlying learning failures they reveal, we examine *what* mistakes individuals make and *when* they make them across rounds. Whenever a participant (all roles included) fails to reach a correct matching outcome in a round, the mistake falls into one of three mutually exclusive categories: (i) *Overmatch*, where the player should remain unmatched but chooses to match; (ii) *Undermatch*, where the player should match but remains unmatched; and (iii) *Wrong Partner*, where the player matches with an incorrect partner. Together, these categories exhaust all deviations from the correct matching outcome. Figure 5 displays the total number of individual mistakes and their composition by market type.

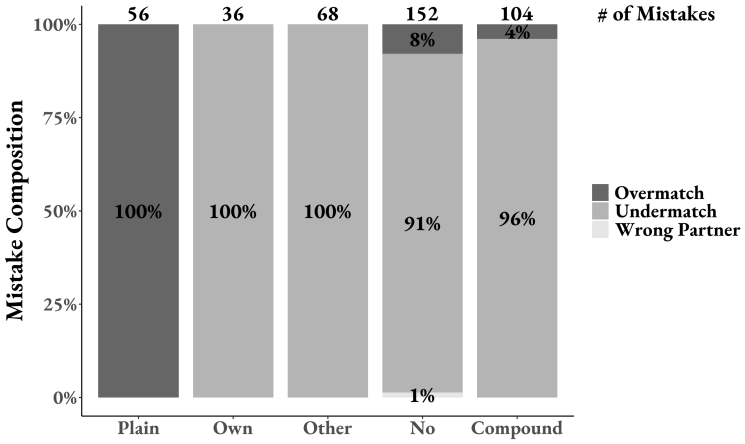


Figure 5: Composition of individual mistake types by market type

We find that participants tend to remain unmatched, the initial status in our markets, although this tendency varies across market types. In markets that require learning, more than 90% of mistakes are Undermatches, a distribution that differs

significantly from *Plain*, where all mistakes are Overmatches. The Undermatch rate is 100% under *Own* and *Other*, 91% under *No*, and 96% under *Compound*. The remaining mistakes are mostly Overmatches, accounting for 8% under *No* and 4% under *Compound*. Wrong-Partner mistakes appear only under *No* and account for about 1% of mistakes. Among the market types that require learning, *Own* and *Other* do not differ significantly from each other, while *No* differs significantly from *Other* but not from *Compound*. We report the corresponding chi-square tests in Table 10.

Table 10: Mistake type distribution: Chi-square tests

	Own	Other	No	Compound
Plain	< 0.001	< 0.001	< 0.001	< 0.001
Own		1.000	0.167	0.233
Other			0.035	0.102
No				0.204

Notes: Each cell reports the p-value from a two-sided chi-squared test at the matching-group level ($n = 16$).

We also find that the distribution of mistakes over the 20 rounds varies strikingly across market types. Figure 6 plots cumulative mistake rates by round. Under *Plain*, 96% of mistakes occur in Round 1, with very few thereafter. Under *Own* and *Other*, 78% and 91% of mistakes occur within the first 10 rounds, respectively. In contrast, under *No* and *Compound*, mistakes persist throughout the experiment: only 50% and 58% of mistakes occur within the first 10 rounds, respectively. These results suggest rapid learning over time under *Plain*, *Own*, and *Other*, but persistent difficulties under *No* and *Compound*.

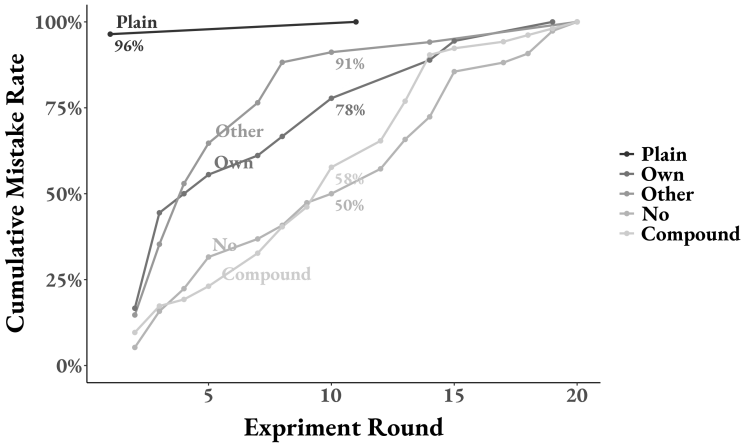


Figure 6: Cumulative mistake rates by round

Result 4 (individual mistake types and evolution across rounds)**.** *Across treatments, both the types of mistake and their evolution over rounds vary strikingly. In Plain, all mistakes are Overmatch, occurring mostly in Round 1. In Own and Other, all mistakes are Undermatch and are concentrated in the first 10 rounds. In No and Compound, over 90% of mistakes are Undermatch, and mistakes are spread more evenly across rounds.*

4.3 Matching dynamics

In this section, we examine the dynamics of correct match formation to better understand how learning requirements affect the decentralized matching process. A correct match is classified as *easy* if the color can match correctly without learning, and *hard* if learning is required. We examine two aspects of matching dynamics: the time required to finalize a correct match and, when both easy and hard matches occur, the order in which they form.

4.3.1 Time required to finalize correct matches

We begin by examining the time required to form correct matches. Figure 7 reports the average matching time by market type and match difficulty, conditional on reaching a

correct match.[11] In *Plain*, all correct matches are easy, while under *No* and *Own+No*, each market contains only one hard match. Under *Own*, *Other*, and *Own+Other*, both easy and hard matches may occur.

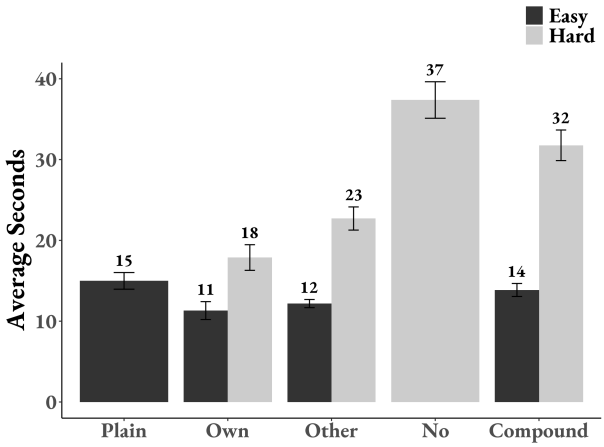


Figure 7: Time required to finalize correct matches by market type and by difficulty

Hard matches take substantially longer to finalize than easy matches, and the delay is greatest under *No* and *Compound*, while easy matches form quickly and at similar speeds across market types. The average time required to finalize a hard match ranges from 18 seconds under *Own* to 23 seconds under *Other*, 32 seconds under *Compound*, and 37 seconds under *No*. In market types where both easy and hard matches occur, hard matches also take significantly longer than easy matches.[12] These results suggest that learning requirements slow correct match formation, particularly when learning involves the absence of another pair's match or multiple learning channels.

Result 5 (time required)**.** *Hard matches take substantially longer than easy matches. The time required to finalize hard matches is ranked as No > Compound > Other > Own.*

[11] Densities of correct match time by match difficulty are reported in Supplemental Appendix C.

[12] All differences are significant at $p < 0.001$ based on two-sided Wilcoxon signed-rank tests.

4.3.2 Order of finalized correct matches

Under *Other* and *Own+Other*, the easy match provides information needed for the hard match and should therefore form first. We examine whether this holds in the experiment, conditional on both correct matches being formed.[13] Figure 8 reports the average easy–hard rate (equal to 1 if the easy match forms before the hard match) by market type and subtype where both easy and hard matches can occur.

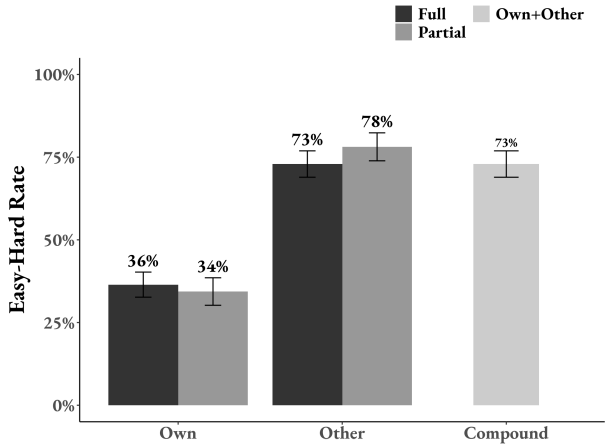


Figure 8: Order of correct matches by market type (easy-hard rate)

We find that easy matches are substantially more likely to precede hard matches under *Other* and *Own+Other* than under *Own*. Under *Own*, the easy–hard rate is 36% in Full and 34% in Partial. It rises to 73% and 78% under *Other*, and to 73% under *Own+Other*. The differences between *Own* and *Other*, and between *Own* and *Own+Other*, are statistically significant.[14] These results suggest that when formation of the hard match depends on information generated by the easy match, the easy match is substantially more likely to form first. This ordering is consistent with participants learning from observed match formation rather than guessing initially and revising afterward.

[13] Nevertheless, we do not require the easy match to precede the hard match for identification, as discussed in Section 2.4.

[14] Both differences have $p < 0.001$ based on two-sided Wilcoxon signed-rank tests.

Result 6 (order of correct matches)**.** *Under Other and Own+Other, easy correct matches precede hard correct matches significantly more often than under Own.*

4.4 Explanations of key findings

The previous sections report a number of findings, several of which call for further interpretation. We discuss possible explanations for our key findings with respect to the two treatment variables: learning channels and the extent of information revelation.

First, Results 1 and 2 show that matching under *No* is more difficult than under the other learning channels and that this difficulty does not diminish with experience. One possible explanation is that making the correct matching decision under *No* requires two steps: (i) recognizing that a fruit–color pair would have formed a match in certain states and (ii) correctly updating beliefs about the realized state after observing that the match did not form. Failure at either step may lead to an incorrect match. Correct matching under *No* is therefore substantially more challenging than under *Own* and *Other*, where colors need only perform the state-inference step based on observed market activities (step (ii)). This explanation also accounts for the higher mistake rates under *Own+No* than under *Own+Other*.

Results 4 and 5 support this explanation in two complementary ways. Under *No*, failure at either step tends to generate Undermatch outcomes, because completing both steps correctly provides the information needed to support matching (Result 4).[15] In addition, identifying the potential match in step (i) increases the time required to match under *No* (Result 5). These findings echo Esponda and Vespa (2014), which shows that subjects often fail to extract information from hypothetical or counterfactual events.[16]

[15]Although the share of Undermatch among all mistakes is slightly lower under *No*, the overall frequency of Undermatch is still the highest under *No*. Moreover, Overmatch mistakes under *No* are made by easy pairs and are therefore unrelated to the relevant inference.

[16]Also see Gui, He and Hu (2026), which investigates agent learning in settings where inaction

Second, Results 1 and 3 indicate that participants make better matching decisions in the *Partial* subtype than in the *Full* subtype under both *Other* and *No*. This suggests that the favorable-state mechanism outweighs the certainty-and-confidence mechanism. Under *Full*, the optimal matching decision requires participants to correctly pin down the realized state. Under *Partial*, by contrast, two of the three states support the same optimal matching decision. As a result, participants may still choose optimally despite imperfect belief updating or reliance on their prior, allowing greater scope for errors or naivety in inference without compromising the final matching decision.

5 Concluding Remarks

Our experimental design of the decentralized markets abstracts from the distinction between strategic and myopic play: because each agent has only one potential partner, the only choice is whether to match, not whom to match with or when. Allowing several acceptable partners would restore that choice, together with the incentives to delay and to reject that arise in structured decentralized games (Pais, 2008; Suh and Wen, 2008; Haeringer and Wooders, 2011; Ferdowsian, Niederle and Yariv, 2025). Whether subjects play the predicted strategies in such environments, and whether outcomes converge to equilibrium, remain open.

A second extension is to two-sided incomplete information (Chen and Hu, 2023; Chen and Ho, forthcoming; Park, 2025), or to the more general cooperative frameworks of Liu (2025). This extension is conceptually subtle because both sides of the market may learn and update simultaneously, so the information structure is no longer determined solely by firms' inferences. Studying how learning and matching evolve in such environments, and how they interact with incentives on both sides of the market, would be an important direction.

arises naturally and introduces a treatment that makes inaction explicitly observable to mitigate the difficulty.

References

Agranov, Marina, Ahrash Dianat, Larry Samuelson, and Leeat Yariv. forthcoming. "Paying to Match: Decentralized Markets with Information Frictions." *The Review of Economic Studies*.

Alston, Max. 2020. "On the non-existence of stable matches with incomplete information." *Games and Economic Behavior*, 120: 336–344.

Bikhchandani, Sushil. 2017. "Stability with One-Sided Incomplete Information." *Journal of Economic Theory*, 168: 372–399.

Chen, Daniel, Martin Schonger, and Chris Wickens. 2016. "oTree—An open-source platform for laboratory, online, and field experiments." *Journal of Behavioral and Experimental Finance*, 9(C): 88–97.

Chen, Yi-Chun, and Gaoji Hu. 2020. "Learning by matching." *Theoretical Economics*, 15(1): 29–56.

Chen, Yi-Chun, and Gaoji Hu. 2023. "A theory of stability in matching with incomplete information." *American Economic Journal: Microeconomics*, 15(1): 288–322.

Chen, Yi-Chun, and Gaoji Hu. 2024. "Bayesian stable states." *Games and Economic Behavior*, 145: 102–116.

Chen, Yi-Chun, and Samuel Cher Sien Ho. forthcoming. "Efficiency of Stable Matchings with Two-Sided Incomplete Information." *Theoretical Economics*.

Dolgopolov, Arthur, Daniel Houser, Cesar Martinelli, and Thomas Stratmann. 2024. "Assignment markets: Theory and experiments." *European Economic Review*, 165: 104738.

Echenique, Federico, Alejandro Robinson-Cortés, and Leeat Yariv. 2025. "An experimental study of decentralized matching." *Quantitative Economics*, 16(2): 497–533.

Esponda, Ignacio, and Emanuel Vespa. 2014. "Hypothetical thinking and information extraction in the laboratory." *American Economic Journal: Microeconomics*, 6(4): 180–202.

Ferdowsian, Andrew, Muriel Niederle, and Leeat Yariv. 2025. "Strategic decentralized matching: The effects of information frictions." *Princeton and Stanford Universities*.

Gale, David, and Lloyd S Shapley. 1962. "College admissions and the stability of marriage." *The American Mathematical Monthly*, 69(1): 9–15.

Gui, Shan, Simin He, and Gaoji Hu. 2026. "Learning from inaction." *SSRN Working Paper 6882699*.

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

He, Simin, Jiabin Wu, Hanzhe Zhang, and Xun Zhu. 2024. "Decentralized matching with transfers: experimental and noncooperative analyses." *American Economic Journal: Microeconomics*, 16(4): 406–439.

Hu, Gaoji. 2026. "Structured Stable Matchings." *SSRN Working Paper 4910284*.

Knuth, Donald Ervin. 1976. *Mariages stables et leurs relations avec d'autres problèmes combinatoires: introduction à l'analyse mathématique des algorithmes.* Presses de l'Université de Montréal.

Lauermann, Stephan, and Georg Nöldeke. 2014. "Stable marriages and search frictions." *Journal of Economic Theory*, 151: 163–195.

Liu, Qingmin. 2020. "Stability and Bayesian Consistency in Two-Sided Markets." *American Economic Review*, 110(8): 2625–2666.

Liu, Qingmin. 2024. "Matching with Incomplete Information." In *Handbook of the Economics of Matching*.

Liu, Qingmin. 2025. "A Cooperative Analysis of Incomplete Information." *Working Paper*.

Liu, Qingmin, George J Mailath, Andrew Postlewaite, and Larry Samuelson. 2014. "Stable Matching with Incomplete Information." *Econometrica*, 82(2): 541–587.

Nalbantian, Haig R., and Andrew Schotter. 1995. "Matching and Efficiency in the Baseball Free-Agent System: An Experimental Examination." *Journal of Labor Economics*, 13(1): 1–31.

Otto, Philipp E., and Friedel Bolle. 2011. "Matching Markets with Price Bargaining." *Experimental Economics*, 14: 322–348.

Pais, Joana. 2008. "Incentives in decentralized random matching markets." *Games and Economic Behavior*, 64(2): 632–649.

Pais, Joana, Ágnes Pintér, and Róbert F Veszteg. 2020. "Decentralized matching markets with (out) frictions: a laboratory experiment." *Experimental Economics*, 23(1): 212–239.

Park, Changwoo. 2025. "Stable Matchings under Two-sided Asymmetric Information." *SSRN Working Paper 5549758*.

Peralta, Esteban. 2025. "Lone wolves just got lonelier." *Games and Economic Behavior*, 152: 55–61.

Pomatto, Luciano. 2022. "Stable matching under forward-induction reasoning." *Theoretical Economics*, 17(4): 1619–1649.

Roth, Alvin E, and John H Vande Vate. 1990. "Random paths to stability in two-sided matching." *Econometrica*, 1475–1480.

Shapley, Lloyd S, and Martin Shubik. 1971. "The assignment game I: The core." *International Journal of Game Theory*, 1(1): 111–130.

Suh, Sang-Chul, and Quan Wen. 2008. "Subgame perfect implementation of stable matchings in marriage problems." *Social Choice and Welfare*, 31(1): 163–174.

Wang, Ziwei. 2023. "Rationalizable Stability in Matching with Incomplete Information." *Working Paper*.

Appendix A Design Details

A.1 Parameters

Table A.1: Payoff tables: parameter set 1 (realized weather: Sunny in all)

	Yellow		Green		
Lemon	7	5	6	−3	Sunny
	4	3	−9	1	Cloudy
	2	8	−2	−8	Rainy
Mango	−1	−5	−8	1	Sunny
	6	−9	5	−4	Cloudy
	−3	4	2	−7	Rainy

(a) Plain (one matched pair)

	Yellow		Green		
Lemon	3	6	5	−4	Sunny
	5	4	−1	3	Cloudy
	9	8	−3	−9	Rainy
Mango	−7	1	3	9	Sunny
	2	−6	6	5	Cloudy
	−4	−5	8	2	Rainy

(b) Plain (two matched pairs)

	Yellow		Green		
Lemon	−6	−3	7	5	Sunny
	−9	1	−4	−3	Cloudy
	2	−8	−2	−8	Rainy
Mango	5	7	−1	5	Sunny
	8	4	2	−9	Cloudy
	6	1	−3	−4	Rainy

(c) Own-Full

	Yellow		Green		
Lemon	−7	−4	8	6	Sunny
	−1	2	5	4	Cloudy
	−3	−9	−3	−9	Rainy
Mango	3	2	−7	1	Sunny
	6	5	2	−6	Cloudy
	9	8	−4	5	Rainy

(d) Own-Partial

	Yellow		Green		
Lemon	2	5	−6	−3	Sunny
	−4	3	−9	1	Cloudy
	−7	8	2	−8	Rainy
Mango	−1	5	5	1	Sunny
	2	−9	8	−4	Cloudy
	−3	−4	6	−7	Rainy

(e) Other-Full

	Yellow		Green		
Lemon	8	6	−7	−4	Sunny
	5	4	−1	2	Cloudy
	−3	9	3	−9	Rainy
Mango	−7	1	3	5	Sunny
	2	−6	6	2	Cloudy
	−4	5	9	−8	Rainy

(f) Other-Partial

	Yellow		Green		
Lemon	−2	5	−6	−3	Sunny
	4	3	−9	1	Cloudy
	7	8	2	−8	Rainy
Mango	−1	5	5	1	Sunny
	2	−9	8	−4	Cloudy
	−3	−4	6	−7	Rainy

(g) No-Full

	Yellow		Green		
Lemon	−3	6	−7	−4	Sunny
	−5	4	−1	2	Cloudy
	8	9	3	−9	Rainy
Mango	−7	1	3	5	Sunny
	2	−6	6	2	Cloudy
	−4	5	9	−8	Rainy

(h) No-Partial

	Yellow		Green		
Lemon	2	5	−6	−3	Sunny
	4	3	−9	1	Cloudy
	−7	8	2	−8	Rainy
Mango	−1	5	6	1	Sunny
	2	−9	−5	−4	Cloudy
	−3	−4	8	−7	Rainy

(i) Own+Other

	Yellow		Green		
Lemon	−3	6	−7	−4	Sunny
	−5	4	−1	2	Cloudy
	8	9	3	−9	Rainy
Mango	−7	1	3	2	Sunny
	2	−6	−6	−5	Cloudy
	−4	5	9	−8	Rainy

(j) Own+No

Table A.2: Payoff tables: parameter set 2 (realized weather: Sunny in all)

	Yellow		Green		
Lemon	3	8	4	−3	Sunny
	7	1	−7	5	Cloudy
	4	2	−6	−9	Rainy
Mango	−2	1	5	−2	Sunny
	6	−5	9	−1	Cloudy
	−8	4	7	−6	Rainy

(a) Plain (one matched pair)

	Yellow		Green		
Lemon	5	−4	6	−5	Sunny
	−9	3	9	−7	Cloudy
	−6	1	−8	2	Rainy
Mango	−8	3	7	−4	Sunny
	−4	−7	2	−3	Cloudy
	−1	6	9	−8	Rainy

(b) Plain (no matched pair)

	Yellow		Green		
Lemon	−4	−3	3	4	Sunny
	−7	−5	−7	−1	Cloudy
	−6	9	−2	−8	Rainy
Mango	−5	−2	−2	1	Sunny
	9	−1	−6	−5	Cloudy
	7	−6	−8	4	Rainy

(c) Own-Full

	Yellow		Green		
Lemon	−6	−5	5	3	Sunny
	−9	7	9	1	Cloudy
	−8	2	−6	−4	Rainy
Mango	−7	4	−8	−3	Sunny
	−2	3	−4	7	Cloudy
	9	−8	−1	6	Rainy

(d) Own-Partial

	Yellow		Green		
Lemon	3	8	−4	3	Sunny
	−7	1	−7	5	Cloudy
	−4	2	−6	−9	Rainy
Mango	−2	1	5	2	Sunny
	6	−5	9	−1	Cloudy
	−8	−4	7	−6	Rainy

(e) Other-Full

	Yellow		Green		
Lemon	5	1	−6	5	Sunny
	9	3	−9	−7	Cloudy
	−6	4	−8	2	Rainy
Mango	−8	−3	2	4	Sunny
	4	−7	7	3	Cloudy
	−6	1	9	−8	Rainy

(f) Other-Partial

	Yellow		Green		
Lemon	−3	2	−4	−3	Sunny
	7	1	−7	5	Cloudy
	4	8	−6	−9	Rainy
Mango	−2	1	5	6	Sunny
	6	−5	2	−1	Cloudy
	−8	−4	7	−9	Rainy

(g) No-Full

	Yellow		Green		
Lemon	−5	4	−6	5	Sunny
	−6	1	−9	−7	Cloudy
	9	3	−8	−2	Rainy
Mango	−8	−3	2	4	Sunny
	1	−7	7	3	Cloudy
	4	−6	9	−8	Rainy

(h) No-Partial

	Yellow		Green		
Lemon	3	8	−4	3	Sunny
	4	1	−7	5	Cloudy
	−7	2	−6	−9	Rainy
Mango	2	−1	7	2	Sunny
	6	−5	−9	−1	Cloudy
	−8	−4	5	−6	Rainy

(i) Own+Other

	Yellow		Green		
Lemon	−5	4	−6	5	Sunny
	−6	3	−9	−7	Cloudy
	9	1	−8	2	Rainy
Mango	−8	−3	9	4	Sunny
	4	−7	−2	−3	Cloudy
	−1	6	7	−8	Rainy

(j) Own+No

A.2 Treatment orders

Each subject experiences all 20 markets (Tables A.1 and A.2), one per round. Market order varies across matching groups according to four pre-randomized treatment orders.

First, to control for order effects across market types, we divide the 20 markets into four blocks of five markets (Blocks A–D). Each block contains one market from each of the five market types. The four blocks are rotated across treatment orders using a Latin-square design, yielding four treatment orders (Table A.3). Within each block, *Plain* always appears first, while the remaining four market types are rotated according to a separate Latin-square design to balance their positions (Table A.4).

Table A.3: Latin-square rotation of blocks

	1	2	3	4
Treatment Order 1	Block A	Block B	Block C	Block D
Treatment Order 2	Block B	Block C	Block D	Block A
Treatment Order 3	Block C	Block D	Block A	Block B
Treatment Order 4	Block D	Block A	Block B	Block C

Table A.4: Within-block market-type sequences

	1	2	3	4	5
Block A	Plain	Own	Other	No	Compound
Block B	Plain	Other	No	Compound	Own
Block C	Plain	No	Compound	Own	Other
Block D	Plain	Compound	Own	Other	No

Second, to control for subtype order effects, we ensure that each subtype appears once in rounds 1–10 and once in rounds 11–20.[17]

[17]The sequence of the four *Plain* markets is held fixed across all treatment orders rather than randomized. Their results may therefore be affected by order effects and should be interpreted with caution. This limitation is unlikely to affect our main analysis, which focuses on comparisons across learning market types and subtypes.

Finally, we pre-randomize payoff-table displays and market assignments to balance visual presentation and role exposure. We randomize the realized weather and vary table layouts by permuting weather rows, fruit row blocks, and color column blocks. We also ensure that each color participant serves equally often in the Hard and Easy roles across learning markets.

Not for Publication:

Supplemental Appendix for "**Learning in Matching**"

Shan Gui Simin He Gaoji Hu

Appendix B Experimental Instructions and Screenshots

In this appendix, we provide the experimental instructions and the experimental screenshots that are translated from the original Chinese version.

B.1 Instructions (All treatments)

[Welcome] Welcome to this experiment! Please read the following introduction carefully.

This experiment will last approximately 70 minutes. During the experiment, please remain silent and do not communicate with others in any way. If you have any questions, please raise your hand, and an experimenter will assist you individually.

Before the experiment begins, you will be randomly assigned to a large group consisting of 12 participants. This group will remain fixed throughout the entire experiment. Each participant will sit alone at a computer terminal, and all decisions will be made on the computer screen. The experiment is anonymous: neither the experimenters nor the other participants will know which participant is seated at which station, nor will they be able to link any decision to you or to anyone else in the room.

Because you arrived on time, you have received a participation fee of 20 yuan. You will earn additional income during the experiment, depending on your own decisions and the decisions of other participants. At the end of the experiment, your total earnings will be paid to you privately.

Throughout the experiment, all earnings will be denominated in "yuan" (RMB).

[Roles and Grouping] Within each large group of 12 participants, you will be randomly assigned to one of two roles: 6 participants will be "Fruits" and 6 participants will be "Colors." These roles remain fixed throughout the entire experiment—once assigned, you will always be either a Fruit or a Color. The experiment consists of 20 rounds, each corresponding to a different "matching game." In each round, the 12 participants in the large group will be randomly divided into three independent subgroups of 4 participants each. Each subgroup of 4 will complete one matching game together. The composition of the 4-person subgroups is fully random and independent across rounds; that is, the participants in your subgroup may vary from round to round. At the beginning of every round, you receive an initial payoff of 10 yuan.

[Matching Game and Matching Payoffs] Each matching game involves four participants: Lemon, Mango, Yellow, and Green. If your role is "Fruit," you will be randomly assigned to be either Lemon or Mango each round. If your role is "Color," you will be randomly assigned to be either Yellow or Green each round. In each round, Fruits and Colors may form "matches," and each Fruit can match with at most one Color, while each Color can match with at most one Fruit. For example, Lemon may match with Yellow, or with Green, or Lemon may remain unmatched. Two Fruits cannot match with each other, and two Colors cannot match with each other. "Weather" affects the payoffs from matching. Each round's weather may be Sunny, Cloudy, or Rainy, with equal probability (33.33% each). The weather is determined independently each round by the computer. At the beginning of each round, the two Fruits will learn the weather, while the two Colors will not. The payoffs for each possible Fruit–Color match are shown in a "Matching Payoff Table." In the example below, you can find the payoff for any Fruit–Color pair at the corresponding cell of the table. Each cell contains three rows: the first row lists the payoffs under Sunny weather, the second row under Cloudy, and the third row under Rainy. Each row contains two numbers: the first number is the Fruit's payoff, and the second number is the Color's payoff. For example:

• If Lemon matches with Yellow and the weather is Sunny, Lemon earns 3 yuan and

Yellow earns 1 yuan.

• If Lemon matches with Green and the weather is Cloudy, Lemon earns –7 yuan and Green earns –6 yuan.

• If Mango matches with Yellow and the weather is Rainy, Mango earns –3 yuan and Yellow earns –4 yuan.

	Yellow		Green		
Lemon	3	1	-6	-7	Sunny
	1	3	-7	-6	Cloudy
	2	4	-8	-9	Rainy
Mango	-1	-5	6	7	Sunny
	-4	-1	9	6	Cloudy
	-3	-4	8	9	Rainy

Figure B.1: Matching Payoff Table (Example)

If a participant does not match, the matching payoff is 0 yuan. For instance, in a 4-person subgroup, if Lemon matches with Yellow, while Mango and Green remain unmatched, and the weather is Sunny, then Lemon earns 3 yuan, Yellow earns 1 yuan, and Mango and Green each earn 0 yuan.

[Matching Process and Final Payoffs] In the matching game, the four participants may freely engage in three types of actions: sending invitations, accepting/rejecting invitations, and breaking a temporary match. Details are as follows.

[Sending Invitations] Each participant may send a matching invitation to any participant of the opposite role (i.e., Fruits may invite any Color, and Colors may invite any Fruit). The invited participant immediately receives a notification showing the inviter's identity; the other participants do not observe this information. Before receiving a response (acceptance or rejection), the inviter cannot send another invitation, but may receive invitations from others and choose whether to accept or reject them.

[Accepting or Rejecting Invitations] Upon receiving an invitation, the invited participant has 15 seconds to respond:

• If the invited participant rejects the invitation, or does not respond within 15 seconds (which counts as a rejection), the invitation expires.

• If the invited participant accepts within 15 seconds, the two participants form a temporary match.

• During these 15 seconds, the invited participant may still send invitations to others.

[Temporary Matches, Final Matches, and Final Payoffs] If an invitation is sent and accepted, a temporary match is formed. If the round has not ended, all matches are temporary. This means:

• Participants in a temporary match may still send new invitations and may accept or reject incoming invitations.

If a participant forms a new temporary match, any previous temporary match is automatically dissolved.

• Either participant in a temporary match may unilaterally break the match at any time; the other party simply receives a notification.

The "Matching Payoff Table" is displayed to everyone, and the currently active temporary matches within the 4-person subgroup are shown by a highlighted indicator at the corresponding cell—this information is public. At the end of the round, all temporary matches become final matches:

• Participants in a final match receive: Final payoff = initial 10 yuan + matching payoff

• Participants without a match receive: Final payoff = initial 10 yuan

[When Does a Round End?] At the beginning of each round, a 60-second public countdown starts (visible in the upper-left corner of the screen):

• If the public matching indicators remain unchanged for 60 consecutive seconds—meaning no temporary match is formed and no existing temporary match is broken—then the round ends.

• Whenever any change occurs (a match is formed or broken), the 60-second countdown restarts. Additionally, once a round has lasted at least 60 seconds, a button

labeled "Agree to Proceed to Next Round" becomes available. If all four participants in the subgroup press this button, the round ends immediately without waiting for the countdown to expire. Note: Even after pressing the "Agree" button, you may continue all matching actions until the round ends, and you may also cancel your agreement.

[Number of Rounds and Total Earnings] You will play 20 matching games in total. All procedures are identical across rounds; the only difference is the Matching Payoff Table, which updates each round. Before the 20 official rounds, there will be one practice round. The practice round is designed to help you become familiar with the procedure, and the earnings from this practice round do not count toward your final payment. After each round ends, the system automatically proceeds to the next round. In the new round, you will not be able to view any payoff tables or outcomes from previous rounds. At the end of the experiment, the computer will randomly select 4 out of the 20 rounds, and your earnings from those 4 rounds will count toward your experiment payment. Your total payment = participation fee (20 yuan) + the sum of the final payoffs from the 4 randomly selected rounds.

B.2 Screenshots

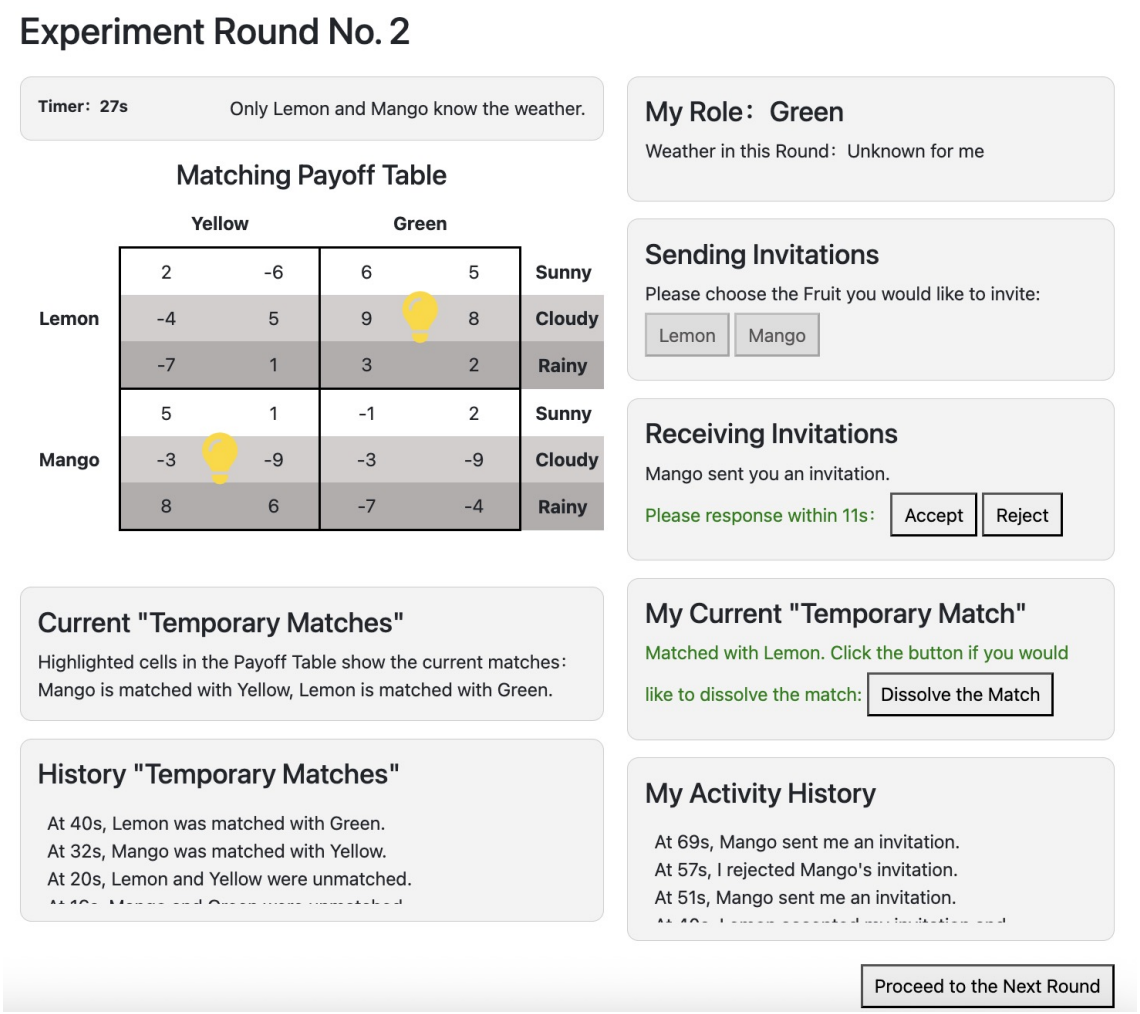
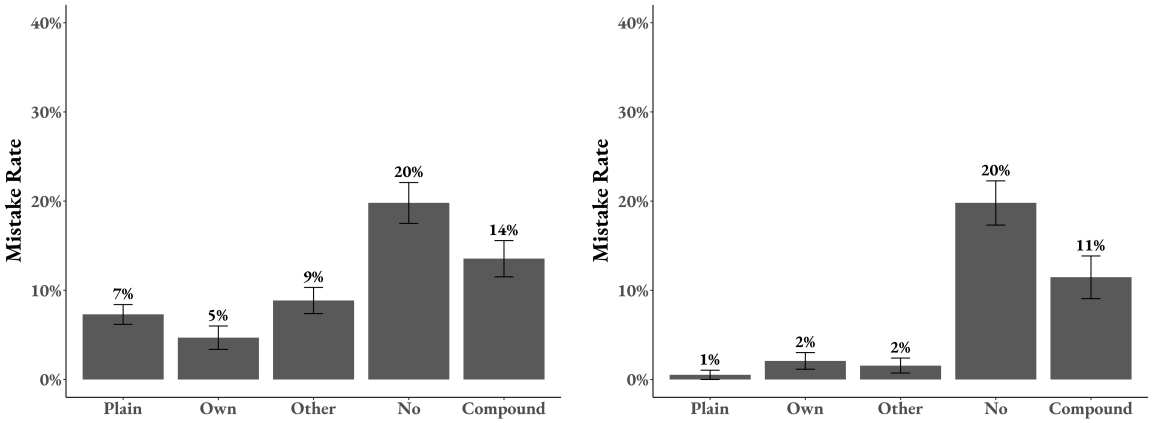


Figure B.2: A translated screenshot of the experiment

Appendix C Supplemental Figures and Tables

This appendix collects figures and tables that complement our main results. Figure C.1 and Table C.1 report individual mistake rates by market type, mirroring the analysis in Section 4.1. Figure C.2 and Table C.2 report them by market subtype, mirroring Section 4.2. Unlike Sections 4.2.1 and 4.2.2 which distinguish participant roles and focus on *Hard* colors, the rates reported here pool all individuals. Figure C.3 reports the time required to reach correct matches by market type and match difficulty, mirroring the analysis in Section 4.3.1.



(a) All 20 rounds (b) Last 10 rounds

Figure C.1: Individual mistake rate by market types

Table C.1: Individual mistake rates: pairwise tests across market types

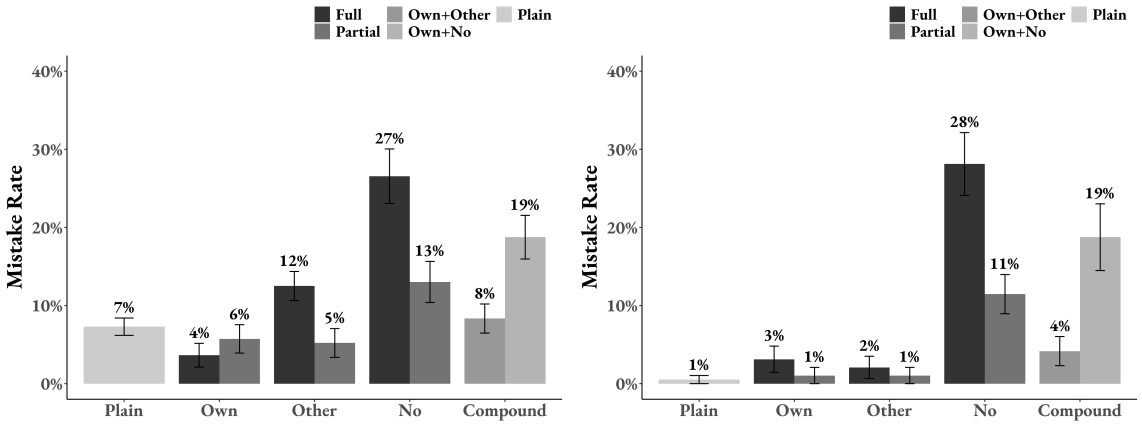
	Own	Other	No	Compound
Plain	0.044	0.657	< 0.001	0.001
Own		0.010	< 0.001	0.001
Other			0.002	0.021
No				0.017

(a) All 20 rounds

	Own	Other	No	Compound
Plain	0.083	0.317	< 0.001	0.002
Own		0.655	< 0.001	0.005
Other			< 0.001	0.001
No				0.022

(b) Last 10 rounds

Notes: Each cell reports the p-value from a two-sided Wilcoxon signed-rank test at the matching-group level ($n = 16$).



(a) All rounds (b) Last 10 rounds

Figure C.2: Individual mistake rate by subtypes

Table C.2: Individual mistake rates: pairwise tests across subtypes

		Within Type	Own+Other	Own+No
Own	Full	0.401	0.081	< 0.001
	Partial		0.110	< 0.001
Other	Full	0.006	0.095	
	Partial		0.127	
No	Full	0.012		0.019
	Partial			0.090
Compound		0.001		

(a) All 20 rounds

		Within Type	Own+Other	Own+No
Own	Full	0.317	0.706	0.006
	Partial		0.180	0.002
Other	Full	0.564	0.317	
	Partial		0.180	
No	Full	0.006		0.045
	Partial			0.302
Compound		0.009		

(b) Last 10 rounds

Notes: Each cell reports the p-value from a two-sided Wilcoxon signed-rank test at the matching-group level ($n = 16$).

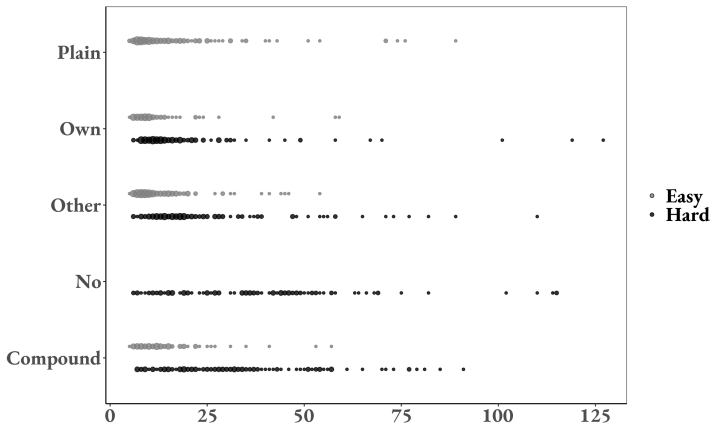


Figure C.3: Density of Final Correct Match Time by Match Difficulty

Appendix D A Stability Theory with Incomplete Information

In the main text, we have identified three canonical learning channels and examined how they shape matching outcomes. We now reflect these ideas in a stability theory, where stable matchings serve both as static solutions and as the steady states of an adaptive matching processes. The materials in Sections D.1 and D.2 largely follow Chen and Hu (2020, 2024), which generalize Liu et al. (2014); Liu (2020). Section D.3 relates the theory to the canonical learning channels in our main text.

D.1 The market environment

Consider a one-to-one matching market with one-sided incomplete information and no monetary transfers. Let $I = \{1, \ldots, n\}$ denote the set of workers (fruits in our experiments) and $J = \{n+1, \ldots, n+m\}$ the set of firms (colors in our experiments), corresponding to fruits and colors in the main text, respectively. A *matching* is a one-to-one function $\mu : I \cup J \to I \cup J$ that pairs up workers and firms such that, for each $i \in I$ and $j \in J$,

(1) $\mu(i) \in J \cup \{i\}$,

(2) $\mu(j) \in I \cup \{j\}$, and

(3) $\mu(i) = j$ if and only if $\mu(j) = i$.

If $\mu(i) = i$ or $\mu(j) = j$, the corresponding agent is said to be *unmatched*. We assume that the matching μ is publicly observable.

A payoff-relevant *state* ω is drawn from a finite set Ω according to a prior distribution $\beta \in \Delta(\Omega)$. Let $a_{ij}(\omega) \in \mathbb{R}$ and $b_{ij}(\omega) \in \mathbb{R}$ denote the ex post payoffs of worker i and firm j, respectively, when they are matched (i.e., $\mu(i) = j$) and the realized state is ω. Normalize the payoff from remaining unmatched to zero; that is, $a_{ii}(\omega) = b_{jj}(\omega) = 0$.

Each worker observes the realized state, whereas firms may not. For convenience, we represent the firms' *information structure* by a partition profile

$$\Pi = (\Pi_{n+1}, \ldots, \Pi_{n+m}),$$

where, for each firm $j \in J$, Π_j is a partition of Ω. If $\omega' \in \Pi_j(\omega)$, then when the true state is ω, firm j cannot distinguish ω from ω'. Agents may also have publicly observable payoff-relevant characteristics, summarized by their indices $i \in I$ and $j \in J$. Finally, the payoff functions

$$a : I \times J \times \Omega \to \mathbb{R} \quad \text{and} \quad b : I \times J \times \Omega \to \mathbb{R}$$

are common knowledge.

An *outcome*, denoted by (μ, ω, Π), consists of a matching μ, a realized state ω (which need not be publicly observable), and an information structure Π. As with μ, we assume Π is publicly observable. That is, everyone knows what each firm can distinguish, but not necessarily which cell the firm is in. Observability of μ follows Liu et al. (2014) and Liu (2020), and that of Π follows Chen and Hu (2020, 2024), who interpret both μ and Π as outcomes of a matching-and-learning process. Because agents' information is shaped by the history of publicly observable matchings, it is itself publicly observable.

D.2 Stability notion

Stability of a matching market consists of three requirements: individual rationality, no blocking pair, and information stability. The first two conditions are standard in classical matching theory, whereas the third is specific to the incomplete information setting.

Definition D.1. *An outcome* (μ, ω, Π) *is* ***individually rational (IR)*** *if agents' expected payoffs are nonnegative, i.e.,*

$$\begin{aligned} a_{i\mu(i)}(\omega) &\geq 0 \quad \textit{for all } i \in I \quad \textit{and} \\ \mathbb{E}\left[b_{\mu(j)j} \middle| \Pi_j(\omega)\right] &\geq 0 \quad \textit{for all } j \in J. \end{aligned}$$

The definition of a blocking pair requires specifying how a firm evaluates a potential deviation. Unlike instandard matching models, worker and firm may possess different information. Consequently, when firm j evaluates a potential match with worker i, only states in which worker i would also benefit from the deviation are relevant. Let $D_{\mu,ij}$ be the set of states under which worker i benefits from rematching with firm j, i.e.,

$$D_{\mu,ij} := \left\{\tilde{\omega} \in \Omega : a_{ij}(\tilde{\omega}) > a_{i\mu(i)}(\tilde{\omega})\right\}.$$

Since a deviation can occur only in those states in $D_{\mu,ij}$, firm j evaluates the deviation using the conditional belief $\beta(\cdot \mid \Pi_j(\omega) \cap D_{\mu,ij})$, which coincides with the off-path belief of Liu (2020).[18]

Definition D.2. *An outcome (μ, ω, Π) is **blocked** by (i, j) if worker i and firm j prefer each other to their assigned partners under μ, i.e.,*

$$\begin{aligned} a_{ij}(\omega) &> a_{i\mu(i)}(\omega) \quad \text{and} \\ \mathbb{E}\left[b_{ij}\middle|\Pi_j(\omega) \cap D_{\mu,ij}\right] &> \mathbb{E}\left[b_{\mu(j)j}\middle|\Pi_j(\omega) \cap D_{\mu,ij}\right]. \end{aligned}$$

The first two conditions are insufficient under incomplete information. Even if an outcome is individually rational and admits no blocking pair, the absence of deviations may itself reveal information to firms. Such information may subsequently create profitable deviations that were previously unavailable. To see this, we recast the example in Table 4 here.

In the market in Table D.1, consider the three outcomes obtained by combining the same matching and information structure with the three possible realized states:

(O) (μ, ω, Π), where $\mu(i_1) = i_1$, $\mu(i_2) = i_2$, $\Pi_{j_1} = \{\{\omega, \omega', \omega''\}\}$, and $\Pi_{j_2} = \{\{\omega, \omega', \omega''\}\}$.

(O') (μ, ω', Π), where $\mu(i_1) = i_1$, $\mu(i_2) = i_2$, $\Pi_{j_1} = \{\{\omega, \omega', \omega''\}\}$, and $\Pi_{j_2} = \{\{\omega, \omega', \omega''\}\}$.

(O") (μ, ω'', Π), where $\mu(i_1) = i_1$, $\mu(i_2) = i_2$, $\Pi_{j_1} = \{\{\omega, \omega', \omega''\}\}$, and $\Pi_{j_2} = \{\{\omega, \omega', \omega''\}\}$.

[18] See Liu (2020, Section IV-E) for alternative ideas of specifying off-path beliefs.

Table D.1: Matching market: learning from absence of deviation

	j_1		j_2		
i_1	-1	2	-10	-10	ω
	1	4	-10	-10	ω'
	1	8	-10	-10	ω''
i_2	-10	-10	1	2	ω
	-10	-10	1	-4	ω'
	-10	-10	1	-8	ω''

Note: The realized state is ω.

Outcome (O) is not blocked by (i_1, j_1) because $a_{i_1,j_1}(\omega) < 0$; it is not blocked by (i_2, j_2) because firm j_2 worries about the large negative payoffs -4 and -8 when $D_{\mu,i_2j_2} = \{\omega, \omega', \omega''\}$. Thus, outcome (O) admits no blocking pair. In contrast, it is straightforward to verify that both outcomes (O') and (O") are blocked by (i_1, j_1) (but not by (i_2, j_2)).

Now suppose the true state is ω. We have argued that there should not be any deviation from outcome (O). However, firm j_2 may make the following inference: If the true state *were* ω' or ω'', then a pairwise deviation by (i_1, j_1) from outcome (O) should have occurred. The absence of such a deviation, therefore, reveals information to firm j_2—that the true state *cannot* be ω' or ω''. Based on her updated belief that the true state must be ω, firm j_2 can then form a blocking pair with worker i_2.

The previous example suggests that firms should also reason from the absence of deviations. We therefore define the set of states under which the current outcome is individually rational and admits no blocking pair:

$$N_{\mu,\Pi} := \{\tilde{\omega} \in \Omega : (\mu, \tilde{\omega}, \Pi) \text{ is IR and not blocked}\}.$$

Suppose (μ, ω, Π) is individually rational and admits no blocking pair. Observing the absence of deviations reveals that the true state must belong to $N_{\mu,\Pi}$. Each firm

should thus refine its information by incorporating this additional knowledge. Let $\mathcal{N}_{\mu,\Pi}$ denote the binary partition generated by $N_{\mu,\Pi}$. The resulting information refinement (i.e., the *join* of Π and $\mathcal{N}_{\mu,\Pi}$) is represented by the operator $H_\mu(\cdot)$.

$$\begin{aligned} H_\mu(\Pi) &= (\Pi_{n+1} \vee \mathcal{N}_{\mu,\Pi}, \ldots, \Pi_{n+m} \vee \mathcal{N}_{\mu,\Pi}), \\ &\quad \text{where } [\Pi_j \vee \mathcal{N}_{\mu,\Pi}](\tilde{\omega}) = \Pi_j(\tilde{\omega}) \cap \mathcal{N}_{\mu,\Pi}(\tilde{\omega}) \text{ for all } \tilde{\omega} \in \Omega \text{ and all } j \in J. \end{aligned}$$

In the example above, $N_{\mu,\Pi} = \{\omega\}$, $\mathcal{N}_{\mu,\Pi} = \{\{\omega\}, \{\omega', \omega''\}\}$, and $[H_\mu(\Pi)]_{j_2} = \{\{\omega\}, \{\omega', \omega''\}\} \neq \Pi_{j_2}$. Therefore, $H_\mu(\Pi)$ represents the information structure after firms learn from the absence of deviations. Information stability requires that this refinement generate no additional information; that is, $H_\mu(\Pi) = \Pi$.[19]

Definition D.3. *An outcome (μ, ω, Π) is **stable** if it satisfies three requirements:*

(1) *(μ, ω, Π) is individually rational.*

(2) *(μ, ω, Π) is not blocked.*

(3) *firms cannot further refine their information from the absence of deviations, i.e.,*

$$H_\mu(\Pi) = \Pi.$$

*Say μ is a **stable matching** at ω if (μ, ω, Π) is a stable outcome for some Π.*

We could write $\omega = (t_1, \ldots, t_n)$ and interpret each t_i as worker i's private type, which contains information about his productivity and/or effort cost. Then there are two kinds of observability to be discussed: (1) whether a worker can observe other workers' types, and (2) within a matched pair, whether a firm can observe her own employee's type. First, in models with one-sided incomplete information, when there is no externality, i.e., $a_{ij}(\omega) \equiv a_{ij}(t_i)$ and $b_{ij}(\omega) \equiv b_{ij}(t_i)$ (Liu, 2020; Chen and Hu, 2024), whether a worker can only observe his own type t_i or can observe the

[19]The requirement here is slightly stronger but simpler than those in Chen and Hu (2020, 2023, 2024), in the sense that their operator $H_\mu(\cdot)$ applies only to the common knowledge event containing ω. We adopt this formulation without losing any learning feature of the matching market, which is the main focus of our paper.

entire worker type profile $\omega = (t_1, \ldots, t_n)$ makes no difference in predicting the market outcome. For simplicity, we assume that each worker can observe the entire state. Second, some papers, e.g., Liu et al. (2014), assume observability within matched pairs, i.e., each matched firm j can observe the true type of her employee $\mu(j)$, while others may not, e.g., Liu (2020). We follow the latter here, yet we notice that imposing the observability assumption is straightforward (Hu, 2026).

D.3 Stability theories and learning channels

This subsection relates the three learning channels of Section 2.2 to the stability notions and matching processes in the literature. Two of the channels (*Own* and *No*) are already visible in the static stability concept of Section D.2. All three appear in dynamic matching processes.

In the static concept, *Own* is built into how a firm evaluates a blocking opportunity. The set $D_{\mu,ij}$ in Definition D.2 restricts attention to states in which the worker would be willing to deviate, so the firm's evaluation already embodies what the worker's willingness reveals. *No* is what separates stability under incomplete information from the standard requirements of individual rationality and the absence of blocking pairs. The information-stability condition $H_\mu(\Pi) = \Pi$ requires that nothing further can be learned from the fact that no deviation has occurred.

These two channels are not artifacts of the particular notion above. The stability concepts developed since Liu et al. (2014) are closely connected and in several cases equivalent (Liu, 2020; Chen and Hu, 2020; Pomatto, 2022; Wang, 2023; Chen and Hu, 2024), and the notion of Section D.2 belongs to that family through Chen and Hu (2020, 2024). Whichever member one adopts, an uninformed agent must take into account a prospective partner's willingness to match and must draw a further inference from the absence of deviations. Therefore, the same learning channels arise across a broad class of matching models with incomplete information, which is why

our design need not commit to any one of them.[20]

In dynamic matching processes, all three learning channels are at work. Definition D.3 evaluates a given outcome, but it admits a dynamic interpretation. Following Chen and Hu (2020, 2024), the matching μ and the information structure Π may both be viewed as endogenous objects generated by a learning-and-blocking process, and the convergence results of those papers carry over to our setting with only notational changes. *Own*, *Other*, and *No* are exactly the canonical channels through which Π evolves along such a process. *Other* is absent from the static concept because that concept evaluates a single outcome, whereas *Other* concerns what an agent learns as the market moves.

Dynamic matching processes are modeled in the literature in two ways. The path-to-stability approach following Knuth (1976) and Roth and Vande Vate (1990) satisfies blocking pairs sequentially and treats agents as myopic, whereas non-cooperative models specify strategies explicitly (Pais, 2008; Haeringer and Wooders, 2011; Lauermann and Nöldeke, 2014; Ferdowsian, Niederle and Yariv, 2025). Our experimental markets make the distinction immaterial. Each agent has a single acceptable partner, so the only decision is whether to match rather than whom to match with, which removes the strategic concerns. Relatedly, our identification strategy depends only on which matches ultimately form and requires no commitment to a particular matching protocol.

[20]Unlike much of this literature, we assume away transfers. Since our focus is on learning and stability rather than matching efficiency, a non-transferable utility environment is better suited to our experimental design because it eliminates the screening role of transfers.