Learning in Matching*

Shan Gui Simin He Gaoji Hu†

December 31, 2025

Abstract

This paper studies how learning shapes outcomes in decentralized matching markets with incomplete information. In a one-to-one worker–firm matching environment where workers privately observe a payoff-relevant state, firms learn from matching behavior. Building on recent Bayesian and prior-free theories of stable matching, we experimentally examine whether stability can emerge through decentralized processes and how learning drives convergence. We identify three fundamental learning patterns, learning from conditional evaluation, learning from blocking, and learning from no blocking, each with fully and partially revealing subtypes, as well as compound learning patterns. Our results show that most markets converge to stable outcomes, demonstrating the predictive power of the theory; however, learning from no blocking and full revelation of the state hinder convergence, and our analysis provides evidence on the mechanisms driving these effects.

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

Keywords: matching, incomplete information, learning.

*We would like to thank Vikram Manjunath for 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. 72394391, No. 72033004). 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, Shanghai, China. **Gui:** guishan@sufe.edu.cn. **He:** he.simin@mail.sufe.edu.cn. **Hu:** hu.gaoji@gmail.com.

1 Introduction

In two-sided markets such as marriage and labor markets, incomplete information is ubiquitous: individuals often have only imprecise knowledge about potential partners, and firms may not know the productivity of their prospective employees. Agents form matches based on what they know, and also learn from observing others' matching behavior. A matching market stabilizes when the interaction between learning and matching settles—that is, when agents' existing knowledge induces no further changes in matches, and the existing matching reveals no new information. Recent matching theory predicts plausible market outcomes through this lens; see Liu et al. (2014) and the literature it pioneers. In this paper, we experimentally examine whether stable matchings emerge in decentralized markets and how learning shapes matching outcomes.

Specifically, we analyze a one-to-one matching environment between workers and firms, where workers are informed of a payoff-relevant *state* that determines their effort costs and productivities (and thus both sides' payoffs), but which may be unknown to firms. Our matching model with one-sided incomplete information resembles both the prior-free framework of Liu et al. (2014) and the Bayesian framework of Liu (2020). As shown in these and subsequent works, incomplete information fundamentally changes the analysis of matching markets relative to the complete-information benchmark of Gale and Shapley (1962), particularly with respect to what agents can infer from various possible observations and how such inferences shape predictions of stable matchings.[1]

We evaluate market performance relative to a theoretical benchmark, referred to as a *stable outcome*, which is consistent with various solution concepts in the recent literature. Informally, a market outcome consisting of a matching and agents' knowledge about the true state is stable if (i) it is individually rational, (ii) it is not blocked by any worker-firm pair, and (iii) the fact of individual rationality and no blocking pair provides no additional information to agents. To the best of our knowledge, this is the first experimental test of such stability-based solution concepts in decentralized matching markets with incomplete information.[2]

[1] See Roth and Sotomayor (1990), and more recently Chiappori, Salanie and Che (2024), for comprehensive surveys of the matching theory with complete information, where agents are fully informed of everyone's characteristics and preferences.

[2] There have been some theoretical progress on this direction; see detailed discussions in Section 5.2. There

When testing these stability concepts, we pay particular attention to the role of *learning from outcomes*, the key new element relative to stability notions under complete information. In decentralized matching markets, three fundamental learning patterns may arise, each associated with a distinct type of matching outcome. First, a firm's (she) hiring decision may depend on the worker's (he) willingness to match. If a bona fide match is formed, the firm refines her knowledge by confirming the worker's willingness. We refer to this as *learning from conditional evaluation* (CE). Second, firms may learn from matches formed between other worker–firm pairs, as such matching behavior can reveal information about the state. We call this *learning from blocking* (LB). Third, firms may learn from the absence of certain worker-firm matches, since non-matching behavior can also be informative about the state. We refer to this pattern as *learning from no blocking* (LNB). These fundamental learning patterns are summarized in Table 1. We also study *compound learning* (Compound) that combine multiple fundamental patterns.

Table 1: Fundamental learning patterns and their subtypes

	Fully revealing	Partially revealing
Learning from conditional evaluation	CE-Full	CE-Partial
Learning from blocking	LB-Full	LB-Partial
Learning from no blocking	LNB-Full	LNB-Partial
Compound learning	CE+LB	-
	CE+LNB	-

Moreover, beyond the learning types, it is important to assess whether full revelation of the state is essential for stable matching. To this end, we introduce two subtypes within each fundamental learning pattern, depending on whether learning fully reveals the true state (*Full*) or only partially reveals it (*Partial*). In the fully-revealing case, firms can, in principle, uniquely pin down the state and match according to ex post payoffs. In contrast, partial-revealing may still suffice for firms' matching decisions; for instance, a firm may prefer to match as long as the potential employee is acceptable, regardless of his exact productivity. Finally, we also consider two compound learning subtypes: CE+LB and CE+LNB.[3]

is related experimental work in Pais, Pintér and Veszteg (2020), but our study differs in the role of incomplete information in bridging theory and experiment; see detailed discussion in Section 5.3.

[3]We do not include a compound LB+LNB as it requires a larger market than our two-by-two market.

Our experiment includes eight treatments listed in Table 1, plus a benchmark treatment, *Plain*, without learning. The primary treatment variable is the learning pattern, while the secondary variable is the subtype. We use a within-subject design with 20 parameter sets covering all market types. Each participant experiences all markets over 20 rounds in randomized order to mitigate order effects. Matching is decentralized: participants may freely propose, accept or reject offers, and leave any matches. Matches are finalized once the market ends, which occurs after a period of inactivity, a similar approach first adopted by Agranov et al. (2025). Markets feature incomplete information with three possible states, and payoffs depend on the realized state. Tentative matches are public, while individual actions are private.

Our results show that, in the absence of learning (Plain), markets achieve nearly complete stable matchings as participants gain experience. Introducing learning patterns has little effect under CE and LB, but substantially impedes stability under LNB and Compound. Most failures to reach stable matchings occur in environments involving LNB (i.e., in LNB or CE+LNB), indicating that learning from the absence of certain matches poses a particularly severe challenge. A plausible explanation is that LNB requires participants to infer the state from the absence of certain matches, which is cognitively more demanding than the inference required in CE or LB; the data support this interpretation.

Furthermore, learning subtypes also matter. Within both LB and LNB, stable matchings arise more often in the Partial subtype than in the Full subtype. This pattern suggests that participants rely on simple heuristics, such as focusing on avoiding large losses or counting favorable states. Because the Partial subtype involves more favorable states under the prior, and because negative payoffs may loom larger due to loss aversion, these heuristics can generate systematically better outcomes in Partial than in Full. Together, these findings highlight that both the structure of learning patterns and the degree of state revelation are critical determinants of whether decentralized matching markets converge to stable outcomes.

Related literature. Our paper complements two streams of literature: (i) matching theory with incomplete information and (ii) the decentralized implementation of stable matchings. We briefly review these literatures here to position our contribution and leave more detailed discussions to Section 5.

Since Gale and Shapley (1962) and Shapley and Shubik (1971), a standard assumption in matching theory has been complete information—that is, agents' characteristics and preferences are commonly known. Liu et al. (2014) made a significant breakthrough by introducing one-sided (worker-side) incomplete information in a prior-free framework that imposes no prior beliefs on the firm side. Their framework has since stimulated a rich line of research, including implementation via adaptive matching processes (Chen and Hu, 2020), epistemic foundations (Pomatto, 2022), two-sided incomplete information (Chen and Hu, 2023; Park, 2025), rural–hospital structures (Peralta, 2025), and efficiency of stable matchings (Chen and Ho, 2025), among others.

In contrast, Liu (2020) introduces a Bayesian framework in which the stability of matchings shall be accompanied by the Bayesian consistency of prior, on-path and off-path beliefs. This approach has likewise been extended to study efficiency of stable matchings and adaptive matching processes (Chen and Hu, 2024), epistemic foundations (Wang, 2023), lattice and rural-hospital structures (Hu, 2025), and general cooperative analysis with (two-sided) incomplete information (Liu, 2023). See also Bikhchandani (2017) and Alston (2020) for an alternative Bayesian approach. These aforementioned papers together build a sound theoretical foundation for matching with incomplete information; see Liu (2024) for a survey.[4]

We contribute to this theoretical literature by experimentally testing stability concepts and identifying learning patterns in decentralized matching markets. Although we present a Bayesian theoretical model, our experiment design fits both the prior-free framework following Liu et al. (2014) and the Bayesian framework following Liu (2020).[5]

The decentralized matching markets in our experimental design relate both to the cooperative "path-to-stability" literature, beginning with Roth and Vande Vate (1990), and to non-cooperative matching theories, such as Pais (2008), Haeringer and Wooders (2011), and Ferdowsian, Niederle and Yariv (2025). The novelty of our design lies in circumventing the

[4]Topics loosely related to this literature include stable mechanisms with uncertain preferences of others (Roth, 1989; Ehlers and Massó, 2007, 2015; Yenmez, 2013), centralized deferred acceptance mechanism with incomplete information Fernandez, Rudov and Yariv (2022), search and matching (Ferdowsian, 2023), information acquisition in matching (Immorlica et al., 2020), etc.

[5]Particularly, our theoretical prediction is consistent with various stability concepts in the recent literature. See Section 5.1 for detailed discussions. One departure is that, unlike Liu et al. (2014) and Liu (2020), we assume away transfers, and thus their role as a screening tool, to focus on learning from match formations.

differences between these two approaches, allowing us to focus directly on learning without relying on ad hoc assumptions about decentralized environments.[6]

There are a few experimental papers testing whether markets achieve stable matchings using a decentralized approach. However, all of these studies test stability concepts derived under the assumption of complete information (Nalbantian and Schotter, 1995; Otto and Bolle, 2011; Pais, Pintér and Veszteg, 2020; Dolgopolov et al., 2024; He et al., 2024; Agranov et al., 2025; Echenique, Robinson-Cortés and Yariv, 2025). These papers vary in two dimensions: (i) whether the information structure is complete information or incomplete information, and (ii) whether the model features transferable or nontransferable utilities.

Most experimental studies focus on matching markets with transferable utility and examine stability, efficiency, and surplus allocation (Nalbantian and Schotter, 1995; Otto and Bolle, 2011; Dolgopolov et al., 2024; He et al., 2024; Agranov et al., 2025). Among them, only Nalbantian and Schotter (1995) and Agranov et al. (2025) introduce incomplete information. In both cases, however, incomplete information is implemented as an experimental friction, while market outcomes are evaluated using theories that assume complete information.

Only Pais, Pintér and Veszteg (2020) and Echenique, Robinson-Cortés and Yariv (2025) study decentralized matching without transfers. Of these, only Pais, Pintér and Veszteg (2020) consider incomplete information, where participants know only their own preferences but not the entire preference profile. Nevertheless, both papers continue to assess outcomes using complete-information stability concepts.

In summary, our paper contributes to the literature in three ways. First, we provide the first experimental test of stability concepts in matching markets with incomplete information, as pioneered by Liu et al. (2014). Second, we show that decentralized interactions can lead to stable matchings in most market environments. Third, we identify two key factors that make stability difficult to achieve: (i) when stability requires learning from the absence of blocking, and (ii) when a greater number of states leads to a loss, rather than a gain, compared to the default of non-matching.

[6] See Section 5.2 for detailed discussions of different approaches to decentralized matching.

2 Matching with Incomplete Information

2.1 The Model

Consider a one-to-one matching market with one-sided incomplete information and without monetary transfer. Specifically, let $I = \{1, \dots, n\}$ be a set of workers, and $J = \{n+1, \dots, n+m\}$ be a set of firms. 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 each $j \in J$,

(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 $\mu(j) = j$, we say that the agent is *unmatched*. Assume that μ is 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 $b_{ij}(\omega) \in \mathbb{R}$ be the ex post *matching values* worker i and firm j receive, respectively, when they are matched, i.e., $\mu(i) = j$, and when the realized state is ω. Normalize the unmatched values to zero, that is, $a_{ii}(\omega) = b_{jj}(\omega) = 0$. Assume that each worker can observe the realized state, whereas firms may not. Instead, the firms' *information structure* is described by a partition profile

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

where for each firm $j \in J$, Π_j is a partition over Ω; $\omega' \in \Pi_j(\omega)$ means that if the true state is ω, firm j cannot distinguish ω from ω'. Each agent can also have publicly observable payoff-relevant attributes that are summarized by their indices $i \in I$ and $j \in J$. Finally, the functional forms of $a : I \times J \times \Omega \to \mathbb{R}$ and $b : I \times J \times \Omega \to \mathbb{R}$ are commonly known.

A *market outcome*, or simply an *outcome*, (μ, ω, Π) specifies a matching μ, a realized state ω which may not be publicly observable, and an information structure Π. As in standard knowledge models, Π is commonly known; we will revisit the interpretations of Π in Section 2.3, after introducing some key concepts. Here, we illustrate a market outcome first.

Example 1 (An illustration of market outcomes)**.**
Consider a matching market with $I = \{i_1, i_2\}$ and $J = \{j_1, j_2\}$. There are three equally

probable states, i.e., $\Omega = \{\omega, \omega', \omega''\}$ and $\beta(\omega) = \beta(\omega') = \beta(\omega'') = 1/3$. The matching values are given in the table below:

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

where the numbers in each box correspond to

$$\begin{array}{|cc|} \hline a_{ij}(\omega) & b_{ij}(\omega) \\ a_{ij}(\omega') & b_{ij}(\omega') \\ a_{ij}(\omega'') & b_{ij}(\omega'') \\ \hline \end{array}.$$ [7]

Consider two market outcomes with the realized state being ω:

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

(B) $(\hat{\mu}, \omega, \hat{\Pi})$, where $\hat{\mu}(i_1) = j_1$, $\hat{\mu}(i_2) = i_2$, $\hat{\Pi}_{j_1} = \{\{\omega, \omega'\}, \{\omega''\}\}$, and $\hat{\Pi}_{j_2} = \{\{\omega, \omega'\}, \{\omega''\}\}$.

In the first outcome, all agents are unmatched and both firms are uninformed; agents' payoffs are all zero. In the second, there is a matched pair (i_1, j_1) and firms are partially informed; worker i_1's payoff is 1, firm j_1's *expected payoff* is 3 conditioning on the event $\Pi_{j_1}(\omega) = \{\omega, \omega'\}$, whereas worker i_2's and firm j_2's payoffs are both zero.

Remark 1 (Observability of the realized state)**.** 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),

[7] For convenience in discussion, we specify matching values directly in the examples, rather than specifying the states and functional forms of a and b. We will present matching values in this format in all examples, unless otherwise specified.

whether a worker can only observe his own type t_i or can observe the 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, 2025).

2.2 Stability of Market Outcomes

Following Chen and Hu (2020, 2023, 2024), and indirectly Liu et al. (2014) and Liu (2020), stability of a market outcome has three requirements: (1) individual rationality, (2) no blocking pair (equivalently, no pairwise deviation) and (3) information stability.

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

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

In Example 1, both outcome (A) and outcome (B) are IR.

To define a blocking pair (i, j) for the outcome (μ, ω, Π), where $\mu(i) \neq j$, we first clarify firm j's belief when evaluating her potential employee i. Let $D_{\mu,ij}$ be the set of states under which worker i benefits from the rematching with firm j, i.e.,

$$D_{\mu,ij} := \left\{\tilde{\omega} \in \Omega : a_{ij}(\tilde{\omega}) > a_{i\mu(i)}(\tilde{\omega})\right\}. \tag{1}$$

Intuitively, for the potential deviation by (i, j) to be viable, worker i and firm j must both expect to benefit from rematching with each other. When firm j calculates her expected payoff, a state $\tilde{\omega}$ is relevant for firm j only when $\tilde{\omega} \in D_{\mu,ij}$; any state that violates the inequality in (1) is irrelevant due to the worker's objection. Based on firm j's initial knowledge $\Pi_j(\omega)$ and the hypothetical knowledge $D_{\mu,ij}$, firm j's belief shall be $\beta(\cdot|\Pi_j(\omega) \cap D_{\mu,ij})$, which is referred to as the "off-path" belief in Liu (2020).[8]

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

Definition 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.,*

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

In Example 1, outcome (A) with all agents unmatched is blocked by (i_1, j_1) since

$$\begin{aligned} a_{i_1j_1}(\omega) = 1 &> 0 = a_{i_1i_1}(\omega) \quad \text{and} \\ \mathbb{E}\left[b_{i_1j_1}\middle|\Pi_{j_1}(\omega)\cap D_{\mu,i_1j_1}\right] = \frac{1}{2}b_{i_1j_1}(\omega) + \frac{1}{2}b_{i_1j_1}(\omega') = 3 &> 0 = \mathbb{E}\left[b_{j_1j_1}\middle|\Pi_{j_1}(\omega)\cap D_{\mu,i_1j_1}\right], \end{aligned}$$

where $\Pi_{j_1}(\omega) = \{\omega, \omega', \omega''\}$ and $D_{\mu,i_1j_1} = \{\omega, \omega'\}$.

Finally, for an outcome to be "stable", we introduce the concept of *information stability*. It captures the intuition that the absence of individual or pairwise deviation does not convey additional information to firms (beyond what is already described in the market outcome—Π). This requirement is crucial because the absence of deviation may indeed provide additional information to firms, which could in turn lead to a deviation based on firms' updated beliefs.

Example 2 (Information updating from the absence of deviation)**.**
Consider the matching market in Example 1, but with the following matching-value table:

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

Consider three market outcomes which differ only in the realized state:

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

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

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

All three outcomes are individually rational. According to Definition 2, outcome (X) is not blocked by (i_1, j_1) because of worker i_1's objection; it is not blocked by (i_2, j_2) because firm

j_2 worries about the terrible matching values -4 and -8 when $D_{\mu,i_2j_2} = \{\omega, \omega', \omega''\}$. Thus, outcome (X) admits no blocking pair. In contrast, it is straightforward to verify that both outcomes (Y) and (Z) 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 (X). 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 state (X) should have occurred. The absence of such 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.

To formalize information stability, we define a set of no-deviation states as follows:

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

Given an outcome (μ, ω, Π) that is IR and not blocked, intuitively, upon noticing *the absence of individual/pairwise deviation*, each firm should refine her partition according to the newly acquired information that *the true state must lie in* $N_{\mu,\Pi}$. For notational convenience, we denote by $\mathcal{N}_{\mu,\Pi}$ the binary partition induced by $N_{\mu,\Pi}$; that is, $\mathcal{N}_{\mu,\Pi} := \{N_{\mu,\Pi},\ \Omega \setminus N_{\mu,\Pi}\}$. Define an operator $H_\mu(\cdot)$ (History) to represent the information refinement (i.e., the *join* between Π and $\mathcal{N}_{\mu,\Pi}$) as follows:

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

In Example 2, $N_{\mu,\Pi} = \{\omega\}$, $\mathcal{N}_{\mu,\Pi} = \{\{\omega\}, \{\omega', \omega''\}\}$, and $[H_\mu(\Pi)]_{j_2} = \{\{\omega\}, \{\omega', \omega''\}\} \neq \Pi_{j_2}$. If $H_\mu(\Pi) = \Pi$, then the fact of individual rationality and no blocking pair provides no further information to firms (in addition to their knowledge Π).

Definition 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) *information stability holds, that is,* $H_\mu(\Pi) = \Pi$.[9]

[9]The requirement here is slightly stronger but simpler than those in Chen and Hu (2020, 2023, 2024), in

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

In Example 1, suppose the true state is ω. Then an outcome (C) $(\check{\mu}, \omega, \check{\Pi})$, where

$$\check{\mu}(i_1) = j_1,\ \check{\mu}(i_2) = j_2 \text{ and } \check{\Pi}_{j_1} = \check{\Pi}_{j_2} = \left\{\left\{\omega, \omega'\right\}, \left\{\omega''\right\}\right\},$$

is stable. Although there are four stable outcomes, they only differ in the information structure (the other three are strict refinements of $\check{\Pi}$). In other words, $\check{\mu}$ is a unique stable matching. Similarly, in Example 2, suppose the true state is ω. Then an outcome (W) $(\bar{\mu}, \omega, \bar{\Pi})$ is stable, where $\bar{\mu}(i_1) = i_1$, $\bar{\mu}(i_2) = j_2$, and $\bar{\Pi}_{j_1} = \{\{\omega\}, \{\omega', \omega''\}\}$, and $\bar{\Pi}_{j_2} = \{\{\omega\}, \{\omega', \omega''\}\}$. Although there are four stable outcomes, they only differ in the information structure (the other three are strict refinements of $\bar{\Pi}$). In other words, $\bar{\mu}$ is a unique stable matching.

2.3 Learning in Dynamic Matching Processes

So far we have introduced a static setting and asked whether or not an exogenously given putative outcome (μ, ω, Π) is stable. However, like in Chen and Hu (2020, 2024), we can also interpret both μ and Π as the result of a "learning-and-blocking" process in the matching market, where they are *endogenously* determined by market dynamics.[10] In what follows, we introduce three kinds of learning in matching with incomplete information. Our analysis shall be clear without formally introducing the matching process of Chen and Hu (2020, 2024).

In Example 1, outcome (A) is the autarky where no one is matched and both firms are uninformed. We have argued that outcome (A) is blocked by (i_1, j_1). Now if i_1 and j_1 indeed form a new match, how would firms update their information? Firm j_1's initial knowledge is $\Pi_{j_1} = \{\{\omega, \omega', \omega''\}\}$. When she evaluates worker i_1, she only considers states in $D_{\mu, i_1 j_1} = \{\omega, \omega'\}$, which is a piece of *hypothetical* information regardless of worker i_1's real preference. Nevertheless, when the new match (i_1, j_1) actually takes place, this hypothetical information is confirmed. Thus, firm j_1's knowledge shall be updated to $\hat{\Pi}_{j_1} = \{\{\omega, \omega'\}, \{\omega''\}\}$. We refer to this kind of learning as *learning from conditional evaluation (CE).*

Since matching is publicly observable, firm j_2 observes the new match (i_1, j_1). She can

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.

[10] Naturally, when we discuss market dynamics where bona fide deviations take place along a path, agents' evaluation of each other shall be based on a state ω (or a partition cell containing it) that is realized and fixed.

infer that the true state must be either ω or ω'. Therefore, firm j_2 shall update her information to $\hat{\Pi}_{j_2} = \{\{\omega, \omega'\}, \{\omega''\}\}$. We refer to this second kind of learning as *learning from blocking (LB)*, which generally includes learning from *other agents'* individual/pairwise deviation. In fact, the new match, together with CE and LB, transitions the market outcome from (A) to (B). With the updated information $\hat{\Pi}_{j_2}$, firm j_2 forms a blocking pair with worker i_2 for outcome (B), which leads to the stable market outcome (C) where two pairs are matched up.

Finally, we have argued in Example 2 that firm j_2, when observing no deviation (potentially by worker i_1 and firm j_1), shall update her information from $\Pi_{j_2} = \{\{\omega, \omega', \omega''\}\}$ to $\bar{\Pi}_{j_2} = \{\{\omega\}, \{\omega', \omega''\}\}$, which leads to a blocking pair with worker i_2 for outcome (X). If they indeed form a new match, the market outcome becomes (W). We refer to this learning pattern as *learning from no blocking (LNB)*.

Remark 2 (Individual "blocking")**.** Although we motivate the learning patterns using pairwise deviations or the absence of them, the idea applies to individual deviations as well. More precisely, suppose the status quo market outcome already includes a matched pair (i, j). If they break up, then agents shall learn from this "individual blocking" (LB). Alternatively, if they do not break up, then agents shall learn from the absence of "individual blocking" or, equivalently, from IR (LNB).

These three learning patterns are arguably fundamental in a matching market. Moreover, they are exhaustive if we maintain the cooperative perspective.[11] In particular, when we focus merely on the evolution of publicly observable matches, it is only possible for a firm to learn from (1) her own matching experience (CE), (2) other agents' new match or separation (LB), and (3) absence of other agents' new match or separation (LNB). Of course, compound learning patterns based on those fundamental ones are possible.

It has been shown in Chen and Hu (2020, 2024) that if agents keep seeking for better matches, and those new matches are formed randomly when there are many, a "learning-and-blocking" process converges to a stable outcome with probability one. Their results provide a dynamic foundation for incomplete-information stability concepts following Liu et al. (2014)

[11] That is, we focus on the change of matches while ignore the details such as "who proposes to whom" or "how long it takes". As Liu (2023) points out, cooperative modeling, as a reduced form, proves useful in analyzing complex markets. It makes no excessive ad hoc assumptions while keeps the flexibility of doing so.

and Liu (2020), which also applies to our setting straightforwardly (with only notational adaption).

In the current paper, we will design matching markets with the following features:

1. There is a unique stable matching.
2. Suppose agents maximize their matching values. There is a unique "learning-and-blocking" path from the autarky (no match and no information) to the unique stable matching.
3. Each learning pattern of CE, LB and LNB, as well as some compound learning patterns, can be uniquely identified by the change of matching in a dynamic market.

Remark 3 (Naive versus sophisticated agents)**.** Chen and Hu (2020, 2024) assume naive/myopic agents when they prove the convergence of their adaptive market processes, i.e., agents opt to execute a blocking opportunity whenever it exists. In decentralized matching markets, e.g., those in Lauermann and Nöldeke (2014) and Ferdowsian, Niederle and Yariv (2025), agents may be sophisticated and strategic interactions make the analysis complex. However, in our setting, this difference is circumvented. Particularly, we design matching markets to make sure that each agent has only one potential partner. Thus, the choice is merely whether to get matched, rather than to which partner.[12]

2.4 Identification of Learning Patterns

We consider eight matching markets in this section. Similar to Examples 1-2, each market has two workers, two firms and three equally probable states. The basic notations remain the same, so are omitted. Suppose the true state is ω and all markets start with the autarky outcome (μ, ω, Π)—that is, $\mu(i_1) = i_1$, $\mu(i_2) = i_2$, and $\Pi_{j_1} = \Pi_{j_2} = \{\{\omega, \omega', \omega''\}\}$. Across the eight markets, we vary matching values to ensure that in each market a unique evolution path leading to stability can identify a specific learning pattern, as long as *agents maximize their matching values*. Two of them identify the CE pattern, two LB, two LNB, and finally two compound learning patterns.

We start with two "CE" markets; the matching value tables are displayed in Figure 1. In market CE-Full (fully revealing), the autarky outcome is blocked by worker i_2 and firm j_2.

[12]The fragility issue pointed out by Rudov (2024) is also circumvented.

When they form a new match (i_2, j_2), firm j_2 learns from the worker's willingness that the true state must be ω (recall that workers can observe the true state). This leads to the unique stable matching $\check{\mu}$, where $\check{\mu}(i_1) = i_1$ and $\check{\mu}(i_2) = j_2$. Market CE-Partial (partially revealing) slightly differs from CE-Full. There, when a new match (i_2, j_2) takes place, firm j_2 learns from the worker's willingness that the true state must be either ω or ω', but never ω''. In these two markets, the learning pattern of CE is uniquely identified by the status of matches, i.e., the match (i_2, j_2) is observed if and only if learning from conditional evaluation occurs, whereas the details such as "who proposes to whom" does not matter.[13]

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

(a) CE-Full

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

(b) CE-Partial

Figure 1: Two markets with learning from conditional evaluation.

Figure 2 shows two "LB" markets. In market LB-Full, the autarky outcome is blocked by worker i_1 and firm j_1, but not by i_2 and j_2 since j_2 worries about the negative payoffs -4 and -8. Forming the new match (i_1, j_1) does not require firm j_1 to conduct conditional evaluation; she is always willing to deviate from the autarky. However, when the match (i_1, j_1) takes place, both firms get fully informed that the true state must be ω. The market outcome transitions to $(\hat{\mu}, \omega, \hat{\Pi})$, where $\hat{\mu}(i_1) = j_1$, $\hat{\mu}(i_2) = i_2$ and $\hat{\Pi}_{j_1} = \hat{\Pi}_{j_2} = \{\{\omega\}, \{\omega', \omega''\}\}$. Now firm j_2 can rule out the negative payoffs and form a blocking pair with worker i_2, leading to the unique stable matching $\check{\mu}$, where $\check{\mu}(i_1) = j_1$ and $\check{\mu}(i_2) = j_2$. Market LB-Partial is similar, only that j_2 learns partially about the true state ω from the match (i_1, j_1). Again, the learning pattern of LB (through firm j_2) is uniquely identified by the status of matches: the match (i_2, j_2) is observed if and only if learning from blocking (by worker i_1 and firm j_1) occurs.

Figure 3 shows two "LNB" markets. The market LNB-Full is identical to Example 2, in

[13] Indeed, firm j_2 may attempt to probe worker i_2's willingness, which is frequently recorded in our experiments.

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

(a) LB-Full

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

(b) LB-Partial

Figure 2: Two markets with learning from blocking.

which the match (i_2, j_2) is observed if and only if learning from no blocking (potentially by worker i_1 and firm j_1) occurs. LNB-Partial differs only in that j_2 learns partially about the true state ω from the *absence* of match (i_1, j_1).

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

(a) LNB-Full

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

(b) LNB-Partial

Figure 3: Two markets with learning from no blocking.

Finally, we consider two compound learning patterns CE+LB and CE+LNB in Figure 4. In market CE+LB, the autarky outcome is blocked by worker i_1 and firm j_1, which does not require firm j_1's conditional evaluation or learning from the outcome. In contrast, worker i_2 and firm j_2 do not form a blocking pair for the autarky, since firm j_2 worries about the negative payoff -8, even if she takes into account the hypothetical information $D_{\mu,i_2j_2} = \{\omega, \omega''\}$. When the match (i_1, j_1) takes place, firm j_2 learns from this blocking that the true state must be either ω or ω'. Therefore, the market outcome transitions from the autarky to $(\hat{\mu}, \omega, \hat{\Pi})$, where $\hat{\mu}(i_1) = j_1$, $\hat{\mu}(i_2) = i_2$ and $\hat{\Pi}_{j_1} = \hat{\Pi}_{j_2} = \{\{\omega, \omega'\}, \{\omega''\}\}$. Now firm j_2 forms a blocking pair with worker i_2, where $D_{\mu,i_2j_2} = \{\omega, \omega''\}$ is confirmed informative. Here, the match (i_2, j_2) is observed if and only if learning from blocking (by worker i_1 and firm j_1) and conditional evaluation simultaneously occur. Similarly, in market CE+LNB, the match (i_2, j_2) is observed if and only if learning from no blocking (potentially by worker i_1 and firm j_1) and conditional evaluation simultaneously occur.

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

(a) CE+LB

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

(b) CE+LNB

Figure 4: Two markets with compound learning patterns.

How about the compound learning pattern of LB+LNB? For experimental tractability, we do not include this case. Indeed, it is impossible to have LB+LNB in a two-worker-two-firm market, for the following reason: First of all, learning from blocking involves a pair to be matched, i.e., the "blocking" part, such as (i_1, j_1). Identifying the effect of learning involves another pair to be matched, such as (i_2, j_2), where firm j_2 learns. Similarly, learning from no blocking involves at least one "off-path pair in mind" to be matched such as (i_3, j_3), which, however, does not actually occur. Identifying the effect of LNB involves another pair to be matched, say (i_2, j_2) or some other pair. Therefore, to have LB+LNB identified, the market shall have at least three firms, which may make other identifications messy.

Remark 4 (Risk attitude)**.** We assumed that firms are risk neutral. However, the matching markets we have designed can identify the learning patterns even if firms are risk averse (whichever extent) or slightly risk seeking. For example, in the market LNB-Partial (Figure 3b), firm j_2 is initially uninformed, which implies $\mathbb{E}\left[b_{i_2 j_2} \mid \Pi_{j_2}(\omega)\right] = 1/3(2+4-8) = -2/3$. Even if firm j_2 is slightly risk seeking and puts more weights on the values 2 and 4, she still prefers not being matched with i_2 without learning from no blocking (potentially by (i_1, j_1)). When firms are extremely risk averse in that they prefer to be matched only if they can clearly guarantee a positive (or better) payoff, we shall adopt the prior-free stability concept of Liu et al. (2014) and Chen and Hu (2020). Other than that, all our experiment design and analysis carry over.

3 Experiment Design

3.1 Treatment variables

According to the theory, we focus on two layers of treatment variables. First, our primary interest lies in the four *learning patterns*: CE, LB, LNB, and Compound, while Plain serves as a baseline without any learning technique. Second, within each of the four learning patterns we introduce two *subtypes*. For CE, LB, and LNB, the subtypes differ in whether state information can be *fully* revealed or *partially* revealed after learning: Full revelation is possible in CE-full, LB-Full, and LNB-Full, but not in CE-Partial, LB-Partial, or LNB-Partial. For the Compound learning pattern, the two subtypes differ in their composition: CE is combined with LB in the first subtype CE+LB, and with LNB in the second subtype CE+LNB.

Across these two treatment dimensions, we employ 20 payoff parameter sets, shown in Figures A.1 and A.2, each corresponding to a specific learning pattern and subtype. In both figures, the first row displays the Plain parameters; the second row presents the CE parameters (CE-Full on the left, CE-Partial on the right); the third row shows the LB parameters (LB-Full on the left, LB-Partial on the right); the fourth row contains the LNB parameters (LNB-Full on the left, LNB-Partial on the right); and the final row reports the Compound parameters (CE+LB on the left, CE+LNB on the right).

3.2 Treatment orders

We use a within-subject design for our treatment variables: each subject plays all 20 markets (Figures A.1 and A.2) over 20 rounds, one per round. Subjects differ only in the order in which these 20 markets are presented. The order is determined by a pre-randomization procedure that is randomized but subject to three sets of constraints.

First, to control for order effects across learning patterns, we divide the 20 markets into four blocks of five markets (Blocks A–D). Each block contains exactly one market from each learning pattern: Plain, CE, LB, LNB, and Compound. The sequence in which these four blocks are played is varied across matching groups using a Latin-square design, generating four distinct treatment orders:

Division of 20 markets/rounds	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

Second, within each block, the five learning patterns follow a structured order. All blocks begin with the Plain market. The remaining four learning patterns appear in different orders across blocks, again determined by a Latin-square design. This ensures variation in the relative positions of CE, LB, LNB, and Compound. The four within-block sequences are:

Five markets/rounds in each block	1	2	3	4	5
Block A	Plain	CE	LB	LNB	Compound
Block B	Plain	LB	LNB	Compound	CE
Block C	Plain	LNB	Compound	CE	LB
Block D	Plain	Compound	CE	LB	LNB

Third, to control for subtype order effects, the pre-randomization additionally ensures that, within each learning pattern, the two subtypes appear an equal number of times in the first 10 rounds and the last 10 rounds. This balances subjects' exposure to subtypes between the early and later stages of the experiment.[14]

3.3 Design details

At the beginning of the experiment, subjects are randomly assigned to matching groups, each of which has size 12 and is fixed throughout the experiment. Each matching group follows one of the four treatment orders for the 20 rounds. Within each group, 6 subjects are randomly

[14] Ideally, the procedure should have been constrained only by the three requirements described above. Due to a coding error, however, the sequence of the four Plain markets was inadvertently fixed across all treatment orders, rather than being randomized. As a result, the Plain markets may be affected by order effects and should be interpreted with caution. Because our primary interest lies in comparing the four learning patterns and their subtypes, this issue does not materially affect our main research question.

assigned to the role of worker and 6 to the role of firm; these roles remain fixed throughout the experiment. Workers are labeled Lemon or Mango, and firms are labeled Yellow or Green. In each round and in each market, a worker is randomly assigned to be Lemon or Mango with equal probability, and a firm is randomly assigned to be Yellow or Green with equal probability.

At the start of each round, 12 subjects in a matching group are randomly partitioned into three matching markets, each consisting of two workers (one Lemon and one Mango) and two firms (one Yellow and one Green). The pre-randomization of role assignment ensures that, across the 20 rounds, each firm-role subject experiences a balanced mix of the easier role (j_1 in the model) and the harder role (j_2). In every round and in every market, a worker may match with a firm or stay unmatched; vice versa for firms.

Each matching market features one of three equally likely states: Sunny, Cloudy, or Rainy. Workers always observe the realized state, whereas firms do not. States are pre-randomized prior to the experiment and independently drawn across matching markets. All players can see the payoff matrix of the current matching market, which displays the worker and firm payoffs for each possible match in each state.[15] Each player receives an initial endowment of 10 points. Remaining unmatched yields this endowment only. A finalized matching generates an additional payoff shown in the payoff matrix, which may be positive or negative.[16]

Within each market, workers and firms may send proposals to one opponent at a time. A proposal must be accepted or rejected within 15 seconds; otherwise, it is automatically rejected. An accepted proposal forms a match. Matches are temporary throughout the market: players may dissolve existing matches and may proposal to or accept proposals from the other opponent to form new ones. Whenever a new match forms, any prior matches involving either player immediately dissolve. Each market runs for at least 60 seconds. A public information area displays a countdown timer above the payoff matrix. When a match forms or dissolves, the corresponding payoff cell lights up or dims, and the 60-second timer restarts. The market ends when the timer reaches zero or when all four players in the market click the "Agree to

[15] The presentation of the payoff matrix is pre-randomized. For each learning pattern, four games are assigned to each of the following payoff-change conditions: no change; row change (swapping i_1 and i_2); column change (swapping j_1 and j_2); and change in both rows and columns (swapping both).

[16] A player's total payoff in a market is non-negative and equals 10 plus the corresponding matching payoff.

Proceed to Next Round" button, which becomes available after the initial 60 seconds. All current temporary matches are finalized at the time the market ends. Subjects do not receive any feedback between rounds. An illustrative screenshot is provided in Figure B.2.

3.4 Procedures

The experiment was conducted at the Shanghai University of Finance and Economics in September and October 2025. Subjects were recruited from the Economics Lab's subject pool through Ancademy, a platform for social science experiments; most participants accessed the experiment using the Ancademy mobile app. We ran eight sessions, each with 24 subjects. In each session, the 24 subjects were randomly divided into two independent matching groups of 12. Treatments were randomized at the matching group level, allowing multiple treatment orders to be implemented within the same session.

In total, 192 subjects participated, forming 16 independent matching groups of 12, which were evenly assigned across the four treatment orders. Each subject participated in only one session. The participant pool consisted primarily of undergraduate students from a variety of majors.

The experiment was computerized using oTree (Chen, Schonger and Wickens (2016)) and conducted in Chinese (English translations of the instructions and screenshots are provided in Appendix B). Upon arrival, subjects were randomly assigned a card indicating their table number and seated in the corresponding cubicles. All instructions were presented on their computer screens, and participants completed a set of control questions to ensure comprehension. The same experimenter oversaw all experimental sessions.

At the end of the experiment, subjects completed a short demographic survey. Four of the 20 rounds were randomly selected for payment. The experimental currency was denominated directly in Chinese yuan (CNY). Average earnings were 70.51 CNY including a 20 CNY participation bonus (approximately 10.53 USD). Each session lasted about 50-60 minutes.

4 Results

In this section, we evaluate experimental performance relative to the predicted stable matching (i.e., the correct matching) along three dimensions. First, at the market level, we examine whether the markets converge to the correct matching. Second, at the individual level, we analyze whether and how participants make mistakes, with particular attention to heterogeneity across firm roles and to different types of mistakes. Finally, we study market dynamics, focusing on the time required to reach the correct matching and the sequence in which correct matches are achieved.

4.1 Market Outcome

4.1.1 Market correct rates by learning pattern

We first examine whether markets reach a stable matching. We classify a 4-player market as *correct* if all four players reach their predicted matches. Figure 5 reports the market correct rate by learning pattern for all 20 rounds (panel (a)) and for the last 10 rounds (panel (b)). Bars report means and error bars indicate ± one standard error, both computed at the 12-player matching-group level.[17]

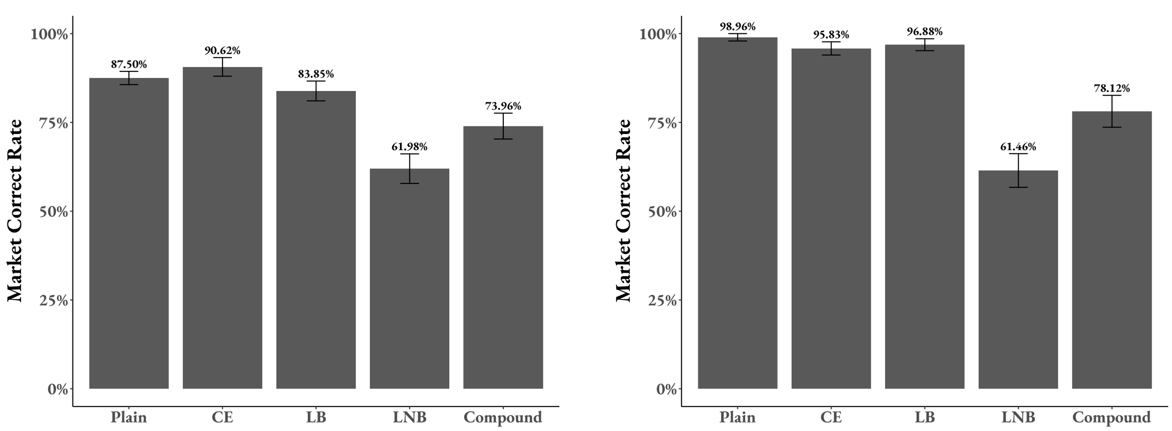

(a) All 20 rounds (b) Last 10 rounds

Figure 5: Market correct rates by learning pattern

[17]Throughout the Results section, bars show means and error bars indicate ± one standard error, both computed at the 12-player matching-group level; statistical tests are also conducted at this level.

As shown in Figure 5a, across all 20 rounds, market correct rates differ substantially across learning patterns, ranging from 61.98% in LNB to 90.62% in CE. For each of the five learning patterns, we reject the hypothesis that markets always reach the stable matching. Comparing Figures 5b and 5a, market correct rates are generally higher in the second half of the experiment, with the exception of LNB. Specifically, during the last 10 rounds, correct rates are close to 100% in Plain, CE, and LB, but remain below 80% in LNB and Compound.

Table 2: Market correct rates: pairwise tests across learning patterns

	CE	LB	LNB	Compound
Plain	0.071	0.293	< 0.001	< 0.001
CE		0.015	< 0.001	0.002
LB			0.001	0.016
LNB				0.013

(a) All 20 rounds

	CE	LB	LNB	Compound
Plain	0.083	0.317	< 0.001	0.002
CE		0.655	< 0.001	0.006
LB			< 0.001	0.001
LNB				0.021

(b) Last 10 rounds

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

Next, we test whether achieving the stable matching is more difficult in CE, LB, LNB, and Compound, compared to Plain. Table 2a reports pairwise Wilcoxon signed-rank tests across learning patterns using data from all 20 rounds. The results show that market correct rates in CE and LB do not differ significantly from Plain, whereas correct rates in LNB and Compound are significantly lower than in Plain. Moreover, correct rates are significantly lower in LB, LNB, and Compound than in CE, lower in LNB and Compound than in LB, and lowest in LNB. Turning to the last 10 rounds, Table 2b reports results that largely mirror those from the full sample, with one exception: the correct rate in LB is no longer significantly lower than in CE, indicating learning over time in LB.

In sum, at the market level, achieving the stable matching is more difficult in LB, LNB, and Compound than in Plain, but not in CE. Moreover, while performance in LB converges to that in Plain as participants gain experience, no comparable improvement is observed in LNB or Compound. Finally, LNB exhibits greater difficulty than Compound.

Result 1 (market-level correct rates by learning patterns)**.** *Across all rounds, no learning pattern achieves full correct rates. The ranking of market correct rates is*

$$\textit{Plain} = \textit{CE} > \textit{LB} > \textit{Compound} > \textit{LNB}.$$

In the last 10 rounds, however, markets under Plain, CE, and LB achieve nearly full correct rates, and the ranking becomes

$$Plain = CE = LB > Compound > LNB.$$

4.1.2 Market correct rates by subtype

In this section, we examine whether market correct rates differ by the secondary treatment variable, subtype, within each learning pattern. For CE, LB, and LNB, the subtypes differ in whether players can uniquely and certainly pin down the true state after learning, which is possible in the subtype *Full* but not in *Partial*. For Compound, the two subtypes differ in their composition, either *CE+LB*, or *CE+LNB*. Plain has no subtypes.

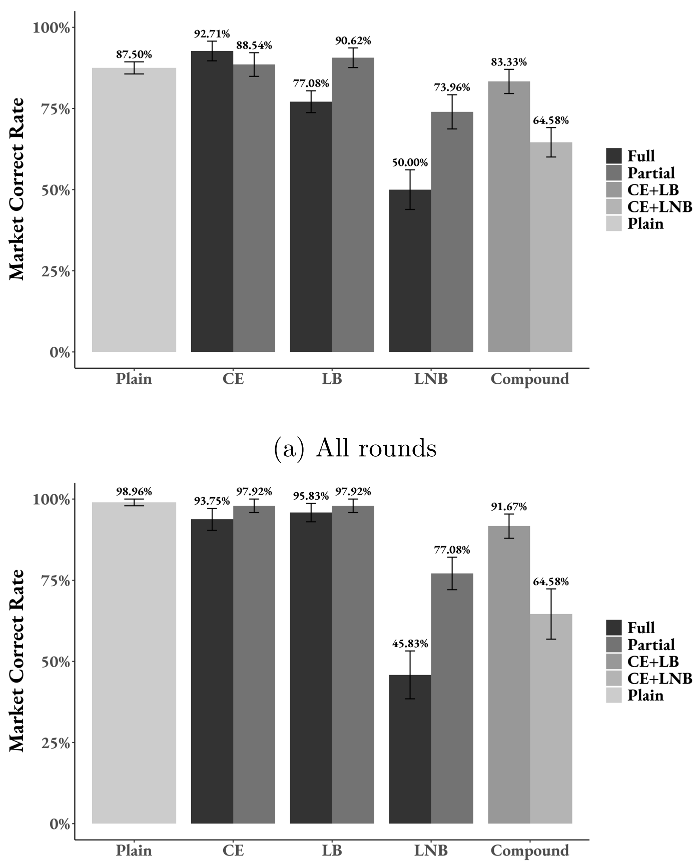

(a) All rounds

(b) Last 10 rounds

Figure 6: Market correct rates by subtypes

Figure 6 shows average market correct rates by subtype, with panel (a) covering all 20 rounds and (b) covering the last 10 rounds; corresponding statistical test results are shown in Table 3. Across all 20 rounds, within LB and LNB, market correct rates are significantly higher in the Partial subtype than in the Full subtype, with gaps of 13.54 percentage points in LB and 23.96 percentage points in LNB. Within Compound, the market correct rate is 18.75 percentage points higher in CE+LB than in CE+LNB, and the difference is statistically significant. This pattern is consistent with the comparison between non-compound LB and LNB.

Turning to the last 10 rounds, Figure 6b shows that market correct rates no longer differ significantly between Partial and Full within LB. In contrast, they remain significantly higher in LNB-Partial than in LNB-Full, as well as in CE+LB than in CE+LNB.

Table 3: Comparisons across subtypes

(a) All 20 rounds

	CE		LB		LNB		Compound
	Full	Partial	Full	Partial	Full	Partial	
Within Type	0.294		0.002		0.011		0.001
CE+LB	0.078	0.199	0.115	0.121			
CE+LNB	< 0.001	0.001			0.011	0.228	

(b) Last 10 rounds

	CE		LB		LNB		Compound
	Full	Partial	Full	Partial	Full	Partial	
Within Type	0.317		0.564		0.005		0.009
CE+LB	0.706	0.180	0.317	0.180			
CE+LNB	0.006	0.002			0.082	0.302	

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

The last two rows of Table 3 report Wilcoxon signed-rank tests comparing the two Compound subtypes with their non-compound counterparts, with panel (a) covering all 20 rounds and panel (b) covering the last 10 rounds. Across all 20 rounds, Table 3a shows that, relative to CE, market correct rates in CE+LB do not differ significantly from either CE subtype, whereas rates in CE+LNB are significantly lower than in both CE subtypes.

Relative to LB, correct rates in CE+LB do not differ from either LB subtype, while rates in CE+LNB are significantly higher than in LNB-Full but not significantly different from LNB-Partial. Turning to the last 10 rounds, results in Table 3b are largely similar to those from the full sample in Table 3a. The only difference is that the market correct rate in CE+LNB is no longer significantly higher than that in LNB-Full. Overall, relative to each non-compound counterparts, only the inclusion of LNB in Compound substantially reduces market performance.

Result 2 (market-level correct rates by subtypes)**.** *Across all rounds, market correct rates in LB and LNB are significantly lower in the Full subtype than in the Partial subtype, although this difference in LB is no longer significant in the last 10 rounds. In both the full sample and the last 10 rounds, market correct rates are significantly lower in CE+LNB than in CE+LB. Compared to the non-compound counterparts, including LNB in Compound reduces market correct rates.*

4.2 Individual Mistakes

Given that markets often fail to reach fully correct matching, we examine whether and how individuals make mistakes. We proceed in three steps. First, we analyze individual mistake rates by firm role, classified as *hard* or *easy* depending on whether learning from others' market activities is required to achieve the correct match.[18] We then focus on hard firms, examining how their performance differs across learning patterns and subtypes. Finally, we study how individuals make mistakes, focusing on the composition of mistake types and their evolution over rounds.

4.2.1 Mistake rates by role

We define an individual as making a *mistake* if her final match differs from the predicted one. In a four-player matching market, two players are workers who know the state, while the other two are firms. Among firms, difficulty differs across CE, LB, LNB, and Compound. A firm is

[18]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 **??**. Moreover, we do not analyze workers' mistake behavior separately, since it is mechanically reflected in firm mistake rates.

classified as a *Hard Firm* if it must update her belief about the state to achieve the correct match, while the other is classified as an *Easy Firm*. In Plain, by contrast, both firms face the same difficulty and neither needs to update beliefs to achieve the correct match. We therefore pool the firms in Plain and use them as a benchmark for comparison with Hard Firms in other treatments.

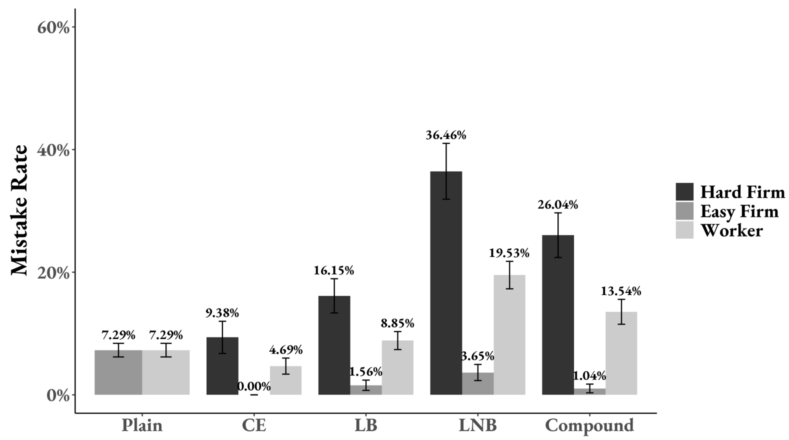

(a) All 20 rounds

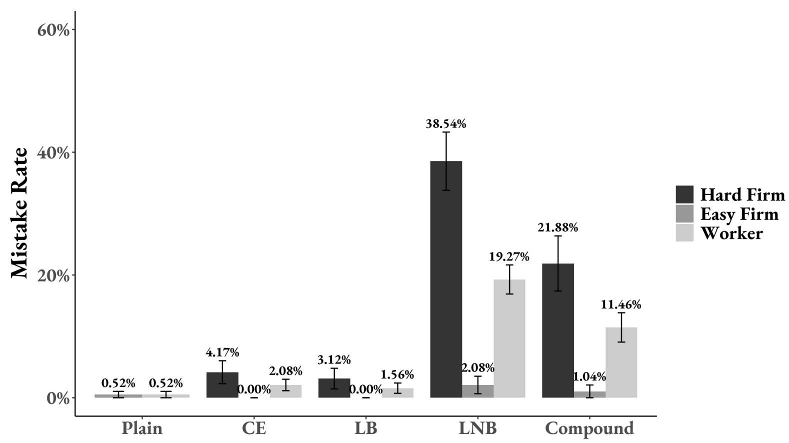

(b) Last 10 rounds

Figure 7: Individual mistake rates by learning pattern and by role

Figure 7 displays individual mistake rates by learning pattern and role, with panel (a) showing all 20 rounds and panel (b) showing the last 10 rounds. Across all learning patterns, Hard Firms exhibit the highest mistake rates, compared to both Easy Firms and workers. As shown in Figure 7a, mistake rates for easy firms fall below 4% across all learning patters in all rounds. In contrast, Hard Firms' mistake rates are lowest in CE (9.38%), followed by LB (16.15%), Compound (26.04%), and highest in LNB (36.46%). The benchmark Plain has a mistake rate of 7.29% for all firms. Figure 7b shows that, in the lats 10 rounds, mistake rates for Easy Firms fall below 3% in all learning patterns. Mistake rates for Hard Firms also decline substantially in CE, and LB, falling below 5%, consistent with learning over time in these environments. By contrast, mistake rates for Hard Firms remain high in LNB and Compound, suggesting persistent learning difficulties in these settings.

Table 4: Mistake rates for Hard Firms: pairwise tests across learning patterns

	CE	LB	LNB	Compound
Plain	0.852	0.009	< 0.001	< 0.001
CE		0.020	< 0.001	0.002
LB			0.002	0.010
LNB				0.069

(a) All 20 rounds

	CE	LB	LNB	Compound
Plain	0.046	0.271	< 0.001	0.001
CE		0.655	< 0.001	0.005
LB			< 0.001	0.001
LNB				0.023

(b) Last 10 rounds

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

Given that Hard Firms account for most individual mistakes, we focus on Hard Firms' mistakes. Table 4 reports the results of paired Wilcoxon signed-rank tests. Across all 20 rounds, Table 4a shows that, compared to Plain, Hard Firms do not make significantly more mistakes in CE ($p = 0.852$), but make significantly more mistakes in LB ($p = 0.009$), LNB ($p < 0.001$), and Compound ($p < 0.001$). Compared to LB, Hard Firms make significantly more mistakes in both LNB ($p = 0.002$) and Compound ($p = 0.010$), while the difference between LNB and Compound is not statistically significant ($p = 0.069$).

Turning to the last 10 rounds, as Hard Firms gain experience, mistake rates remain low in CE (4.17%) and LB (3.12%), neither of which differs significantly from Plain ($p = 0.046$ and $p = 0.271$). In contrast, mistake rates remain high in LNB (38.54%) and Compound (21.88%), with the rate in Compound significantly lower than that in LNB ($p = 0.023$). These results indicate that hard firms are able to learn over time in CE and LB, whereas persistent learning difficulties remain in LNB and Compound.

Result 3 (Hard-Firm mistake rates by learning patterns)**.** *For each learning pattern, Hard Firms account for most mistakes. Across all 20 rounds, the ranking of mistake rates for Hard Firms is*

$$\textit{Plain} < \textit{CE} < \textit{LB} < \textit{Compound} = \textit{LNB}.$$

In the last 10 rounds, this ranking changes to

$$\textit{Plain} = \textit{CE} = \textit{LB} < \textit{Compound} < \textit{LNB}.$$

4.2.2 Mistake rates for Hard Firms by subtype

The previous section shows that the mistake-rate patterns for Hard Firms closely mirror the market-level results, which is expected given that Hard Firms account for most mistakes. We now investigate whether mistake rates for Hard Firms also differ by subtype. Figure 8 plots average mistake rates by subtype, with panel (a) showing all 20 rounds and panel (b) showing the last 10 rounds. Corresponding statistical tests are reported in Table 5.

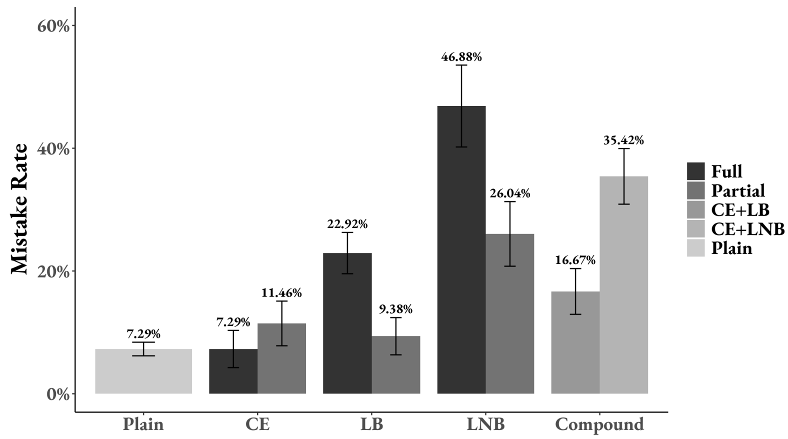

(a) All rounds

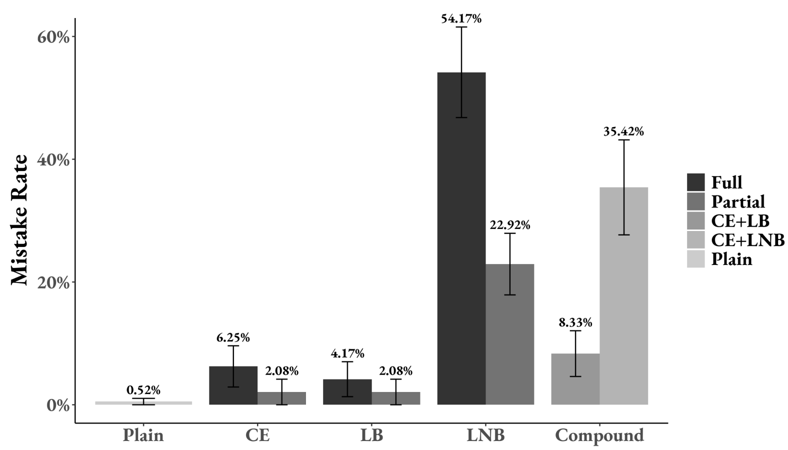

(b) Last 10 rounds

Figure 8: Mistake rates for Hard Firms by subtype

Subtype differences for Hard Firms again closely parallel the market level subtype differences discussed in Section 4.1.2. Across all 20 rounds, mistake rates for Hard Firms do not differ between the Full and Partial subtypes within CE ($p = 0.401$), but are significantly higher in the Full subtype than in the Partial subtype in LB and LNB ($p = 0.003$ and $p = 0.024$, respectively). Within Compound, Hard Firms make significantly more mistakes in CE+LNB than in CE+LB ($p = 0.001$). Results for the last 10 rounds are similar, with one exception. There is no longer a significant difference between Full and Partial within LB ($p = 0.564$).

Table 5: Mistake rates for Hard Firms: pairwise tests across subtypes

(a) All 20 rounds

	CE		LB		LNB		Compound
	Full	Partial	Full	Partial	Full	Partial	
Within Type	0.401		0.003		0.024		0.001
CE+LB	0.081	0.110	0.220	0.107			
CE+LNB	0.001	< 0.001			0.054	0.150	

(b) Last 10 rounds

	CE		LB		LNB		Compound
	Full	Partial	Full	Partial	Full	Partial	
Within Type	0.317		0.564		0.005		0.009
CE+LB	0.706	0.180	0.317	0.180			
CE+LNB	0.006	0.002			0.061	0.342	

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

Comparing the two Compound subtypes to their non-compound counterparts over all 20 rounds, we find the following. First, the mistake rate for Hard Firms in CE+LB does not differ significantly from either subtype of CE or from either subtype of LB. Second, Hard Firms in CE+LNB make significantly more mistakes than in CE-Full and CE-Partial (both $p < 0.001$), marginally fewer mistakes than in LNB-Full ($p = 0.054$), and do not differ significantly from LNB-Partial. The last 10 rounds exhibit the same pattern.

Result 4 (Hard-Firm mistake rates by subtypes)**.** *Mistake rates for Hard Firms in LB and LNB are significantly higher in the Full subtype than in the Partial subtype, although the difference in LB is no longer significant in the last 10 rounds. In both the full sample and the last 10 rounds, mistake rates for Hard Firms are significantly higher in CE+LNB than in CE+LB. Compared to the non-compound counterparts, including LNB in Compound increases the mistake rates.*

4.2.3 Mistake Types

In this section, we study *how* individuals make mistakes. Whenever a player fails to reach a predicted match 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 predicted outcome.

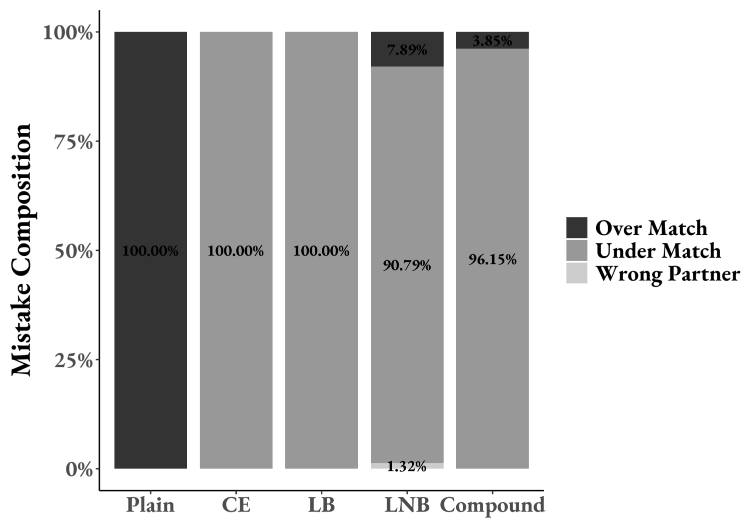

Figure 9: Composition of individual mistake types by learning pattern

Individual mistakes vary sharply across learning patterns. Figure 9 displays the composition of mistake types by learning pattern. In Plain, all mistakes are Overmatch; in CE and LB, all are Undermatch. Wrong-Partner mistakes appear only in LNB (1.32%). In both LNB and Compound, most mistakes are Undermatch (90.79% and 96.15%), with the remainder being Overmatch.

The chi-square tests in Table 6 confirm that the distribution of mistake types differs across treatments. Relative to Plain, all other treatments exhibit statistically different mistake-type distributions ($p < 0.001$). CE and LB exhibit identical distributions ($p = 1.000$). CE and LNB do not differ from each other ($p = 0.167$). By contrast, LNB differs significantly from LB ($p = 0.035$) but not from Compound ($p = 0.204$).

Table 6: Mistake Type Distribution: Chi-square Tests

	CE	LB	LNB	Compound
Plain	< 0.001	< 0.001	< 0.001	< 0.001
CE		1.000	0.167	0.233
LB			0.035	0.102
LNB				0.204

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

The timing of mistakes also varies systematically across learning patterns. Figure 10 plots cumulative mistake rates by round. In Plain, nearly all mistakes occur in Round 1 (96.43%), with very few thereafter. In CE and LB, mistakes are concentrated in the first 10 rounds (77.78% and 91.18%, respectively). In contrast, LNB and Compound exhibit mistakes throughout the experiment: in LNB, exactly 50% occur in the first 10 rounds, and in Compound, 57.69% occur in the first 10 rounds. These patterns indicate rapid learning over time in Plain, CE, and LB, but not in LNB and Compound.

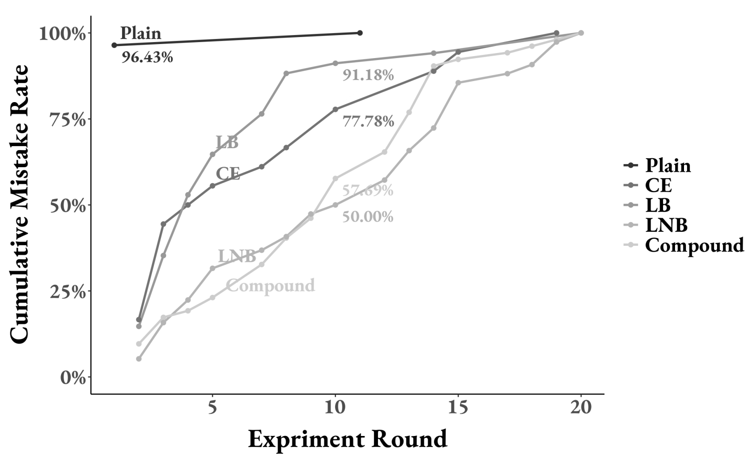

Figure 10: Cumulative mistake rates by round

Result 5 (individual mistake types)**.** *Individual mistake types and their timing differ sharply across treatments. In Plain, all mistakes are Overmatch, occurring mostly in Round 1. In CE and LB, all mistakes are Undermatch and are concentrated in the first 10 rounds. In LNB and*

Compound, over 90% mistakes are Undermatch, and mistakes are spread more evenly across rounds.

4.3 Matching Dynamics

In this section, we investigate the matching dynamics, that is, decentralized market activities. Studying these dynamics allows us to assess more precisely the difficulty imposed by different learning patterns.

Across learning environments, each matching market can generate up to two correct matches. Each match can be classified as either an *easy match*, where the firm does not need to infer the state to match correctly, or a *hard match*, where the firm must update its belief about the state in order to match correctly. In Plain, all correct matches are easy. In LNB and CE+LNB, each market contains exactly one match, and it is hard. In CE, LB, and CE+LB, both easy and hard matches may occur.

We focus exclusively on correct matches and examine two aspects of matching dynamics. First, we study the time required to finalize a correct match. Second, we examine the order in which correct matches occur within a market.

4.3.1 Time required to form correct matches

We begin by analyzing the time required to achieve easy and hard matches. Figure 11 plots the average matching time for each treatment, conditional on reaching a correct match. A clear pattern emerges. Easy matches occur quickly and at similar speeds across treatments. In contrast, hard matches take much longer, with average matching time rising from CE (17.88 seconds) to LB (22.71 seconds), Compound (31.77 seconds), and peaking in LNB (37.37 seconds). In treatments with both easy and hard matches (CE, LB, and Compound), hard matches also take significantly longer than easy matches (two-sided Wilcoxon signed-rank tests, $p < 0.001$). These results indicate that greater learning difficulty slows decentralized matching.

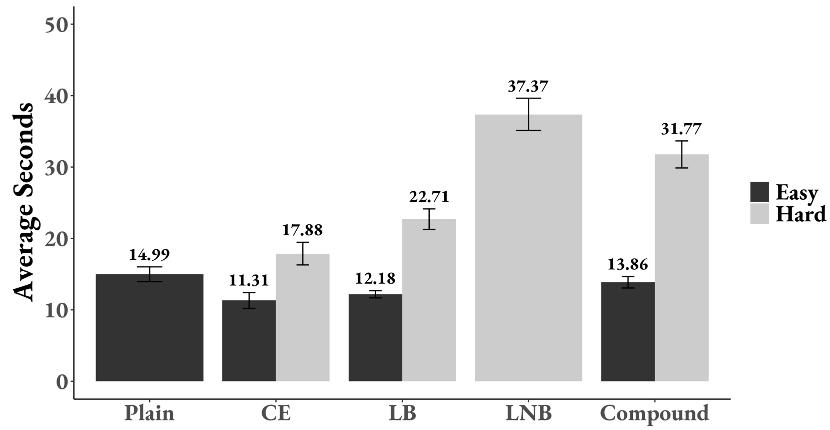

Figure 11: Time required to form correct matches by learning pattern and by difficulty

Result 6 (matching dynamics)**.** *Hard matches take substantially longer than easy matches. The time required to form hard matches is ranked as*

$$\text{LNB} > \text{Compound} > \text{LB} > \text{CE}.$$

4.3.2 Order of correct matches

In CE, LB, and CE+LB, both easy and hard matches can occur. Moreover, in LB and CE+LB, theory predicts that easy matches should precede hard matches, as hard matches rely on information revealed by easy matches. We now examine whether this holds in the experiment.

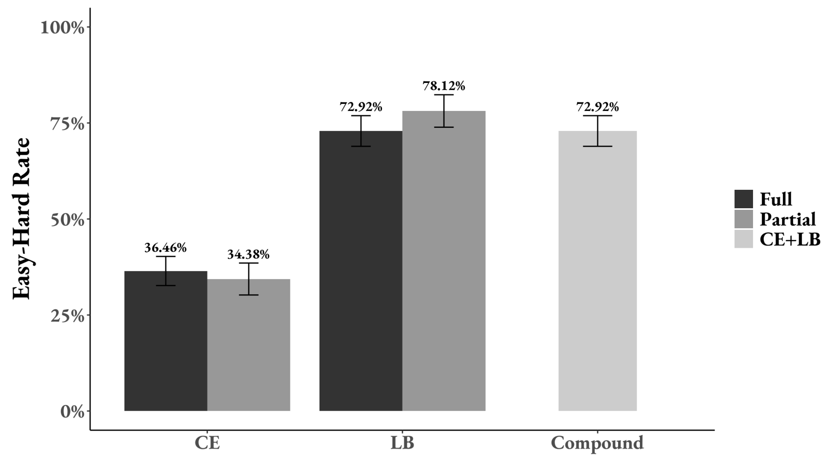

Figure 12: Order of correct matches by learning pattern (easy-hard rate)

Figure 12 reports the average easy–hard rate (equal to 1 if the easy match forms strictly before the hard match) by learning pattern and subtypes. The easy–hard rate is relatively low in CE (36.46% in Full and 34.38% in Partial), but substantially higher in LB (72.92% in Full and 78.12% in Partial) and in CE+LB (72.92%). The differences between CE and LB, and between CE and CE+LB, are statistically significant (both $p < 0.001$, two-sided Mann–Whitney tests). These results indicate that in LB and CE+LB, hard matches are considerably more difficult to form, making the hard match more likely to occur after the easy match.

Result 7 (order of correct matches)**.** *In LB and CE+LB, easy correct matches precede hard correct matches significantly more often than in CE.*

4.4 Explanations of key findings

In this section, we discuss possible explanations for our two key findings with respect to the two treatment variables: learning patterns and subtypes.

First, Results 1 and 3 show that LNB is more difficult than the other learning patterns, and that this difficulty does not diminish with experience. One possible explanation is that making a correct matching decision in LNB requires the firm to go through three steps: (i) identifying which worker-firm pair can potentially match in some states, (ii) waiting sufficiently long to verify the absence of the match, and (iii) correctly updating beliefs about the state based on observed non-matching outcomes. Failure at any of these steps may lead to an incorrect match. As a result, correct matching in LNB is substantially more challenging than in CE and LB, where firms only need to perform a single state-inference step based on market activities (step (iii)). Results 5 and 6 support this explanation in two complementary ways: first, failures at any step are likely to generate Undermatch outcomes (Result 5); while the waiting requirement in step (ii) leads to a longer time to match in LNB (Result 6).

Second, Results 2 and 4 indicate that the Partial subtype is easier than the Full subtype in both LB and LNB. Recall that the benchmark payoff (the worst case payoff from staying unmatched) is zero. One potential explanation stems from *loss aversion*, which may be particularly pronounced in our design, where small positive payoffs coexist with relatively large negative ones. When a firm makes an effort to draw inferences, they may focus more on avoiding negative payoffs rather than on confirming positive ones, even though these considerations are theoretically equivalent. For example, it may be more difficult to rule out the two negative outcomes (-4 and -8), as in the potential payoff vector $(2, -4, -8)$ under Full, than to rule out a single negative outcome (-8), as in $(2, 4, -8)$ under Partial. A related but distinct explanation focuses on the state complexity rather than on payoff magnitudes. Given a uniform prior over the three states, $(1/3, 1/3, 1/3)$, a match is more likely to be correct (i.e., to generate positive payoffs) in the Partial subtype $(2, 4, -8)$ than in the Full subtype $(2, -4, -8)$. By effectively reducing the number of adverse states that must be ruled out, the Partial subtype simplifies the inference problem faced by firms, thereby increasing the likelihood of correct matching.

5 Discussions

5.1 Alternative stability notions

In matching theory with one-sided incomplete information, there have been a few stability notions since Liu et al. (2014), including those of Chen and Hu (2020), Liu (2020), Pomatto (2022), Wang (2023), and Chen and Hu (2024). They are all closely connected. Our theoretical prediction belongs to this class via Chen and Hu (2020, 2024).

Liu et al. (2014) formalized the first cooperative stability notion in a prior-free model. More precisely, when a firm faces uncertainty over the type profile of workers and evaluates a potential employee, she is willing to form a new match with that potential employee only if the match yields a higher ex post payoff than the firm's status quo match under *all* possible (reasonable) type profiles. Liu et al. (2014) start with the set of all individual rational matchings and iteratively exclude matchings that are clearly blocked in the above sense. They define the limit of this procedure—*iterative elimination of clearly blocked matchings*—as the set of incomplete information stable matchings. This approach resembles the idea of rationalizability (Osborne and Rubinstein, 1994, Chapter 4) and circumvent the analysis of firm-specific beliefs. Pomatto (2022) adopts an epistemic approach and proposes a stability notion that is equivalent to Liu et al.'s notion.

Chen and Hu (2020) extends Liu et al. (2014) to incorporate flexible firm knowledge, represented by information partitions. Unlike Liu et al. (2014) which concerns a stable set of matchings, they focus on whether an individual matching is stable with the companion information structure being stable as well. Nevertheless, their stability notion is also equivalent to that of Liu et al. (2014) if we restrict attention to the set of stable matchings. This equivalence is established in the prior-free setting.

Introducing a prior belief into Chen and Hu (2020) in a straightforward way yields the model of Chen and Hu (2024). Their Bayesian stability notion turns out to be equivalent to that of Liu (2020). Particularly, Liu (2020) defines stability of matching functions, each of which specifies a matching (observable) for each type profile of workers (unobservable), which shares the spirit of rational expectation equilibrium (Radner, 1979). In contrast, Chen and Hu (2024) focus on the Bayesian stability of individual matchings (together with the companion information structure). They show that Liu's stable matching function can be translated into a series of their stable matchings; conversely, any stable matching in their sense can be induced from some stable matching function of Liu (2020). Wang (2023) establishes a similar equivalence in a more general model.

Our theoretical prediction in Section 2.2 is clearly akin to Chen and Hu (2020, 2024). The only major difference is that we assumed away transfers. Since our focus is on the stability criterion and the learning patterns, instead of matching efficiency, a nontransferable utility setting is more suitable for experiment design (avoiding the screening role played by transfers).

5.2 Alternative approaches to decentralized matching

Other than the decentralized matching market in our experimental design, there are at least two alternative approaches to decentralized matching: the "path-to-stability" approach following Knuth (1976) and Roth and Vande Vate (1990), and the non-cooperative models of Haeringer and Wooders (2011), Pais (2008), Suh and Wen (2008), Ferdowsian, Niederle and Yariv (2025), Lauermann and Nöldeke (2014) and Ferdowsian (2023). Our experimental design balances the advantages of these approaches.

To be precise, consider a complete-information matching market where blocking pairs are sequentially satisfied, without specifying the details of who proposes to whom and so on. Knuth (1976) demonstrates that this dynamic matching process may generate indefinite cycles, whereas Roth and Vande Vate (1990) shows that there is always at least one blocking path that leads to a stable matching. An immediate corollary is that the process which randomly satisfies blocking pairs (each with positive probability if there are many) will converge to a stable matching with probability one.

When there is incomplete information, Chen and Hu (2020, 2024) argue that information will be updated along a matching process. They replicate the convergence result of Roth and Vande Vate (1990) by showing that there is always a "learning-and-blocking" path that leads to an incomplete-information stable matching. This line of research maintains the cooperative feature of blocking pairs and circumvents the strategic behavior of agents. A specific assumption is that agents are myopic—they opt to form a new matching as long as a blocking opportunity exists.

In contrast, agents' strategies are explicitly modeled in Haeringer and Wooders (2011); Pais (2008); Suh and Wen (2008); Ferdowsian, Niederle and Yariv (2025) (non-cooperative decentralized matching), Lauermann and Nöldeke (2014) (search with random meetings) and Ferdowsian (2023) (multi-armed bandit problem). However, as Liu (2023, 2024) has argued, cooperative models should have been more attractive as analytical tools for complex economic applications involving incomplete information. One advantage is that it makes little ad hoc (and thus diverse) assumptions about agents' behavior.

The decentralized matching markets in our experimental design make as little assumption as possible while uniquely identifying all learning patterns. First, we construct the matching environment so that each agent has only one potential partner. Choices therefore reduce to whether to form a match or remain unmatched, rather than targeting which partner. This feature rules out the complicated strategic behavior studied in the aforementioned papers. Second, each matching market in our experiment has a unique incomplete-information stable matching and there is a unique "learning-and-blocking" path from the autarky (no matches and no information) to this stable matching. As a result, our design is robust to the specific details of how matches are formed and how blocking pairs are selected.

5.3 Other decentralized implementations of stability notions

Most experimental studies of stability in matching markets test stability notions derived under the assumption of complete information; see these notions in Roth and Sotomayor (1992). While some experiments introduce incomplete information, it typically enters as an information friction layered onto theories that assume complete information.

A first strand of the literature studies decentralized matching with transferable utility. Under complete information, Otto and Bolle (2011), Dolgopolov et al. (2024), and He et al. (2024) examine how agents share matching surplus in decentralized environments. Otto and Bolle (2011) study repeated price bargaining in a two-by-two market and find frequent inefficiency and instability. Dolgopolov et al. (2024) analyze convergence to the core under different trading institutions, namely double auction, posted prices, and decentralized communication, in a three-by-three assignment market, finding strong support for core allocations but weaker evidence for core prices. He et al. (2024) investigate how assortativity and "whether equal split is in the core" affect stability and efficiency, showing that both—especially the latter—significantly influence market performance.

Two papers study transferable-utility markets under incomplete information. Nalbantian and Schotter (1995) examine decentralized matching where pairs negotiate salaries via phone without knowing their match-specific surpluses. Agranov et al. (2025) compare complete- and incomplete-information environments and show that incomplete information reduces both efficiency and stability. In their incomplete-information treatments, agents know their own surplus profile but not which partners generate which surpluses, and they learn match values only after an offer is accepted. Relative to our study, these papers focus primarily on stability concepts and surplus allocation under complete information, rather than on testing stability concepts derived for incomplete-information environments.

A second strand examines decentralized matching without transfers, closer to our setting. Pais, Pintér and Veszteg (2020) study both complete- and incomplete-information environments, where under incomplete information agents know only their own preferences, and find no significant effect of incomplete information on stability or efficiency. Echenique, Robinson-Cortés and Yariv (2025) focus on complete-information markets and study equilibrium selection among multiple stable matchings. Although these papers consider nontransferable-utility markets, they continue to test stability notions based on complete-information theory, with incomplete information introduced only as an experimental friction. In contrast, our experiment studies nontransferable-utility matching markets in which incomplete information arises from uncertainty about states that determine agents' preferences, and we explicitly test stability concepts developed for matching under incomplete information.

5.4 Future directions

In existing decentralized implementations of stable matchings, experiments have mostly adopted either free negotiation or communication (common in matching markets with transfers), or free proposals, acceptances, and rejections with temporary matches that evolve until the market unravels. While stability notions are well defined in theory, these decentralized processes are largely outside the formal theoretical frameworks. A key limitation of such decentralized approaches is the limited control over agents' interim strategies prior to the formation of final matches. A natural direction for future research is therefore to study structured dynamic games, such as those proposed by Pais (2008), Suh and Wen (2008), and Haeringer and Wooders (2011), and possibly incorporate some form of incomplete information as in Pomatto (2022) or Ferdowsian, Niederle and Yariv (2025), in order to examine whether agents follow the predicted strategies and achieve equilibrium outcomes in these formal settings.

References

Agranov, Marina, Ahrash Dianat, Larry Samuelson, and Leeat Yariv. 2025. "Paying to Match: Decentralized Markets with Information Frictions." *forthcoming at 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. 2025. "Efficiency of Stable Matchings with Two-Sided Incomplete Information." *Available at SSRN: https://ssrn.com/abstract=3496454*.

Chiappori, P.A., B. Salanie, and Y.K. Che. 2024. *Handbook of the Economics of Matching. Handbooks in Economics*, Elsevier Science.

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.

Ehlers, Lars, and Jordi Massó. 2007. "Incomplete information and singleton cores in matching markets." *Journal of Economic Theory*, 136(1): 587–600.

Ehlers, Lars, and Jordi Massó. 2015. "Matching markets under (in) complete information." *Journal of Economic Theory*, 157: 295–314.

Ferdowsian, Andrew. 2023. "Learning through transient matching in congested markets." *Working Paper.*

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

Fernandez, Marcelo Ariel, Kirill Rudov, and Leeat Yariv. 2022. "Centralized matching with incomplete information." *American Economic Review: Insights*, 4(1): 18–33.

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

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. 2025. "The Structure of Bayesian Stable Matchings." *Available at SSRN 4910284.*

Immorlica, Nicole, Jacob Leshno, Irene Lo, and Brendan Lucier. 2020. "Information acquisition in matching markets: The role of price discovery." *Available at SSRN 3705049.*

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. 2023. "Cooperative Analysis of Incomplete Information." *Working Paper.*

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

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.

Osborne, Martin J, and Ariel Rubinstein. 1994. *A course in game theory.* MIT press.

Otto, Philipp E., and Friedel Bolle. 2011. "Matching Markets with Price Bargaining." *Experimental Economics*, 14: 3222–348.

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." *Available at SSRN 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.

Radner, Roy. 1979. "Rational expectations equilibrium: Generic existence and the information revealed by prices." *Econometrica: Journal of the Econometric Society*, 655–678.

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

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

Roth, Alvin E, and Marilda A Oliveira Sotomayor. 1990. *Two-sided matching: A study in game-theoretic modeling and analysis.* Cambridge University Press.

Roth, Alvin E, and Marilda Sotomayor. 1992. "Two-sided matching." *Handbook of game theory with economic applications*, 1: 485–541.

Rudov, Kirill. 2024. "Fragile stable matchings." *arXiv preprint arXiv:2403.12183.*

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.*

Yenmez, M. Bumin. 2013. "Incentive-compatible matching mechanisms: consistency with various stability notions." *American Economic Journal: Microeconomics*, 5(4): 120–141.

Appendix A Parameters

	j_1		j_2	
	7	5	6	-3
i_1	4	3	-9	1
	2	8	-2	-8
	-1	-5	-8	1
i_2	6	-9	5	-4
	-3	4	2	-7

(a) Plain (one pair)

	j_1		j_2	
	3	6	5	-4
i_1	5	4	-1	3
	9	8	-3	-9
	-7	1	3	9
i_2	2	-6	6	5
	-4	-5	8	2

(b) Plain (two pairs)

	j_1		j_2	
	-6	-3	7	5
i_1	-9	1	-4	-3
	2	-8	-2	-8
	5	7	-1	5
i_2	8	4	2	-9
	6	1	-3	-4

(c) CE-Full

	j_1		j_2	
	-7	-4	8	6
i_1	-1	2	5	4
	-3	-9	-3	-9
	3	2	-7	1
i_2	6	5	2	-6
	9	8	-4	5

(d) CE-Partial

	j_1		j_2	
	2	5	-6	-3
i_1	-4	3	-9	1
	-7	8	2	-8
	-1	5	5	1
i_2	2	-9	8	-4
	-3	-4	6	-7

(e) LB-Full

	j_1		j_2	
	8	6	-7	-4
i_1	5	4	-1	2
	-3	9	3	-9
	-7	1	3	5
i_2	2	-6	6	2
	-4	5	9	-8

(f) LB-Partial

	j_1		j_2	
	-2	5	-6	-3
i_1	4	3	-9	1
	7	8	2	-8
	-1	5	5	1
i_2	2	-9	8	-4
	-3	-4	6	-7

(g) LNB-Full

	j_1		j_2	
	-3	6	-7	-4
i_1	-5	4	-1	2
	8	9	3	-9
	-7	1	3	5
i_2	2	-6	6	2
	-4	5	9	-8

(h) LNB-Partial

	j_1		j_2	
	2	5	-6	-3
i_1	4	3	-9	1
	-7	8	2	-8
	-1	5	6	1
i_2	2	-9	-5	-4
	-3	-4	8	-7

(i) CE+LB

	j_1		j_2	
	-3	6	-7	-4
i_1	-5	4	-1	2
	8	9	3	-9
	-7	1	3	2
i_2	2	-6	-6	-5
	-4	5	9	-8

(j) CE+LNB

Figure A.1: Parameter 1.

	j_1		j_2	
	3	8	4	-3
i_1	7	1	-7	5
	4	2	-6	-9
	-2	1	5	-2
i_2	6	-5	9	-1
	-8	4	7	-6

(a) Plain (one pair)

	j_1		j_2	
	5	-4	6	-5
i_1	-9	3	9	-7
	-6	1	-8	2
	-8	3	7	-4
i_2	-4	-7	2	-3
	-1	6	9	-8

(b) Plain (no pair)

	j_1		j_2	
	-4	-3	3	4
i_1	-7	-5	-7	-1
	-6	9	-2	-8
	-5	-2	-2	1
i_2	9	-1	-6	-5
	7	-6	-8	4

(c) CE-Full

	j_1		j_2	
	-6	-5	5	3
i_1	-9	7	9	1
	-8	2	-6	-4
	-7	4	-8	-3
i_2	-2	3	-4	7
	9	-8	-1	6

(d) CE-Partial

	j_1		j_2	
	3	8	-4	3
i_1	-7	1	-7	5
	-4	2	-6	-9
	-2	1	5	2
i_2	6	-5	9	-1
	-8	-4	7	-6

(e) LB-Full

	j_1		j_2	
	5	1	-6	5
i_1	9	3	-9	-7
	-6	4	-8	2
	-8	-3	2	4
i_2	4	-7	7	3
	-6	1	9	-8

(f) LB-Partial

	j_1		j_2	
	-3	2	-4	-3
i_1	7	1	-7	5
	4	8	-6	-9
	-2	1	5	6
i_2	6	-5	2	-1
	-8	-4	7	-9

(g) LNB-Full

	j_1		j_2	
	-5	4	-6	5
i_1	-6	1	-9	-7
	9	3	-8	-2
	-8	-3	2	4
i_2	1	-7	7	3
	4	-6	9	-8

(h) LNB-Partial

	j_1		j_2	
	3	8	-4	3
i_1	4	1	-7	5
	-7	2	-6	-9
	2	-1	7	2
i_2	6	-5	-9	-1
	-8	-4	5	-6

(i) CE+LB

	j_1		j_2	
	-5	4	-6	5
i_1	-6	3	-9	-7
	9	1	-8	2
	-8	-3	9	4
i_2	4	-7	-2	-3
	-1	6	7	-8

(j) CE+LNB

Figure A.2: Parameter 2.

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

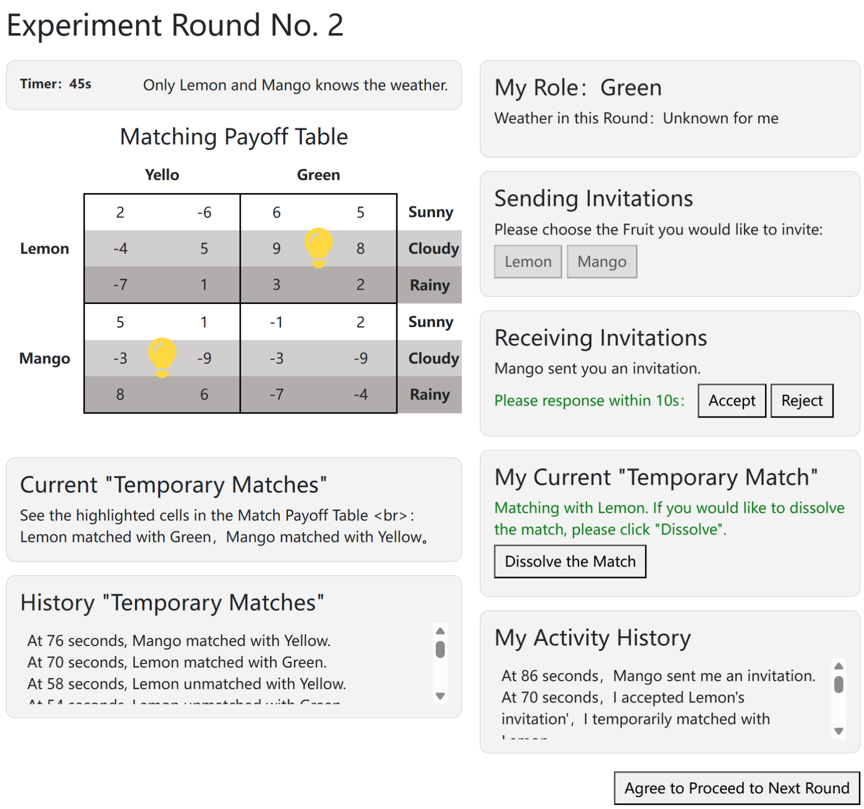

Figure B.2: A translated screenshot of the experiment

Appendix C Supplemental figures and tables

In this appendix, we provide supplemental figures and tables, which are useful for understanding the experimental results.

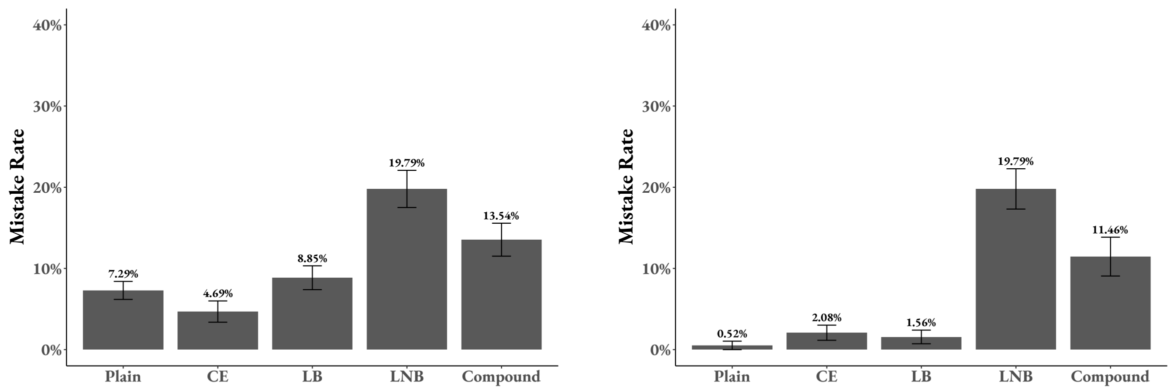

(a) All 20 rounds

(b) Last 10 rounds

Figure C.1: Individual mistake rate by learning patterns

Table C.1: Individual mistake rates: pairwise tests across learning patterns

	CE	LB	LNB	Compound
Plain	0.044	0.657	< 0.001	0.001
CE		0.010	< 0.001	0.001
LB			0.002	0.021
LNB				0.017

(a) All 20 rounds

	CE	LB	LNB	Compound
Plain	0.083	0.317	< 0.001	0.002
CE		0.655	< 0.001	0.005
LB			< 0.001	0.001
LNB				0.022

(b) Last 10 rounds

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

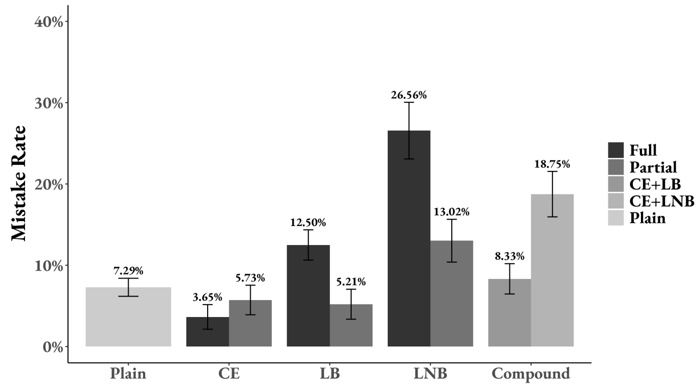

(a) All rounds

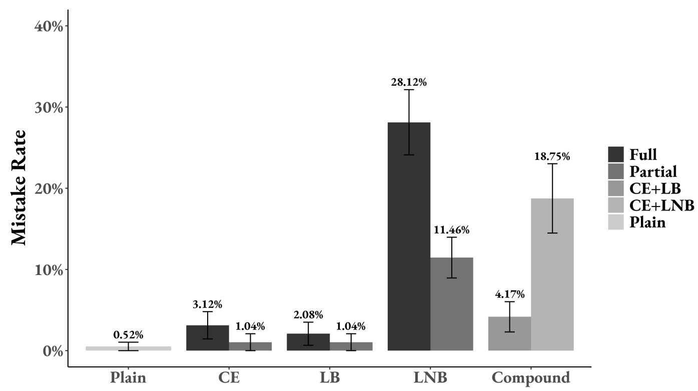

(b) Last 10 rounds

Figure C.2: Individual mistake rate by subtypes

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

	CE		LB		LNB		Compound
	Full	Partial	Full	Partial	Full	Partial	
Within Type	0.401		0.006		0.012		0.001
CE+LB	0.081	0.110	0.095	0.127			
CE+LNB	< 0.001	< 0.001			0.019	0.090	

(a) All 20 rounds

	CE		LB		LNB		Compound
	Full	Partial	Full	Partial	Full	Partial	
Within Type	0.317		0.564		0.006		0.009
CE+LB	0.706	0.180	0.317	0.180			
CE+LNB	0.006	0.002			0.045	0.302	

(b) Last 10 rounds

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

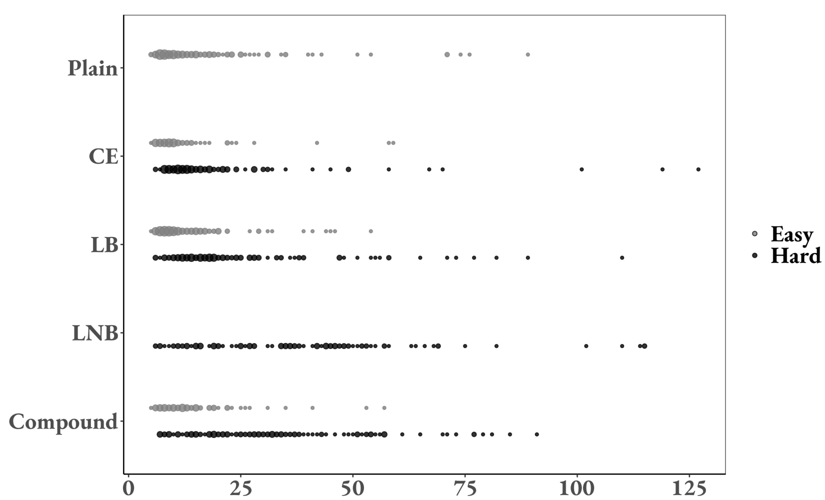

Figure C.3: Density of Final Correct Match Time

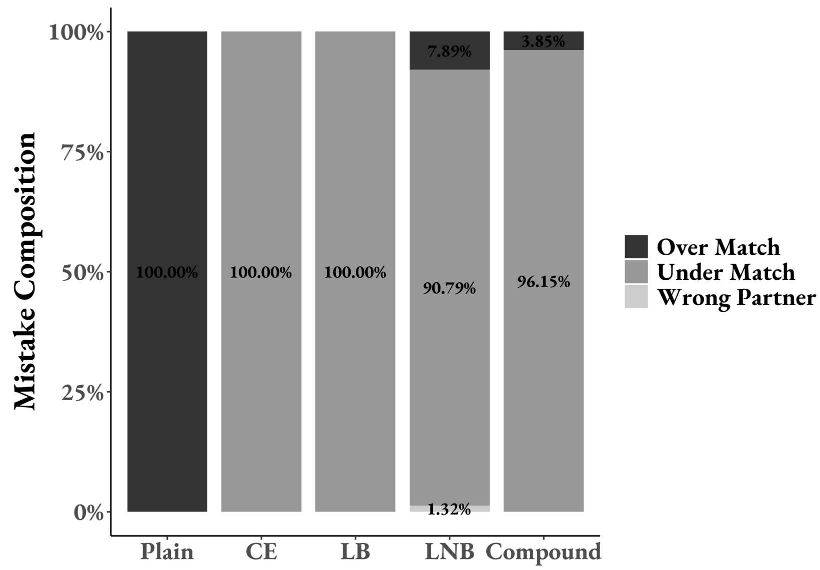

Figure C.4: Hard-firm mistake type composition by learning pattern

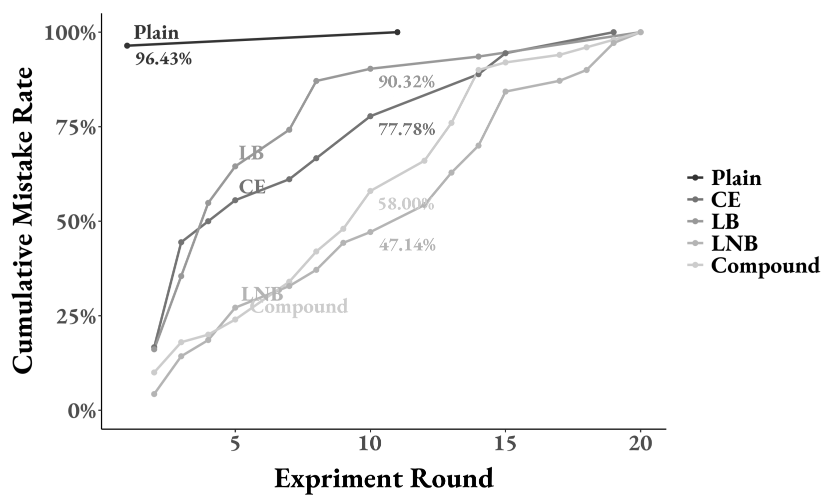

Figure C.5: Hard-firm cumulative mistake rate by round and by learning pattern