CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

Discussion Paper No. 2025-03

Lawrence Choo,
Todd R. Kaplan and
Xiaoyu Zhou

December 2025

Intermediaries and Information Aggregation in Fragmented Markets: Experimental Evidence

CeDEx Discussion Paper Series
ISSN 1749 - 3293

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit http://www.nottingham.ac.uk/cedex for more information about the Centre or contact

Samantha Stapleford-Allen
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 74 86214
Samantha.Stapleford-Allen@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx

Intermediaries and Information Aggregation in Fragmented Markets: Experimental Evidence*

Lawrence Choo† Todd R. Kaplan‡ Xiaoyu Zhou§

December 7, 2025

Abstract

Markets have the ability to aggregate information even under difficult conditions. What is not known is whether fragmented markets can also achieve information aggregation when differently-informed traders do not participate in the same market. We experimentally test this by extending the Arrow-Debreu version of the Plott and Sunder (1988) framework to two separate markets, adding uninformed market intermediaries who can trade in both markets. We find that indeed markets are able to aggregate information, but more experience is necessary when there is less competition among the intermediaries. The rate of convergence does not differ systematically across treatments once traders gain sufficient experience. Convergence arises through intermediaries' cross-market trading in competitive settings, whereas in less competitive settings it relies more on non-price signals. However, despite price convergence, intermediaries earn profits similar to those of informed traders in competitive markets and higher profits than informed traders in less competitive markets.

Keywords: *Information aggregation, Market fragmentation, Prediction markets, Financial market, Experiments.*

JEL code: C90, D80, G14, G40

*We are grateful to Jürgen Huber, Nobi Hanaki, Jörg Oechssler, Peiran Jiao, Ben Greiner, Felix Fattinger, and Christoph Huber for helpful comments and discussions. We also thank seminar participants at Singapore Management University, WU Vienna, Duke-Kunshan University, Shanghai Jiao Tong University, Nanyang Technological University, and the University of Nuremberg, as well as participants at the Experimental Finance Conferences 2021 and 2022, the ESA World Meeting 2021, the RCEA Conference, and the ESA Meeting in Bologna for valuable feedback. We acknowledge the funding support of the University of Macau grant: SRG2024-00062-FSS.

†Department of Economics, University of Macau. Email: `lawrencechoo@um.edu.mo`.

‡Department of Economics, University of Exeter, and Department of Economics, University of Haifa. Email: `toddrkaplan@gmail.com`.

§Business School, University of Nottingham. Email: `Xiaoyu.Zhou@nottingham.ac.uk`.

1. INTRODUCTION

The ability of markets to aggregate dispersed information into prices is central to economic thinking (e.g., Hayek, 1945; Fama, 1970). Controlled laboratory experiments (e.g., Plott and Sunder, 1988; Forsythe and Lundholm, 1990; Choo et al., 2019, 2022) show that market prices can reveal the true state of nature even when participants hold diverse and imprecise private information about the true state, although the effectiveness of information aggregation may depend on the underlying market structure (see Anderson et al., 2020; Corgnet et al., 2023). However, this evidence largely comes from *centralised markets*, in which all informed traders interact in a single venue. Real-world markets are often more fragmented, raising the question of whether information can still aggregate when trade is dispersed across multiple, only imperfectly connected venues in which price transparency is limited.

Institutional factors are a natural driver of market fragmentation. For example, the coexistence of multiple exchanges with different listing rules, the separation between on-exchange and over-the-counter trading, and the use of dark pools or other alternative trading systems that restrict price transparency. However, segmentation can also arise from non-institutional forces. Behavioural biases such as familiarity bias (e.g., Fox and Tversky, 1995; Fox and Weber, 2002; Chew et al., 2008) and home-market bias (e.g., Lewis, 1999) may lead traders to focus on a narrow subset of assets that are salient or familiar to them, thereby overlooking potentially informative prices from other markets. In practice, professional traders frequently specialise in particular asset classes (e.g., fixed income, equities, or real estate), which can limit their ability to interpret or react to price signals outside their domain of expertise. Limited attention and cognitive constraints may reinforce such patterns, as monitoring and processing information across multiple venues is costly. As a result, information relevant to the true state may remain dispersed across markets even when no formal barriers to integration exist. This motivates our interest in whether, and through which mechanisms, fragmented markets can still reveal the true state.

Against this backdrop, we use a laboratory experiment to study information aggregation where markets are fragmented. Our setting, deliberately stark, considers two *separated markets* (left and right), each trading a series of Arrow-Debreu assets whose payoffs depend on the same underlying state of nature. Traders who are *native* to each market—henceforth, **side traders**—receive imprecise private information about this common state but can only trade and observe market activity within their own market.

The information structure is similar to that in Plott and Sunder (1988), where there are two groups of traders and each group member receives the same signal about the true state, but the two groups have different signals. However, in our setting, the groups of informed traders (which we call side traders) are in separated markets. Thus, unlike Plott and Sunder, there is no single market that contains all the private information required to identify the true state.

We allow the left and right markets to be "connected" by a group of uninformed *intermediaries*—the **middle traders**—who are permitted to trade in both markets. Intermediaries here are analogous to real-world cross-market participants—such as broker-dealers, arbitrageurs, and multi-venue liquidity providers—who, despite not possessing private information themselves, may facilitate the diffusion of information across markets through their cross-market activities (e.g., Menkveld, 2013; Makarov and Schoar, 2020). The experiment thus places traders in a setting where information is deliberately split across fragmented markets, enabling us to examine whether, and through which mechanisms, information aggregation can emerge in the absence of a fully centralised market. In doing so, we combine a classic Plott-Sunder information environment with a fragmented market structure and uninformed intermediaries, offering a new laboratory setting for studying information aggregation across separated markets.

An analogy for our design is to imagine two national markets (e.g., countries C and U) where traders possess private information about local demand for a new battery technology that can influence both countries. This local information is partially reflected in each country's domestic market prices. However, traders in country C do not know which country U's prices are informative for them, and country U's traders face the same difficulty with country C's prices. In this setting, intermediaries are akin to traders operating in a third country, country E, which although they do not know the underlying demand conditions in either countries C and U, they know the relevant prices to observe and can trade in both countries, enabling them to infer which price movements in one market are relevant for the other and thereby transmit information across the two otherwise disconnected venues.

In our design, intermediaries are the only channel through which information can move from one separated market to the other, since the informed side traders are restricted to trading and observing in their own markets. As such, the uninformed intermediaries should learn about the true state before the side traders.

Competition among intermediaries to exploit cross-market price discrepancies—and thereby make use of their information advantage—is a potential mechanism through which information about the true state may eventually be reflected in prices in both separated markets. We therefore vary the degree of competition among intermediaries in order to examine how the strength of this cross-market competitive force affects information aggregation.

To do so, we consider two treatments in which the intermediaries are, relative to the side traders, a *minority* (the M424 treatment) or a *majority* (the M464 treatment) within each separated market. We use the continuous double auction mechanism to facilitate trade in both treatments. When intermediaries are a minority relative to side traders (M424), each middle trader faces little competitive price pressure to reveal, through their trades, the information they have inferred from the other market—recall that side traders in each separated market cannot directly observe the opposing market. In contrast, when intermediaries form a majority (M464), competition among them is stronger, creating greater pressure to react quickly to price movements in the other market and thereby transmit information across markets more effectively. We formalise our main research question as follows:

> **Research Question:** *Can fragmented markets aggregate information, and how does the success of aggregation depend on the degree of competition among intermediaries?*

To study information aggregation, we benchmark market prices against the fully-revealing equilibrium (Radner, 1979), in which prices incorporate all privately held information. Using the myopic-reasoning model of Choo et al. (2019), we show that market prices in both M424 and M464 can converge towards this benchmark, although achieving full revelation is non-trivial in fragmented environments.

Our findings can be summarised as follows. Information aggregation across the separated markets becomes successful once traders gain some experience, and this holds in both treatments—whether intermediaries are a majority (M464) or a minority (M424) within each separated market. Furthermore, the pace of convergence towards the fully-revealing equilibrium does not differ significantly across the treatments.

To better understand the drivers of information aggregation, we analyse trader behaviour in each treatment. The evidence reveals distinct patterns across treatments. In the M464 treatment, where intermediaries form a majority, they play a prominent role in steering prices towards the fully-revealing equilibrium. By contrast, the influence of intermediaries is naturally more limited in the M424 treatment, where they are a minority. In M424, we find evidence that side traders use non-transaction price signals to form inferences about the true state. Finally, we find that the profitability of intermediaries depends on market structure: when they are numerous (M464), competition erodes their informational rents and yields payoffs similar to side traders, whereas when they are fewer (M424), limited competition allows them to earn substantially higher returns from their information advantage.

We contribute to several strands of the economic literature, with the most immediate being the experimental literature on information aggregation (e.g., Sunder, 1995; Plott, 2000; Deck and Porter, 2013; Galanis et al., 2024; Corgnet et al., 2021, 2023). This is the first laboratory study of information aggregation to extend the centralised-market paradigm to a fragmented multi-market setting. We show that Arrow-Debreu markets can successfully aggregate information across two separated venues even when no single venue observes all state-relevant information. Furthermore, we shed light into the mechanisms through which uninformed intermediaries facilitate the diffusion of information across markets. Finally, the experimental framework we introduce is flexible and can be extended in multiple directions to study information aggregation under alternative market structures; we outline several such extensions in the discussion section.

Beyond the information-aggregation experimental literature, our study also connects to the theoretical and empirical (e.g., Froot and Dabora, 1999; Foucault et al., 2023; O'Hara, 2015; Chen and Duffie, 2021; Babus and Parlatore, 2022) literature on market fragmentation and multi-venue trading. This literature typically focuses on *trading fragmentation* and examines how the dispersion of order flow across venues affects liquidity and price discovery when traders have full access to price information and can freely choose where to trade. Our experiment departs from this framework by allocating private information to venue-restricted traders—the side traders—and allowing information to flow across markets only through uninformed intermediaries. In doing so, we incorporate not only trading fragmentation but also *information fragmentation*, creating a setting in which state-relevant information is deliberately split across markets. In such an environment, intermediaries are the only traders capable of linking the markets, making their behaviour central to the diffusion of information across venues.

In our experiment, intermediaries play a critical role in facilitating the diffusion of information across markets: by attempting to arbitrage cross-market price discrepancies, they exploit their informational advantage from observing activity in both venues. This mechanism connects directly to the empirical literature documenting how cross-market arbitrageurs help integrate fragmented markets and complements the limited experimental evidence on arbitrage activity (Angerer et al., 2023; Neugebauer et al., 2023). By varying whether intermediaries form a majority (M464) or a minority (M424), our design provides, to our knowledge, the first suggestive evidence on how intermediary competition affects information transmission across markets. Somewhat surprisingly, the degree of intermediary competition has little effect on the pace of convergence. To some extent, the fact that intermediaries in M464 do not earn significantly higher profits than side traders despite their information advantage, whereas intermediaries in M424 do, suggests that competition among arbitrageurs is an important factor in shaping their ability to extract rents. This pattern is consistent with the broader limits-of-arbitrage (e.g., Shleifer and Vishny, 1997; Gromb and Vayanos, 2010) literature, which emphasises how constraints and competitive pressures can restrict arbitrageurs' capacity to exploit mispricing .

The between-treatment comparison provides some insights as to how traders adjust their behaviours depending on the degree of competition among intermediaries. More generally, the comparison contributes to the empirical (e.g., Biais et al., 1995) and experimental (e.g., Camerer and Weigelt, 1991; Sunder, 1992) discussions on the "informativeness" of non-transaction market signals. In the M424 setting, where intermediaries exert only limited influence on convergence, we find that side traders can extract meaningful information from non-transaction signals within their own venue. This observation also sheds light on static rational-expectation equilibrium models (e.g., Grossman, 1976; Radner, 1979; Hellwig, 1980) of information aggregation, which specify the equilibrium outcome but not the dynamic path toward it. The ability of traders to learn from non-transaction signals implies that prices need not transact at off-equilibrium levels for markets to move toward the equilibrium outcome.

For side traders in our experiment, the two-market setting is analogous to a lit-dark structure: while they can observe prices in their own ("lit") market, the other market is effectively "dark" to them, as they cannot observe its price path or trading activity. Much of the lit-dark literature examines informational efficiency by studying whether and how the presence of dark venues impairs or supports price discovery in lit markets. In this regard, our study is related to experimental work on lit-dark and fragmented trading venues. For example, Halim et al. (2025) examine how transparency (lit vs. dark) and the distribution of informed traders affect price and allocative efficiency when traders receive private signals.

Similarly, Hendershott et al. (2022) study how "pre-trade transparency" (i.e., the full order-book) influences trading behaviour and market quality in fragmented limit-order markets. Both studies, however, consider a setting where all traders observe lit prices, so that information asymmetry arises only across traders and not across venues. In contrast, our design introduces information fragmentation across markets: informed side traders are restricted to a single venue, cannot observe prices in the other market, and information can flow across venues only through uninformed intermediaries.

The over-the-counter (OTC) markets can be viewed as an extreme form of fragmentation, where trading is bilateral and information *percolates* slowly across a decentralised network (Duffie et al., 2009, 2014). Consistent with this, Asparouhova and Bossaerts (2018) show experimentally that bilateral OTC-style trading leads to slow and incomplete information aggregation. In this regard, our experiment can be seen as a midpoint between centralised and OTC markets: information is local and cross-market visibility is limited, as in OTC settings, but trading within each venue is organised and information flows across venues only through intermediaries.

The rest of the paper is organised as follows. Section 2 describes the experimental design. Section 3 outlines the theoretical framework. Section 4 presents the results, and Section 5 concludes.

2. EXPERIMENT DESIGN

The experiment involved the M464 (9 matching-groups, 126 student subjects) and M424 (9 matching-groups, 90 student subjects) treatments. The sessions were conducted between December 2019 and April 2021 at the Shanghai Jiaotong University and the experiment was programmed with zTree (Fischbacher, 2007).

Each session had two parts. Part I served as a five-round "training stage" and Part II, the main experiment, comprised 12 rounds (one M464 session ran only nine). Payoffs were recorded in "points."

Each session lasted roughly 120 minutes and used fixed matching groups. After the experiment, payoffs from one randomly chosen training-stage round and three randomly chosen main-stage rounds were converted to cash at 100 points = 9 yuan; paying for three rounds in Part II allows subjects to smooth earnings, following Charness et al. (2016). Including a 10-yuan show-up fee, average earnings were 128 yuan ($\approx$ US\$19.40) in the M464 treatment and 131 yuan ($\approx$ US\$20.60) in the M424 treatment.

Figure 1: **Illustrative example of market design and information for all traders.** The example assumes that the true state is Y. The arrow denote the markets that the traders can participate in. Assets are labelled true, fake, or neutral according to the true state and the set of states that side traders know to be impossible.

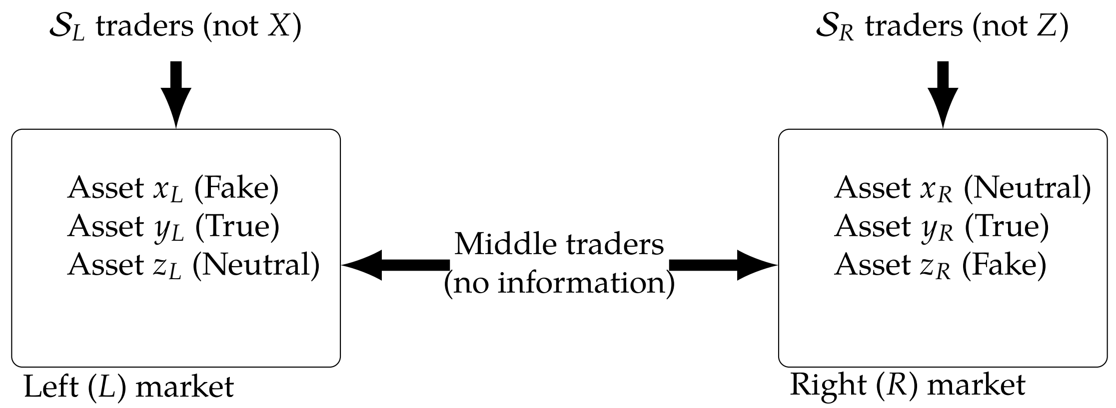

The following details the design of main experiment (part II) and thereafter discusses the training stage (part I).

2.1. DESIGN

At the start of the experiment, subjects are randomly assigned either the role of a *side trader* or a *middle trader*—role fixed over all rounds. The treatments vary the number of middle traders. Throughout the paper, we use the terms middle traders and intermediaries interchangeably. We now detail the proceedings of each round.

Information. At each round, the side traders will be randomly separated into two equal groups, the *left* (L) and *right* (R) groups, and nature selects one of three equiprobable states X, Y and Z. Denote a $j = L, R$ side trader by $\mathcal{S}_j$. Each side-trader group learns which one of the two non-realised states cannot have occurred. For instance, if nature chooses state Y (see Fig. 1), the left group $\mathcal{S}_L$ is told that X is impossible, while the right group $\mathcal{S}_R$ learns that Z is impossible—the groups each learn a different impossible state. Middle traders, by contrast, know only that the three states X, Y, and Z are equally likely.

Market. Trade takes place in two independent Arrow-Debreu markets—left (L) and right (R)—-each listing three state-contingent assets (x_L, y_L, z_L) or (x_R, y_R, z_R). Denote asset i_j where $i \in \{x, y, z\}$ and $j \in \{L, R\}$. Side traders $\mathcal{S}_j$ may trade and observe prices only in their own market $j \in \{L, R\}$; middle traders can trade and observe prices in both.

Markets operate as 180-second continuous double auctions with no short sales. Each side trader starts the market with 200 points in cash and six units of every asset in her market, while each middle trader starts with 200 points and three units of every asset in both markets.

At the end of trading, each asset pays a dividend of 20 points if its state is the true state, and zero otherwise. Thus, if the true state is Y, only assets y_L and y_R pay 20 points. A trader's payoff is

$$\begin{aligned} \text{Side trader } \mathcal{S}_j &: \quad L + \sum_i d(i_j)e^{i_j} \\ \text{Middle trader } \mathcal{M} &: \quad L + \sum_j \sum_i d(i_j)e^{i_j} \end{aligned}$$

where $L \geq 0$ is the trader's end-of-round cash balance, $e^{i_j} \geq 0$ is her inventory of asset i_j, and $d(i_j) \in \{20, 0\}$ is the dividend paid by that asset.

Belief elicitation. We elicit traders' post-market belief about the true state: a trader either selects a specific state (X, Y, or Z) or selects "I don't know" ($\emptyset$). The $\emptyset$ earns 5 points, a correct prediction earns 20 points, and an incorrect one incurs a 20 points penalty. Thus, a risk-neutral trader submits a prediction only if she believes her chance of being correct exceeds 62.5%.

2.2. TREATMENTS

Although the left and right markets operate independently—side traders in one never observe orders or prices in the other—they are connected in two crucial ways: (i) middle traders can trade in both markets, and (ii) all assets, though non-interchangeable across markets, are contingent on the same underlying true state. The information needed to identify that state is deliberately split between the markets—each holds only half of the puzzle—so full revelation is possible only if middle traders transmit signals across the two markets. The two treatment differ as follows:

- **M464** (9 matching-groups): $4\times\mathcal{S}_L$ traders, $4\times\mathcal{S}_R$ traders, and 6 middle traders: intermediaries are a *majority* within each market.
- **M424** (9 matching-groups): $4\times\mathcal{S}_L$ traders, $4\times\mathcal{S}_R$ traders, and 2 middle traders: intermediaries are a *minority* within each market.

We chose to vary the number of middle traders (4 vs. 6) and fix the number of side traders (4 in each market) to hold constant the "degree of information" across treatments. Had we held the total head-count constant and merely reassigned roles, faster convergence might instead reflect the larger supply of informed signals rather than the intermediaries themselves. That said, aggregation may still be harder in M424—where only ten traders participate, versus fourteen in M464—because the thinner order flow reduces liquidity and slows the diffusion of price signals.

2.3. TRAINING STAGE (PART I)

Part I, identical across treatments, serves only as a five-round training stage to familiarise participants with the information structure, trading rules, and payoff scheme (see also Choo et al., 2022). In each round, nature chooses the true state; one half of the traders learn that a particular state is impossible, the other half learn a different state is impossible, and then all participants trade the six Arrow-Debreu assets ($x_L, x_R, y_L, y_R, z_L, z_R$). This means that each trader can trade and observe prices of all six assets. The trading mechanism, dividend rules, and incentives mirror those used later in the main experiment. The purpose of the training stage is to familiarise subjects, especially the middle traders, with the trading platform and the information structure.

3. THEORY

For each market j, we label the *true* state as nature's choice, the *fake* state as the one the local side traders know is impossible, and the remaining as the *neutral* state. Because the side traders on the left and right receive different signals, their fake and neutral labels differ across markets. We adopt the same terminology when referring to the corresponding assets (see Figure 1).

To evaluate how well prices aggregate information, we benchmark them against two standard models: the Fully-Revealing Rational-Expectations equilibrium (FRE) (Radner, 1979) and the Prior-Information equilibrium (PIE) (Plott and Sunder, 1982, 1988; Choo et al., 2019, 2022). Both models impose the usual market-clearing condition but differ in how beliefs are formed: in PIE, traders' posteriors are fixed by their priors; in FRE, posteriors adjust endogenously to prices. Table 1 lists the corresponding equilibrium prices.

Table 1: **The Fully-Revealing Rational-Expectation equilibria (FRE) and Prior-Information equilibria (PIE) prices in the M424 and M464 treatments.** The example assumes that the true state is Y.

	Left market ($\mathcal{S}_L$ see not X)			Right market ($\mathcal{S}_R$ see not Z)		
asset type	True (y_L)	Fake (x_L)	Neutral (z_L)	True (y_R)	Fake (z_R)	Neutral (x_R)
FRE	20	0	0	20	0	0
PIE	10	20/3	10	10	20/3	10

PIE. Risk-neutral traders use Bayes' rule to update their posteriors from their private signals, set demands accordingly, yet do not infer anything from market prices. In the long run, therefore, prices in market j converge to the highest valuations implied by those fixed posteriors: side traders value the true and neutral assets at 10 points and the fake asset at 0, while middle traders value every asset at 20/3 points.

FRE. Traders not only update their posteriors with their private signals but also continually revise them in response to observed prices. This iterative learning process continues until prices convey all information held anywhere in the market—an outcome Radner (1979) shows is equivalent to traders openly pooling their signals. In our setting, the FRE therefore corresponds to the benchmark in which every trader effectively becomes fully informed of the true state by the end of trading.

The PIE and FRE differ only in what information prices reveal: under PIE, information about the true state is only partly aggregated, whereas under FRE it is fully aggregated.

3.1. CONVERGENCE TO FRE? A MYOPIC REASONING MODEL ANALYSIS.

Whether convergence to FRE is even feasible in our split-market setting is unclear. As Biais and Pouget (2000) note, with the information asymmetry used here, trade should occur only at FRE prices in equilibrium—but without trade at other prices traders cannot discover the true state in the first place. Equally important is the role of intermediaries: can uninformed middle traders drive prices toward FRE, and does their impact depend on whether they form a minority or majority within each market?

Dynamic models of information aggregation are typically complex and focus on a single-market (e.g., Dubey et al., 1987; Hellwig, 1982; Ostrovsky, 2012; Radner, 1979). To study our two-market setting, we adapt the myopic reasoning model (MRM) introduced by Choo et al. (2019, 2022). The MRM shows that prices can, in principle, reach FRE in both the M424 and M464 treatments, but also highlights how the difficulty of aggregation differs between them.

The model assumes that trading duration in each market j can be partitioned into $t = 1, 2, 3, 4$ hypothetical periods. Traders are risk-neutral, face no capital constraints, and behave myopically—that is, they act in every period as if it were the last. Each period unfolds in three steps:

> **Step 1** (Belief updating). Traders observe prices from period $t-1$ and revise their beliefs about the true state.
>
> **Stage 2** (Demand submission). They place orders based on these updated beliefs.
>
> **Stage 3** (Short-run clearing). The market clears at prices that equate current supply and demand, with each asset purchased by the traders who value it most.[1]

Panels A and B of Table 2 list the MRM-predicted prices for each period t in the M464 and M424 treatments. The example assumes the true state is Y: $\mathcal{S}_L$ know state X is impossible, $\mathcal{S}_R$ know state Z is impossible, and middle traders receive no information. Cell shading indicates who ultimately holds each asset at the period-t clearing price: blue for side traders, red for middle traders, and green when ownership is undetermined.

M424. In period 1 the middle traders value every asset at 20/3, while each side trader values the true and neutral assets at 10 and the fake at 0, so the clearing prices at each j market are 10-0-10 (true-fake-neutral); side traders hold the true and neutral assets, middle traders the fake.[2] By period 2 middle traders infer that Y is the true state and now value the true asset at 20 and the others at 0. Side traders' beliefs are unchanged, and the j market clearing prices remain 10-0-10 as middle traders are a minority in either market; the true asset is owned by middle traders who value them more, the fake is owned by side traders, and the neutral

[1]As in the M464 and M424 designs, every trader starts with a positive endowment of each asset they are allowed to trade.

[2]Recall that at the market-clearing price, all assets will be purchased by those traders that value them more.

Table 2: **Application of the Myopic Reasoning Model in the M464 and M424 treatments.** Each cell reports the model's predicted market-clearing price. The cells' colour report the predicted owners of the assets at the market clearing prices. The example assumes that the true state is Y.

	Left market ($\mathcal{S}_L$ see not X)			Right market ($\mathcal{S}_R$ see not Z)		
asset type	True (y^L)	Fake (x^L)	Neutral (z^L)	True (y^R)	Fake (z^R)	Neutral (x^R)
Panel A. The M424 treatment						
$t = 1$	10	0	10	10	0	10
$t = 2$	10	0	10	10	0	10
$t = 3$	20	0	10	20	0	10
$t = 4$	20	0	0	20	0	0
Panel B. The M464 treatment						
$t = 1$	20/3	20/3	20/3	20/3	20/3	20/3
$t = 2$	10	20/3	10	10	20/3	10
$t = 3$	20	0	10	20	0	10
$t = 4$	20	0	0	20	0	0

Blue cells: All assets are owned by the Side ($\mathcal{S}$) traders.
Red cells: All assets are owned by the Middle traders.
Green cells: Assets can be owned by the Middle traders or Side ($\mathcal{S}$) traders.

shows no clear ownership. In period 3 only middle traders trade the true asset, pushing its price to 20. The j market clearing price will be 20-0-10. By period 4 side traders, seeing the 20-point price, also deduce the true state, at the j market clearing price correspond to the FRE.[3]

M464. In contrast to the M424 case, the j market clearing price in period 1 will be 20/3-20/3-20/3 as the uninformed middle traders are more numerous than the side traders in each market: side traders hold the true and neutral assets, middles the fake. Because prices reveal no new information, the period 2 prices will be 10-20/3-10; only side traders trade the true and neutral assets. In period 3, only the middle traders can determine the true state from prices in both markets. The resulting j market clearing prices will be 20-0-10. In period 4, the side traders now learn the true state from prices, and the market-clearing price corresponds to the FRE.

The MRM assumes traders revise beliefs solely from observed prices. In practice, continuous-double-auction participants also infer information from order

[3]Intuitively, the side traders reason that an asset is 20 only because the middle traders have learnt that it is the true state.

flow and unexecuted bids and asks (e.g., Choo et al., 2019); actual price paths may therefore diverge from those in Table 2. Even so, the model shows that prices can reach the FRE in both treatments: middle traders learn the true state first, and competition among them propagates that information across markets. At the same time, the model underscores the hurdles of convergence to FRE:

- In the M424 treatment the sparse ranks of middle traders relative to side traders dampen competition among intermediaries, slowing price movement toward the FRE benchmark. Middle traders can exploit this by buying the true asset below its realised value and offloading the neutral asset above its payoff, capitalising on side traders' uncertainty about the true state.

- In M464, a larger pool of intermediaries intensifies competition: once they infer the true state, their trading pushes prices rapidly toward FRE. Yet when the intermediaries are unaware of the true state, the same competitive pressure can drown out side-trader signals and slow aggregation.

For these reasons, we proceed in three steps. First, we test whether prices in each treatment—M464 and M424—ultimately converge to the FRE benchmark despite information being split across markets. If both succeed, we then ask whether the speed of that convergence differs when intermediaries are a majority (M464) versus a minority (M424). Finally, we examine traders' behaviour in each treatment and whether intermediaries exploit their cross-market position to earn excess profits. While the MRM provides intuition about information aggregation in the M424 and M464 treatments, it is silent on which treatment should converge faster to the FRE benchmark and which factors drive that convergence. Addressing these questions forms the basis of our empirical strategy.

4. RESULTS

To analyse prices in the continuous double auction, we partition each 180-second trading round into six consecutive 30-second blocks, indexed by t. For each asset i_j we define:

- **Block-t market price**: the average transaction price within block t. The asset i_j block t market price is undefined if no trades occur within the block.

- **Closing market price**: the market price recorded in the last block in which the asset is traded. The closing prices account for 38% of all transactions in the M464 treatment and 43% in the M424 treatment.[4]

The block-t market price offers a snapshot of transactions within each 30-second interval, whereas the closing market price provides a terminal measure of valuation, reflecting the final point at which information is incorporated into trade. For example, if the asset never trades after the first 30 seconds (i.e., block 1), its closing price for that round is the block 1 market price.

Figure 2 plots the average block-t market prices for the true (top row), fake (centre), and neutral (bottom) assets across successive rounds (columns) in the M464 and M424 treatments—we aggregate data over the left and right markets. Trading volumes for each asset i_j average between one and two transactions per block. We observe that the true asset reliably trades at higher prices than the fake and neutral assets, clustering around its realised value, while the latter two assets seem to converge towards zero—we expand upon this in the next subsection.

To capture potential learning dynamics, we split the data into EARLY (rounds 1-6) and LATE (rounds 7-12) periods.[5] This split is particularly relevant for M424, where side traders may gradually become better at anticipating the behaviour of middle traders. Also, we classify a trader's belief as *correct* if his post-market prediction about the true state is accurate—otherwise classified as an *incorrect* belief.

4.1. INFORMATION AGGREGATION

Information aggregation, if it occurs, is expected to be a dynamic process in which traders continually update their beliefs based on observed market activity. Consequently, prices at the end of the round should lie closer to the FRE prices than to the PIE prices. Building on this, we compare the **mean absolute deviations** (MAD) of the **closing market prices** from the FRE (MAD^{FRE}) and PIE (MAD^{PIE}) benchmarks (e.g., Plott and Sunder, 1988; Choo et al., 2019, 2022; Corgnet et al., 2023).[6,7] Here, a smaller value indicates closer alignment with the

[4]In the M464 treatment, 17% of closing market prices are drawn from block prices in the first 60 seconds (blocks 1-2), while 58% are drawn from block prices in the final 60 seconds (blocks 5-6). The corresponding proportions in the M424 treatment are 21% and 52%, respectively.

[5]Group 1 (M464) contains only nine rounds and thus rounds 6-9 are treated as the LATE rounds.

[6]Plott and Sunder (1988) and Corgnet et al. (2023) measure deviations using the *last transaction price* of each asset relative to its equilibrium benchmark; adopting their method leaves our conclusions unchanged.

[7]For every group-round combination we compute the mean absolute deviation from a given benchmark $B \in \{\text{FRE}, \text{PIE}\}$:

$$MAD^B = \frac{\sum_j \sum_i |p^{i_j} - m^{i_j,B}|}{\text{num}},$$

Figure 2: **Mean market prices and transaction volumes by block and treatment.** Panels are arranged with assets (rows) and rounds (columns). The solid line and dashed lines detail the average block market price in the M464 and M424 treatments, respectively. Transaction activity is depicted with bars: hollow (M464) and blue (M424) bars. Prices are aggregated across the left and right markets; session-level series appear in the Appendix.

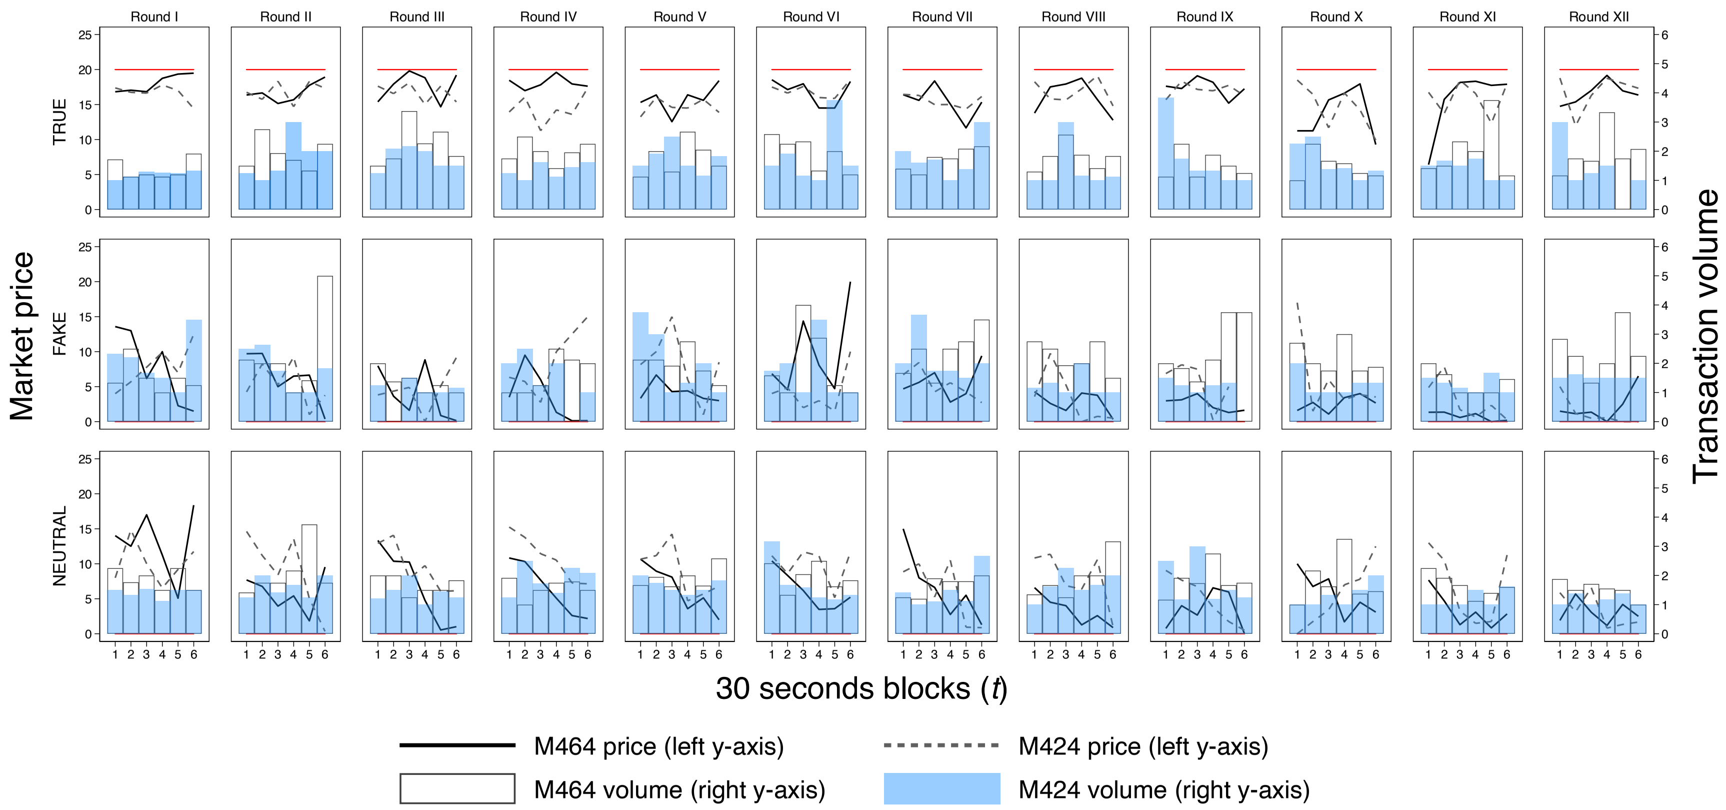

Note. The average number of transactions (over all assets and markets) per round is 22 in the M464 treatment and 13 in the M424 treatment.

Figure 3: **Mean and 95% confidence intervals of the mean absolute deviation (MAD) of closing market prices from the fully-revealing rational-expectations (MAD^{FRE}) and prior-information equilibrium (MAD^{PIE}) benchmarks.** Within-treatment differences between the MAD^{FRE} and MAD^{PIE} are tested with the fixed-effects panel regression. Between-treatment differences in MAD^{FRE} are assessed with the random-effects panel specification. Significance levels are denoted by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

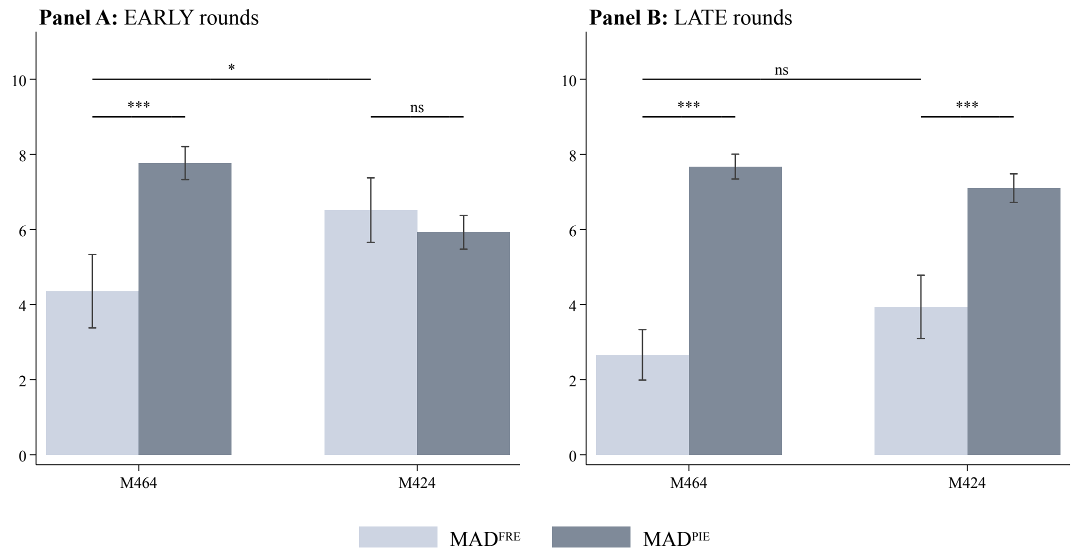

Note. Each bar pools observations from all groups across all trading rounds.

relevant benchmark. Given the above, we define the markets to be successful at aggregating information with prices when MAD^{FRE} is lower than MAD^{PIE}.

Figure 3 details the mean MAD^{FRE} and MAD^{PIE} (with 95% confidence intervals) for EARLY and LATE rounds of the respective treatments. The fixed-effects regressions (within-treatment) finds MAD^{FRE} to be significantly smaller than MAD^{PIE} in both EARLY and LATE rounds ($p < .001$ in each) of the M464 treatment.[8] In contrast, MAD^{FRE} is only significantly smaller that MAD^{PIE} in the LATE rounds ($p < .001$) of the M424 treatment—no significant differences in the EARLY rounds ($p = .480$).[9] This leads us to our first result.

Result 1. *In the M464 treatment, where uninformed intermediaries are a majority in each market, information aggregation across the two separated markets is successful with prices in both EARLY and LATE rounds. By contrast, in M424—where uninformed intermediaries are a minority—information aggregation is only successful with prices in the LATE rounds where traders have some experience.*

Further to Result 1, the random-effects panel regression finds no significant ($p = .398$) between-treatment difference in MAD^{FRE} during the LATE rounds. This suggests that when information aggregation succeeds, closing market prices in both M464 and M424 cluster equally tightly around the FRE benchmark. In contrast, the EARLY rounds MAD^{FRE} are marginally ($p = .051$) lower in the M464 relative to M424 treatment.

We also find the correctness of traders' post-market beliefs to be consistent with Result 1. In M464, side traders' beliefs are correct 85% in EARLY rounds and 79% in LATE rounds, while middle traders achieve 88% and 84%, respectively. In M424, side traders improve from 74% to 84% across EARLY and LATE rounds, and middle traders from 85% to 95%. In all cases, these proportions are significantly higher than the respective prior probabilities of correct beliefs—50% for side traders and 33% for the middle traders—indicating that traders do learn from prices.

Robustness. We perform three robustness checks. First, Result 1 is robust to different block interval specifications: block intervals of 10-, 20-, and 45-seconds.

where p^{i_j} is the closing market price of asset i_j, $m^{i_j,B}$ is the benchmark price, and num is the number of assets in that round whose closing prices are observed. For example num = 4 if only the true and neutral assets are traded in each market.

[8]We regressed $\Delta = MAD^{FRE} - MAD^{PIE}$ on a constant: a negative estimate indicates prices are closer to the FRE benchmark, while a positive estimate indicates the opposite.

[9]The same conclusion is obtained when computing mean absolute deviations using only closing market prices from the left market or only from the right market, confirming that the above observations are not driven by activities in either sides of the market.

Second, we focus only on those closing market prices that occur in the last 90 seconds of the round in computing the relevant mean absolute deviations—i.e., we exclude assets that do not attract traders' attention later in the round. Re-estimating all comparisons with this restriction leaves our main findings unchanged.[10]

Third, we perform the *sign-flip* permutation test. For each matching-group and round, we compute the paired difference ($\Delta = MAD^{FRE} - MAD^{PIE}$). Under the null hypothesis that closing market prices are, on average, equally close to the two benchmarks, the sign of every Δ is exchangeable. We therefore create a null distribution by randomly multiplying each Δ by ± 1 with 50% probability and repeating the procedure 10,000 times to obtain 10,000 "randomise" sample means. The observed M464 treatment mean (average Δ over all rounds) lies in the lower tail of this null distribution, confirming that closing market prices are significantly closer to the FRE benchmark than to the PIE benchmark in both EARLY and LATE rounds of the M464 treatment ($p < .001$ in each case). In the M424 treatment, however, a significant difference emerges only in the LATE rounds ($p < .001$)—the EARLY rounds show no such effect ($p = .512$).

4.2. RATE OF CONVERGENCE TO FRE BENCHMARK

Building on Result 1 we now study whether the speed of price convergence to the FRE benchmark depends on the relative number of intermediaries. The analysis is primarily concerned with the LATE rounds where information aggregation with prices is successful in both treatments—we report the EARLY rounds for completeness.

To do so, we track the block-by-block dynamics of *mispricing*, defined as the deviation of asset prices from the FRE benchmark. For every 30-second block t, we compute the mean absolute deviation of block-t market prices of all assets from their FRE benchmark, denoted as MAD_t^{FRE}. Unlike the closing-price analysis in the previous subsection, MAD_t^{FRE} is calculated only from assets that traded within block t, providing a contemporaneous snapshot of mispricing at that point in the round, whereas closing prices reflect the terminal valuations reached by the end of trading.[11]

Although traders can update their beliefs continuously as new price signals

[10]Using this restricted sample, we find MAD^{FRE} to be significantly lower than MAD^{PIE} in the EARLY ($p < .001$) and LATE ($p < .001$) rounds of the M464 treatment. In contrast, MAD^{FRE} is only significantly smaller in the LATE ($p < .001$) rounds of the M424 treatment—no significant differences in EARLY ($p = .332$) rounds.

[11]If there are no transactions in block t, MAD_t^{FRE} is not defined for that block.

Table 3: **Estimated exponential-decay parameters for mispricing (MAD_t^{FRE}) in EARLY and LATE rounds**. λ is the baseline decay rate in the M464 treatment and δ is the incremental effect in the M424 treatment. Standard errors (clustered by matching-groups) are in parentheses. Significance levels are denoted by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Dependent variable: MAD_t^{FRE}		
	EARLY	LATE
λ	0.473^{***}	0.434^{***}
	(0.119)	(0.147)
δ	-0.326^{**}	-0.129
	(0.131)	(0.177)
n	467	511
Estimated $\hat{\lambda}_{\text{M464}}$	0.473	0.434
Estimated $\hat{\lambda}_{\text{M424}}$	0.147	0.304

arrive, the informational value of each additional trade—and hence the scope for "learning from prices"—diminishes as the round progresses. Indeed, prices typically move rapidly toward the FRE benchmark in the first few transactions, after which adjustments become smaller. To capture this pattern, we model the evolution of MAD_t^{FRE} with an exponential decay process:

$$MAD_t^{FRE} = \alpha e^{-\lambda(t-0.5)}, \qquad t = 1, \ldots, 6$$

where α is the initial deviation and $\lambda > 0$ is the decay rate. Each 30-second block reduces the remaining deviations by the factor of $e^{-\lambda}$, so a larger λ indicates faster convergence toward the FRE benchmark over the round. Note there is a $t - 0.5$ in the formula since the average time in block t is halfway between the time at end of block $t - 1$ (which is the time at the start of block t) and the time at the end of block t.

Given the above specification, we estimate the following log-linear fixed-effects model:

$$\ln(MAD_t^{FRE}) = \mu - \lambda t - \delta(t \times \mathbb{1}_{M424}) + \epsilon_t$$

where μ is the matching-group × round fixed-effects, λ is the baseline decay rate for the M464 treatment, and δ measures how that rate changes in the M424 treatment.[12]

[12]In the panel setup, each matching-group × round combination is treated as a pseudo-group observed over the six time-varying blocks t. We then estimate a fixed-effects regression and cluster the standard errors at the matching-group level.

Figure 4: **Cumulative probability functions of $\overline{MAD}^{FRE}$ by treatment.** A leftward shift indicates smaller deviations and hence faster convergence toward the FRE benchmark.

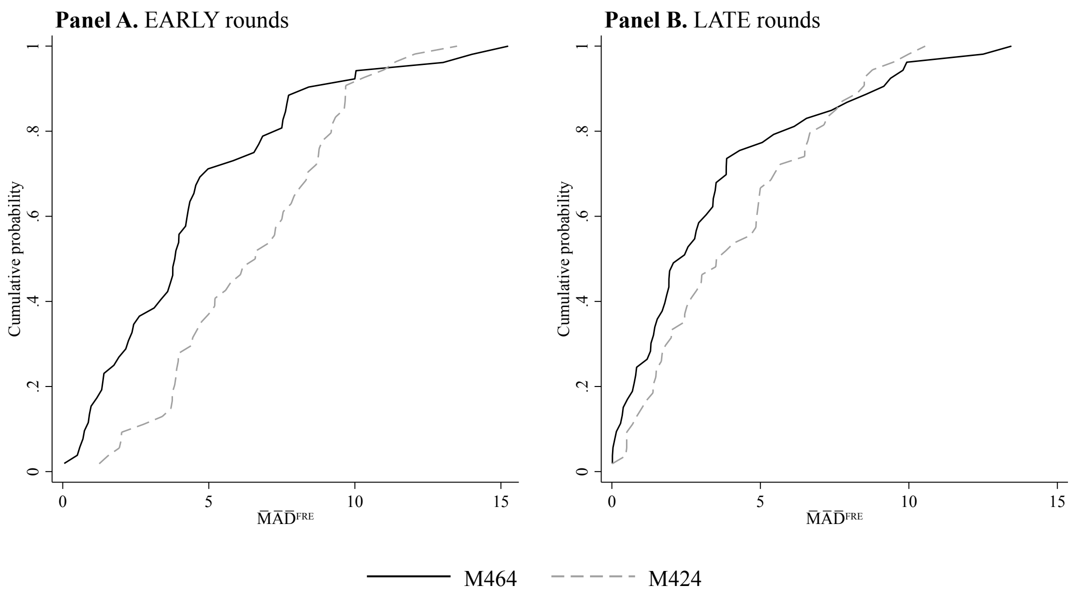

The estimates are reported on Table 3 and can be summarised as follows. In the EARLY rounds, convergence is markedly faster in M464 than in M424—$\hat{\lambda}_{\text{M464}} = 0.473$ ($p = .005$) versus $\hat{\lambda}_{\text{M464}} = 0.147$ ($p = .033$)—and the between-treatment gap is significant ($p = .024$). By contrast, in the LATE rounds the rates narrow to $\hat{\lambda}_{\text{M464}} = 0.434$ ($p = .021$) and $\hat{\lambda}_{\text{M424}} = 0.304$ ($p = .017$), and the difference is no longer significant ($p = .433$).[13] This leads us to our second result.

Result 2. *In the LATE rounds, the rate of convergence towards the Fully-Revealing Rational-Expectation prices do not differ significant depending on whether the uninformed intermediaries are a majority (M464) or minority (M424) within each separated market.*

Robustness. We do two robustness checks. First, Result 2 is robust to different block interval specifications: block intervals of 10-, 20-, and 45-seconds.

Second, because information aggregation succeeds in the LATE rounds of both treatments and their convergence rates do not differ significantly, we should likewise expect the cumulative mispricing—averaged over all blocks—to show

[13] We can also compute the corresponding half-life, $t_{1/2} = \ln 2/\hat{\lambda}$, that is the number of 30-second blocks required to halve the remaining deviations. EARLY rounds: M464 (1.5 blocks ≈ 45s) and M424 (4.7 blocks ≈ 140s). LATE rounds: M464 (1.6 blocks ≈ 48s) and M424 (2.3 blocks ≈ 70 s).

no meaningful difference between treatments. We therefore compute, for each group-round cell, the average mispricing (i.e., MAD_t^{FRE}) over all blocks. This single measure of average mispricing—denoted by $\overline{MAD}^{FRE}$—summarises how far the round's prices, on average, deviate from the benchmark; the smaller the value, the closer the market has come to full information. Figure 4 details the cumulative probability function of $\overline{MAD}^{FRE}$ in both treatments. The curve for M464 often lies to the left of that for M424, indicating smaller cumulative mispricing overall. However, a random-effects regression finds no significant between-treatment differences in $\overline{MAD}^{FRE}$ during the LATE rounds (p = .422); the same test finds significantly lower average mispricing in M464 during the EARLY rounds (p = .028).

4.3. PATTERNS OF TRADE

To better understand the drivers of information aggregation, we study the patterns of buying and selling within each treatment.

4.3.1. Who is buying and who is selling?

For each asset in market $j \in \{L, R\}$, trade can be classified into four *flows*: middle-to-middle trader (M→M), middle-to-side (M→S), side-to-middle (S→M), and side-to-side (S→S). For example, S→M indicates that a middle trader buys from a side trader. Figure 5 shows, for every t block, the share of transactions accounted for by each flow (left y-axis). The scatterplot (right y-axis) indicates the median transaction price for a specific flow. Finally, the numeral on the top of each bar details the total number of transactions. For example, we see that the total number of true asset transactions in block 6 of the M424 LATE rounds (aggregating over all groups) is 29. Of these, 79% are S→M flow with the median price of 19.

Side traders initially hold an informational advantage over the fake asset: they can profit by selling the fake asset to uninformed middle traders at any positive price. Indeed, most fake asset transactions involve the S→M flow—though at low prices—and occur more frequently in M464 than in M424, due to the larger number of middle traders. We also observe some S→S flows, especially in LATE rounds of M424, again at low prices. Since side traders know the asset is worthless, such trades cannot reflect informational differences.

Because middle traders can operate in both the left and right markets, they should eventually gain an informational advantage over the side traders for the true and neutral assets. Consistent with this, S→M flows account for a large share

Figure 5: **Trade flows between side and middle traders, by asset type.** Each stacked bar reports the proportion of transactions involving a given asset type that follow one of four flows: middle-to-middle (M⟶M), middle-to-side (M⟶S), side-to-middle (S⟶M), and side-to-side (S⟶S). The numeral on the top of each bar denotes the total number of transactions for that given condition. Overlaid scatter points indicate the median transaction price for each flow (right y-axis).

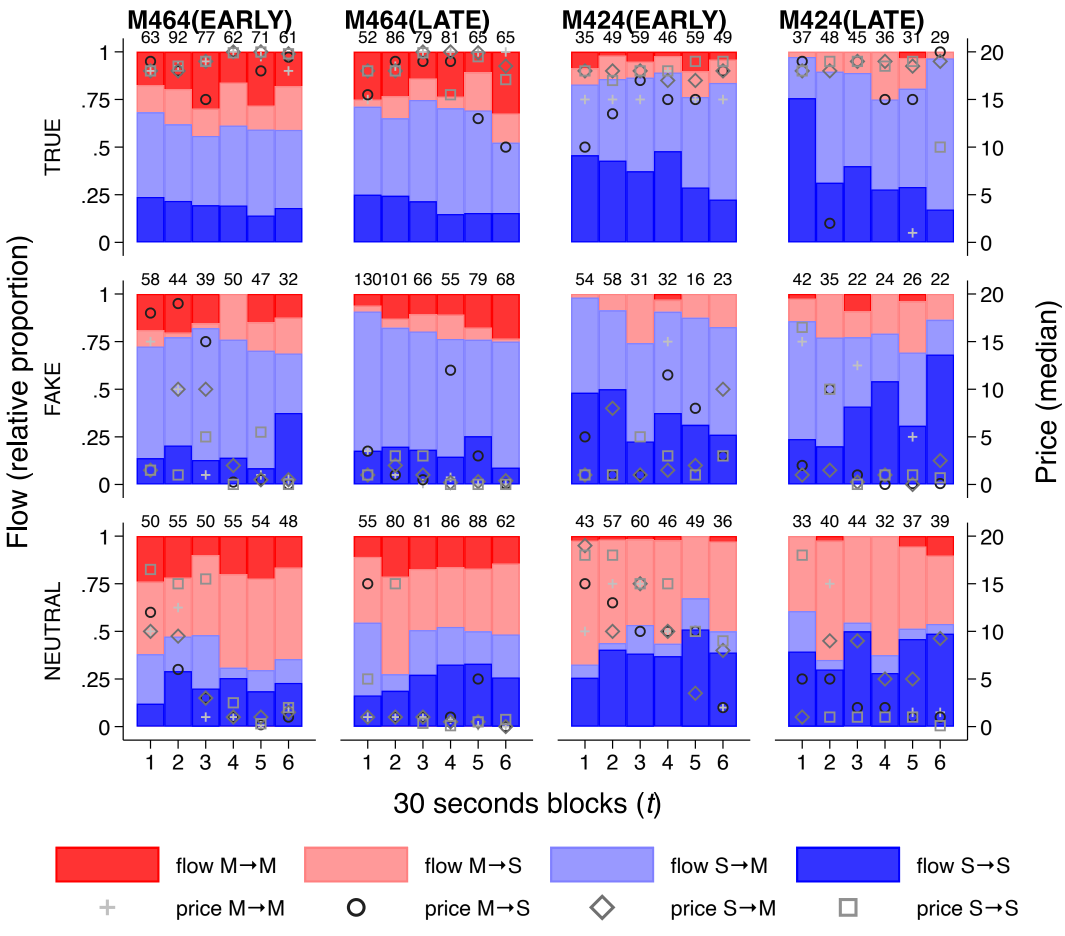

of true asset trades. Likewise, M→S flows account for much of the trading in the neutral asset. Taken together, these patterns lead us to the following observation.

Observation 1. *Trading activity is consistent with traders exploiting their informational advantages: side traders predominantly sell the fake asset to uninformed middle traders, while middle traders predominantly sell the neutral asset and buy the true asset from the side traders.*

The astute reader may inquire whether the fake and neutral flows are due to arbitrage by the middle trader exploiting the price differences across the left and right markets. For instance, if x^L is the fake asset in the left market then the price would be lower than x^R in the right market where it is the neutral asset. Arbitraging would involve the middle trader will buying x^L (fake asset) in left market, S→M flow, and selling x^R (neutral asset) in the right market, M→S flow—recall that assets are not transferable across markets. However, arbitrage alone cannot account for flows in the true asset since the same logic would imply that the S→M flows will be similar to the M→S, which we do not observe. This suggests that the middle trader ought to be learning about the true state from prices in both markets.

4.3.2. Price asymmetries across the left and right markets

Suppose the true state is Y, and side traders in the left and right markets are endowed with the not-X and not-Z information, respectively. The model in Section 3.1 predicts that middle traders deduce the true state from the fake states (X in the left market and Z in the right) in each market. Given that side traders know which asset is worthless, this should result in lesser demand for that asset and potentially predictable price asymmetries across markets: the *True-Fake* (assets y_L and x_L) price gap in the left market should exceed the corresponding *True-Neutral* (assets y_R and x_R) gap in the right, while the *True-Neutral* gap in the left market should be smaller than the *True-Fake* gap in the right. Middle traders use these cross-market asymmetries to infer the true state. As their market activities transmit information between markets, the asymmetries gradually diminish, and price gaps converge as the markets approach the fully revealing equilibrium.

In continuous time, it is difficult to pinpoint when middle traders learn the true state, as beliefs evolve dynamically and are not directly observable. Nevertheless, this diffusion process should leave a distinct price signature. For each matching-group and round, let P_i^j be the average transaction price of the $i \in \{\text{true}, \text{fake}, \text{neutral}\}$ asset in market $j \in \{L, R\}$. The above reasoning yields the following prediction:

$$\left(P^{j}_{\text{true}} - P^{j}_{\text{fake}}\right) \geq \left(P^{j'}_{\text{true}} - P^{j'}_{\text{neutral}}\right) \quad \text{for } j \neq j'.$$

Empirically, this inequality holds in 64% of the 122 *eligible markets* under the M464 treatment and in 65% of the 71 eligible markets under M424—market j is eligible if it records at least one transaction in both the true and fake assets within that market and at least one transaction in the true and neutral assets in the opposite market $j \neq j$.[14] A two-sided proportion test rejects the null of equal likelihood ($p \leq .013$ in both treatments), indicating that such cross-market price asymmetries occur significantly more often than would be expected by chance.[15] The above lead us to the next observation:

Observation 2. *The pattern of cross-market price asymmetries is consistent with middle traders learning about the true state indirectly—by first identifying the fake states in each separated market—and with information diffusing across markets.*

4.4. DRIVERS OF INFORMATION AGGREGATION

Observations 1 and 2 together suggest that information aggregation in the experiment arises through two complementary mechanisms. First, side traders incorporate private information into prices by exploiting their initial informational advantage in the fake asset. Second, the middle traders learn the true state by identifying the fake states and transmitting this information across markets. These findings motivate a closer examination of the drivers of information aggregation—specifically, how the trading behaviour of side and middle traders contributes to price convergence toward the FRE.

4.4.1. Influence of the middle traders

As a first step, we examine the influence of middle traders on price convergence relative to that of side traders. Because middle traders learn the true state before side traders, they should exert direct pressure on the price of the true asset, pushing it toward the FRE, whereas their influence on the fake and neutral assets is limited since they assign these assets a value of zero and therefore have little

[14]Eligible markets account for 58% and 30% of all markets in the M464 and M424 treatments, respectively.

[15]When we further restrict our attention to rounds containing exactly two eligible markets, both inequalities hold jointly in 34% of the 44 rounds under M464 and 55% of the 20 rounds under M424. A two-sided proportion test again rejects the null of a 0.25 likelihood for the M424 treatment ($p < .001$) but not for M464 ($p = .167$).

incentive to purchase them. In each separated market, the relative prices of the true and neutral assets therefore help side traders deduce the true state.

Intuitively, the relative true-asset influence of the middle traders, over the side traders, should be especially pronounced in the M464 treatment, where six informed middle traders compete to exploit their informational advantage. By contrast, this mechanism is less clear in the M424 treatment, where only two middle traders are present and competitive forces among informed middle traders may be weaker.

To study the relative influence of the middle traders, we focus on traders' purchasing decisions.[16] For each market $j \in \{L, R\}$, we define a *convergent purchase* as a purchase that brings the transaction price of an asset weakly closer to its FRE price relative to the preceding trade.[17] For each trader, we define the *convergent purchase ratio* (CPR) as the number of asset-specific convergent purchases in market j made by that trader as a proportion of the total number of that asset's transactions in the same market—the CPR is computed separately for each asset type (true, fake, or neutral). For middle traders, who may trade in both markets, we compute their CPR as the average over the Left and Right markets. For example, if a middle trader makes 20 true asset convergent purchases in the left market (out of 100 true asset transactions by all traders in that market) and 5 true asset convergent purchases in the right market (out of 50 true asset transactions by all traders in that market), her true asset CPR will be computed as $[20/100 + 5/50]/2 = 0.15$. A higher CPR for a given asset type indicates a greater influence by that trader in moving that asset's price toward its FRE.

The random-effects regression estimates in Table 4 show that the true-asset CPR is significantly higher for middle traders than for side traders in both the EARLY ($p = .032$) and LATE ($p = .015$) rounds of the M464 treatment. In the M424 treatment, we find only a weakly higher true-asset CPR for middle traders in the EARLY round ($p = .075$) and no difference in the LATE round ($p = .966$). For the fake and neutral assets, however, we find no significant differences between middle and side traders ($p \geq .159$ in all comparisons).

[16]We focus on purchases rather than sales because purchases directly reflect traders' updated valuations and drive prices toward the fully revealing equilibrium, whereas sales often occur when traders liquidate positions at prices above the FRE, reflecting profit-taking motives.

[17]For instance, with an FRE price of 20 and the sequence of trade prices of 17, 19, 23, 22, 22, and 25, the second, fourth, and fifth transactions are convergent purchases—they bring prices weakly closer the equilibrium value.

Table 4: **Random-effects estimates of the Convergent Purchase Ratio (CPR) by trader and asset types.** Standard errors (clustered by matching group) are shown in parentheses. Significance levels are denoted by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Dependent variable. Convergence purchase ratio (CPR).

	EARLY rounds			LATE rounds		
Asset type:	True	Fake	Neutral	True	Fake	Neutral
Panel A. M464 treatment						
Middle trader	0.060**	-0.172	0.047	0.083**	-0.032	0.004
	(0.028)	(0.113)	(0.141)	(0.034)	(0.091)	(0.073)
Constant	0.245***	0.479***	0.373***	0.291	0.431	0.360
	(0.032)	(0.118)	(0.054)	(0.036)	(0.066)	(0.037)
n	158	83	103	156	119	157
Panel B. M424 treatment						
Middle trader	0.145*	0.008	-0.150	0.003	0.074	-0.067
	(0.081)	(0.128)	(0.107)	(0.080)	(0.106)	(0.153)
Constant	0.346***	0.564***	0.466***	0.428***	0.385***	0.507***
	(0.049)	(0.098)	(0.054)	(0.081)	(0.083)	(0.038)
n	99	44	96	85	51	87

Result 3. *Middle traders exert significantly greater influence than side traders in driving the true asset's price toward the fully revealing equilibrium in the M464 treatment through their purchasing decisions, whereas this relative influence is substantially weaker in the M424 treatment.*

4.4.2. The influence of non-transaction signals

Result 3 raises the question about how side traders in the M424 treatment learn about the true state. One possibility is that they infer the true state from non-transaction signals—such as observing order submissions, cancellations, or unexecuted bids and asks—which provide additional cues about others' beliefs and intentions even when no trade occurs.

It is difficult to observe the influence of non-transactional signals directly. Nevertheless, if side traders in the M424 treatment rely more on such signals—relative to their counterparts in M464—then their deduction of the true state should be less sensitive to transaction-based price signals. In other words, side traders in M424 may be "*doing more with less*", inferring the true state from weaker transaction price signals by relying more heavily on non-transactional cues. We test this by regressing the proportion of side traders whose ex-post market beliefs are correct on the mean absolute deviation of closing market prices from the FRE

Table 5: **Random-effects estimates of the proportion of side traders whose ex-post beliefs are correct by treatment.** Robust Standard errors are in parentheses. Significance levels are denoted by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Dependent variable. Proportion of side traders whose ex-post beliefs are correct.		
	EARLY	LATE
MAD^{FRE}	-0.027** (0.010)	-0.060*** (0.004)
M424	-0.063 (0.082)	-0.018 (0.050)
M424 × MAD^{FRE}	0.001 (0.016)	0.029** (0.013)
Constant	0.968*** (0.059)	0.972*** (0.029)
n	106	107

benchmark (MAD^{FRE}).

As expected, the random-effects estimates in Table 5 show that MAD^{FRE} has a negative and highly significant effect ($p \leq .014$ in both EARLY and LATE rounds) on the proportion of side traders who correctly identify the true state. In the EARLY rounds, there is no significant between-treatment difference in this effect ($p = .993$), whereas in the LATE rounds the negative influence of MAD^{FRE} is significantly weaker in the M424 treatment ($p = .033$).

Result 4. *Larger deviations of closing prices from the fully revealing equilibrium significantly reduce side traders' ability to identify the true state in both treatments. However, this price sensitivity is similar across treatments in the EARLY rounds and becomes significantly weaker in M424 during the LATE rounds. This LATE round evidence is consistent with the M424 side traders relying on non-transaction price signals to infer the true state.*

To some extent, Result 4 also helps explain why information aggregation succeeds in both the EARLY and LATE rounds of M464, but only in the LATE rounds of M424. When middle traders form a majority in each separated market, as in M464, information diffuses quickly through their price-improving trades, generating rapid convergence early in trading (Result 3). In contrast, when middle traders are a minority in each market, as in M424, side traders must gradually learn to interpret non-transactional signals to infer the true state. As they become more adept at extracting information from these indirect cues, aggregation in M424 improves over time, producing convergence only in the LATE rounds.

Table 6: **Random-effects estimates of traders market payoff in EARLY and LATE rounds.** Matching-group clustered standard errors are in parentheses. Significance levels are denoted by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Dependent variable: Market payoff (end of round)				
	M464		M424	
	EARLY	LATE	EARLY	LATE
Middle Trader	-2.133	1.783	10.205*	4.869**
	(2.909)	(1.168)	(5.656)	(2.134)
Constant	320.914***	319.235***	317.958***	319.026***
	(1.246)	(0.501)	(1.131)	(0.426)
n	728	742	540	540

4.5. CAN INTERMEDIARIES PROFIT FROM THEIR POSITION?

Having shown that middle traders play a central role in transmitting information across markets, a natural question arises: do intermediaries profit from their position? In principle, middle traders could earn informational rents by exploiting their unique access to both markets—buying underpriced assets and selling overpriced ones as information diffuses. However, such profits may depend on the market structure: when middle traders are numerous (M464) competition among them may quickly erode these rents; when they are fewer (M424) limited competition could allow intermediaries to sustain higher gains.

We study the difference in end-of-round *market payoffs* between middle and side traders within each treatment.[18] The within-treatment analysis isolates the effect of trader role under a common market structure, holding constant the liquidity conditions, trading intensity, and information environment specific to each treatment—all of which may vary substantially across treatments. The random-effects regression estimates on Table 6 find that end-of-round market are not significantly different (EARLY: $p = .463$; LATE: $p = .127$) for middle and side traders in the M464 treatment. In contrast, middle traders earn significantly higher payoffs than side traders (EARLY: $p = .071$; LATE: $p = .023$) in the M424 treatment.

[18]The market payoff for each trader is calculated as the sum of their end-of-round cash holdings and the value of their remaining asset inventory, based on their realised values.

Result 5. *The ability of intermediaries to profit from their position depends on market structure. When intermediaries are numerous (M464), competition erodes informational rents, yielding similar payoffs for middle and side traders. When they are fewer (M424), middle traders earn significantly higher payoffs, indicating that limited competition allows them to profit from their informational advantage.*

Building on Result 5, we next ask *when* middle traders in the M424 treatment are able to profit from their position. Do they capture informational rents early in the round—when price discrepancies across markets are largest—or do profits accumulate later, as information gradually diffuses and prices converge? Understanding the timing of these gains provides insight into how intermediaries exploit their informational advantage.

To study the above, we define a trader's block-t portfolio value PortValue_t as the sum of her cash holdings and the market value of her asset inventory at the end of that block—asset prices are evaluated at their realised values. The initial portfolio value, PortValue_0, corresponds to the common initial endowment and is identical across all traders. We define the one-period in portfolio value as

$$\Delta\text{PortValue}_t = \text{PortValue}_t - \text{PortValue}_{t-1}, \qquad \text{for } t = 1, 2, .., 6$$

We define a trader as *active* if she purchases or sells at least one asset within a given block. For each active trader, a positive $\Delta\text{PortValue}_t$ indicates a gain in portfolio value during block t resulting from trading activity, whereas a negative $\Delta\text{PortValue}_t$ indicates a loss in portfolio value over the same period.[19]

We use the multilevel (random-effects) model to study the within-treatment difference in $\Delta\text{PortValue}_t$ for *active* middle and side traders across the different blocks. The resulting predicted margins are detailed in Figure 6, where a positive value indicates that $\Delta\text{PortValue}_t$ is higher for middle traders relative to side traders. In the M464 treatment, we do not observe any consistent or coherent differences between middle and side traders in $\Delta\text{PortValue}_t$ across either the EARLY or LATE rounds. In the M424 treatment, we observe $\Delta\text{PortValue}_t$ to be, on average, consistently higher for middle traders relative to side traders from block 1 onwards.[20] This suggests that middle traders in the M424 treatment are able to profit from their informational advantage throughout the trading round.

[19]We focus on changes in portfolio value across blocks rather than portfolio levels, as the change captures the timing and source of traders' gains or losses, isolating new value created within each block from cumulative effects of earlier trading.

[20]M424 EARLY rounds: $\Delta\text{PortValue}_t$ are weakly higher for active middle traders in block 3 ($p = .098$), block 4 ($p = .068$) and block 5 ($p = .098$). M424 LATE rounds $\Delta\text{PortValue}_t$ are weakly higher in block 2 ($p = .052$), block 3 ($p = .021$) and block 6 ($p = .016$).

Figure 6: **Predicted differences in ΔPortValue$_t$ between side and middle traders.** Predicted margins are based on multilevel model estimates in which *blocks* are nested within *traders*, and traders are nested within *matching groups*. Each line shows the estimated difference in the change in portfolio value between middle and side traders across blocks, with 95% confidence intervals. Positive values indicate greater portfolio gains for middle traders, while negative values indicate greater gains for side traders.

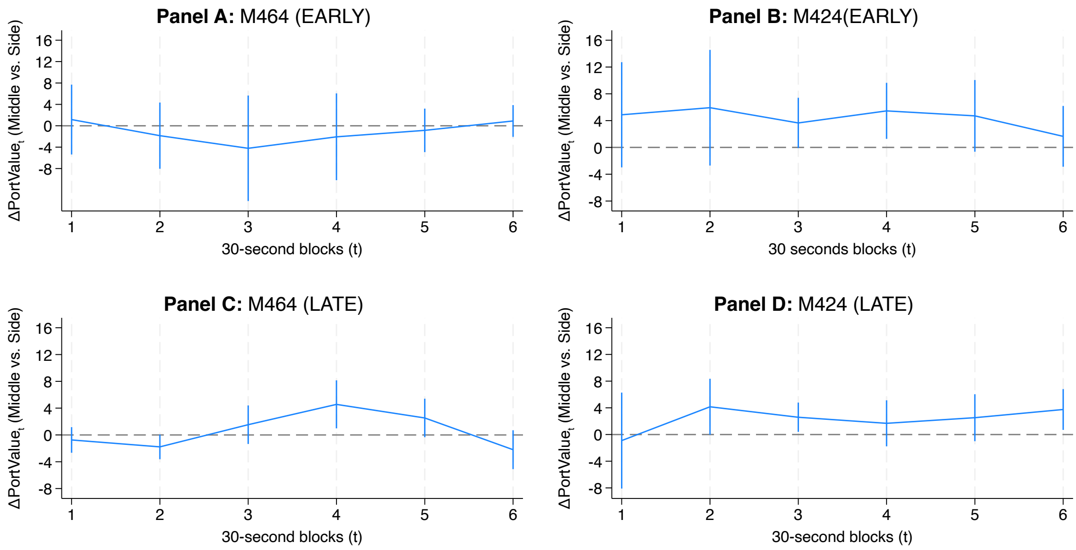

This leads us to our final observation.

Observation 3. *The extraction of informational rents by intermediaries in the M424 treatment does not appear to be concentrated at the beginning or end of the trading round.*

5. CONCLUSION

This paper asks whether fragmented Arrow-Debreu markets can aggregate dispersed private information when no single venue observes all state-relevant signals and when cross-venue information flows only through uninformed intermediaries whose competitive pressure may vary across markets. Our design embeds the classic Plott and Sunder (1988) information structure—where each group of side traders receives a distinct noisy signal about the same underlying state—but places these signals into separated markets that are connected only through uninformed intermediaries. In our setting, fragmentation arises along two distinct dimensions. First, there is market fragmentation: trading takes place in two physically separated venues. Second, there is information fragmentation: side traders are restricted to a single venue and can observe only local market activity. Hence, no individual market contains all the private information needed to identify the true state, and information aggregation in both markets can only be facilitated through the market cross-market activities of uninformed intermediaries (middle traders). Finally, to vary the extent of intermediary competition, the two treatments differ only in the relative number of middle and side traders within each separated market: in M424 the intermediaries are a minority, whereas in M464 they form a majority.

In M464, where uninformed intermediaries are a majority within each market, information aggregation succeeds in both EARLY and LATE rounds. In M424, where intermediaries are a minority, aggregation emerges only in the LATE rounds once traders have learned the environment. Despite these differences, the rate at which prices converge toward the fully-revealing equilibrium benchmark in the LATE rounds does not differ significantly across treatments. The mechanisms underlying convergence also differ: middle traders exert substantially greater influence on true-asset prices in M464, whereas their influence is much weaker in M424. Side traders in both treatments are less able to identify the true state when prices deviate further from the fully-revealing equilibrium, but this sensitivity diminishes in M424 during the LATE rounds, consistent with increased reliance on non-transaction signals. Finally, intermediaries' ability to earn informational rents depends on market structure—competition among numerous intermediaries in M464 eliminates excess returns, while limited competition in M424 allows intermediaries to profit from their cross-market informational advantage.

Taken together, our results show that the ability of markets to aggregate diverse and imprecise private information—well established in centralised settings—can extend to less centralised, fragmented multi-venue environments in which information is distributed unevenly and intermediaries are the only link between venues. The experimental framework introduced here also opens the door to studying far more complex fragmented environments, providing a flexible platform for analysing how information diffuses, interacts, and potentially becomes distorted in richer multi-venue market settings.

Future studies. A natural extension concerns the informational environment faced by intermediaries. In our M424 and M464 design, intermediaries know that the two separated markets embed different noisy signals about the true state. In many real-world settings, however, markets may sometimes draw on the same underlying signal, creating the possibility of *information mirages* (e.g., Camerer and Weigelt, 1991; Noussair and Xu, 2015; Morone and Nuzzo, 2019), where redundant noise is mistakenly treated as independent information.

Our design also lends itself to studying *daisy-chain* market structures. In the current treatments, intermediaries can access both left and right markets and therefore observe information from each venue directly. In practice, intermediaries often cannot observe all trading activity. Instead, information diffuses gradually through a sequence of partially connected venues, much like supply chains in which regional and local wholesalers pass information along step by step—other relevant examples include OTC markets or interbank FX. This motivates extending the design to a M4334-type framework, where three uninformed intermediaries each connect to either the Left or the Right markets—both markets will remain populated by informed side traders. There is an additional intermediate venue (middle market) that is open solely to uninformed intermediaries. Such a structure would allow the study of information diffusion over a chain of markets in which intermediaries themselves hold no private information. Furthermore, this pushes the analysis of fragmented markets toward questions of how information *percolates* through multi-market networks.

The exploration of chains of markets naturally extends to questions about how the topology of the chain (e.g., star vs. circular network) affects the speed of information aggregation. Insights from such structures may be directly relevant to research on information aggregation in spatial economics and economic geography, where the pattern of connections between locations shapes how quickly information and innovations may diffuse across space.

REFERENCES

Anderson, Ryan, Peter Bossaerts, and Felix Fattinger, "How information structure impacts aggregation of information in markets," *Available at SSRN 3618878*, 2020.

Angerer, Martin, Tibor Neugebauer, and Jason Shachat, "Arbitrage bots in experimental asset markets," *Journal of Economic Behavior & Organization*, 2023, *206*, 262–278.

Asparouhova, Elena and Peter Bossaerts, "Experiments on percolation of information in dark markets," *Economic Journal*, 2018, *128* (613), 2281–2302.

Babus, Ana and Cecilia Parlatore, "Strategic fragmented markets," *Journal of Financial Economics*, 2022, *145* (3), 876–908.

Biais, Bruno and Sébastien Pouget, "Microstructure, incentives, and the discovery of equilibrium in experimental financial markets," IDEI Working Paper 2000.

__ , Pierre Hillion, and Chester Spatt, "An empirical analysis of the limit order book and the order flow in the Paris Bourse," *Journal of Finance*, 1995, *50* (5), 1655–1689.

Camerer, Colin and Keith Weigelt, "Information mirages in experimental asset markets," *Journal of Business*, 1991, pp. 463–493.

Charness, Gary, Uri Gneezy, and Brianna Halladay, "Experimental methods: Pay one or pay all," *Journal of Economic Behavior & Organization*, 2016, *131*, 141–150.

Chen, Daniel and Darrell Duffie, "Market fragmentation," *American Economic Review*, 2021, *111* (7), 2247–2274.

Chew, Soo Hong, King King Li, Robin Chark, and Songfa Zhong, *Source preference and ambiguity aversion: models and evidence from behavioral and neuroimaging experiments*, Vol. 20 of *Advances in Health Economics and Health Services Research*, Emerald Group Publishing Limited,

Choo, Lawrence, Todd R. Kaplan, and Ro'i Zultan, "Information aggregation in Arrow-Debreu Markets: An experiment," *Experimental Economics*, 2019, *22*, 625–652.

_ , _ , and _ , "Manipulation and (mis)trust in prediction markets," *Management Science*, 2022, *68* (9), 6716–6732.

Corgnet, Brice, Cary Deck, Mark DeSantis, Kyle Hampton, and Erik O Kimbrough, "When do security markets aggregate dispersed information?," *Management Science*, 2023, *69* (6), 3697–3729.

_ , Mark DeSantis, and David Porter, "Information aggregation and the cognitive make-up of market participants," *European Economic Review*, 2021, *133*, 103667.

Deck, Cary and David Porter, "Prediction markets in the laboratory," *A Collection of Surveys on Market Experiments*, 2013, pp. 219–234.

Dubey, Pradeep, John Geanakoplos, and Martin Shubik, "The revelation of information in strategic market games: A critique of rational expectations equilibrium," *Journal of Mathematical Economics*, 1987, *16* (2), 105–137.

Duffie, Darrell, Semyon Malamud, and Gustavo Manso, "Information percolation with equilibrium search dynamics," *Econometrica*, 2009, *77* (5), 1513–1574.

_ , _ , and _ , "Information percolation in segmented markets," *Journal of Economic Theory*, 2014, *153*, 1–32.

Fama, Eugene F, "Efficient capital markets: A review of theory and empirical work," *Journal of Finance*, 1970, *25* (2), 383–417.

Fischbacher, Urs, "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental economics*, 2007, *10* (2), 171–178.

Forsythe, Robert and Russell Lundholm, "Information aggregation in an experimental market," *Econometrica*, 1990, pp. 309–347.

Foucault, Thierry, Marco Pagano, and Ailsa Röell, *Market liquidity: theory, evidence, and policy*, Oxford University Press, 2023.

Fox, Craig R. and Amos Tversky, "Ambiguity aversion and comparative ignorance," *Quarterly Journal of Economics*, 1995, *110* (3), 585–603.

Fox, Craig R and Martin Weber, "Ambiguity aversion, comparative ignorance, and decision context," *Organizational behavior and human decision processes*, 2002, *88* (1), 476–498.

Froot, Kenneth A and Emil M Dabora, "How are stock prices affected by the location of trade?," *Journal of Financial Economics*, 1999, *53* (2), 189–216.

Galanis, Spyros, Christos A Ioannou, and Stelios Kotronis, "Information aggregation under ambiguity: theory and experimental evidence," *Review of Economic Studies*, 2024, *91* (6), 3423–3467.

Gromb, Denis and Dimitri Vayanos, "Limits of arbitrage," *Annual Review Financial Economics*, 2010, *2* (1), 251–275.

Grossman, Sanford, "On the efficiency of competitive stock markets where trades have diverse information," *The Journal of finance*, 1976, *31* (2), 573–585.

Halim, Edward, Yohanes E Riyanto, Nilanjan Roy, and Yan Wang, "How dark trading harms financial markets," *Economic Journal*, 2025, *135* (669), 1711–1733.

Hayek, FA, "The use of knowledge in society," *American economic review*, 1945, *35* (4), 519–530.

Hellwig, Martin F, "On the aggregation of information in competitive markets," *Journal of economic theory*, 1980, *22* (3), 477–498.

_, "Rational expectations equilibrium with conditioning on past prices: A mean-variance example," *Journal of Economic Theory*, 1982, *26* (2), 279–312.

Hendershott, Terrence, Marvin Wee, and Yuanji Wen, "Transparency in fragmented markets: Experimental evidence," *Journal of Financial Markets*, 2022, *59*, 100732.

Lewis, Karen K., "Trying to explain home bias in equities and consumption," *Journal of Economic Literature*, 1999, *37* (2), 571–608.

Makarov, Igor and Antoinette Schoar, "Trading and arbitrage in cryptocurrency markets," *Journal of Financial Economics*, 2020, *135* (2), 293–319.

Menkveld, Albert J, "High frequency trading and the new market makers," *Journal of Financial Markets*, 2013, *16* (4), 712–740.

Morone, Andrea and Simone Nuzzo, "Market efficiency, trading institutions and information mirages: evidence from a laboratory asset market," *Journal of Economic Interaction and Coordination*, 2019, *14* (2), 317–344.

Neugebauer, Tibor, Jason Shachat, and Wiebke Szymczak, "A test of the Modigliani-Miller theorem, dividend policy and algorithmic arbitrage in experimental asset markets," *Journal of Banking & Finance*, 2023, *154*, 106814.

Noussair, Charles and Yilong Xu, "Information mirages and financial contagion in an asset market experiment," *Journal of Economic Studies*, 2015, *42* (6), 1029–1055.

O'Hara, Maureen, "High frequency market microstructure," *Journal of financial economics*, 2015, *116* (2), 257–270.

Ostrovsky, Michael, "Information Aggregation in Dynamic Markets With Strategic Traders," *Econometrica*, 2012, *80* (6), 2595–2647.

Plott, Charles R., "Markets as information gathering tools," *Southern Economic Journal*, 2000, *67* (1), 2–15.

_ **and Shyam Sunder**, "Efficiency of experimental security markets with insider information: An application of rational-expectations models," *Journal of Political Economy*, 1982, *90* (4), 663–698.

_ **and** _ , "Rational expectations and the aggregation of diverse information in laboratory security markets," *Econometrica*, 1988, *56* (5), 1085–1118.

Radner, Roy, "Rational expectations equilibrium: Generic existence and the information revealed by prices," *Econometrica*, 1979, *47* (3), 655–678.

Shleifer, Andrei and Robert W Vishny, "The limits of arbitrage," *Journal of finance*, 1997, *52* (1), 35–55.

Sunder, Shyam, "Market for information: Experimental evidence," *Econometrica*, 1992, pp. 667–695.

_ , "Experimental Asset Markets: A Survey," in John H. Kagel and Alvin E. Roth, eds., *The Handbook of Experimental Economics*, Princeton, NJ: Princeton University Press, 1995, chapter 6, pp. 445–500.

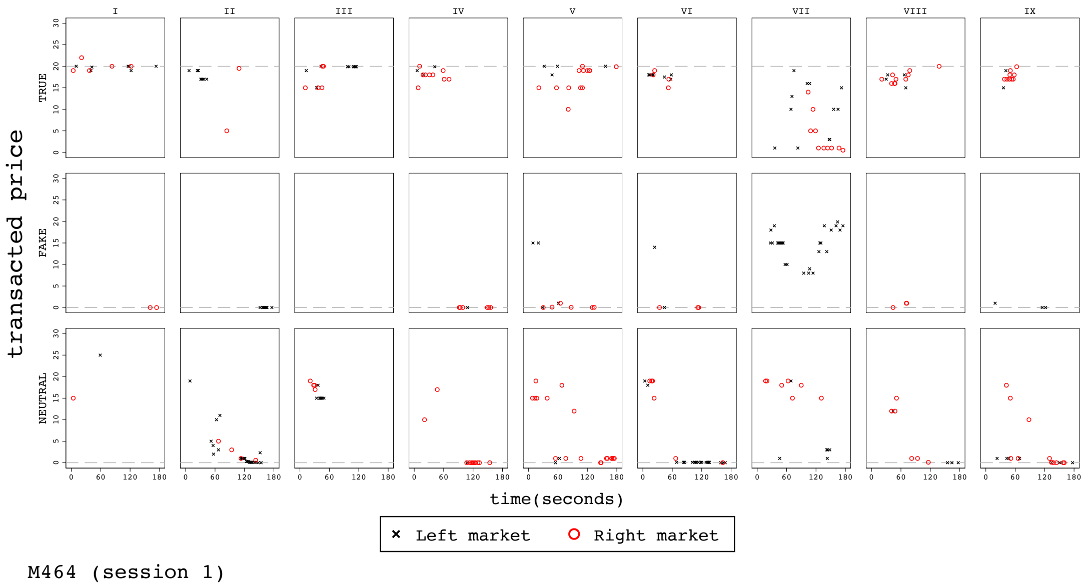

Figure A1: Session 1 (M464 treatment).

Online Appendix

Appendix A. DATA AND ANALYSIS

A.1. EXPERIMENTAL DATA

Figures A1-A18 present the full set of experimental data for each matching group. Within each figure, panels are arranged by round (columns) and by asset type (top row: true asset; middle row: fake asset; bottom row: neutral asset). Each panel reports all transaction prices (y-axis) and transaction times (x-axis), with trades in the left market shown as crosses and trades in the right market shown as hollow circles. The fully revealing equilibrium prices are detailed by the horizontal dashed lines. Finally, Figure A19 details the proportion of side and middle traders who submitted correct ex-post predictions (about the true state) in the respective matching groups and rounds.

Figure A2: Session 2 (M464 treatment).

Figure A3: Session 3 (M464 treatment).

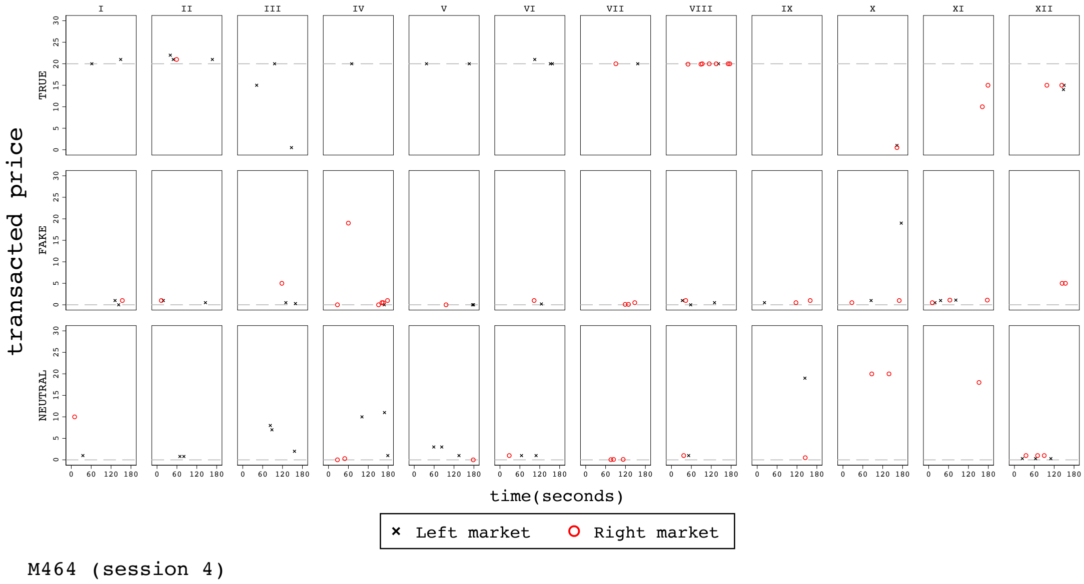

Figure A4: Session 4 (M464 treatment).

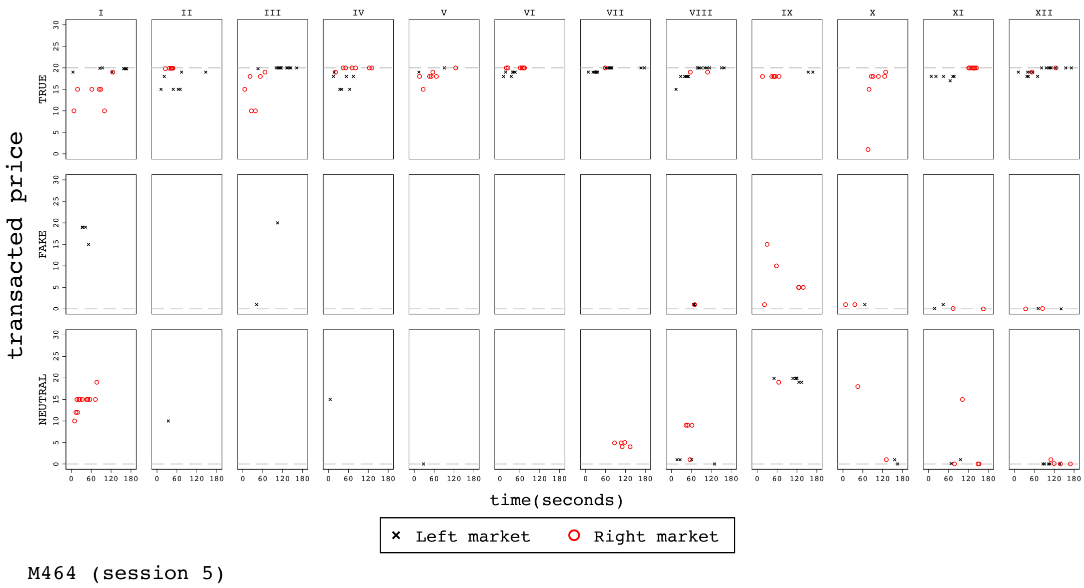

Figure A5: Session 5 (M464 treatment).

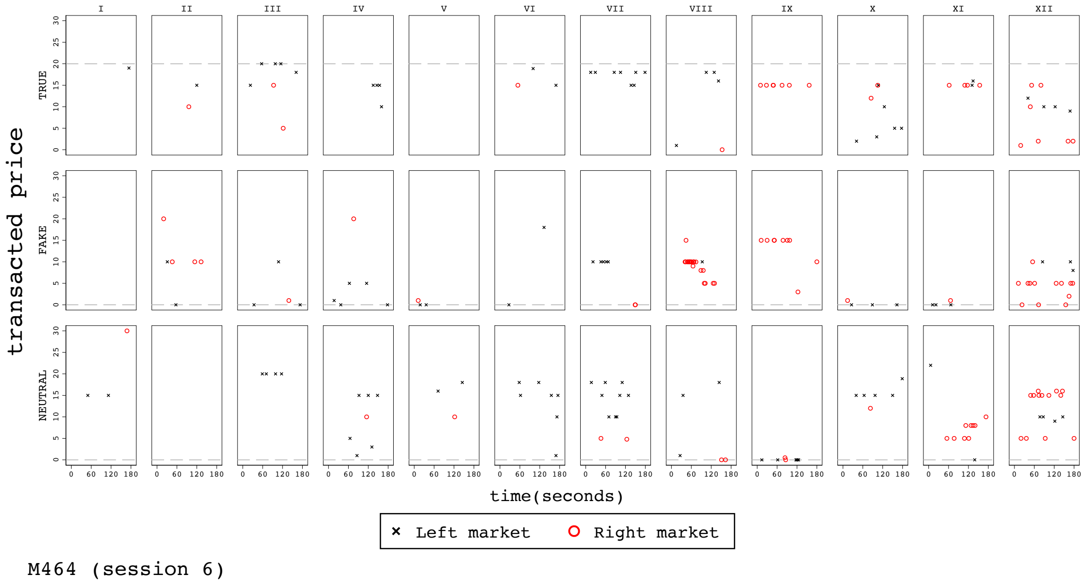

Figure A6: Session 6 (M464 treatment).

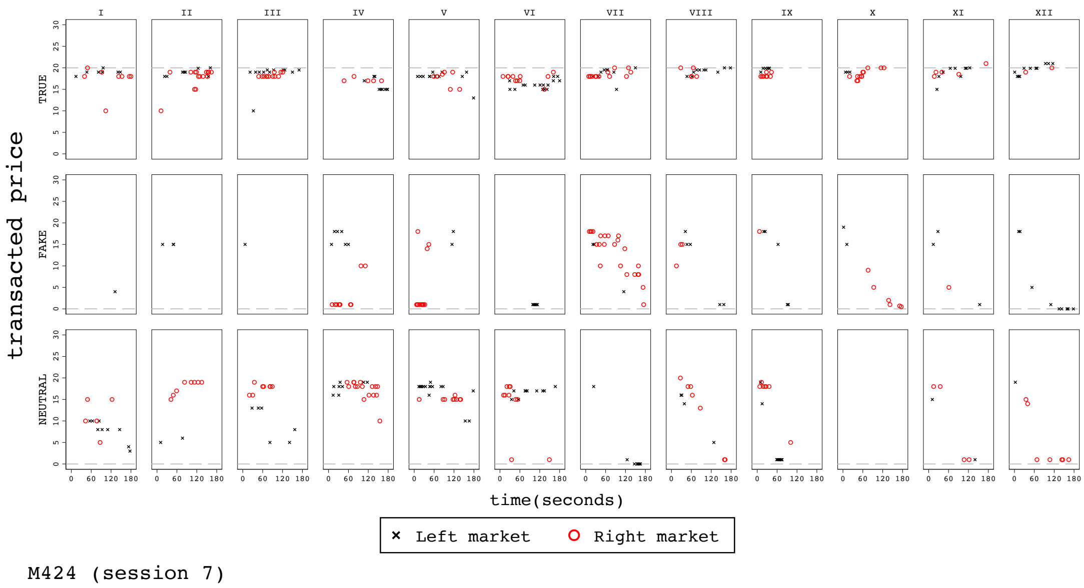

Figure A7: Session 7 (M424 treatment).

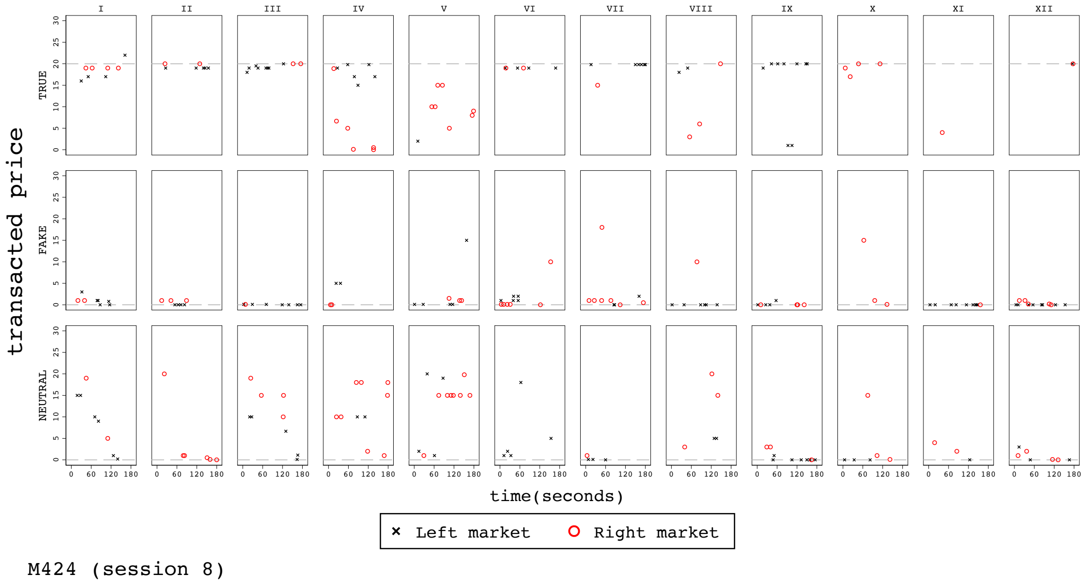

Figure A8: Session 8 (M424 treatment).

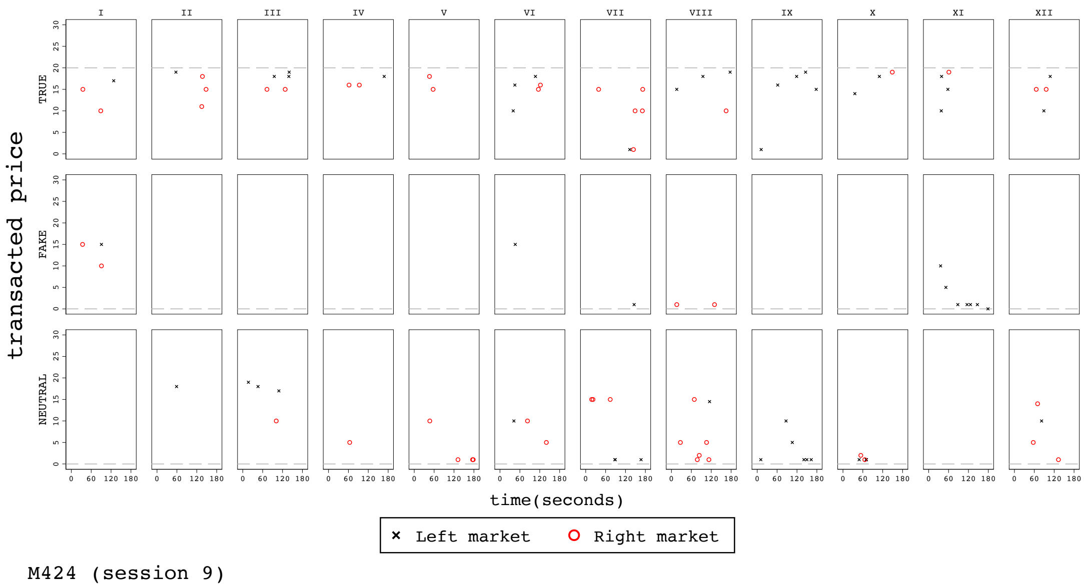

Figure A9: Session 9 (M424 treatment).

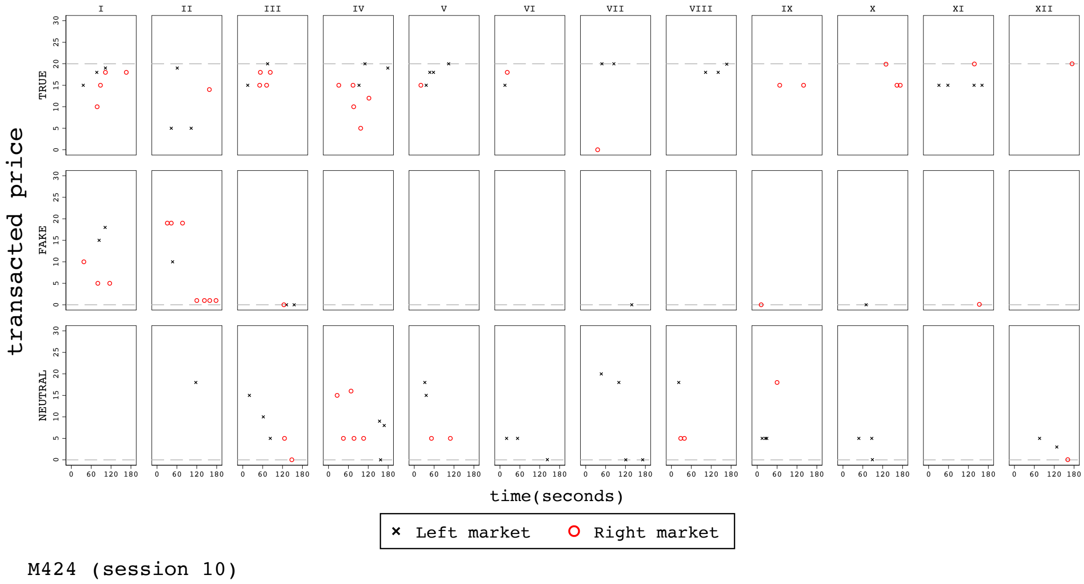

Figure A10: Session 10 (M424 treatment).

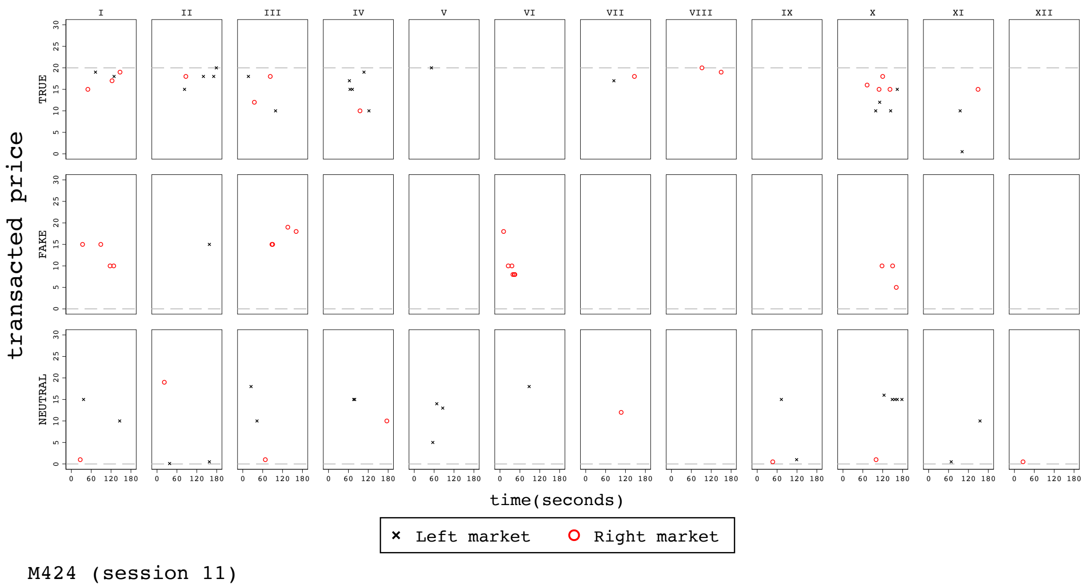

Figure A11: Session 11 (M424 treatment).

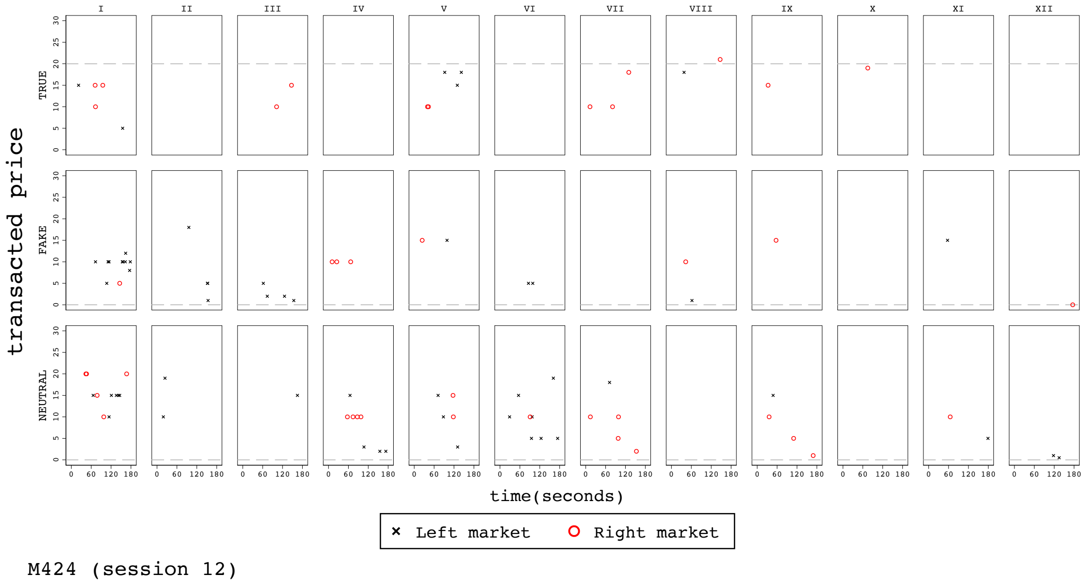

Figure A12: Session 12 (M424 treatment).

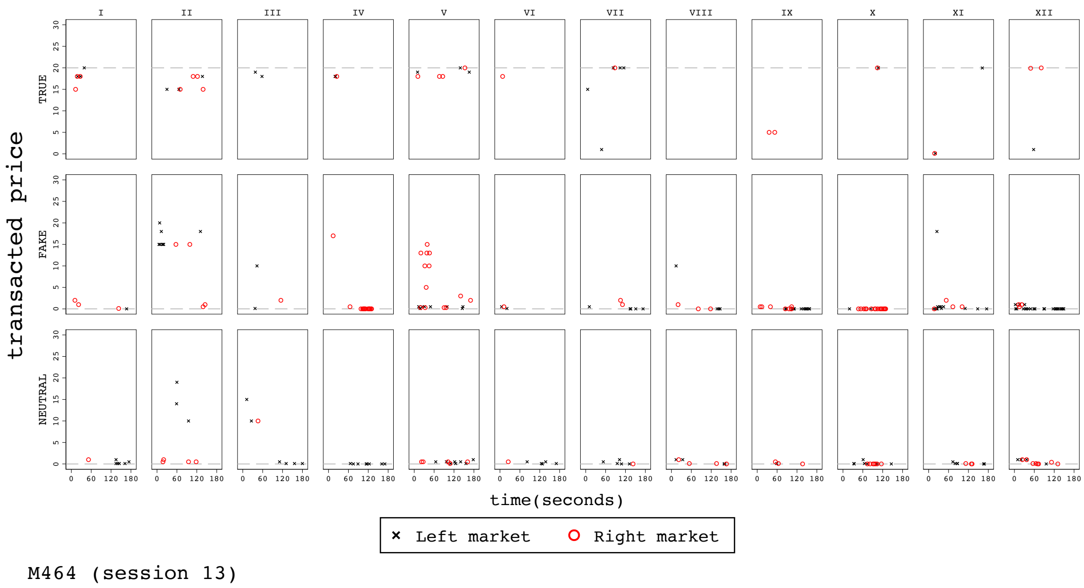

Figure A13: Session 13 (M464 treatment).

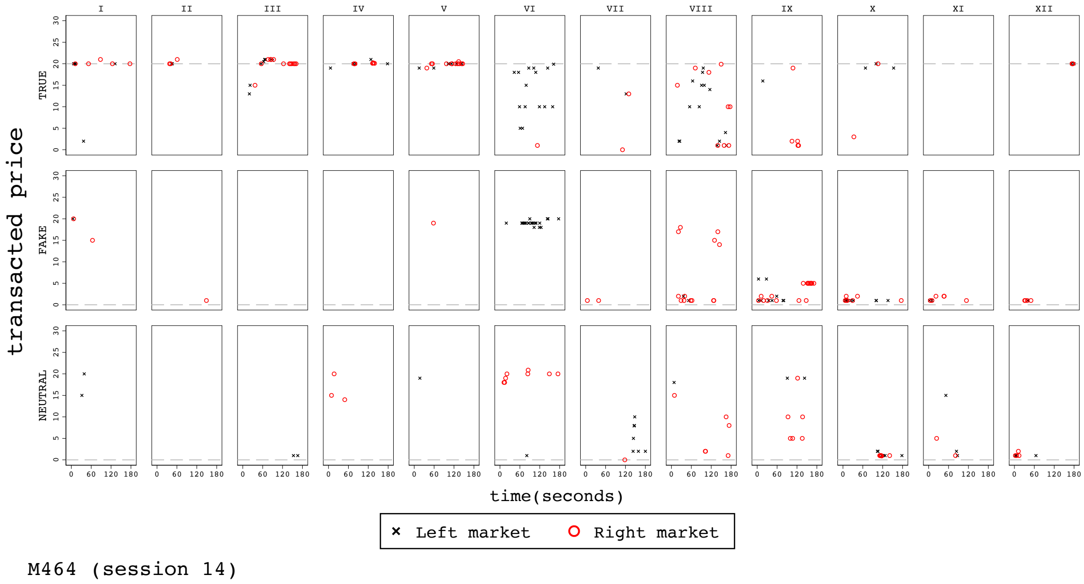

Figure A14: Session 14 (M464 treatment).

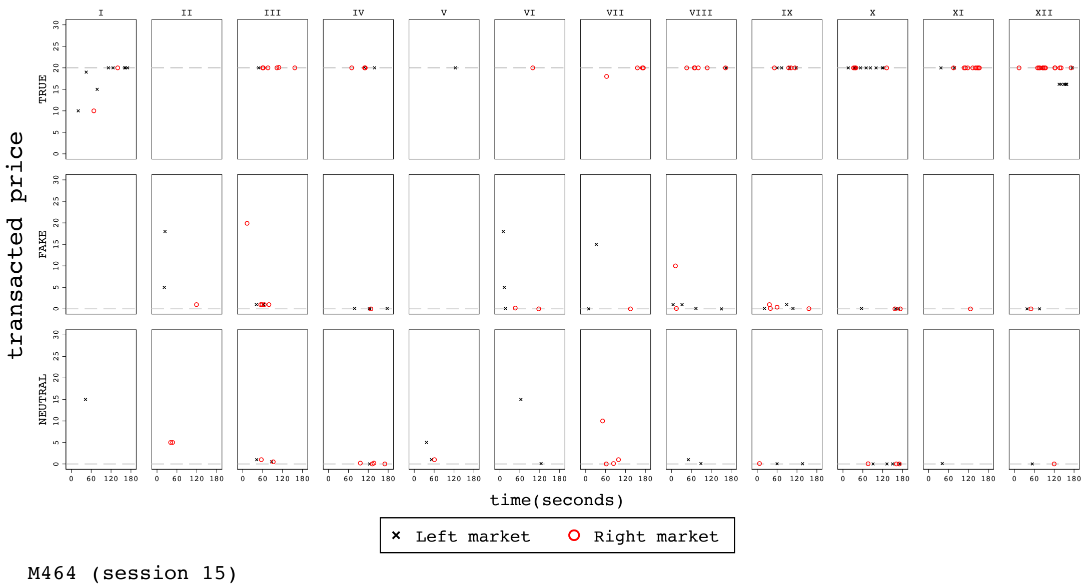

Figure A15: Session 15 (M464 treatment).

Figure A16: Session 16 (M424 treatment).

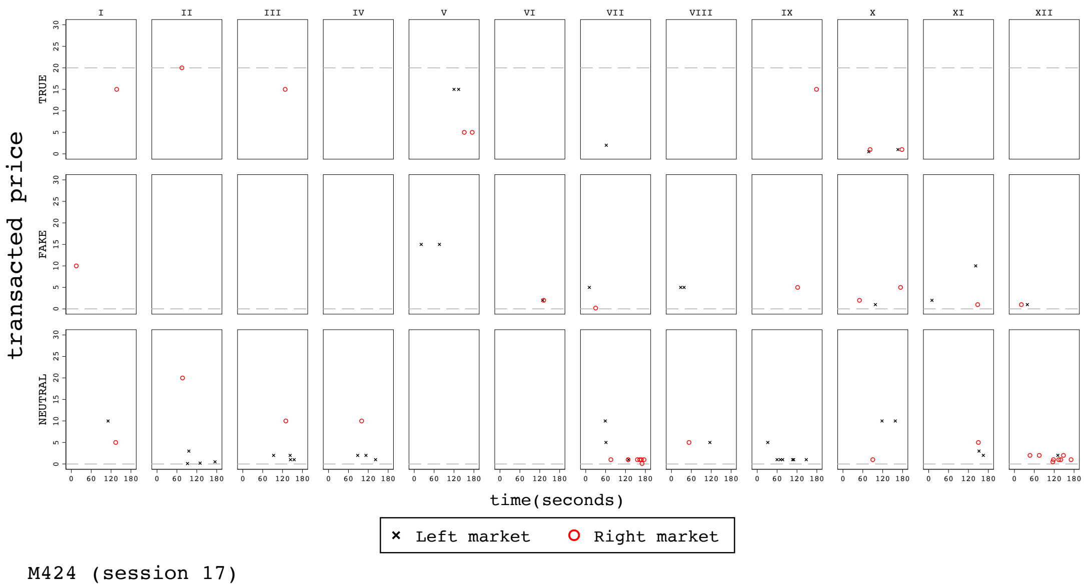

Figure A17: Session 17 (M424 treatment).

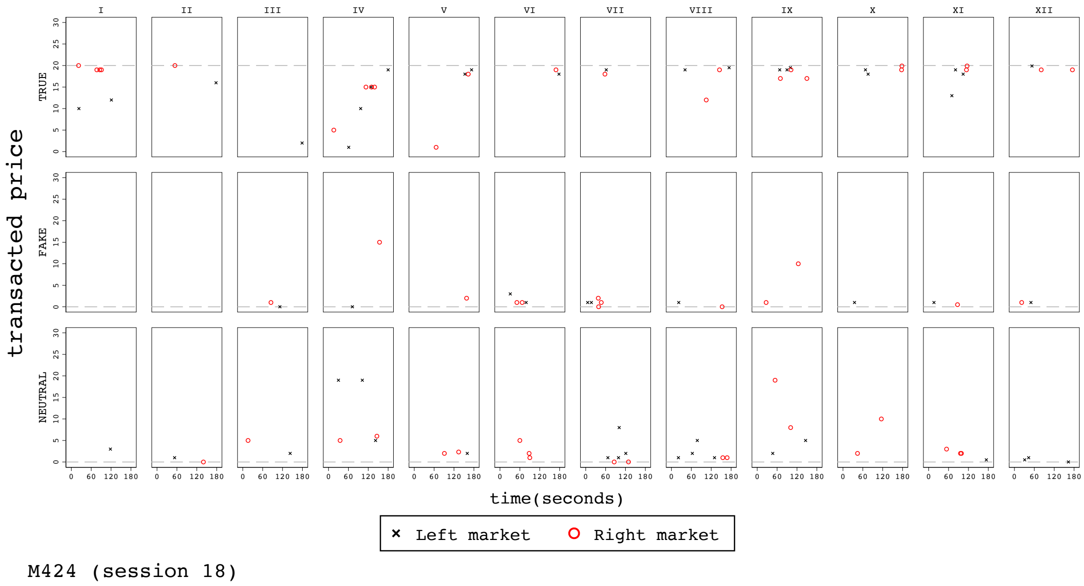

Figure A18: Session 18 (M424 treatment).

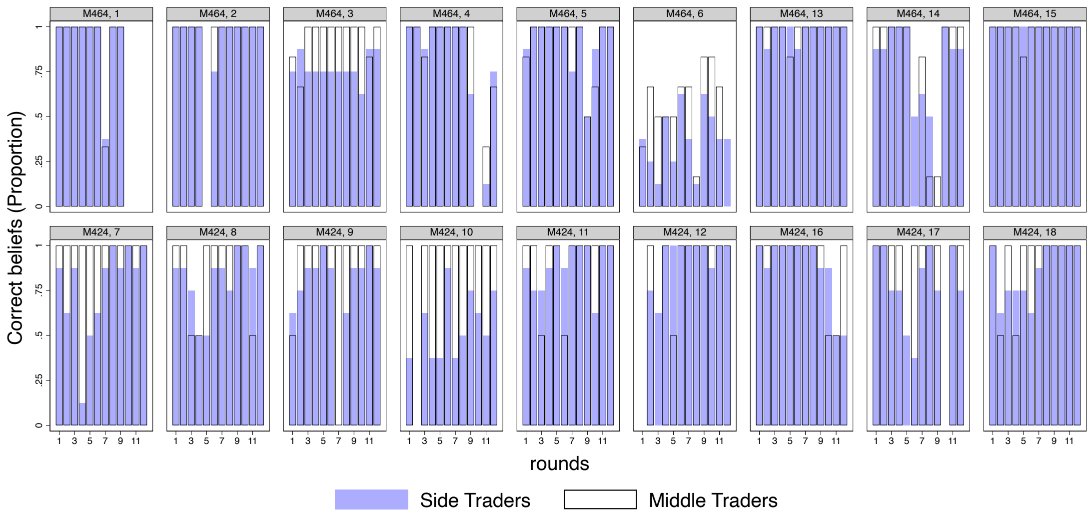

Figure A19: Proportion of side and middle traders that submitted correct ex-post market beliefs.

Dependent variable: $\Delta = MAD^{FRE} - MAD^{PIE}$

	EARLY rounds				LATE rounds			
Block length	45sec	30sec	20sec	10sec	45sec	30sec	20sec	10sec
Panel A: M464 treatment								
Constant	-3.414^{***}	-3.429^{***}	-3.539^{***}	-3.592^{***}	-4.394^{***}	-4.648^{***}	-4.662^{***}	-4.730^{***}
	(0.548)	(0.548)	(0.545)	(0.548)	(0.552)	(0.543)	(0.546)	(0.527)
n	52				53			
Panel B: M424 treatment								
Constant	0.343	0.388	0.347	0.325	-3.418^{***}	-3.496^{***}	-3.532^{***}	-3.539^{***}
	(0.532)	(0.545)	(0.559)	(.552)	(0.587)	(0.576)	(0.577)	(0.576)
n	54				54			

Note. Each column reports a fixed-effects regression of Δ on a constant. Significance levels are denoted by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Table A7: Are closing market pricing closer to the FRE or PIE benchmark? Fixed-effects regression estimates (within-treatment)

A.2. SUPPLEMENTARY DATA ANALYSIS

The supplementary data analysis documents empirical results that are not reported in full in the main manuscript.

A.2.1. Information aggregation

Given the experimental data, we consider the block interval specifications of 10-, 20-, 30-, and 45-seconds. For each interval specification, we compute mean absolute deviation (MAD) of **closing market prices** from the FRE (MAD^{FRE}) and (MAD^{PIE}) benchmarks. To compare whether closing market prices are closer to the FRE or PIE, we $\Delta = MAD^{FRE} - MAD^{PIE}$ on a constant (within-treatment analysis). Here, a negative Δ implies that closing prices are closer to the FRE. The fixed-effects regression estimates are reported on Table A7 and show that closing market prices, for all interval specifications, are significantly closer to the FRE for the EARLY and LATE rounds of the M464 treatment. For the M424 treatment, closing prices (all specifications) are significantly closer to the FRE in the LATE rounds—no significant differences in the EARLY rounds. We also make between-treatment comparison of MAD^{FRE} for each interval specification. The random-effects regression estimates on Table A8 find no significant between-treatment differences ($p \geq .389$) in the LATE rounds. In contrast, MAD^{FRE} are often lower in the M464 ($p \leq .051$) in the EARLY rounds. In all of the above analysis, we note that the estimates do not vary substantially with changes in the interval specification.

Dependent variable: MAD^{FRE}

	EARLY rounds				LATE rounds			
Block length	45sec	30sec	20sec	10sec	45sec	30sec	20sec	10sec
Reference group: M464								
M424	2.025**	2.035*	2.133**	2.151**	0.601	0.677	0.686	0.706
	(1.010)	(1.041)	(1.041)	(1.044)	(0.832)	(0.801)	(0.810)	(0.806)
Constant	4.280***	4.277***	4.245***	4.209***	3.131***	3.023***	3.024***	2.984***
	(0.717)	(0.739)	(0.739)	(0.741)	(0.590)	(0.568)	(0.574)	(0.572)
n	106				107			

Note. Significance levels are denoted by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Table A8: Between-treatment comparison of MAD^{FRE}. Random-effects regression estimates.

Dependent variable: MAD^{FRE}

	EARLY rounds				LATE rounds			
Block length	45sec	30sec	20sec	10sec	45sec	30sec	20sec	10sec
$\hat{\lambda}_{464}$	0.715	0.473	0.343	0.172	0.758	0.434	0.322	0.160
$\hat{\lambda}_{424}$	0.346	0.146	0.122	0.045	0.421	0.304	0.201	0.100
Difference p-value	0.116	0.024**	0.028**	0.009***	0.216	0.473	0.262	0.213

Note. The above cells reported the estimate λ for the M464 and M424 treatment. The final row details the fixed-effects regression between-treatment difference in λ. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Table A9: Rate of convergence in the M424 and M464 treatments.

A.2.2. Rate of convergence

We study the rate of convergence for the different interval specifications. We model the decay rate with

$$MAD_t^{FRE} = \alpha e^{-\lambda(t-0.5)}, \qquad t = 1, \ldots, T$$

where T here refers to the final block in the interval specification, $\alpha = MAD_0^{FRE}$, and λ is the rate of decay. We use a fixed-effects regression to estimate to between-treatment comparisons of λ. The first two rows of Table A9 report the estimated λ by treatments and interval specifications. The final row details the p-value for the between-treatment comparison of λ. In the LATE rounds, we do not see any significant between-treatment differences ($p \geq .213$) for each interval specification.

Appendix B. EXPERIMENT INSTRUCTIONS

The experiment was conducted in Chinese. Below is the English translated version of the instructions. Parts of the instructions that are unique to the **M464** treatment are denoted in *text*, and those unique to the **M424** treatment are denoted using text.

Part I of the experiment is identical across both treatments, and Part II only differs on the number of middle traders. Throughout the instructions we use the terms blue and red traders to denote the side and middle traders, respectively.

B.1. INTRODUCTION

First, thank you for taking part in this experiment. Please remain quiet during the experiment and do not talk to other participants. If you have any questions, please raise your hand and a member of the experiment team will come to assist you.

To avoid disturbing others or being disturbed, please switch off your mobile phone and put it in your bag. Please also put away any items unrelated to today's experiment, including books, notebooks, and other electronic devices. Do not use any other functions on the computer, as this may cause it to crash or the programme to malfunction. Please note: if you violate these rules, you will lose the opportunity to participate in today's study and will not receive any payment.

All of your decisions in today's experiment are anonymous. You will not receive any information about other participants, and they will not receive any information about you. In addition, we will pay you individually at the end of the experiment. This means that other participants will not know what decisions you made, nor how much you earned.

Please read the instructions carefully, as your payment will depend on how well you understand the experiment. If you read and understand the instructions carefully, you can expect to earn a good amount in today's study. In addition, you will receive a 10 RMB "show-up fee".

Today's experiment consists of two parts (Part 1 and Part 2). Your total earnings today will depend on the decisions you make in both parts.

We will first distribute the instructions for Part 1. After Part 1 is finished, we will distribute the instructions for Part 2.

If you have any questions while reading the instructions, please raise your hand and we will help you.

B.2. INSTRUCTIONS FOR PART 1

Part 1 of the experiment consists of one practice decision and four paid decisions. The practice decision is designed to help you understand today's experiment; the choice you make in the practice decision will have no effect on your earnings today. Your earnings from Part 1 will depend on the choices you make in the four paid decisions.

In each round, you will face the same type of problem. In addition, each round is independent. This means that the decision you make in one round will not affect any other part of the experiment.

In Part 1, you will be placed in a group with 13 other randomly selected participants (so there are 14 participants in each group). The group composition in Part 1 is fixed. That is, you will complete all of Part 1 together with the same 13 participants.

In each round, you will take part in a simulated trading market. Your earnings will be expressed in experimental points. The number of points you earn in each round depends on the decisions you make in that round.

At the end of Part 1, the computer will randomly select one of the four paid rounds to determine your earnings from Part 1. The points you earn in that round will be converted into RMB according to the following exchange rate:

$$1 \text{ point} = 0.09 \text{ RMB}$$

Next, we will describe the details of each round. Please read them carefully. You will also be asked to complete a short comprehension quiz at the end.

B.2.1. Section 1: Random Ball Draw

There are three balls in an opaque box: an X ball, a Y ball, and a Z ball.

- The computer will randomly draw one ball. Each ball has an equal probability of being drawn, i.e. 1/3.
- The 14 members of your group will then be randomly divided into two subgroups, each containing 7 participants.
- After the grouping, the computer will send each 7-person subgroup a hint about which ball was drawn. Note that the hints received by the two subgroups are always different.

Example 1: Suppose the computer draws the X ball. One subgroup receives the hint "Not Y" (so the ball could be X or Z). The other subgroup receives the hint "Not Z" (so the ball could be X or Y).
Example 2: Suppose the computer draws the Y ball. One subgroup receives the hint "Not X". The other subgroup receives the hint "Not Z".
Example 3: Suppose the computer draws the Z ball. One subgroup receives the hint "Not X". The other subgroup receives the hint "Not Y".

From the examples above, you can see that the two subgroups always receive different hints. In addition, the hint you receive will help you make an informed guess. A hint rules out one incorrect "answer". As a result, the remaining two possibilities each become equally likely (50%).

B.2.2. Section 2: Trading Stage

Once both 7-person subgroups have received their hints, the trading stage begins. In this stage, you may trade six different asset certificates: X1, Y1, Z1, and X2, Y2, Z2.

Asset certificates X1, Y1, Z1 can be traded only in Market 1. Similarly, asset certificates X2, Y2, Z2 can be traded only in Market 2. Each participant may trade in both Market 1 and Market 2 simultaneously.

- At the start of trading, each participant receives the following initial endowment:
  - 200 points
  - 3 units of X1
  - 3 units of Y1
  - 3 units of Z1
  - 3 units of X2
  - 3 units of Y2
  - 3 units of Z2
- You may buy and sell asset certificates on the trading platform. The trading period lasts 3 minutes. You may use your points to buy additional certificates, or sell certificates that you hold. The next section will explain in detail how to trade on the platform.
- After the three-minute trading window ends, trading stops. The system then redeems all asset certificates you hold (X1, Y1, Z1, and X2, Y2, Z2). In other words, the system "buys back" your certificates at specific redemption prices. The redemption price of each certificate depends on which ball the computer drew.
- Table B1 below shows the exact redemption prices for each asset certificate under each possible ball draw. Naturally, during the trading period, you will not know these redemption prices. The system will reveal them at the end of each round.

Table B1: Redemption values of asset certificates

Market	Asset	Ball X drawn	Ball Y drawn	Ball Z drawn
Market 1	X1	20 points	0 points	0 points
	Y1	0 points	20 points	0 points
	Z1	0 points	0 points	20 points
Market 2	X2	20 points	0 points	0 points
	Y2	0 points	20 points	0 points
	Z2	0 points	0 points	20 points

Example 1: Suppose the computer draws the X ball. Then the redemption value of asset certificates X1 and X2 is 20 points. The redemption value of all other certificates—Y1, Y2, Z1, and Z2—is 0.
Example 2: Suppose the computer draws the Y ball. Then the redemption value of asset certificates Y1 and Y2 is 20 points. The redemption value of all other certificates—X1, X2, Z1, and Z2—is 0.
Example 3: Suppose the computer draws the Z ball. Then the redemption value of asset certificates Z1 and Z2 is 20 points. The redemption value of all other certificates—X1, X2, Y1, and Y2—is 0.

In summary, your earnings from the trading stage are calculated as follows:

Trading-stage earnings = (Remaining funds) + (Number of certificates held × Corresponding redemption value)

Example 2: Suppose that at the end of the trading stage, you still have 180 points and hold the following quantities of asset certificates:

	Market 1			Market 2		
Asset type	X1	Y1	Z1	X2	Y2	Z2
Quantity held	3	1	4	2	5	0

- If the drawn ball is X, your earnings from the trading stage are: 180 + (3×20) + (2×20) = 280
- If the drawn ball is Y, your earnings from the trading stage are: 180 + (1×20) + (5×20) = 300
- If the drawn ball is Z, your earnings from the trading stage are: 180 + (4×20) + (0×20) = 260

Below, we describe in detail how the trading platform works and how to buy and sell asset certificates on it.

Figure B1: Your current funds are shown in the upper-right corner of the screen. This amount updates continuously as you trade. In addition, six equally sized "boxes" are displayed on the screen. Each box shows the trading prices for a specific asset certificate and the quantity you currently hold.

Figure B2: If you want to buy or sell an asset certificate, you must enter a bid price (buy order) or an ask price (sell order) in the corresponding box.

- When you enter a bid price, you indicate to other market participants how many points you are willing to pay to buy that certificate. The system always displays your most recent bid. The minimum bid is 0, the maximum is 30, and prices may be entered with up to two decimal places.
- When you enter an ask price, you indicate how many points you are willing to accept to sell that certificate. The system always displays your most recent ask. The minimum ask is 0, the maximum is 30, and prices may be entered with up to two decimal places.

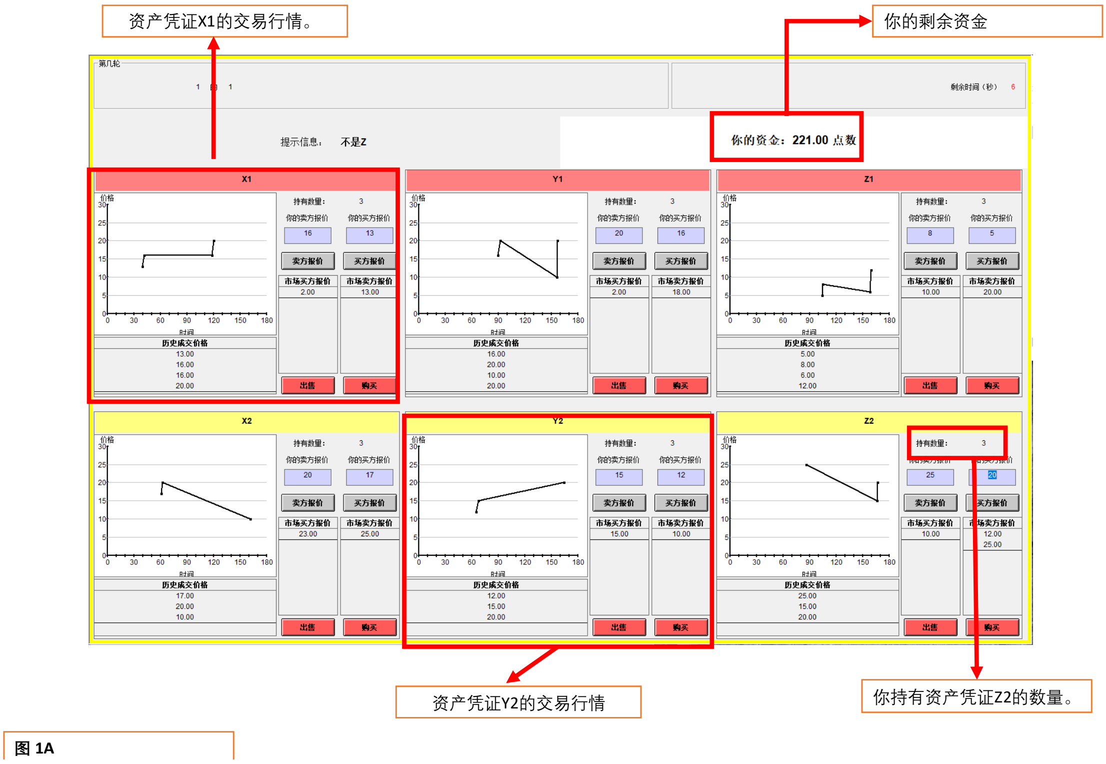

Figure B1: Screenshot.

Figure B2: Screenshot.

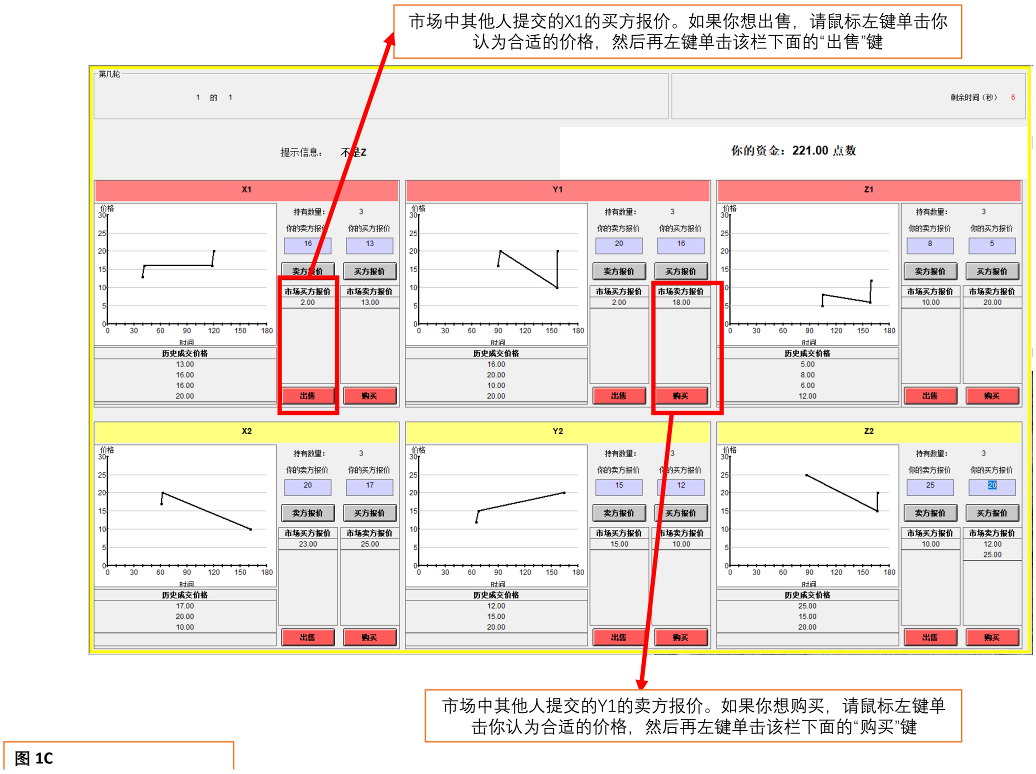

Figure B3: Screenshot.

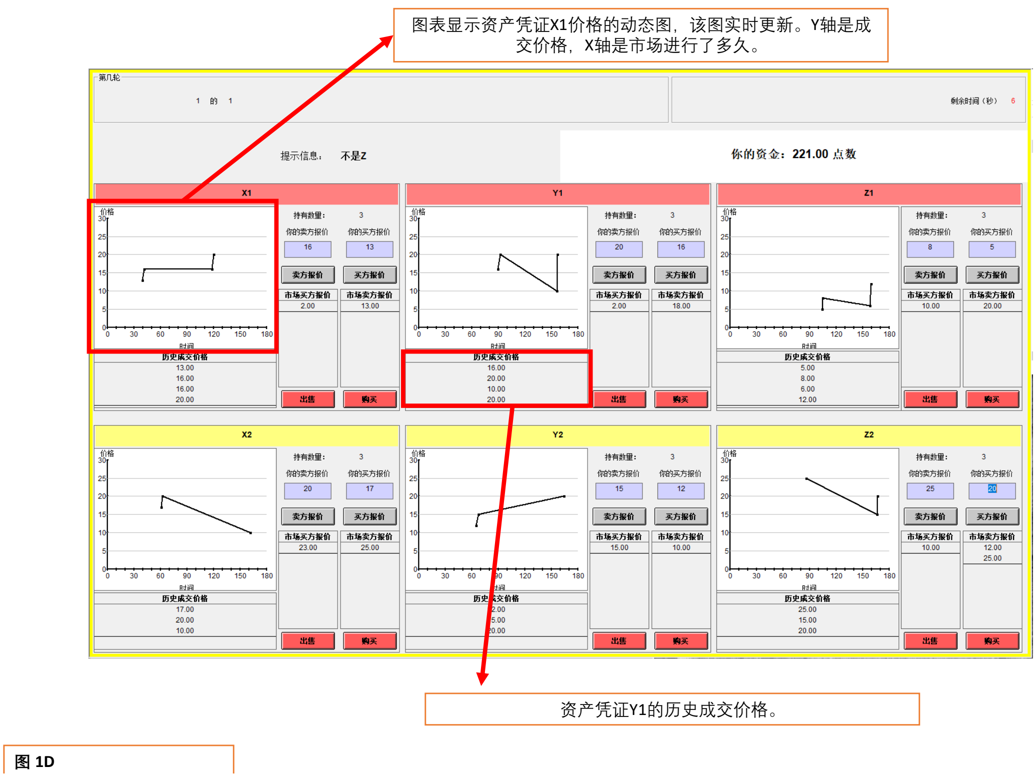

Figure B4: Screenshot.

Figure B3: The bid prices and ask prices submitted by other participants are displayed in two separate lists: "Market Bid Price" and "Market Ask Prices".

- How to buy a certificate: Left-click a price you find acceptable in the "Market Ask Prices" list, then left-click the "Buy" button below that list.
- How to sell a certificate: Left-click a price you find acceptable in the "Market Bid Prices" list, then left-click the "Sell" button below that list.

Figure B4: A real-time price chart displays the transaction prices of the certificate. You may also view the same information in the "Past Transaction Prices" list below the chart.

Please note the following trading rule: You cannot trade with yourself. Your own bid and ask prices appear in blue in the market price lists.

B.2.3. Section 3: Guessing Stage

After the trading stage ends, and before the computer reveals the redemption values of the asset certificates, you must guess which ball the computer actually drew. The screen will display the following question: "Which ball did the computer randomly select?"

You will have four options:

- I don't know
- The computer selected the X ball
- The computer selected the Y ball
- The computer selected the Z ball

Your earnings from the guessing stage are determined as follows:

- If you choose "I don't know", you receive 5 points.
- If you make the correct guess, you receive 20 points.
- If you make the wrong guess, you lose 20 points.

Examples:

- If the computer drew the X ball and you chose "The computer selected the X ball", then your earnings from the guessing stage are 20 points.
- If the computer drew the X ball and you chose "The computer selected the Z ball", then your earnings from this stage are -20 points.
- If the computer drew the X ball but you chose "I don't know", then your earnings from this stage are 5 points.

B.2.4. Total Earnings for Each Round

Your total earnings in each round are calculated as follows:

Total earnings = (Earnings from the trading stage) + (Earnings from the guessing stage)

B.2.5. Additional Information

To ensure that all participants fully understand the experiment, you must complete the short comprehension quiz below. The experiment will begin only after all participants have answered the quiz correctly. Part 1 consists of five rounds in total. Throughout these five rounds, the 14 members of your group remain fixed. In other words, you will complete all of Part 1 together with the same 13 participants.

B.2.6. Comprehension Quiz

1. If the hint you receive is "Not Z", what is the probability that the computer selected the Y ball? ___

   Ans: 50%

2. If the computer selected the Y ball, both 7-person subgroups receive the same hint. (True/False)

   Ans: False

3. If the computer selected the Y ball and your 7-person subgroup receives the hint "Not Z", then what hint will the other subgroup receive?

   Ans: Not X

4. If the computer selected the X ball, what are the redemption values of asset certificates X1 and X2?

   Ans: 20 points

5. If the computer selected the Y ball, what are the redemption values of asset certificates X1 and X2?

   Ans: 0 points

6. If the computer selected the X ball, and at the end of the trading stage you have 200 points remaining and hold the following quantities of asset certificates as in Table B2. So, how many points do you earn in the trading stage?

Table B2: Porfolio

	Market 1			Market 2		
Asset type	X1	Y1	Z1	X2	Y2	Z2
Quantity held	2	3	4	1	2	3

Ans: 260 points

7. If the computer selected the Y ball, and at the end of the trading stage you have 200 points remaining and hold the following quantities of asset certificates as in Table B3. So, how many points do you earn in the trading stage?

Ans: 240 points

Table B3: Porfolio

	Market 1			Market 2		
Asset type	X1	Y1	Z1	X2	Y2	Z2
Quantity held	2	1	6	3	1	2

8. If the computer selected the X ball and you choose option "I don't know" in the guessing stage, how many points do you earn in the guessing stage?

Ans: 5 points

9. If the computer selected the X ball and you choose option "The computer selected the X ball" in the guessing stage, how many points do you earn in the guessing stage?

Ans: 20 points

10. If the computer selected the X ball and you choose option "The computer selected the Y ball" in the guessing stage, how many points do you earn in the guessing stage?

Ans: -20 points

B.3. PART 2

Part 2 of the experiment consists of one practice decision and twelve paid decisions.

In each round, you will face the same type of problem. In addition, each round is independent. This means that the decision you make in one round will not affect any other part of the experiment.

In Part 2, you will be grouped with the same 13 participants who were in your group in Part 1. Thus, your group members are the same in both Part 1 and Part 2. Moreover, the same 14 participants will complete all twelve rounds of Part 2 together; there will be no change in group composition during the experiment.

In each round, you will again take part in a simulated trading market. Your earnings will be expressed in experimental points. The number of points you earn in each round depends on the decisions you make in that round.

At the end of Part 2, the computer will randomly select three of the twelve paid rounds to determine your earnings from Part 2. The points you earn in those rounds will be converted into RMB according to the following exchange rate:

$$1 \text{ point} = 0.09 \text{ RMB}$$

Next, we will describe the details of each round. Please read them carefully. You will also be asked to complete a short comprehension quiz at the end.

Some stages in Part 2 are the same as in Part 1, while others are different. Please read the instructions carefully.

B.3.1. Section 1: Assignment of Roles

Each group member will be randomly assigned one of the following roles: a) Blue Trader, or b) Red Trader.

In each round of the experiment, there are:

- 8 Blue Traders, and
- *6 Red Traders* (2 Red Traders)

At the beginning of each round, the system will inform you which role you have been assigned.

Important: Your role remains fixed throughout Part 2. This means that if you are assigned the role of Blue Trader in Round 1, you will remain a Blue Trader for all remaining 11 rounds. Similarly, if you are assigned the role of Red Trader in Round 1, you will remain a Red Trader for all remaining 11 rounds.

B.3.2. Section 2: Random Ball Draw

There are three balls in an opaque box: an X ball, a Y ball, and a Z ball. The computer will randomly draw one ball. Each ball has an equal probability of being drawn.

- After the programme draws a ball, the 8 Blue Traders will be randomly divided into two groups (4 traders per group).
- The programme will then send each group a hint.
- The two Blue groups always receive different hints.
- The *6 Red* (2 Red) Traders receive no hint.

Example 1: Suppose the programme draws the X ball. One Blue group receives the hint "Not Y". The other Blue group receives the hint "Not Z". Red traders receive no information.
Example 2: Suppose the programme draws the Y ball. One Blue group receives the hint "Not X". The other Blue group receives the hint "Not Z". Red traders receive no information.

From the examples above, you can see that the two Blue groups always receive different hints. In addition, the hint you receive helps you update your beliefs. A hint rules out one incorrect option, so the remaining two options each become 50% likely. In other words, before receiving the hint, all three balls are equally likely from your perspective.

Note: Red Traders receive no hint. Therefore, for Red Traders, the probabilities of X, Y, and Z remain equally likely.

B.3.3. Section 3: Trading Stage

After all participants have received their information (if any), the trading stage begins. In this stage, you may buy and sell the assets you hold. However, the types of assets you can trade and the markets you may enter (Market 1 or Market 2) depend on your role.

As described earlier, the 8 Blue Traders are randomly divided into two groups of 4.

- One Blue group may trade only in Market 1,
- and the other Blue group may trade only in Market 2.

The *6 Red* (2 Red) Traders may trade in both Market 1 and Market 2.

- Blue Traders (Market 1 Group): You may trade only X1, Y1, and Z1 in Market 1—See FigureB5. Each participant in this group starts with:
  - 200 points
  - 6 units of X1
  - 6 units of Y1
  - 6 units of Z1
- Blue Traders (Market 2 Group): You may trade only X2, Y2, and Z2 in Market 2—See FigureB6. Each participant in this group starts with:
  - 200 points
  - 6 units of X2
  - 6 units of Y2
  - 6 units of Z2
- Red Traders (Market 1 and Market 2): You may trade: X1, Y1, Z1 in Market 1, and X2, Y2, Z2 in Market 2—See FigureB7. Each Red Trader starts with:

Figure B5: Screenshot.

- 200 points
- 3 units of X1
- 3 units of Y1
- 3 units of Z1
- 3 units of X2
- 3 units of Y2
- 3 units of Z2

Important Notes

- Members of the Market 1 Blue group all receive the same hint. Similarly, members of the Market 2 Blue group receive the same hint.
- Red Traders do not receive any information.
- Blue Traders in Market 1 cannot participate in Market 2 and cannot see prices in Market 2.
- Likewise, Blue Traders in Market 2 cannot participate in Market 1 and cannot see prices in Market 1.

Figure B6: Screenshot.

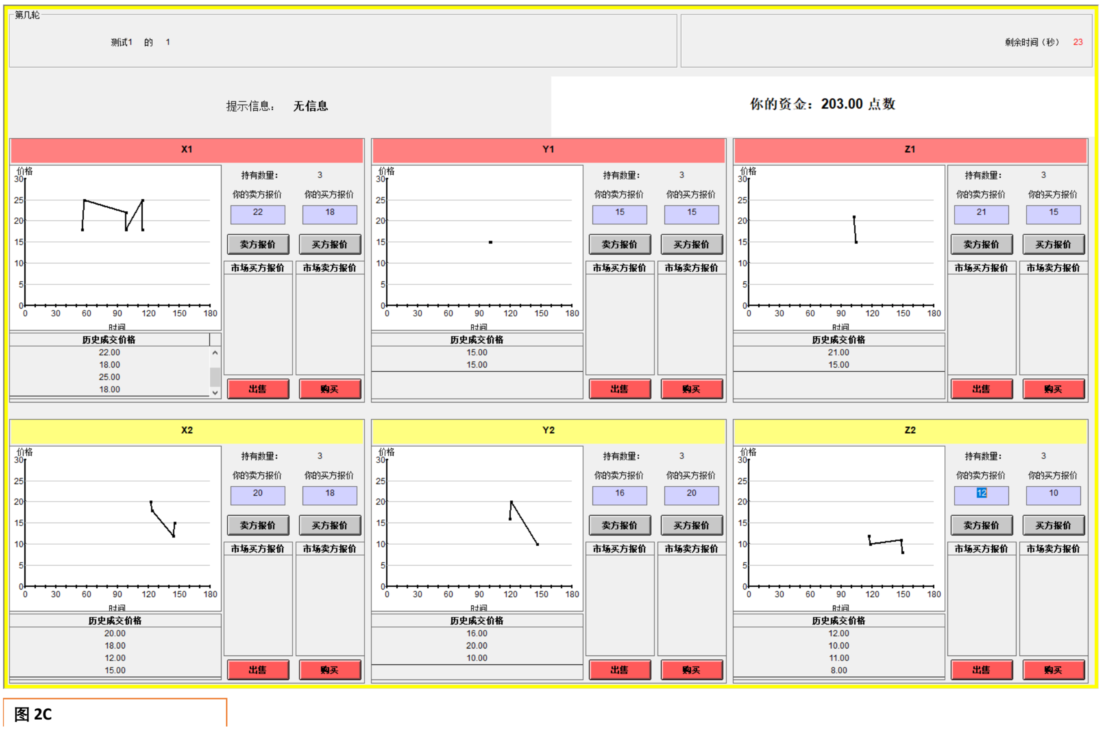

Figure B7: Screenshot.

B.4. TRADING PROCESS

You will trade on the same trading platform used in Part 1. You may sell certificates to obtain funds and use funds to buy certificates. Each round provides 180 seconds (3 minutes) for trading.

After three minutes, trading stops. The system will redeem all certificates you hold (X1, Y1, Z1, X2, Y2, Z2). The redemption value of each certificate depends on which ball the programme drew.

Table B4 lists the redemption values corresponding to each possible draw. Naturally, you will not know the redemption values during trading. The system reveals them only at the end of each round. Example 1: Suppose the computer draws the X ball. Then the redemption value of asset certificates X1 and X2 is 20 points. The redemption value of all other certificates—Y1, Y2, Z1, and Z2—is 0.

Table B4: Redemption values of asset certificates

Market	Asset	Ball X drawn	Ball Y drawn	Ball Z drawn
Market 1	X1	20 points	0 points	0 points
	Y1	0 points	20 points	0 points
	Z1	0 points	0 points	20 points
Market 2	X2	20 points	0 points	0 points
	Y2	0 points	20 points	0 points
	Z2	0 points	0 points	20 points

Example 2: Suppose the computer draws the Y ball. Then the redemption value of asset certificates Y1 and Y2 is 20 points. The redemption value of all other certificates—X1, X2, Z1, and Z2—is 0.

Example 3: Suppose the computer draws the Z ball. Then the redemption value of asset certificates Z1 and Z2 is 20 points. The redemption value of all other certificates—X1, X2, Y1, and Y2—is 0.

In summary, your earnings from the trading stage are calculated as follows:

Trading-stage earnings = (Remaining funds) + (Number of certificates held × Corresponding redemption value)

B.5. SECTION 4: GUESSING STAGE

After the trading stage ends, and before the computer reveals the redemption values of the asset certificates, you must guess which ball the computer actually drew. The screen will display the following question: "Which ball did the computer randomly select?"

You will have four options:

- I don't know
- The computer selected the X ball
- The computer selected the Y ball

- The computer selected the Z ball

Your earnings from the guessing stage are determined as follows:

- If you choose "I don't know", you receive 5 points.
- If you make the correct guess, you receive 20 points.
- If you make the wrong guess, you lose 20 points.

B.5.1. Additional Information

To ensure that all participants fully understand the experiment, you must complete the comprehension quiz below. The experiment will begin only after all participants have answered the quiz correctly.

Part 2 consists of 12 rounds. Throughout these 12 rounds, the 14 members of your group remain fixed. In other words, you will complete all of Part 2 with the same 13 participants.

Note:

- The *6 Red* (2 Red) Traders remain the same throughout Part 2.
- The 8 Blue Traders also remain the same participants, but in each round the system randomly divides them again into two Blue groups of 4.
- The composition of these 4-person Blue groups may differ from round to round.

B.5.2. comprehension quiz

1. If your hint is "Not Z", what is the probability that the programme selected the Y ball?

   Ans: 50%

2. If the programme selected the Y ball, the two Blue 4-person groups receive the same hint. (True/False)

   Ans: False

3. If the programme selected the Y ball and your Blue 4-person group receives the hint "Not Z", then what hint will the other Blue 4-person group receive?

   Ans: Not X

4. All participants freely trade all six types of asset certificates. (True/False)

   Ans: False