Information Aggregation in Arrow-Debreu Markets: An Experiment*

Lawrence Choo   Todd R. Kaplan   Ro'i Zultan[†]

Abstract

Studies of experimental and betting markets have shown that markets are able to efficiently aggregate information dispersed over many traders. We study information aggregation in Arrow-Debreu markets using a novel information structure. Compared to previous studies, the information structure is more complex, allows for heterogeneity in information among traders – which provides insights into the way in which information is gradually disseminated in the market – and generates situations in which all traders hold identical beliefs over the traded assets' values, thus providing a harsh stress test for belief updating. We find little evidence for information aggregation and dissemination in early rounds. Nonetheless, after traders gain experience with the market mechanism and structure, prices converge to reveal the true state of the world. Elicited post-market beliefs reveal that markets are able to efficiently aggregate dispersed information even if individual traders remain uninformed, consistent with the marginal trader hypothesis.

1 Introduction

The idea that markets can aggregate dispersed private information into prices is central to much of economic thinking.[1] When some traders are fully informed about the true state of the world, competitive bidding will push prices to reflect the true state, thus revealing the information to uninformed traders. A more challenging case is when each trader only holds partial information about the true state. If traders understand how the different *pieces* of information affect behaviour, the movement of prices eventually reveals the information held by other traders. Furthermore, if traders continuously condition their market demands on their own information, then market prices should continuously move until no trader can learn any more from the prices.[2] Naturally, whether equilibrium market prices fully reveal the true state depends on the information structure. Radner (1979) showed that prices converge to an equilibrium as if each trader openly and truthfully communicates his own private information. This result, however, requires that traders hold diverse or *heterogeneous* partial private information and that the true state is in the unique element within the intersect of at least two different private information.[3] This analysis equates the states of the world with beliefs over the traded assets' possible valuations. It remains an open question, therefore, whether market prices are able to fully reveal the true state even when traders hold identical or *homogeneous* partial private information, in the sense that it induces the same beliefs over asset valuations. This paper studies the information aggregation properties of markets in a novel design that allows us to clearly distinguish between situations where traders hold homogeneous and heterogeneous partial private information about the true state of the world.

Empirical studies suggest that the intricate process of information aggregation (e.g., learning about the information of others through price movements) postulated by theory often places too many demands on the abilities of the traders (see De-Bondt and Thaler, 1985; Thaler, 1988; Hirshleifer, 2001; Shleifer, 2000). However, identifying 'who knows what'

---

[1] The idea often builds upon the early discussions by Hayek (1945) and was subsequently development further by Grossman (1976). It is also the central theme of the "efficient market hypothesis" (Fama, 1970).

[2] Alternatively, some traders could seek to learn more about the true event at a cost. However, Grossman and Stiglitz (1980) challenged the notion that traders will be willing to incur such costs if market prices will eventual reveal what they learnt.

[3] For example, consider the state space $\{1, 2, 3, 4, 5, 6\}$, where the true state is 4. Suppose that trader A knows that the true state is in $\{1, 2, 3, 4\}$ and trader B knows that the true state is in $\{4, 5, 6\}$. Combining both pieces of information through open communication, both traders will realise that the true state must be 4. Conversely, if trader B's partition is such that she knows the true state to be in $\{3, 4, 5\}$, then the combined information can only reveal that the true state is either 3 or 4.

is often challenging in empirical market research. Controlled experimental environments complete the picture by studying market behaviour in a highly structured environment, where the experiment designer controls the information structure. Laboratory experiments study the general principles underpinning economic behavior with extraordinary control over independent variables, and generally generalise well to various field environments (Camerer, 2015). Market experiments, in particular, yield similar results with student and businessmen populations (e.g., Porter and Smith, 1994).

In a study of information dissemination in experimental single-period markets, Plott and Sunder (1982) found that asset prices converge to fully reveal the true state even when only a subset of the traders had full information about the true state.[4] Markets where traders have *partial* information pose a greater challenge to successful convergence, as market prices have a double role of information *aggregation* and *dissemination* (e.g., Plott and Sunder, 1988; Forsythe and Lundholm, 1990; Copeland and Friedman, 1987; Camerer and Weigelt, 1991; O'Brien and Srivastava, 1991).[5] Plott and Sunder (1988) pioneered the laboratory study of information aggregation in single-period markets structured as event-contingent assets. In their experiments, the true state of the world can either be X, Y or Z with some known probabilities. For each state other than the true state, half of the traders are informed of the false state. For example, let the true state be Y. In this case, half of the traders know the true state to be either X or Y (i.e., not Z), and the other traders know the true state to be either Y or Z (i.e., not X). Thus, traders have heterogeneous information that, if aggregated efficiently, is enough to fully reveal the true state. In the treatments closest to our design, the market consists of three types of assets x, y and z, which pay a positive dividend for the states X, Y, and Z, respectively (e.g., asset x pays a positive dividend if and only if the true state is X).[6] The results of Plott and Sunder (1988) show that market prices converge to the true value of the asset – thereby revealing the true state – even when there is no immediate incentive for any trader to reveal their private information. However, the information structure used in the experiments greatly facilitates information dissemination and aggregation. For example, if the true event is Y, then the distribution

---

[4] Developments in uses of predictions markets (e.g., Wolfers and Zitzewitz, 2004; Chen and Plott, 2002) are often inspired by the success of markets in this respect.

[5] For a broader summary of market experiments on information aggregation, see surveys by Sunder (1995); Duxbury (1995); Noussair and Tucker (2013).

[6] Subsequent studies have used a similar design to study price manipulation (e.g., Hanson et al., 2006; Veiga and Vorsatz, 2010) and the interplay between trading performances and overconfidence (Biais et al., 2005).

of private information implies that demand for asset y will easily outstrip those for assets x and z.[7] As such, the market demand for asset y relative to those of assets x and z will immediately reveal the true event to be Y.

In this paper, we study a more complex and asymmetric information structure. This structure allows us to test an environment where information aggregation requires counterfactual reasoning, and to test gradual aggregation and dissemination of information. The information structure we employ differs from that of Plott and Sunder (1988) in two important ways. First, it relaxes the one-to-one relation of states of the world to profitable assets (Arrow-Debreu securities) by allowing different states to map onto the same asset. Second, while the states of the world are symmetric and equiprobable, the different possible asset evaluations are not. This asymmetry in information translates to asymmetry with regards to the information held by different types of traders.

These two design innovations raise two new issues, respectively. First, as different states of the world map to the same asset value, we generate a situation where, in some states, all of the traders hold exactly the same beliefs over the assets' values despite having diverse information regarding the underlying state of the world. In previous studies, such as Plott and Sunder (1988), traders always potentially observe some market activity (asks, bids, or transactions) that is inconsistent with some state of the world that they consider possible. Such market activity reveals that there exist other traders who hold different information, and thus conveys new information about the true state. In our setting, in contrast, the different information held by different traders translates to the same beliefs and the same market activity, hence the observed market activity does not provide any new information regarding the true state. The diverse information held by different traders, if aggregated successfully, is nonetheless sufficient to fully reveal the true state. To do so requires counterfactual reasoning: traders should realize that all observed market activity is consistent with their own beliefs about the assets' values, and consequently eliminate the states of the world in which some of the other traders hold differing beliefs. Our experiment tests whether this feature of the market impedes, or even prevents, successful aggregation of information.

Second, in states of the world where different traders do hold different beliefs, we can identify a minority of traders who hold different beliefs to those held by the majority traders. Since the majority traders dominate the market and the price setting, the minority traders stand to better learn new information from the initial movement of prices. Accordingly, we

---

[7] The not X traders will have positive demands for assets z and y, whilst the not Z traders will have positive demands for assets x and y.

test whether these minority traders become better informed about the true state of the world and about the true value of the assets, and whether they can leverage this knowledge for profit.[8]

Finally, previous studies have focused on aggregate outcomes in the market. Our experimental design introduces yet another new feature. To test *individual* belief updating and information dissemination, we elicit traders' post-market beliefs about the true asset-relevant event. This allows us to test whether individual revelation of the true state is a necessary condition for information revelation in market prices.

We present a simple *fully revealing rational expectations equilibrium* (FRE) model, which shows that information aggregation and full revelation of the true state is possible in our setup both when traders' beliefs over the assets' values are heterogeneous and homogeneous. Furthermore, when traders hold heterogeneous private information, the model predicts that a minority subset of traders (the *minority* traders) will learn about the true state before the others and will be instrumental in driving prices towards the FRE. Against this backdrop, three questions are investigated:

RQ1: Can Arrow-Debreu markets be successful at aggregating private information about the true event into prices, even when traders are endowed with homogeneous beliefs, such that information aggregation is challenging and requires traders to employ sophisticated counterfactual reasoning?

RQ2: Are *minority* traders better informed about the true event, can they leverage this knowledge for profit, and are they instrumental in driving prices to reveal the true state?

RQ3: Can individuals in the market learn the true event from the movement of prices, and is individual learning a necessary condition for information aggregation in the prices?

To measure success at information aggregation in the experimental setting, we benchmark the FRE model against a *prior-information equilibrium* (PIE) model.[9] The PIE model

---

[8] Plott and Sunder (1988) included several treatments where a single asset was traded, which valuations depend on the state of the world. This induces some asymmetries, not in the information structure as in the present study, but in the asset valuations. As (with a limited supply of the asset) the price reveals the highest valuation in the market, traders who hold lower valuations are potentially better informed. Plott and Sunder (1988) did not explore the implications of such differences between trader types in their data.

[9] This builds on prior experimental research on the issue of information aggregation (e.g., Banks, 1985;

assumes that the traders update their beliefs given their private information, but fail to take market activity into account. Thus, we test the ability of the market to aggregate dispersed information by comparing the observed final market prices to the FRE and PIE predictions. The post-market beliefs allow us to test the role of individual beliefs in market prices convergence to the true values and the dissemination of information throughout the market. By repeating the market game, we are able to test the role of market experience in information aggregation and dissemination.[10]

Our findings can be summarised as follows. Markets are successful at aggregating dispersed information in the prices when traders are experienced, but not in early rounds. Although market prices reflect the true state, traders largely remain ignorant in the post-market stage. This finding is consistent with the *marginal trader hypothesis* (e.g., Forsythe et al., 1992), which posits that efficiency in market prices only require a small population of sophisticated traders, who are influential in driving prices to the competitive equilibrium. We observe successful price convergence even in the more challenging case of homogeneous beliefs. That is, we do not find any significant differences in the information aggregation properties of markets when traders are endowed with homogeneous or heterogeneous private information. Finally, we find some evidence suggesting that minority traders learn about the true event before others and are instrumental in driving prices towards the FRE. There is no evidence, however, that the better-informed minority traders are able to extract a higher share of the surplus.

2 Experimental Design

We conducted 7 experimental sessions involving 12 subjects each. Participants were recruited from the subject pool of the experimental economics laboratory at Ben-Gurion University of the Negev using the recruitment software ORSEE (Greiner, 2015). All sessions were programmed with the z-Tree (Fischbacher, 2007) software. Subjects in each session were first given 45 minutes to familiarise themselves with the instructions, complete an accompanying set of control questions and participate in two non-paying practice rounds (see the appendix for the translated instructions). The experiment itself included 10 rounds. All payoffs were

---

Barner et al., 2005; Camerer and Weigelt, 1991; Copeland and Friedman, 1987; Forsythe and Lundholm, 1990; O'Brien and Srivastava, 1991; Plott and Sunder, 1982, 1988).

[10] Plott (2000) found that prices converge fairly efficiently to the FRE when subjects have some experience (i.e., after three rounds or more).

stated in ECU (Experimental Currency Unit), and converted to New Israeli Shekels (NIS) at the end of the experiment at a conversion rate of 15 ECU = 1 NIS. The final payoff consisted of the payoff of one randomly chosen round in addition to a show-up fee of 25 NIS. The average payment including show-up fee was 73.33 NIS (approximately 21 USD). Average session duration was approximately 120 minutes.

2.1 Market structure and procedure

The twelve traders are randomly allocated into three groups of four traders each. Each group $k \in \{1, 2, 3\}$ is randomly assigned to be of type X or type O with equal probabilities. Thus, there are eight equally probable states of the world $s \in S \equiv \times_k H_k$, where $H_k \in \{X, O\}$ denotes group k's type.[11] Let the event $x(s) \in \{0, 1, 2, 3\}$ denote the total number of type X groups at state s (e.g., $x(s) = 1$ iff $s \in \{XOO, OXO, OOX\}$). At each state $s \in S$, traders observe the type of the two other groups, but not their own group's type. For example, in the case where $s = XXO$, the first two groups observe one X and one O, hence only events 1 and 2 are possible. The third group observes two X's, hence only events 2 and 3 are possible.

For the true state $s^* \in S$, we define the event $x(s^*)$ as the corresponding *true event*. This distinction is necessary as market prices will convey information about the true *event* rather than the true *state*. The information structure is commonly known and endows traders with private and partial information over the true event. For example, if $s^* = XXX$, then all traders will observe two other type X groups, implying that they each know the true event $x(s^*)$ to be either 2 or 3.

Trade involves four "types" of Arrow-Debreu assets, indexed by $\alpha \in \{0, 1, 2, 3\}$. The value of asset α given the true event $x(s^*)$ is given by

$$d(x(s^*), \alpha) = \begin{cases} 100 & \text{if } \alpha = x(s^*), \\ 0 & \text{otherwise.} \end{cases}$$

Each trader is endowed with a working capital loan of 5,000 ECU and six units of each of the four assets. There are four simultaneous markets, one for each asset. Trade is conducted using the Continuous Double Auction (CDA) mechanism for a duration of 300 seconds. Short sales are prohibited, and a price ceiling of 200 ECU is imposed on the bid and ask prices to

---

[11] The eight states are $XXX, OXX, XOX, XXO, XOO, OXO, OOX$ and OOO, where the first, second and third coordinates of each state denote the type of groups 1, 2 and 3, respectively.

limit potential negative profits. Note that traders at the start of the markets do not know the value of each asset (i.e., they don't know which assets are valued at 0 or 100). Instead, they only know that the value of each asset can be either 100 or 0, depending on the true event $x(s^*)$.

After 300 seconds, the CDA market closes and traders enter the *post-market* stage. Here, traders are asked to guess how many type X groups were in the market. The response $b_i \in \{0, 1, 2, 3\}$ is taken to be participant i's elicited belief.[12] The payment for beliefs is given by

$$g(x(s^*), b_i) = \begin{cases} 200 & \text{if } b_i = x(s^*), \\ -200 & \text{otherwise.} \end{cases}$$

At the end of the round, traders receive feedback on the true state, the unit value and number of units owned for each of the four assets, and their detailed final round payoff. The round payoff consist of the total value of owned assets and remaining capital, minus the initial loan of 5,000 ECU, plus payment for the submitted beliefs.[13] Denote by $L_i \geq 0$ and $e_i^\alpha \geq 0$ trader i's inventory of capital and asset α, respectively. Thus, each trader i's payoff is given by:

$$\Pi_i = \underbrace{(L_i - 5,000) + \sum_{\alpha=0}^{3} d(x(s^*), \alpha_i) e_i^\alpha}_{\text{market payoff}} + \overbrace{g(x(s^*), b_i)}^{\text{post-market payoff}},$$

where the first term captures the consequences of the trader's market decisions and the second his submitted beliefs.

---

[12] Eliciting a (modal) point belief is sufficient for the purpose of assessing individuals' ability to learn the true state from the observed market prices. The alternative of eliciting a full belief distribution would require an incentive compatible scoring rule, which is likely to be confusing for experimental subjects.

[13] Because the loan of 5,000 ECU is to be returned at the end of the round, there were some instances of bankruptcy (13% of all observations). In the first five rounds, these appeared to be distributed rather uniformly across all subjects. In the last five rounds, bankruptcies were concentrated in a handful of subjects, suggesting that the initial bankruptcies were part of the learning process. See the appendix for more details. Note that the limited liability doesn't impact the theoretical analysis, as it doesn't bias trading decisions with respect to the true asset, but may impede information aggregation by increasing risky and/or noisy investments. Thus, if anything, we are underestimating the market's true ability to aggregate information without limited liability.

3 Theoretical analysis

3.1 Preliminaries

We shall refer to the events $x(s^*) \in \{1, 2\}$ as the *heterogeneous* events, given that they induce different beliefs in different traders. That is, the private information held by traders in one group, the *minority traders*, induces different beliefs over the assets' valuations than for the other traders. For example, if $s^* = XXO$, groups 1 and 2 know that the true event is either 1 or 2, while group 3 (the minority traders) knows that the true event is either 2 or 3. In the complementary events $x(s^*) \in \{0, 3\}$, all traders observe the same number of X's, and therefore hold the same beliefs over the events $x(s)$. We therefore refer to these events as the *homogeneous* events. For example if $s^* = XXX$, all traders know that the true event is either 2 or 3. Naturally, it follows that there are no minority traders in homogeneous events.

3.2 The FRE and PIE models

The experimental design endows traders with partial and private information about the true event. Because the complexity of the CDA mechanism makes it extremely difficult to analyse with standard game theoretical methods, we evaluate the market's success at aggregating information by comparing market prices against the *fully revealing rational expectations equilibrium* (FRE) and the *prior–information equilibrium* (PIE) described in full below. Table 1 details the FRE and PIE predicted prices of each asset α in the market and the corresponding implied proportion of submitted beliefs in the post-market stage.[14]

The PIE describes the market clearing prices when risk-neutral traders use Bayes' rule to update their posteriors about the true event given their own private information and condition their demands upon such posteriors. However, there is no belief updating following trade, and consequently the true event is only partially revealed in prices. For example if $s^* = XXO$, then traders in groups 1 and 2 will assign equal posteriors to the events 1 and 2, whilst those in group 3 will assign equal posteriors to the events 2 and 3. As traders condition their demand for assets on their posteriors, the corresponding market clearing prices are 0, 50, 50 and 0 for assets $\alpha = 0$, $\alpha = 1$, $\alpha = 2$ and $\alpha = 3$, respectively.[15] Uniformed of the true

---

[14] The models assume that prices are set only by the expected values of the assets given the traders' beliefs. That is, we assume that traders are risk neutral and have no other motivations for trading other than divergent beliefs.

[15] Traders from groups 1 and 2 are willing to sell asset $\alpha = 3$ at zero. Given group 3 traders are the

Table 1: Predicted Prices and Beliefs.

	Market Prices				Submitted Beliefs			
True Event	p^0	p^1	p^2	p^3	B^0	B^1	B^2	B^3
Panel A: Fully revealing equilibrium rational expectations (FRE)								
$x(s^*) = 0$	100	0	0	0	1.00	0	0	0
$x(s^*) = 1$	0	100	0	0	0	1.00	0	0
$x(s^*) = 2$	0	0	100	0	0	0	1.00	0
$x(s^*) = 3$	0	0	0	100	0	0	0	1.00
Panel B: Prior-information equilibrium (PIE)								
$x(s^*) = 0$	50	50	0	0	0.50	0.50	0	0
$x(s^*) = 1$	0	50	50	0	0.17	0.50	0.33	0
$x(s^*) = 2$	0	50	50	0	0	0.33	0.50	0.17
$x(s^*) = 3$	0	0	50	50	0	0	0.50	0.50

p^α denotes the predicted price of asset α and B^j denotes the predicted proportion of traders submitting beliefs $b = j$.

event, traders in groups 1 and 2 will uniformly randomise between the responses $b = 1$ and $b = 2$. Similarly, those in group 3 will randomise between the responses $b = 2$ and $b = 3$.

In contrast to the PIE, the FRE describes an outcome where risk-neutral traders not only use Bayes' law to update their posteriors but also continuously learn about the private information of others from market prices. This endogenous process continues until no trader can learn any more from prices and the true event is fully revealed in the prices. Since all traders are fully informed about the true event by definition, this should be reflected in their submitted beliefs.

In many respects, both equilibria describe some form of market aggregation. Where they differ is on the issues as to whether the true event is partially (PIE) or fully (FRE) revealed in prices. The expected market payoff in both PIE and FRE is 600. The expected post-market payoff is 0 and 200 for the former and latter, respectively.

How prices might actually converge to the FRE remains an open question. There is an extensive literature that discusses the dynamic behaviours of traders (e.g., Dubey et al., 1987; Hellwig, 1982; Ostrovsky, 2012; Radner, 1979). However, the complexity of such models are largely beyond the scope of our experimental design. Furthermore, Biais and Pouget (1999) show that given the information asymmetry in our design, there should be no trade in equilibrium, except at the FRE prices. Yet without trade at non FRE prices, it is difficult to see how traders learn about the true state.

In the following subsection, we present a simple *reasoning model* describing the convergence process supporting the FRE prices.[16] As we later elaborate, counterfactual reasoning by traders will play in integral role in the convergence process.

3.3   The reasoning model

The model builds on Jordan's (1982) discussion of the way in which prices may first converge to some 'pre-equilibrium' on the path to the FRE.[17] Trade takes place over $t = 1, 2, 3, 4$ hypothetical periods. Traders are risk-neutral, not capital constrained, always seek to maximise

---

minority in the population, the only market clearing price for asset 3 is zero.

[16] We introduce the reasoning model to highlight the challenges and complexities in convergences to the FRE. The model is not based on a game theoretic analysis of double auction market structure, does not incorporate risk-aversion and does not include other motivations to trade other than differences in opinion.

[17] Jordan's (1982) model is sometimes known as the 'Jordan Path'. Copeland and Friedman (1987) and Plott and Sunder (1988) have noted that the convergence of asset prices to the FRE in their data is consistent with the Jordan Path.

payoffs and are non-strategic in the sense that they trade in each period as if it were the last. Denote p_t^α as the period t price of asset α. Given the above, the model assumes that behaviours at each period t are described by the following four stages:

Stage 1. Traders observe the period $t-1$ prices of each asset (i.e., p_{t-1}^α).

Stage 2. Traders update their beliefs regarding the true event.

Stage 3. Traders condition their demands for assets given their updated posteriors or beliefs.

Stage 4. The market clears and period t prices are established.

To see how convergence to the FRE occurs in the reasoning model, first consider the case where $s^* = XXX$ or $x(s^*) = 3$ (homogeneous event). Here, all traders observe two type X groups and therefore assign equal posteriors to the events 2 and 3. At period $t = 1$, traders condition their asset demands on such posteriors. This results in the market clearing prices: $p_1^2 = p_1^3 = 50$ and $p_1^0 = p_1^1 = 0$ (the PIE prices). At period $t = 2$, traders observing the previous period prices will reason the true event to be 3. The reasoning is as follows. If the true event were 2, some traders would have observed one X and one O. Such traders would assign a posterior probability of 0.5 to the event 1, driving p_1^1 up to 50. Since $p_1^1 = 0$, this must mean that no trader observes only one other X ruling out event 2. Conditioning their demands on the updated beliefs now results in the market clearing prices: $p_2^3 = 100$ and $p_2^0 = p_2^1 = p_2^2 = 0$ (the FRE prices). Since the true event is fully revealed in $t = 2$, there should be no revisions to the prices in periods $t = 3$ and $t = 4$. This process of information aggregation holds by symmetry for the $s^* = OOO$ case.

Now consider the case where $s^* = XXO$ or $x(s^*) = 2$ (heterogeneous event). Here, traders in groups 1 and 2 observe one type X group and therefore assign equal posteriors to the events 1 and 2. Conversely, the minority traders in group 3 observe two type X groups and therefore assign equal posteriors to the events 2 and 3. At period $t = 1$, traders condition their asset demands on such posteriors, resulting in the market clearing prices: $p_1^1 = p_1^2 = 50$ and $p_1^0 = p_1^3 = 0$ (the PIE prices).[18] At period $t = 2$, the minority traders, having observed non-zero prices for asset 1, can deduce that some of the other traders observed a type O group. This rules out event 3 from being the true event, hence the true

---

[18] Although group 3 traders value asset 3 at 50, those in groups 1 and 2 value asset 3 at 0. The excess supply therefore drives the market clearing price of asset 3 towards 0.

event must be $x(s^*) = 2$. In contrast, the non-minority traders (those in groups 1 and 2), who already knew that there is at least one type O group receive no new information from the prices, and therefore remain uninformed. Given the updated beliefs, there is still excess demand for asset 1 by groups 1 and 2, leading to market clearing prices of $p_2^1 = p_2^2 = 50$ and $p_2^0 = p_2^3 = 0$.[19] However, given their higher valuations, the market mechanism will allocate asset 2 to the minority traders. At period $t = 3$, groups 1 and 2 again remain uninformed of the true event. Given that the trade of asset 2 will now mostly involve group 3 traders, though, the subsequent market clearing prices will be: $p_3^0 = p_3^3 = 0$, $p_3^1 = 50$ and $p_3^2 = 100$.[20] At period $t = 4$, the non-minority traders observing that $p_3^2 > p_3^1$, should conclude that the true event is $x(s^*) = 2$. As traders condition their asset demands on their updated beliefs, we arrive at the market clearing prices: $p_4^0 = p_4^1 = p_4^3 = 0$ and $p_4^2 = 100$, the FRE prices.[21] This process of information aggregation holds by symmetry for all states where $x(s^*) \in \{1, 2\}$.

The main results of the reasoning model are as follows. Firstly, convergence to the FRE prices is possible for all states and events. Secondly, on the *path* of convergence towards the FRE, prices go through the PIE prices. Finally, minority traders (heterogeneous event) learn about the true event before the other traders and are 'instrumental' in driving prices towards the FRE.

Some comments

The model assumes that traders only update their beliefs on observed prices. In CDA markets, traders can also learn from the volume of transactions or the unfilled bid and ask prices. Will these additional elements influence the information aggregation process? Suppose that traders only update their beliefs on observed transaction volumes. When $s^* = XXX$, transactions should be concentrated on assets 2 and 3. So again, traders will only learn about the true event through reasoning about why transactions for asset 1 is relatively low. When $s^* = XXO$, transactions should be concentrated on assets 1 and 2.[22] As such, minority traders should still learn about the true state before the others.

---

[19] If leveraged short-sales are permitted, then p_2^1 may converge to 0 as traders in group 3 act to short the asset.

[20] The intuition here is that because asset 2 is all held by the group 3 traders, 'Bertrand' like competition between informed traders drive the price of asset 2 up to 100.

[21] An interesting corollary of the model is that if there is only one minority trader, as opposed to group of traders, prices may never converge to the FRE.

[22] The population of group 3 traders (minority traders) who will purchase asset 3 is substantially smaller than those in groups 1 and 2, who will mainly transact in assets 1 and 2.

Consequently, transaction volumes for asset 2 will increase relative to those for asset 1 as minority traders concentrate their purchases on the true asset. Non minority traders will thus learn the true event to be 2. Thus, the additional elements do not seem to alter the information aggregation process in any substantial way.

The model suggests that convergence to the FRE might be more challenging for the heterogeneous information events given that minority subjects must first be informed about the true event and be instrumental at driving prices towards the FRE.[23] At the same time, convergence to the FRE also requires agents to employ counterfactual reasoning. In heterogeneous information events, agents make inferences from actual observed trade. In the homogeneous information events, in contrast, agents deduce the true event from not observing trade on an asset, which would have been traded in some counterfactual event. Recognising that empirical (experimental and otherwise) data are often noisy, and some trade is likely to happen, it is not clear whether low volumes or prices are enough for agents to be confident at ascertaining the true event. Even complete market inactivity is somewhat ambiguous, as it is not clear how much time should be allowed to elapse before concluding that there is no demand for the asset. Lack of trade may be explained, for example, by agents concentrating on other parts of the market. For these reasons, convergence to FRE prices may depend on the information structure, with homogeneous beliefs leading to slower convergence.

4 Experimental results

To study prices in dynamic CDA markets, we partition the market data into time windows of 30 seconds, each indexed by $t = (1, 2, 3, .., 10)$. Denote by p_t^α the mean transaction price for asset α in window t. For example, p_2^3 corresponds to the mean transaction price of asset 3 within the 30–60 seconds time window. Where p_t^α cannot be computed as there are no transactions in the time window, we employ the *carry forward rule*, substituting p_t^α with the last known computed price.[24] As such, p_t^α will always reflect the most recent price for

---

[23]Note that because t refers to hypothetical periods, it does not necessarily imply that convergence to the FRE is slower for the heterogeneous relative to homogeneous information states.

[24] The carry forward rule is performed recursively until a price is determined. That is, we first set $p_t^\alpha = p_{t-1}^\alpha$. If this cannot be determined, we then set $p_t^\alpha = p_{t-2}^\alpha$ and so forth. If no transaction of the asset occurred since the beginning of trader, we set $p_1^\alpha = 0$. It was rather rare that an asset was not traded at all in $t = 1$, occurring for only 24 of 280 (8.6%) of traded assets, 15 instances of which pertain to an asset that all traders can rule out as the true asset at the beginning of trade (e.g., $\alpha = 0$ in $s^* = XXO$).

asset α. Where appropriate, we will also consider the prices without the carry forward rule, denoted by $\bar{p}_t^\alpha$. That is, $\bar{p}_t^\alpha = 0$ if no transactions are observed.

Since all events of the form $x(s)$ are a one-to-one mapping to assets, we will sometimes use both terms interchangeably. For any true event $x(s^*)$, define the *alternative* event as the non-true event to which traders assign a positive posterior given their private information at the start of the market. Since this will differ amongst the minority and non-minority traders, we use Alt-M and Alt-NM to denote the alternative event for the former and latter, respectively. Finally, for any true event, define the *implausible* events as those which all traders assign zero posterior to given their private information at the start of the market. For example if $x(s^*) = 2$, then assets 0, 1, 2 and 3 are the implausible, Alt-NM, true, and Alt-M assets, respectively. If $x(s^*) = 0$, then assets 0 and 1 are the true and Alt-NM assets, respectively, while assets 2 and 3 are the implausible assets.

Figures 1 and 2 present the median p_t^α and submitted beliefs for homogeneous and heterogeneous private information events, respectively. The rows in each figure correspond to prices for the true, Alt-NM, Alt-M, and implausible assets. The header of each panel details the round, followed by the number of sessions which data is pooled from and the corresponding mean submitted belief in parenthesis. The solid and dash horizontal lines on each panel detail the FRE and PIE prices, respectively. For example, the top left panel of Figure 1 corresponds to the the median p_t^α of the true asset over the first round. The header shows that the round-1 true event was homogeneous in three of the seven sessions, and that 42% of subjects' submitted beliefs corresponded to the true event.
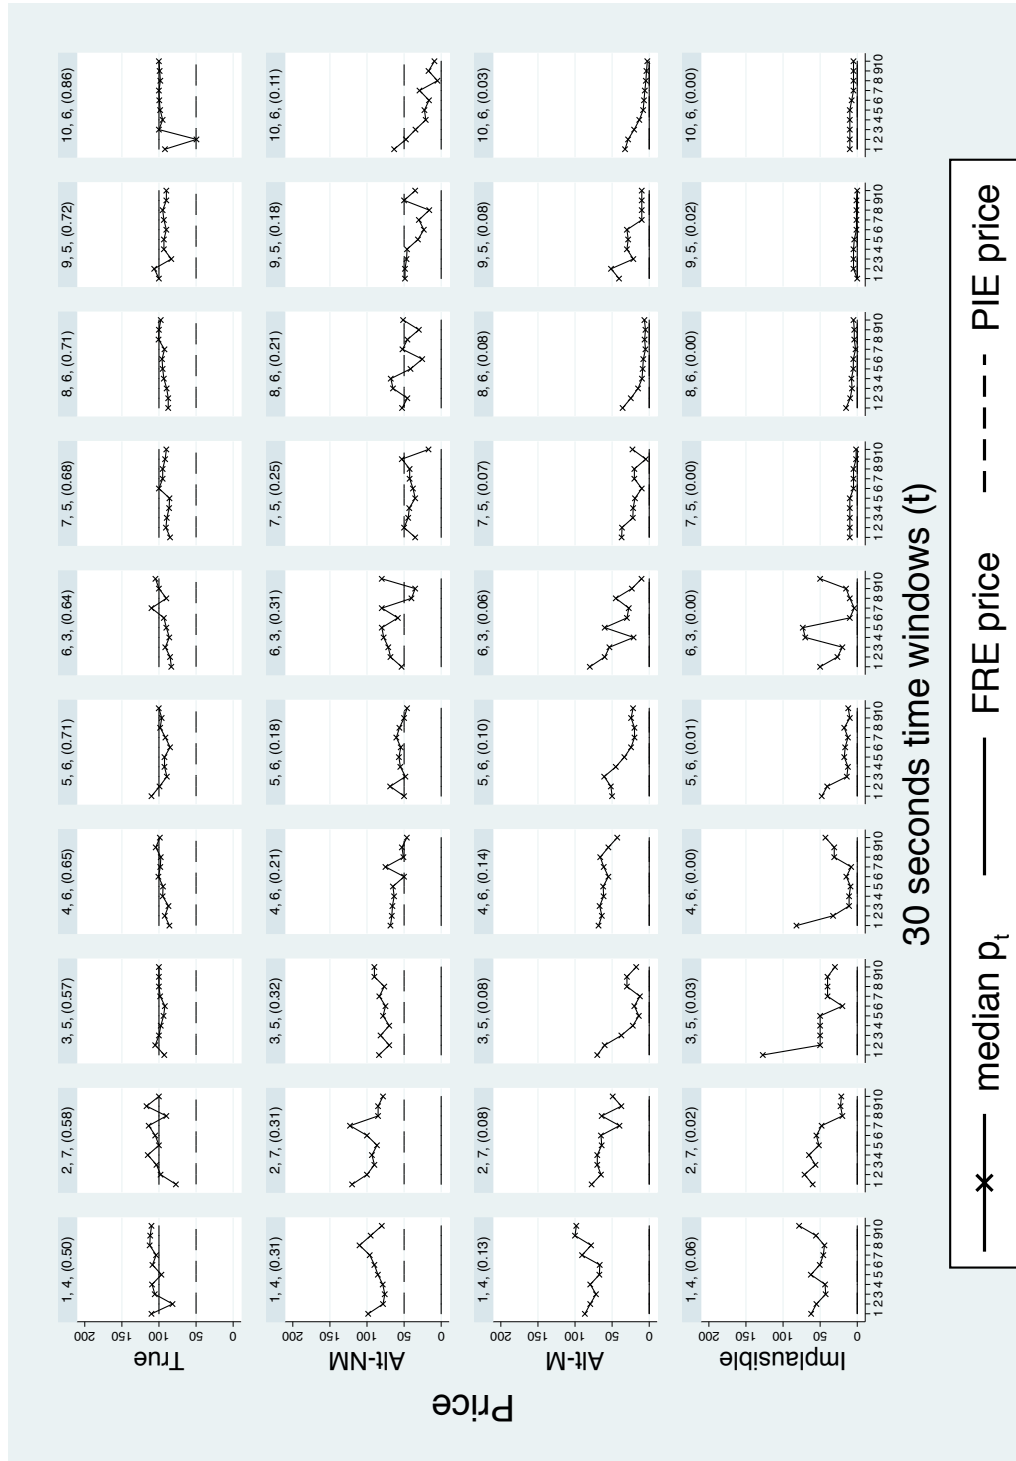
There are several indications of learning over the ten rounds. We first consider purchases at a price of above 100 ECU. Such purchases might seem 'irrational' given that the maximum value of any asset is 100 ECU. Note, however, that this is not necessarily the case. For example, a trader might be willing to buy an asset she knows to be worthless at a price of 120 ECU, if she believes that this will influence the beliefs of other traders and enable her to sell five units of the same asset at a price of 30 ECU, or if she believes that she will be able to resale the asset at a higher price. We find that 19% of all purchases in the first five rounds were at a price of above 100 ECU, decreasing to 9% of all purchases in the last five rounds. Furthermore, whilst 79% of subjects made a seemingly irrational purchase at least once in the first five rounds, only 35% did so in the last five rounds. See the appendix for more details.[25] Importantly, the persistence of such purchases suggests that there are still

---

[25] In fact, one of these subjects – call her subject A – was responsible for nearly 25% of all purchases above 100 ECU in the last 5 rounds. However, all of such purchases for subject A pertained to the true

Figure 1: Prices and submitted beliefs for homogeneous information events.
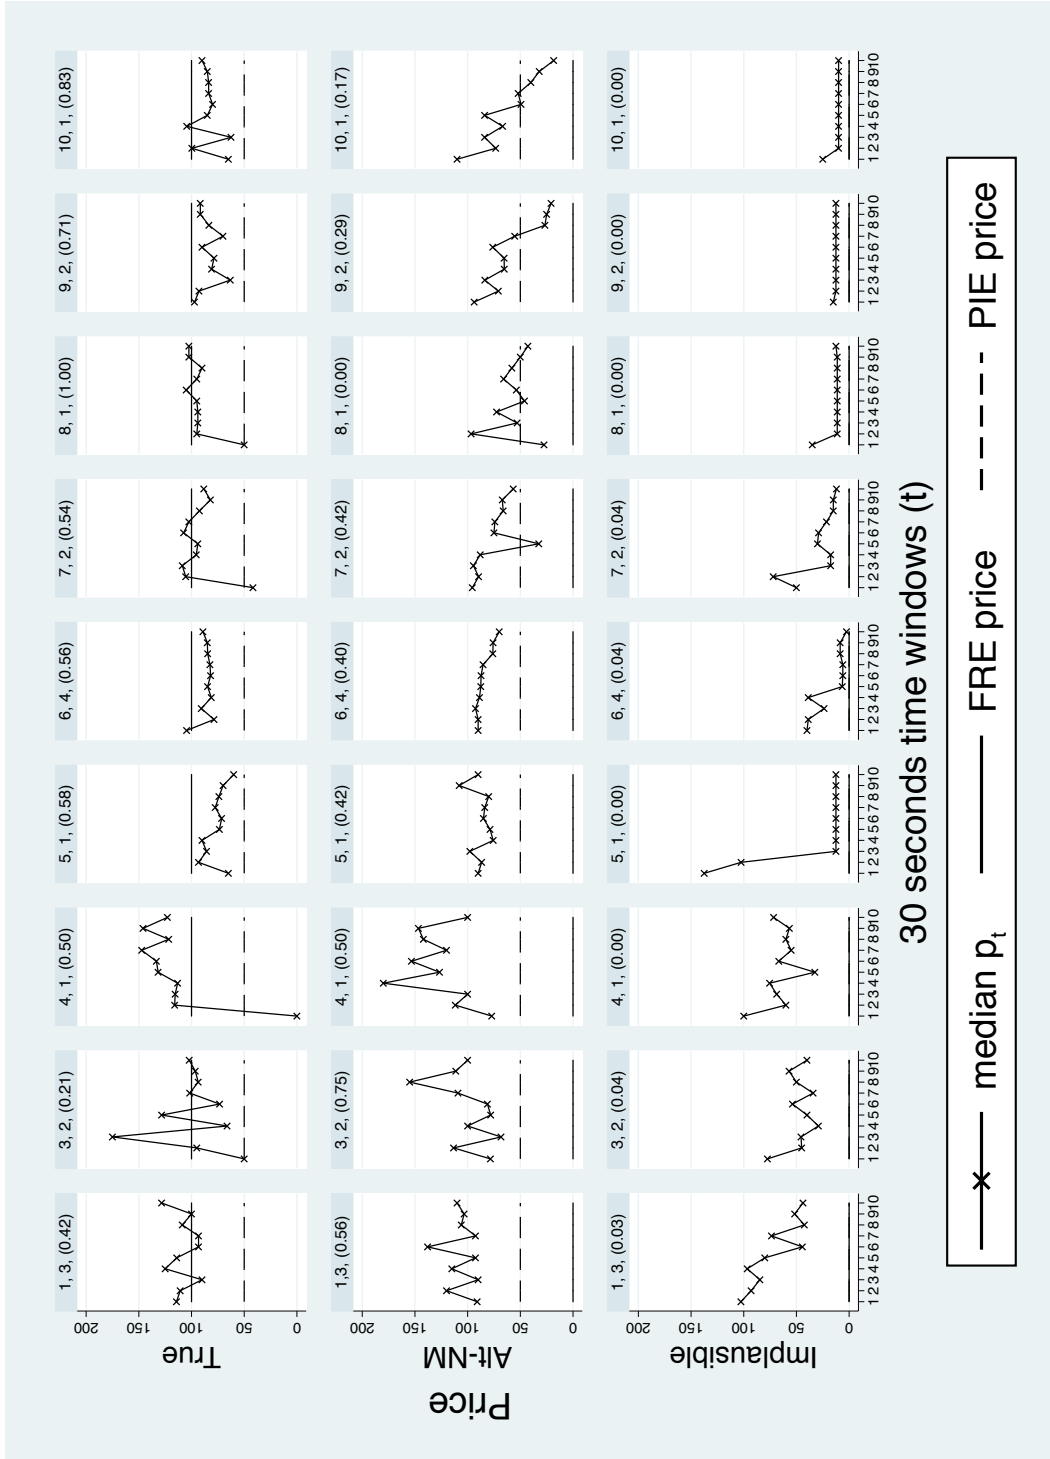
The first, second and third rows correspond to median prices for the True, Alt-NM and Implausible assets, respectively, aggregating over all sessions. The header of each panel details the round, followed by the number of sessions and finally, the corresponding submitted beliefs in parenthesis. The solid and dash horizontal lines denote the FRE and PIE prices, respectively.

Figure 2: Prices and submitted beliefs for heterogeneous information events.

The first, second, third and fourth rows correspond to median prices for the True, Alt-NM, Alt-M and Implausible assets, respectively, aggregating over all sessions. The header of each panel details the round, followed by the number of sessions and finally, the corresponding submitted beliefs in parenthesis. The solid and dash horizontal lines denote the FRE and PIE prices, respectively.

some subjects who behave like *noise traders* (e.g., De-Long et al., 1990).

Turning our attention to price dynamics, we observe that the true asset often transacted at close to 100 ECU. Consistent with the FRE and PIE predictions, prices of the implausible and Alt-M assets are converging towards zero, but only in rounds 6–10. Finally, prices for the Alt-NM asset seem to converge towards 100 ECU in rounds 1–4, 50 ECU in rounds 5–6 and finally, 0 ECU in rounds 7–10. To acknowledge the role of experience, we separate the data in the following analyses into 'Inexperienced' (rounds 1–5) and 'Experienced' (rounds 6–10) subjects.[26]

In the following, we will first address RQ1, investigating whether Arrow-Debreu markets can be successful at aggregating dispersed private information into prices and revealing the true event. Orthogonal to this, we will also investigate whether the nature of information aggregation might be dependent on the information structure, more specifically, whether price and belief convergence differ between homogeneous and heterogeneous events. Finally, we study the behaviours of minority subjects to investigate RQ2.

4.1 Information aggregation

Our approach seeks to investigate the question of information aggregation from three different angles. Since relative asset prices can be interpreted as the markets' prediction about the true event, we will first investigate whether the markets are often correct (i.e., the true asset is priced higher than the other assets). Thereafter, we will focus on the issue of arbitrage opportunities, such as when prices of the various assets do not sum up to 100 ECU. The intuition here is that subjects may learn about the true event at different t, resulting in disagreements and consequently arbitrage opportunities. However, such arbitrage opportunities should also be eliminated if information aggregation is successful. Finally, we will benchmark the FRE against the PIE to examine if prices are relatively closer to either equilibria.
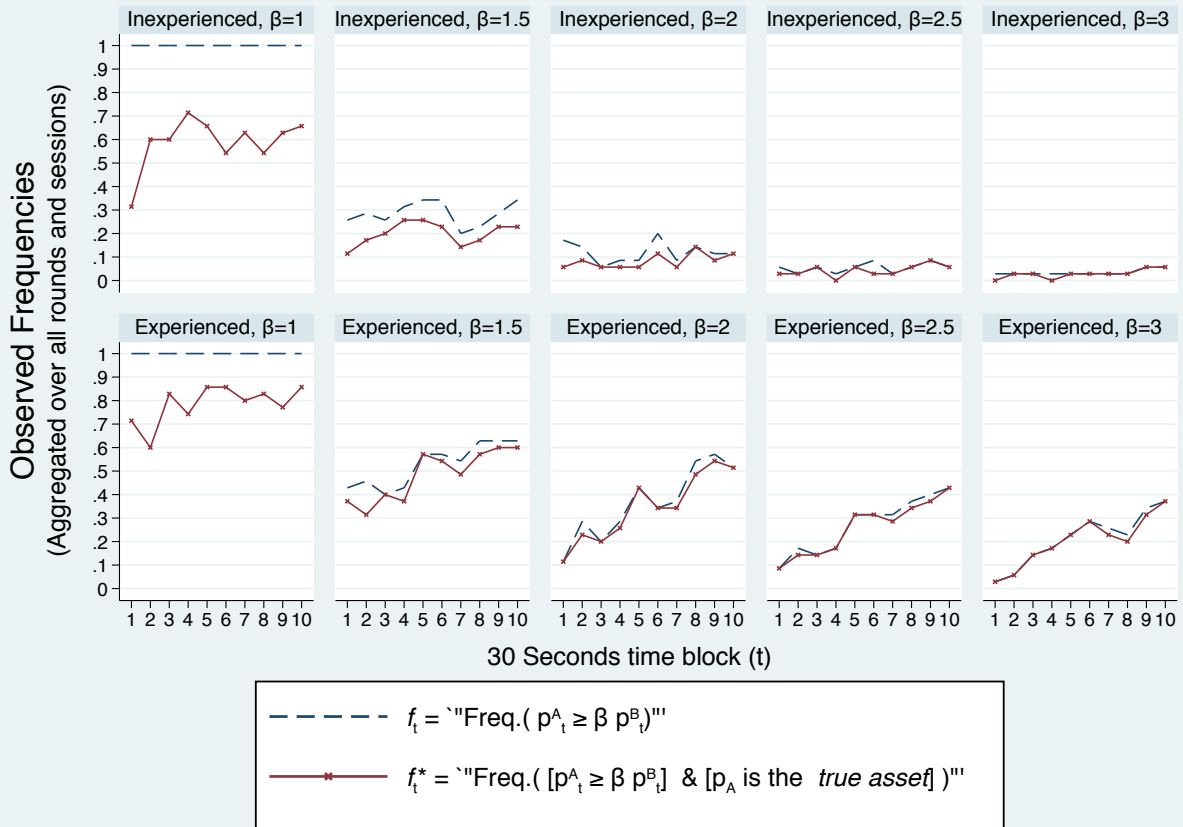
Assets A and B are the highest and second highest price assets at window t, respectively, with correspond prices p_t^A and p_t^B. The constant $\beta \geq 1$ captures the extent to which p_t^A is higher than p_t^B. For a given β, f_t shows the frequency to which the market has "picked a true event" and f_t^* the frequency to which the market is correct. Specifically, for $\beta = 1$, f_t^* is the proportion of time in which the prices reveal the underlying true event.

Figure 3: Highest and second highest price assets.

4.1.1 Relative prices

From Figures 1 and 2, we note that prices for the Alt-M and implausible assets are consistently lower than those for the Alt-NM and true assets. Indeed, transactions of the former two assets are also relatively less frequent. These observations are helpful since they suggest that markets often excluded the Alt-M and implausible assets as candidates for the true event. In this case, our analysis can focus on the relative prices of the Alt-NM and true assets.

To make such comparisons, we will define assets A and B as the highest and second highest priced assets at time window t, respectively, with the corresponding prices p_t^A and p_t^B.[27] We derive the following two frequencies for some $\beta \geq 1$:

f_t: The frequency to which $p_t^A \geq \beta p_t^B$ at window t.

f_t^*: The frequency to which $p_t^A \geq \beta p_t^B$ conditional on A being the true asset at window t.

For $\beta = 1$, f_t^* reflects the proportion of markets in which the true asset was priced higher than the other assets. Thus, f_{10}^* measures the proportion of times in which the final prices reveal the true event. This happens overall in 52 of the 70 (74.3%) markets included in the experiment, or 29 of the 35 (82.86%) markets with experienced subjects. In each of the final four rounds, final prices revealed the true event in six of the seven sessions.

Values of β larger than 1 capture the extent to which asset A is priced higher than asset B. Intuitively, f_t captures the markets' candidate true asset at time t given the criterion β. For example when $\beta = 2$, then f_t details the ratio of observations which the price of asset A is at least twice that of asset B in window t – data aggregated over all sessions. In this case, $f_t^* \leq f_t$ can be interpreted as the frequency to which the markets are correct (e.g., the frequency to which the true asset is priced at least β times higher than the Alt-NM asset at time t). Inversely, $(f_t - f_t^*)$ details the frequency to which the markets are incorrect.[28]

Figure 3 details the f_t and f_t^* for $\beta = 1, 1.5, 2, 2.5, 3$.[29] A few observations are apparent from the figure. First, for $\beta > 1$, f_t^* is very close to f_t, implying that traders are sometimes

---

asset. Furthermore, irrational purchases were often between the range of 101-120. Perhaps, subject A might be using the prices to signal the true event to others.

[26] As noted by Forsythe and Lundholm (1990), experience with the trading mechanism or experimental design is crucial in giving markets a better chance at succeeding.

[27] Prices for assets A and B here are derived applying the carry forward rule.

[28] The FRE predicts that f_t and f_t^* are unity for all $\beta \geq 1$. In contrast, the PIE predicts that f_t and f_t^* are zero for all $\beta > 1$.

[29] See the appendix for f_t and f_t^* for $\beta = 1, 1.1, 1.2, 1.3, 1.4, 1.5$.

incorrect in identifying the true asset, but rarely by a large margin. Alternatively, when markets clearly favour one asset over the others, they are almost always correct.

Second, experience has a strong effect on information aggregation. When subjects are inexperienced, f_{10} is approximately 0.35 and 0.10 when $\beta = 1.5$ and $\beta = 2$, respectively, suggesting that final trading prices do not discriminate between the assets by much. Furthermore, the plots are rather flat, indicating that discrimination doesn't improve by much across the trading period. For experienced subjects, in contrast, f_t and f_t^* are very close to each other. Discrimination clearly develops as trade proceeds, as reflected in the upwards trends in the plots, culminating in f_{10}^* at approximately 0.60 and 0.51 when $\beta = 1.5$ and $\beta = 2$, respectively. The markets were, therefore, quite efficient in discriminating between the true asset and the other assets, with the true asset priced at at least twice the other asset in 51% of all observations.

4.1.2 Arbitrage opportunities

To study arbitrage opportunities, we define $SUM_t = \sum_\alpha \overline{p}_t^\alpha$ as the sum of al prices at window t.[30] The thick black line on Figure 4 plots the median SUM_t over all sessions. Here, we clearly observe SUM_t to be decreasing with t for both inexperienced and experienced subjects, indicating that arbitrage opportunities are reduced with t.[31] Are they 'eliminated' at the end of the market?

To study this, we focus on the prices at window $t = 10$. Panel A of Table 2 details the mean SUM_{10} apportioned by sessions and experience. We clearly observe SUM_{10} to be higher than 100 for inexperienced subjects. Indeed, aggregating over all sessions, we find SUM_{10} to be significantly greater from 100 (Signrank, $\rho < 0.001$, $n = 35$ observations) for inexperienced subjects. This clearly suggests that arbitrage opportunities are still apparent at the end of the market.

The picture is nuance for experienced subjects. Firstly, we observe some heterogeneity amongst the sessions, with SUM_{10} be noticeably different from 100 for sessions 1, 4 and 5 and less so for sessions 3, 6 and 7. Aggregating over all sessions, we do not find SUM_{10} to be significantly different from 100 for experienced subjects (Signrank, $\rho = 0.13$, $n = 35$).[32] We

---

[30] Neglecting the carry forward rule here to avoid double counting.

[31] Notice that SUM_t for inexperienced subjects are frequently above 100 ECU. This is mainly due to the activities of those subjects who purchased assets at prices above 100 ECU.

[32] For further robustness checks, we also regressed SUM_{10} on a constant – fixed effect regression model. For inexperienced subjects, the constant is positive and significantly different from 100 at the 1% level. For
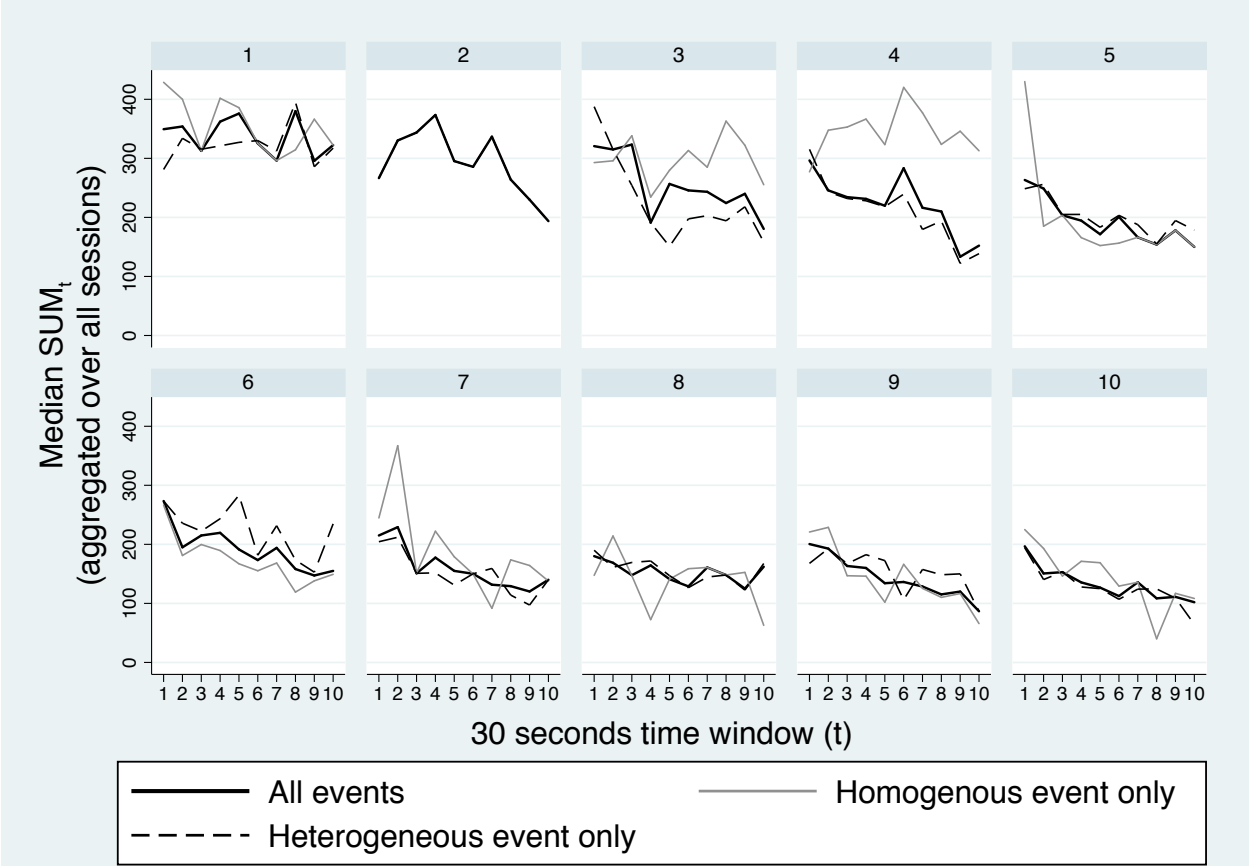
The header of each panel details the round. The variable $SUM_t = \sum_\alpha \bar{p}_t^\alpha$ describes the sum of all prices at window t. Arbitrage opportunities are defined to arise when $SUM_t \neq 100$.

Figure 4: Median SUM_t over all sessions.

Table 2: Summary statistics for SUM_{10} and $DIFF_{10}$.

Variables: $SUM_{10} = \sum_\alpha \overline{p}_{10}^\alpha$ and $DIFF_{10} = MAD_{10}^P - MAD_{10}^F$

Variable	Inexperienced		Experienced	
	SUM_{10}	$DIFF_{10}$	SUM_{10}	$DIFF_{10}$
Panel A: Data aggregated by sessions				
Session 1 (mean)	248	-1	143	4
Session 2 (mean)	134	6	80	13
Session 3 (mean)	98	3	86	4
Session 4 (mean)	353	-4	183	-6
Session 5 (mean)	314	0	141	-1
Session 6 (mean)	190	-3	103	13
Session 7 (mean)	175	-16	98	10
[i]All sessions (mean)	216*** $(n=35)$	-3 $(n=35)$	119 $(n=35)$	5** $(n=35)$
Wilcoxon Signrank				
Null: $SUM_{10} = 100$	$\rho < 0.01$		$\rho = 0.13$	
Null: $DIFF_{10} = 0$		$\rho = 0.09$		$\rho = 0.02$
Panel B: Data aggregated by homogeneous and heterogeneous events				
Homogeneous (mean)	277 $(n=7)$	-3 $(n=7)$	112 $(n=10)$	0 $(n=10)$
[ii]Heterogeneous (mean)	201 $(n=28)$	-2 $(n=28)$	122 $(n=25)$	7 $(n=25)$
Mann-Whitney (Homo vs. Hetero)	$\rho = 0.13$	$\rho = 0.83$	$\rho = 0.66$	$\rho = 0.12$

The variable $SUM_t = \sum_\alpha \overline{p}_t^\alpha$ describes the sum of all prices at window t. The variable $DIFF_t = MAD_t^P - MAD_t^F$ describes the difference between the mean absolute deviation of prices from the PIE (MAD^P) and FRE (MAD^F).

[i] Aggregating over all sessions, we use the fixed effect model to regress SUM_{10} ($DIFF_{10}$) over a constant, where *** and ** imply that the constant is significantly different from 100 (0) at the 1% and 5% levels, respectively.

[ii] Aggregating over all sessions, we use the fixed effect model to regress the dependent variable (SUM_{10} or $DIFF_{10}$) over a constant and a situation dummy for Heterogeneous information events. ** and * imply that the situation dummy is significantly different from zero at the 1% and 5% levels, respectively.

thus cannot exclude the possibility that arbitrage opportunities are eliminated at the end of the market.

4.1.3 Deviations of prices from the FRE and PIE

To benchmark the FRE against the PIE, we compute the mean absolute deviations of p_t^α from the former and latter, henceforth known as MAD_t^F and MAD_t^P, respectively.[33] Note that the symmetrical nature of the FRE and PIE imply that MAD_t^F and MAD_t^P will often be close and at times even identical.

Given the above, we define $DIFF_t = MAD_t^P - MAD_t^F$ as the difference between both mean absolute deviations at window t. Interpretations here should be straightforward: A positive (negative) $DIFF_t$ implies that prices at window t are relatively closer to the FRE (PIE). This relationship also implies that $DIFF_t$ will converge to 25 (-25) as prices converge to the FRE (PIE). Therefore, any relative measure between the FRE and PIE only needs to focus on the sign of $DIFF_t$. Information aggregation is also expected to be a dynamic process from traders continuously about the true event from prices. Furthermore, the reasoning model in Section 3.3 predicts that the path of convergence to the FRE will also go through the PIE. In this case, we should typically expect $DIFF_t$ to start in the negative region and increase with t towards the positive region. However, the rate of increase is also expected to decrease with t.

The construct of $DIFF_t$ is also convenient given that variations will only depend on the prices of the true and Alt-NM assets, the two most frequently traded asset types. That is, for fixed prices of the true and Alt-NM assets, variations to the implausible and Alt-M assets will not influence $DIFF_t$ given that such variations shift the MAD_t^F and MAD_t^P by the same magnitude.

The thick solid line on Figure 5 details the median $DIFF_t$, aggregating over all sessions. For inexperienced subjects, we observe $DIFF_t$ to often be zero (i.e., prices are not getting closer to the FRE nor the PIE with t.) This is consistent with the observations in Section 4.1.1, where we find that prices for the true and Alt-NM assets are often not very different. For experienced subjects, we often observe $DIFF_t$ to be negative at the start of the round and increase progressively over t. To see the evolution of prices more formally, we regress $DIFF_t$ on the time trend t and t^2. The fixed-effects coefficient estimates for experienced

experience subjects, the constant is positive but not significantly different from 100 at the 5% level.

[33] More formally, $MAD_t^F = \sum_\alpha |p_t^\alpha - p^\alpha|/4$ and $MAD_t^P = \sum_\alpha |p_t^\alpha - \hat{p}^\alpha|/4$, where p^α and $\hat{p}^\alpha$ are the FRE and PIE prices, respectively.

24
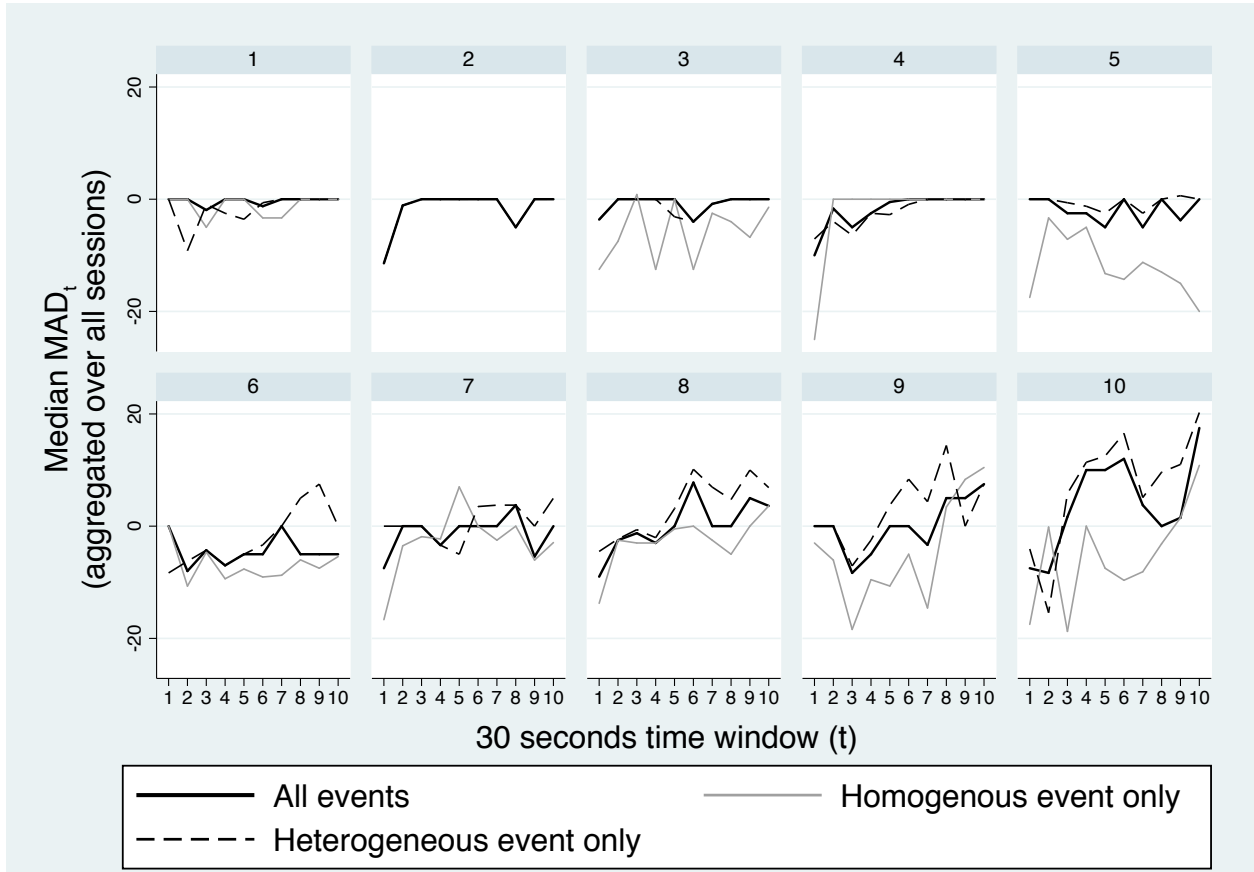
The header of each panel details the round. The variable $DIFF_t = MAD_t^P - MAD_t^F$ describes the difference between the mean absolute deviation of prices from the PIE (MAD^P) and the mean absolute deviation of prices from the FRE (MAD^F). A positive and negative $DIFF_t$ indicate that prices are relatively closer to the FRE and PIE, respectively.

Figure 5: Median $DIFF_t$ over all sessions.

25

and inexperienced subjects are reported as follows (robust standard errors in parenthesis):[34]

$$D\hat{I}FF_t = \underset{(1.49)}{-6.31} + \underset{(0.64)}{0.92t} - \underset{(1.49)}{0.06t^2} \qquad n = 350 \text{ (Inexperience)}$$

$$D\hat{I}FF_t = \underset{(1.73)}{-7.48} + \underset{(0.72)}{1.69t} - \underset{(0.06)}{0.05t^2} \qquad n = 350 \text{ (Experience)}$$

The negative intercept term in both regressions imply that prices towards the start of the round are significantly closer to the PIE. For inexperienced subjects, the coefficient estimates for t and t^2 are not significant at the 10% level. In contrast, the t coefficient for experienced subjects is both significant ($\rho = 0.02$) and positive. The t^2 coefficient for experience subjects is negative but not significant ($\rho = 0.42$). The estimates therefore suggest that only for experienced subjects do prices seem to move away from the PIE towards to the FRE.

To test whether markets are indeed successful at aggregating dispersed private information and revealing the true event through the equilibrium prices, we now focus on the end of the markets (i.e., window $t = 10$). Panel A of Table 2 details the mean $DIFF_{10}$ apportioned by sessions and experience. Aggregating over all sessions, we do not find $DIFF_{10}$ to be significantly different from zero (Signrank, $\rho = 0.093$, $n = 35$). In contrast, we find $DIFF_{10}$ to be positive and significantly (Signrank, $\rho = 0.027$, $n = 35$) different from zero for experienced subjects.[35] Together with the observations from Sections 4.1.1 and 4.1.2, the evidence here leads us to our first result.

Result 1 *For experienced subjects we find that (i) Markets frequently price the true asset significantly higher than the other assets, (ii) There is no significant evidence of arbitrage opportunities at the end of the round, (iii) Prices move towards the FRE and away from the PIE within a round, and (iv) End of round prices are significantly closer to the FRE relative to the PIE. These observations do not hold for inexperienced subjects. The evidence therefore suggests that markets can be successful at information aggregation, even when the task is challenging and requires market participants to employ counterfactual reasoning, but only when market participants have experience with the trading mechanism or market design.*

---

[34] To setup our panel data model, we defined each round as an "independent" market, with the time variable t. Therefore, we estimated the fixed-effects model, clustering the standard errors at the session level. We also considered the estimates from GLLAMM (Rabe-Hesket et al., 2005) linear regression model, a class of Multilevel mixed models, which allows for a hierarchy of random effects for windows nested in rounds nested in sessions. The conclusions are similar.

[35] For further robustness checks, we also regressed $DIFF_{10}$ on a constant – fixed effect regression model. For inexperienced subjects, the constant is negative but not significantly different from zero at the 5% level. For experience subjects, the constant is positive and significantly different from zero at the 5% level.

For further robustness checks, we also employ the approach of Plott and Sunder (1988) and Forsythe and Lundholm (1990) where only the last ($P1$ price) and last 5 ($P5$ prices) transactions prices of each asset are considered. For inexperienced subjects, we find $P1$ (Signrank, $\rho = 0.859$, $n = 35$) and $P5$ (Signrank, $\rho = 0.022$, $n = 35$) prices to be somewhat closer to the PIE. In contrast, $P1$ (Signrank, $\rho = 0.044$, $n = 35$) and $P5$ (Signrank, $\rho = 0.025$, $n = 35$) prices are significantly closer to the FRE for experienced subjects.

4.1.4 Homogeneous vs. heterogeneous information events

Building on Result 1, we now investigate whether the information aggregation properties of markets depend on the information structure (i.e., *homogeneous* or *heterogeneous* events).[36]

First consider those instances where the true and Alt-NM assets are the highest (A) and second highest (B) priced assets at window t. Focusing on window $t = 10$, we compute the ratio p_{10}^A/p_{10}^B, the relative price of the true over the Alt-NM asset. This ratio does not differ significantly between homogeneous and heterogeneous events for inexperienced (Mann-Whitney, $\rho = 0.412$, $n = 35$) and for experienced (Mann-Whitney, $\rho = 0.372$, $n = 35$) subjects.

Turning our attention to arbitrage opportunities, Figure 4 details the median SUM_t for homogeneous and heterogeneous events, apportioned by rounds. With the exception of round 4, there does not seem to be any obvious differences in SUM_t for both events. Restricting our analysis to prices at the end of the round, Panel B of Table 2 details the mean SUM_{10} apportioned by events and experience. Aggregating over all sessions, we do not find SUM_{10} for homogeneous and heterogeneous events to be significantly different amongst inexperienced (Mann-Whitney, $\rho = 0.13$, $n = 35$) and experienced (Mann-Whitney, $\rho = 0.66$, $n = 35$) subjects. Similarly, we do not find SUM_{10} for experienced subjects to be significantly different from 100 for homogeneous (Signrank, $\rho = 0.54$, $n = 10$) and heterogeneous (Signrank, $\rho = 0.20$, $n = 23$) events.

Finally, Figure 5 details the $DIFF_t$ for homogeneous and heterogeneous events, apportioned by rounds. Now with the exception of round 5, there does not seem to be any differences in $DIFF_t$ amongst both events. Panel B of Table 2 also provides some summary statistics for $DIFF_{10}$ over the homogeneous and heterogeneous events. We again do not find any significant differences in $DIFF_{10}$ for inexperienced (Mann-Whitney, $\rho = 0.83$, $n = 35$)

---

[36] For inexperienced subjects, there were 7 and 28 occurrences of homogeneous and heterogeneous information events, respectively. For experienced subjects, the corresponding occurrences are 10 and 23, respectively.

and experienced (Mann-Whitney, $\rho = 0.12$, $n = 35$) subjects. This leads us to the second result.

Result 2 *For inexperienced and experienced subjects, we find that homogeneous and heterogeneous events do not result in significant end of round differences in (i) The ratio of highest and second highest priced assets, (ii) Arbitrage opportunities, and (iii) Differences in the mean absolute deviation of prices from the PIE and FRE. This suggests that the information aggregation properties of markets are independent of whether traders are endowed with homogeneous or heterogeneous private information about the true event.*

4.2   Post-market stage

Turning our attention to the post-market stage, we define a subject to be *correct* if her submitted beliefs correspond to the true event, i.e., $b = x(s^*)$. The corresponding proportions of correct inexperienced and experienced subjects are 0.51 and 0.71, respectively.

Given that both PIE and FRE make precise predictions as to the submitted beliefs of minority and non-minority traders, it seems natural to investigate how such predictions match subjects' submitted beliefs. To do so, we compute the mean absolute deviations of the submitted beliefs from the PIE and FRE beliefs.[37]

Submitted beliefs for inexperienced subjects are significantly closer to the PIE prediction (Signrank, $\rho < 0.01$, $n = 35$). This may not be surprising given that Result 1 suggests that the true event is unlikely to be reflected in prices. It is perhaps more surprising that beliefs submitted by experienced subjects are not significantly closer to either the FRE or the PIE prediction (Signrank, $\rho = 0.79$, $n = 35$) despite the fact that subjects could learn the true event by looking at the closing market prices. This result reveals that individual learning is not a necessary condition for successful information aggregation in prices. That is, it possible for market prices to reflect the true event even when a sizeable proportion of traders remain uninformed of the true event.

This result is consistent with the *marginal trader hypothesis* (e.g., Forsythe and Lundholm, 1990), which posits that it is sufficient to have a few *sophisticated* traders that understand how market prices are related to the states of the world in order to drive prices

---

[37] For example if $x(s^*) = 1$, then the FRE predicts that all 4 minority subjects and all 8 non-minority subjects will submit beliefs $b = 1$. In contrast, the PIE predicts that beliefs $b = 0$ and $b = 1$ will be each submitted by 2 minority subjects. In addition, beliefs $b = 1$ and $b = 2$ will each be submitted by 4 minority subjects.

towards the FRE. We explore this possibility by using the post-market stage to identify the 'sophisticated' traders independent of their market behaviour. We define expert traders to be those who submitted a correct belief in at least nine out of the ten markets they participated in. This categorization yields 24% (20 of 84) expert traders.[38] Table 3 presents mixed-effects regressions comparing the expert traders with the other traders. Columns (1) and (2) reveal that expert traders made larger profits, and held higher share of the true asset at the end of trade. Looking at the expert traders' purchase decisions, we see in Columns (3) and (4) that experts were more judicious in their purchase decisions, making fewer purchases, and being more likely to purchase the true asset compared to non-experts. The MAD in column (6) of the table, defined as the mean absolute deviation of the trader's purchases from the FRE, shows that the experts' mean purchase price is indeed significantly closer to the FRE price. To further test whether expert traders were more likely to push prices closer to the FRE, we look at how prices move with each purchase made by the trader.[39] We define *convergent purchases* to be purchases made at a price that is at least as close to the FRE price as in the preceding transaction of the same asset.[40] Column (5) in Table 3 details the share of convergent purchases out of the trader's total purchases in the round. Close to 61% of non-experts' purchases do not move the price away from the FRE compared to 65% for expert traders. This difference increases to nine percent points and becomes significant at the 10% level if we restrict the analysis to experienced traders only.

Returning to the comparison of homogeneous and heterogeneous events, we find that inexperienced subjects are more likely to submit a correct belief in a heterogeneous event than in a homogeneous event (0.69 vs. 0.39, Fisher exact, $\rho < 0.001$). With experience, however, we see a substantial improvement even in the homogeneous events. Experienced subjects are roughly equally likely to be correct in both event types, submitting correct beliefs in 0.73 and 0.65 of instances for the heterogeneous and homogeneous events, respectively (Fisher exact, $\rho = 0.152$).[41] It thus appears that homogeneous events pose greater difficulty for

---

[38] Only four traders submitted correct beliefs in all ten rounds. The results are robust to setting the criterion at 7, 8, or 10 correct beliefs.

[39] Selling decisions are more difficult to analyse since even traders who are informed about the true event should seek to sell assets at prices away from the FRE.

[40] Suppose that the sequence of purchase prices for an asset is 75, 79, 122, 60, 89 and 89 with the FRE price being 100. Then only the second, fifth and sixth purchases will be defined as convergent as they are no further away from 100 than the preceding purchase. If the FRE price is instead 0, then only fourth and sixth can be defined as convergent purchases.

[41] Alternatively, we can also consider the mean absolute deviations of submitted beliefs from the FRE

Table 3: Expert traders.

	(1) Profits	(2) True in inventory	(3) Total purchases	(4) Proportion true	(5) Proportion convergent	(6) *MAD*
EXPERT	196.048**	9.814**	-2.429*	0.161***	0.044	-8.651***
	(86.193)	(3.946)	(1.325)	(0.040)	(0.029)	(2.240)
Round		-0.038	-0.481***	0.008**	0.004	-4.650***
		(0.157)	(0.107)	(0.004)	(0.004)	(0.290)
Constant	553.322***	22.254***	12.940***	0.352***	0.584***	73.898***
	(24.561)	(0.836)	(1.333)	(0.040)	(0.031)	(3.099)
N	840	840	840	553	768	768

Mixed-effects regressions with random effects for subjects and robust standard errors clustered at the session level. Note that profits do not change across rounds as trade is zero-sum. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

traders to discover the true state of the world. Notwithstanding, this does not impede price convergence. This conclusion should be taken cautiously, however, as the post-market beliefs may be influenced by the *ex-ante* low probability of the homogeneous events.[42] Result 3 summarises:

Result 3 *Inexperienced traders remain uninformed about the true state of the world. Although with experience market prices converge to the FRE prices, traders remain largely*

predictions. For inexperienced subjects, submitted beliefs are closer to the FRE predictions in heterogeneous relative to homogeneous information events (Mann-Whitney, $\rho < 0.001$, $n = 35$). In contrast, we find no significant differences for experienced subjects (Mann-Whitney, $\rho = 0.328$, $n = 35$).

[42] This explanation implies a failure to properly apply Bayes' rule. To see why, first consider the case where $s^* = OOO$ (homogeneous event). By Bayes' rule, each uninformed subject should assign equal probability to the events 0 and 1. However, the prior probability of nature choosing event 1 is also three times more than 0. Thus, the failure to properly apply Bayes' rule will lead to uninformed subjects incorrectly assigning greater weights to the event 1. Now consider the case where $s^* = XXO$ (heterogeneous event). Traders in groups 1 and 2 should believe that events 1 and 2 are possible. Since nature chooses events 1 and 2 each with probability 3/8, such subjects should not weight one event more than the other, even when they fail to properly apply Bayes rule. However, for traders in group 3, the failure to properly apply Bayes' rule will lead them to put greater weights on event 2, given that the prior probability on event 3 is 1/8. Thus, the proportion of correct subjects will be more than 1/2.

Table 4: Minority traders.

	(1) Correct beliefs	(2) Profits	(3) True in inventory	(4) Total purchases	(5) Proportion true	(6) Proportion convergent	(7) MAD
Minority	0.106*	-6.006	2.204**	2.928***	0.002	-0.009	-6.672**
	(0.045)	(30.259)	(0.815)	(0.441)	(0.020)	(0.023)	(2.033)
Round	0.031**		0.022	-0.462***	0.006	0.003	-4.835***
	(0.009)		(0.189)	(0.080)	(0.004)	(0.005)	(0.246)
Constant	0.459***	602.002***	23.789***	11.653***	0.401***	0.594***	72.497***
	(0.047)	(10.086)	(1.125)	(0.414)	(0.024)	(0.026)	(1.507)
N	636	636	636	636	405	581	581

OLS regressions with subject fixed effects and robust standard errors clustered at the session level. Only data from the heterogeneous events is included. Note that profits do not change across rounds as trade is zero-sum. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

uninformed. There is some evidence that a few sophisticated traders both succeed in discovering the true state and are instrumental in driving prices towards the FRE.

4.3 Minority Report

The reasoning model of Section 3.3 predicts that minority traders in the heterogeneous events learn about the true state before the non-minority traders, and are influential at driving prices towards the FRE. In this section we test these predictions using the market behaviour and submitted beliefs of minority traders.

Table 4 provides some evidence supporting the model's predictions. Comparing behaviour in the heterogeneous events within subjects, we find that when (exogenously) placed in the minority role, traders are better informed, as measured by post-market beliefs ($p < 0.05$ for the marginal effect in the regression reported in Column (1)). Column (4) shows that minority traders are more active, making 9% more purchases than non-minorities. They are, however, not more likely to purchase the true asset or to make a convergent purchase, as seen in Columns (5) and (6). Notwithstanding, Column (7) reveals that minority traders buy at closer to the FRE price, suggesting that when minority traders do make a convergent

purchase, their purchase price moves a larger distance from the preceding traded price towards the FRE price. Finally, Columns (2) and (3) show that, at the end of trade, minority traders hold a larger share of the true asset in their inventory, but do not make a larger profit overall compared to rounds in which they are cast in the non-minority role. Result 4 summarises these observations.

Result 4 *Relative to non-minority subjects, minority subjects (i) Are significantly more likely to be correct, (ii) Hold a significantly greater proportion of the true asset in their inventory and (iii) Purchase additional assets at prices significantly closer to the FRE. This provides some tentative evidence that minority subjects are informed about the true event before the others and are influential at driving prices towards the FRE.*

If minority traders are better informed, why are they unable to leverage their knowledge to purchase the true asset at a discount and thus improve their payoffs? A central claim of the *efficient market hypothesis* (e.g., Fama, 1970) is that if markets are strong-form efficient, then any trader seeking to exploit his private information immediately reveals the same information from his demand of assets, thus negating any possible gains. Furthermore, competitive bidding between minority subjects might also reduce the scope for such gains.

The above suggests that whilst minority traders may initially be able to profit their knowledge of the true event, such gains will eventually be eroded over time. To study the development of minority gains over time , we compute the *liquidation value* (i.e., how much the portfolio of cash and assets is worth after the loan of 5,000 ECU is returned) of the trader's portfolio at window t. We then regress for each window t, the liquidation value on the minority subjects, using fixed subject effects with session level clustering. The minority dummy coefficient, $\beta_{1,t}$, thus captures the expected difference in liquidation value between the minority and non-minority roles at window t. Figure 6 plots the estimate for $\beta_{1,t}$ at each window t. The hump shape distribution of β_1 suggests that the liquidation value of a minority trader increases during the early periods of the round (i.e., between windows 2–5), with the gap disappearing as trade proceeds. However, the differences are not significant at the 5% level.

5   Conclusion

This study is motivated by the question as to whether markets are able to aggregate dispersed private information about the true state of the world into prices. We consider an
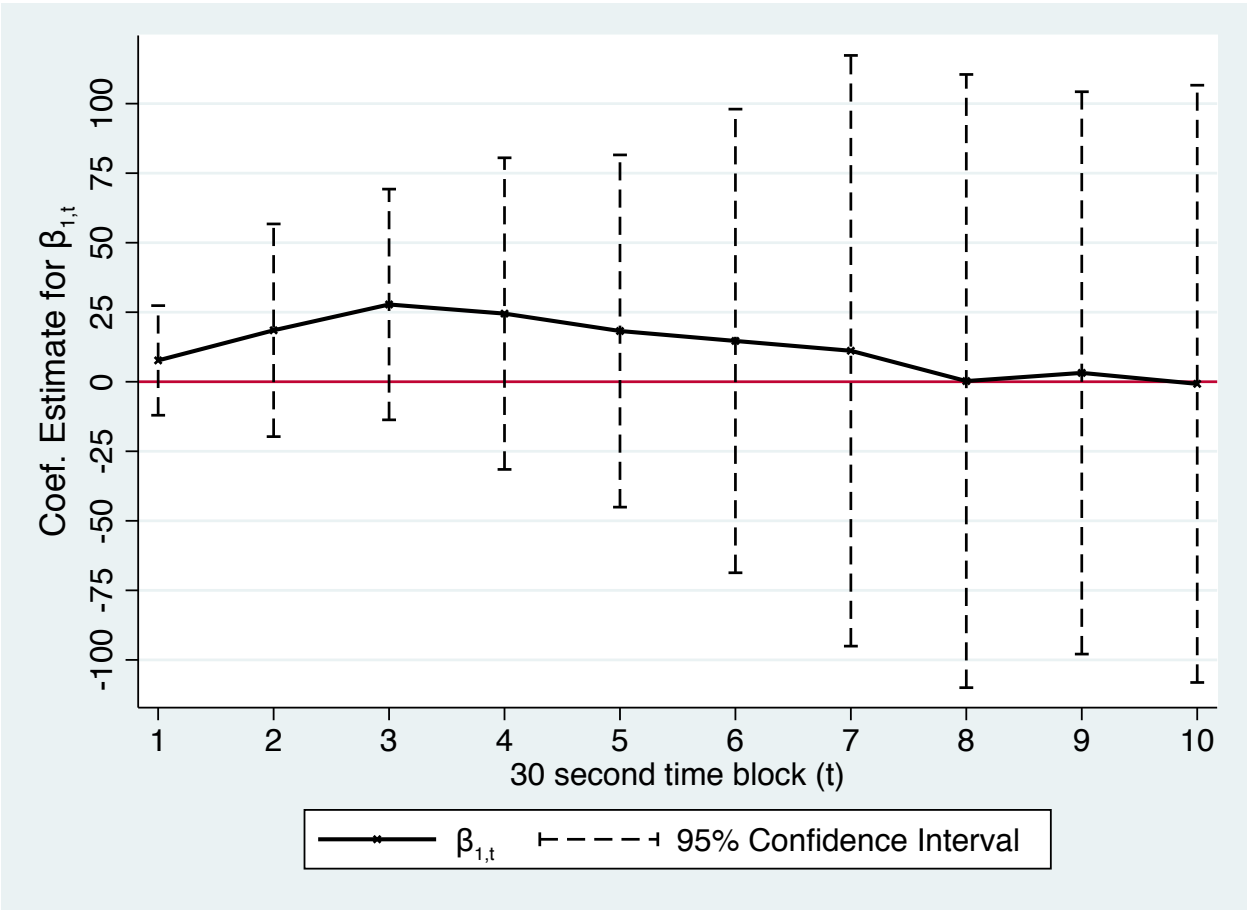
We regress subjects' liquidation value at window t on a constant and situation dummy for minority subjects ($\beta_{1,t}$ coefficient). The above details the estimate of β at each t.

Figure 6: Coefficient estimate of $\beta_{1,t}$.

experimental design which allows us to endow traders with homogeneous and heterogeneous private information about the true event. With this framework, we advance a simple model which shows that markets converge to the fully revealing rational expectations equilibrium for both homogeneous and heterogeneous cases, indicating that the true event will be reflected in prices. However, the process of convergence to the FRE is challenging and requires traders to perform counterfactual reasoning. A post-market stage allows us to estimate individual learning independent of market activity, and to correlate learning and market behaviour.

We find that markets can be successful at aggregating information, but only when market participants have experience with the trading mechanisms. When traders are experienced, the market prices successfully aggregate and reflect the true state, even when traders start out with induced homogeneous beliefs. The requirement for counterfactual reasoning does not appear to impede market convergence. Although the prices reveal the true state, individual learning is strongly constrained. Consistent with the marginal trader hypothesis, we find that a limited number of sophisticated traders is sufficient to drive market prices towards the rational expectations equilibrium. The role traders play in price convergence depends both on endogenous characteristics and on information asymmetries. The last point suggests that even when some traders gain insider knowledge, markets can efficiently price this knowledge.

Taken together, this paper extends the pioneering work by Plott and Sunder (1988) on information aggregation in an environment which we believe to be significantly more complicated. The new design allows us to explore questions of information structure complexity and symmetry, and compare market results to directly elicited post-market beliefs. We conclude that markets composed of experienced traders are able to successfully aggregate dispersed information even in a challenging environment, and even if only a minority of the traders become fully informed in the process.

References

Jeffery S. Banks. Price-conveyed information versus observed insider behavior: A note on rational expectations convergence. *Journal of Political Economy*, 93(4):807–815, 1985.

Martin Barner, Francesco Feri, and Charles R. Plott. On the microstructure of price determination and information aggregation with sequential and asymmetric information arrival in an experimental asset market. *Annals of Finance*, 1(1):73–107, 2005.

Bruno Biais and Sébastian Pouget. Microstructure, incentives and convergence to equilibrium in experimental financial markets. Working Paper, Toulouse University, 1999.

Bruno Biais, Denis Hilton, Karine Mazurier, and Sébastian Pouget. Judgemental overconfidence, self-monitoring, and trading performance in an experimental financial market. *The Review of Economic Studies*, 72(2):287–312, 2005.

Colin Camerer. The promise and success of lab-field generalizability in experimental economics: A critical reply to levitt and list. In G. Frechette and A. Schotter, editors, *Handbook of Experimental Economic Methodology.* Oxford, UK, Oxford University Press, 2015.

Colin F. Camerer and Keith W. Weigelt. Information mirages in experimental asset markets. *Journal of Business*, 64(4):463–493, 1991.

Kay-Yut Chen and Charles R. Plott. Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Social Science Working Paper 1131, California Institute of Technology, 2002.

Thomas E. Copeland and Daniel Friedman. The effect of sequential information arrival on asset prices: An experimental study. *The Journal of Finance*, 42(3):763–797, 1987.

Werner F.M De-Bondt and Richard Thaler. Does the stock market overreact? *The Journal of Finance*, 40(3):793–805, 1985.

J. Bradford De-Long, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. Noise trader risk in financial markets. *The Journal of Political Economy*, 98(4):703–738, 1990.

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

Darren Duxbury. Experimental asset markets within finance. *Journal of Economic Surveys*, 9(4):331–371, 1995.

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

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

Robert Forsythe and Russell J. Lundholm. Information aggregation in an experimental market. *Econometrica*, 58(2):309–347, 1990.

Robert Forsythe, Forrest Nelson, George R. Neumann, and Jack Wright. Anatomy of an experimental political stock market. *The American Economic Review*, 82(5):1142–1161, 1992.

Ben Greiner. Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125, 2015.

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

Sanford J. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

Robin Hanson, Ryan Oprea, and David Porter. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization*, 60(4):449–459, 2006.

Friedrich A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4): 519–530, 1945.

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

David Hirshleifer. Investor psychology and asset pricing. *The Journal of Finance*, 56(4): 1533–1597, 2001.

James S. Jordan. A dynamic model of expectations equilibrium. *Journal of Economic Theory*, 28(2):235–254, 1982.

Charles N. Noussair and Steven Tucker. Experimental research on asset pricing. *Journal of Economic Surveys*, 27(3):554–569, 2013.

John W. O'Brien and Sanjay Srivastava. Dynamic stock markets with multiple assets: An experimental analysis. *The Journal of Finance*, 46(5):1811–1838, 1991.

Michael Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2647, 2012.

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

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

Charles R. Plott and Shyam Sunder. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*, 56(5):1085–1118, 1988.

David P Porter and Vernon L Smith. Stock market bubbles in the laboratory. *Applied Mathematical Finance*, 1(2):111–128, 1994.

Sophia Rabe-Hesket, Anders Skrondal, and Andrew Pickles. Maximum likelihood estimation of limited and discrete dependent variable models with nested random effects. *Journal of Econometrics*, 128(2):301–323, 2005.

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

Andrei Shleifer. *Inefficient Markets*. Oxford University Press, Oxford, 2000.

Shyam Sunder. Experimental asset markets: A survey. In John H. Kagel and Alvin E. Roth, editors, *The Handbook of Experimental Economics*, volume 1, chapter 6, pages 445–500. Princeton University Press, 1995.

Richard H. Thaler. Anomalies: The winner's curse. *Journal of Economic Perspectives*, 2(1): 191–202, 1988.

Helena Veiga and Marc Vorsatz. Information aggregation in experimental asset markets in the presence of a manipulator. *Experimental Economics*, 13(4):379–398, 2010.

Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.

Appendix A: Experimental Instructions

(The experimental instructions were written in Hebrew. Below is the English Translation version)

In this experiment, we are going to simulate a stock market environment where you are able to buy and sell certificates. All payoffs in this experiment will be denoted in Experimental Currency (ECU). The experiment will consist of 2 practice rounds (non-paying) followed by 10 experimental rounds. At the end of the experiment, the computer will randomly select one of the 10 experimental rounds, and your earnings in that round will be paid to you in cash at the exchange rate of 15 ECU to 1 Shekels. In addition, we will also pay you a 25 (15 in Session 1) Shekel show up fee.

We will now describe each experimental round. At the beginning of each round, you will be randomly assigned into either Group A, Group B or Group C. Each group will consist of exactly 4 subjects. Thereafter, the computer will choose a type for each group, which could either be X or O with equal chances. This implies that there is a 1 in 2 chance of your group type being X and a 1 in 2 chance of your group type being O. The table below describes all possible type combinations for the three groups and the chance for each combination. You will not know your own group type but will be able to observe the group type of all other groups.

Group A	X	O	X	X	X	O	O	O
Group B	X	X	O	X	O	X	O	O
Group C	X	X	X	O	O	O	X	O
Probability of combination	1 in 8	1 in 8	1 in 8	1 in 8	1 in 8	1 in 8	1 in 8	1 in 8
Total number of Type X groups	3	2	2	2	2	1	1	0

You will first enter the trading stage, where you have the opportunity to buy and sell 4 classes of certificates labelled S0, S1, S2 and S3. Each participant begins the trading stage with 6 units of each class of certificate and a loan of 5,000 ECU cash, which you will pay back at the end of the round. You buy and sell certificates through a continuous double auction mechanism and the trading period will last for exactly 300 seconds (5 minutes). We shall now describe the mechanism.

The Double Auction Trading Mechanism

Each participant can buy or sell only one class of certificate at each time. Participates will first select the class of certificate they would like to buy or sell and thereafter decide submit the "Bid" and "Ask" price to the market.

- Ask price (between 0 and 200 ECU) = How much you are willing to sell a unit of certificate for.

- Bid price (between 0 and 200 ECU) = How much you are willing to buy a unit of certificate for.

The "Bid" and "Ask" prices of all participants will be instantaneously reflected in the columns: "Market Ask Prices" and "Market Bid Prices"

- Buy: To buy a certificate, simply select the price in the "Market Ask Price" column and click the "Buy" Button

- Sell: To sell a certificate, simply select the price in the "Market Bid Price" column and click the "Sell" Button

The column "Latest Transaction Price" describes the history of transaction prices for each class of certificate. Please note that you are able to buy and sell any class of certificates.

At the end of round, your certificates Ð S0, S1, S2 and S3 Ð will be redeemed at a specific value that depends on the total number of type X groups in the round.

- If there was no type X group, security S0 will be redeemed at 100 ECU and the others at 0 ECU

- If there was only one type X group, security S1 will be redeemed at 100 ECU and the others at 0 ECU

- If there were only two type X groups, security S2 will be redeemed at 100 ECU and the others at 0 ECU

- If there were only three type X groups, security S3 will be redeemed at 100 ECU and the others at 0 ECU

For example, if Group A & Group B & Group C were all type X, then certificates of class S0, S1 and S2 will be redeemed at 0 ECU, whilst certificate of class S3 will be redeemed at 100 ECU. Alternatively, if Group A & B were type X but Group C was type O, then

certificates S0, S1 and S3 will be redeemed at 0 ECU, whilst certificate S2 will be redeemed at 100 ECU.

Submitting your Decision

Before you will know the redemption value of your certificates, you will be presented with the following question, which you must answer.

Qns: "How many type X groups are there in this round?" to which you must answer with either (a) 0 Type X groups, (b) 1 Type X groups, (c) 2 Type X groups or (d) 3 Type X groups

Please note that the question refers to the "total number of type X groups" and not the type of your group. You will be awarded 200 ECU if you are correct and lose 200 ECU if you are wrong.

End of Round Payoffs

After you have completed the decision stage, your payoff for the round are computed as: Payoffs = (Money at the end of Trading Stage) + (Money from Certificate Redemption) + (Payoffs from Decision) – (5,000 ECU loan)

Please feel free to clarify any questions or doubts that you might have with regards to the instructions or the experimental design.

Control questions

Q1 If there is only one type-X group, for how much will each asset (S0, S1, S2, S3) be redeemed?

Q2 If there is only two type-X group, for how much will each asset (S0, S1, S2, S3) be redeemed?

Q3 What is the probability that your group is of type X?

Appendix B: Screenshots
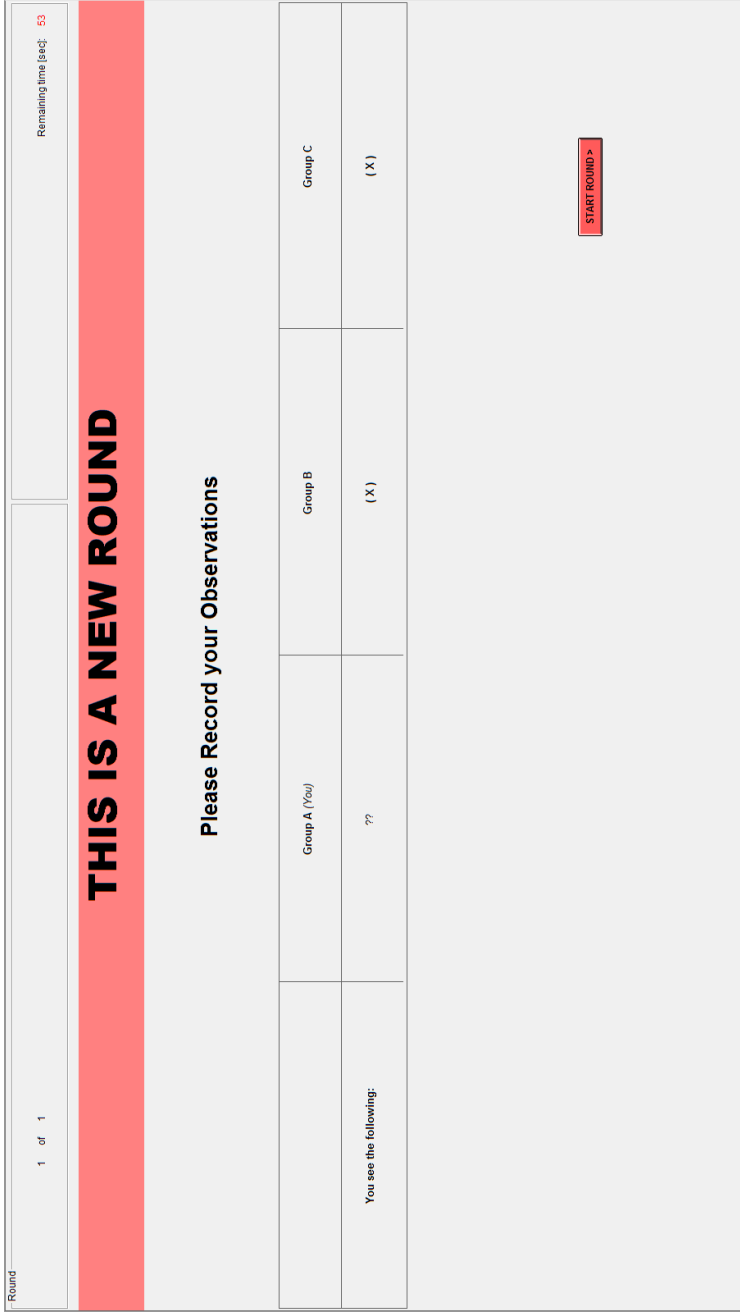
Figure B1: Screen shot, information at the start of a round.

TRADING STAGE

Your Money (ECU): 5000

Cert. S0

No. of S0 Certs. Owned: 6

Your "ASK" Price — ASK

Your "BID" Price — BID

Latest Transaction Price (S0)

Market "ASK" Prices — Buy

Market "BID" Prices — Sell

Cert S1

No. of S1 Certs Owned: 6

Your "ASK" Price — ASK

Your "BID" Price — BID

Latest Transaction Price (S1)

Market "Ask" Prices — Buy

Market "Bid" Prices — Sell

Cert S2

No. of S2 Certs Owned: 6

Your "ASK" Price — ASK

Your "BID" Price — BID

Latest Transaction Price (S2)

Market "Ask" Prices — Buy

Market "Bid" Prices — Sell

Cert S3

No. of S3 Certs Owned: 6

Your "ASK" Price — ASK

Your "BID" Price — BID

Latest Transaction Price (S3)

Market "Ask" Prices — Buy

Market "Bid" Prices — Sell

Figure B2: Screen shot, market trading.

42

Remaining time [sec]: 57

DECISION STAGE

Your Money (ECU): 5000

	Group A	Group B (You)	Group C
You see the following:	(O)	??	(X)

Cert S0		Cert S1	Cert S2	Cert S3
No. of S0 Certs Owned:	6	No. of S1 Certs Owned: 6	No. of S2 Certs Owned: 6	No. of S3 Certs Owned: 6
Latest Transaction Price (S0)		Latest Transaction Price (S1)	Latest Transaction Price (S2)	Latest Transaction Price (S3)

Average Trade Price in Market (ECU):	0	Average Trade Price in Market (ECU): 0	Average Trade Price in Market (ECU): 0	Average Trade Price in Market (ECU): 0
Volume of Trade (units):	0	Volume of Trade (units): 0	Volume of Trade (units): 0	Volume of Trade (units): 0

How many type X groups are there in this Round?
- ○ 0 Type X Group
- ○ 1 Type X Group
- ○ 2 Type X Groups
- ○ 3 Type X Groups

SUBMIT ANSWERS

Figure B3: Screen shot, post-market submission of beliefs.

Appendix C: Market prices

C.1 Session Level Data

The experimental data from the Market and Post-market stages of sessions 1 to 7 are presented in Figures C1–C7, respectively. The panels in each figure are arranged with rounds (1–10, increasing from left to right) and assets (increasing from top to bottom) along the columns and rows, respectively. For example, the top left panel of Figure C1 details the transacted prices, unfilled bids and asks prices of asset 0 in round 1 of session 1. The FRE and PIE prices are depicted in each panel with the solid and dashed horizontal lines, respectively.

The header of each panel presents the round, followed by the asset type and the proportion of submitted beliefs corresponding to the asset. For example, 8% subjects in round 1 of session 1 submitted the belief that the true event is $x(s^*) = 0$.
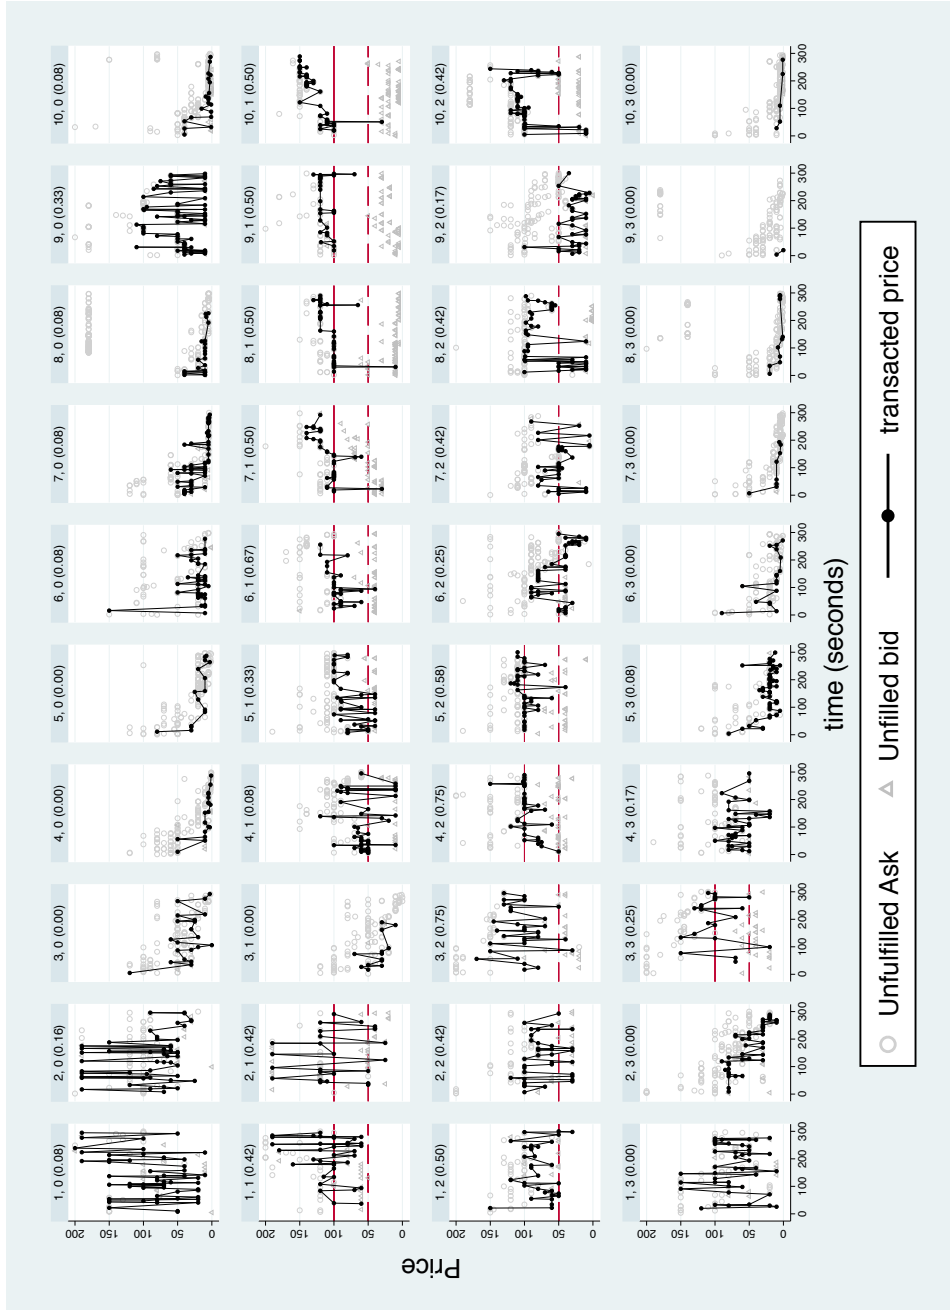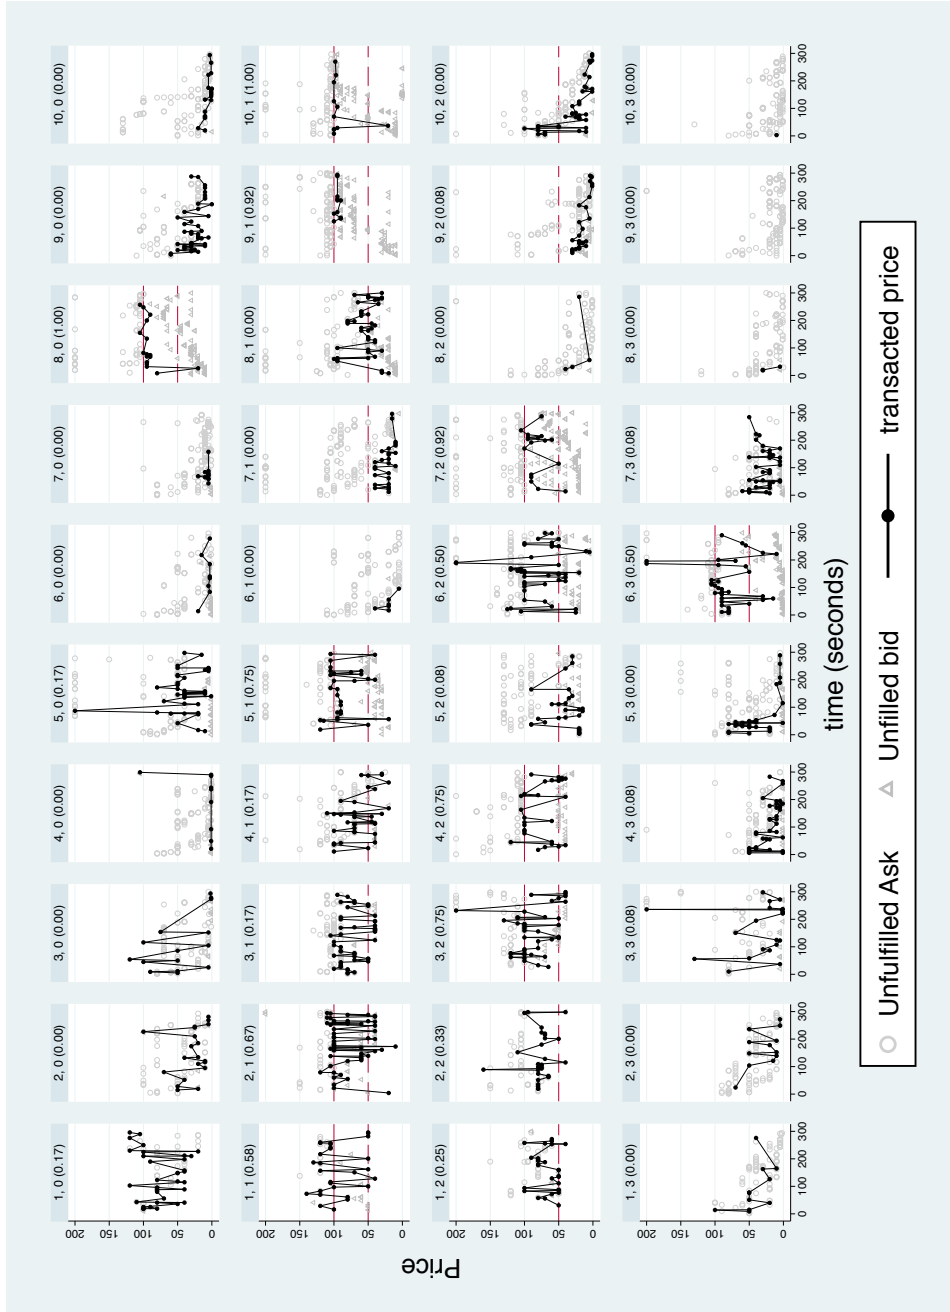
Figure C1: Session 1 Data.

The header of each panel details the round, followed by the asset type and finally, the proportion of subjects who submitted beliefs $b = \alpha$ (in parenthesis).

45
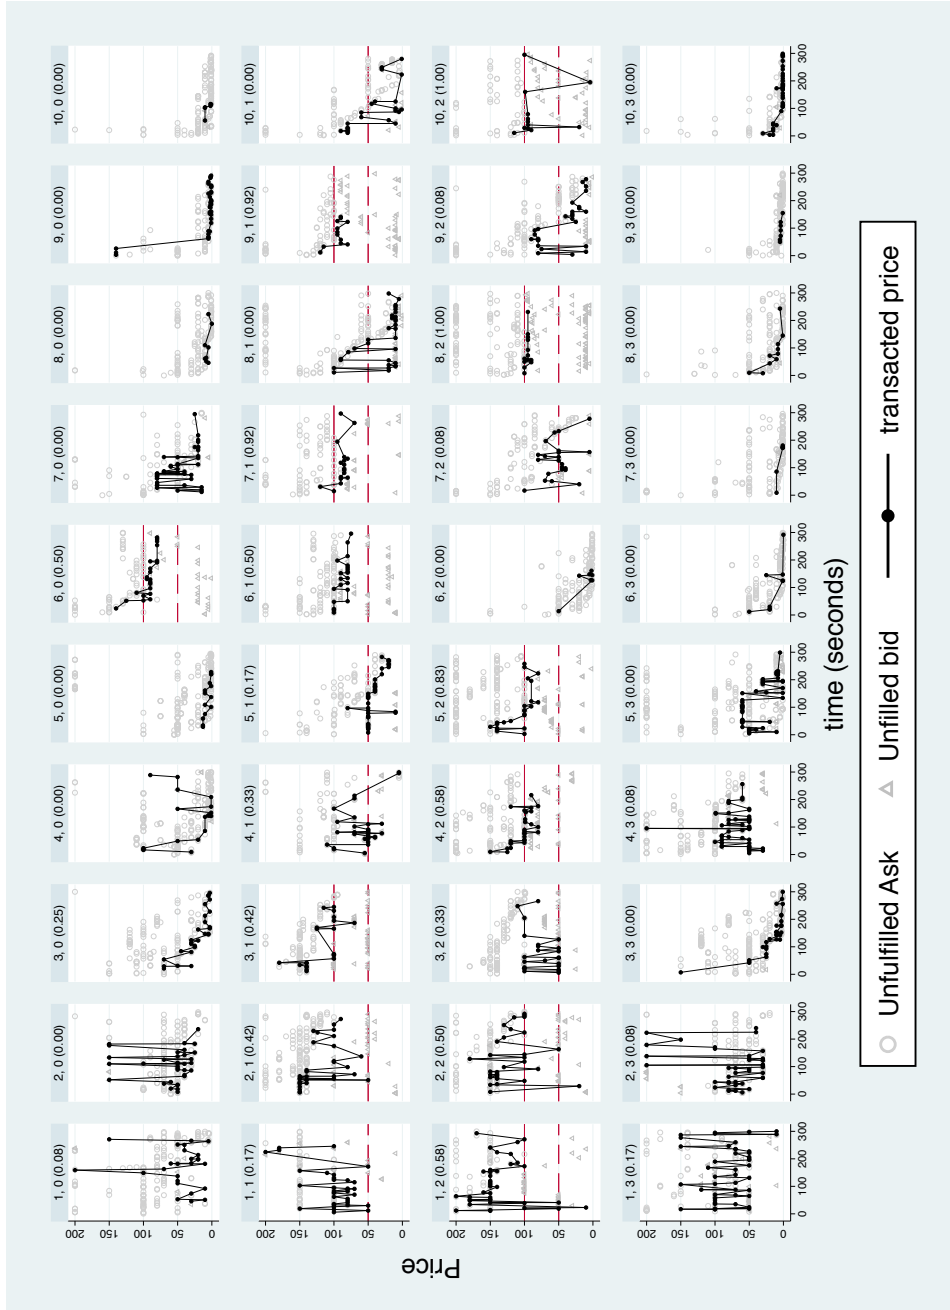
Figure C2: Session 2 Data.
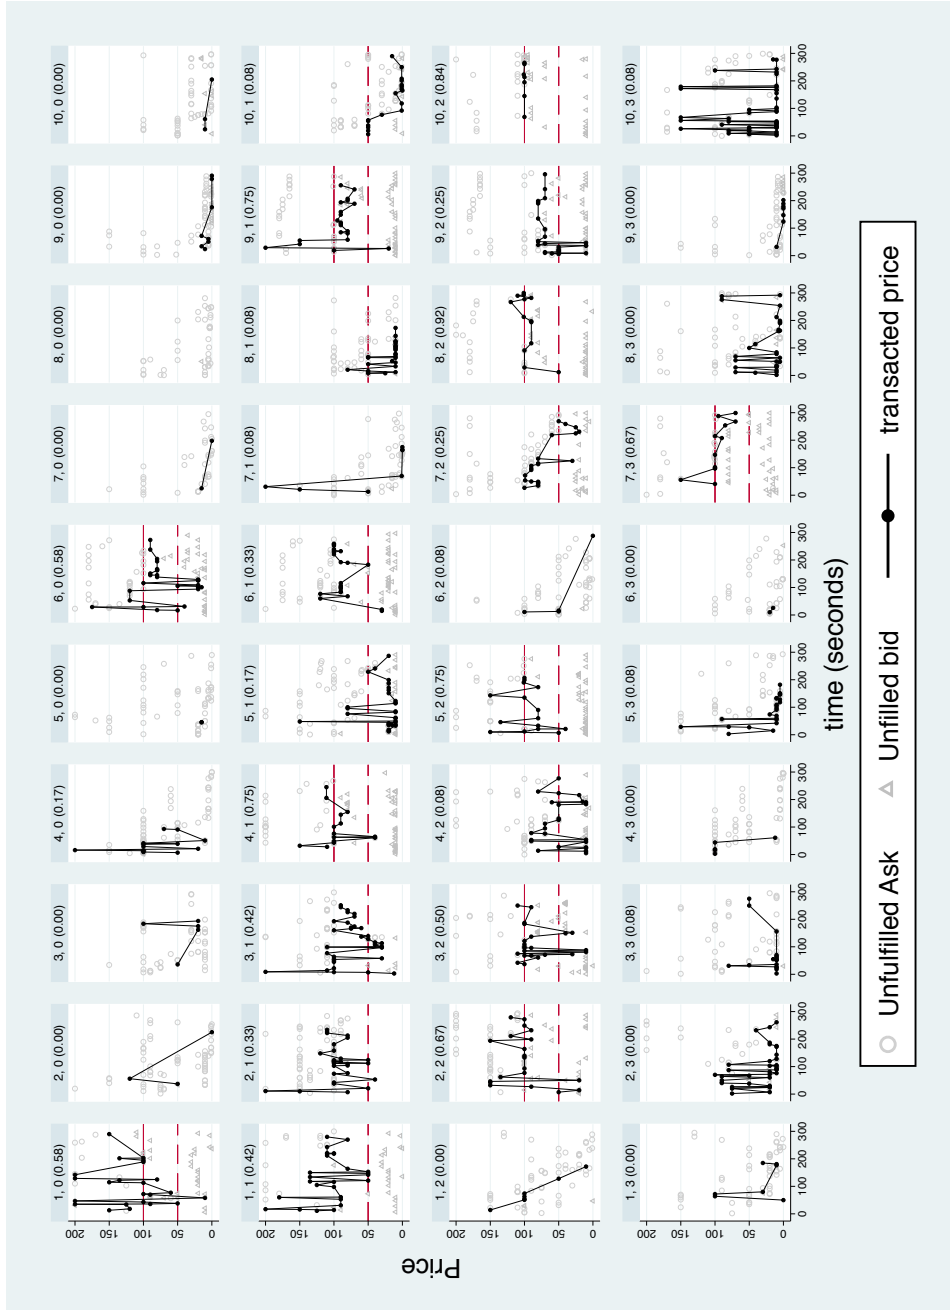
The header of each panel details the round, followed by the asset type and finally, the proportion of subjects who submitted beliefs $b = \alpha$ (in parenthesis).

Figure C3: Session 3 Data.
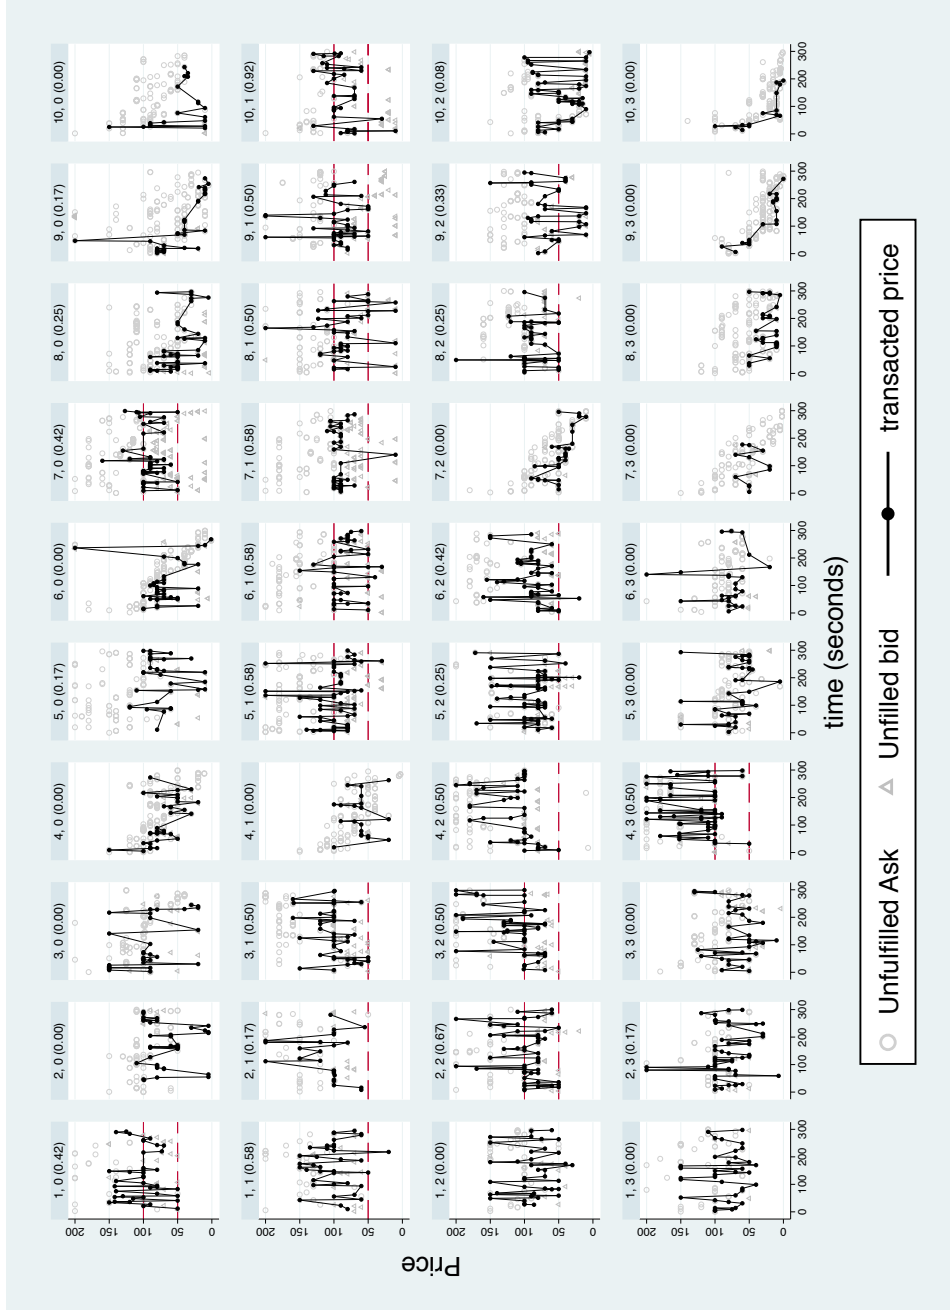
The header of each panel details the round, followed by the asset type and finally, the proportion of subjects who submitted beliefs $b = \alpha$ (in parenthesis).

47

Figure C4: Session 4 Data.
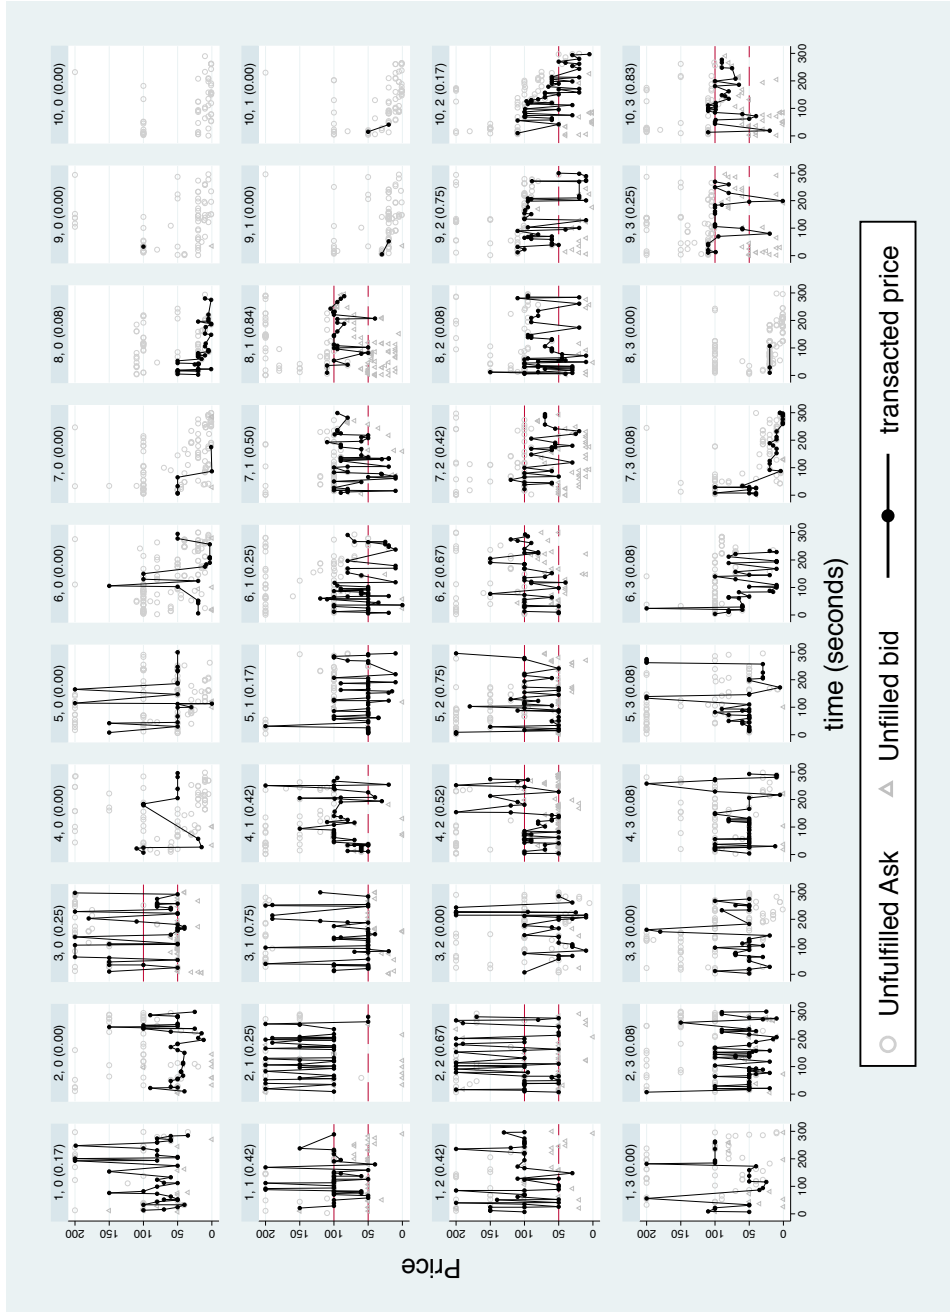
The header of each panel details the round, followed by the asset type and finally, the proportion of subjects who submitted beliefs $b = \alpha$ (in parenthesis).

Figure C5: Session 5 Data.
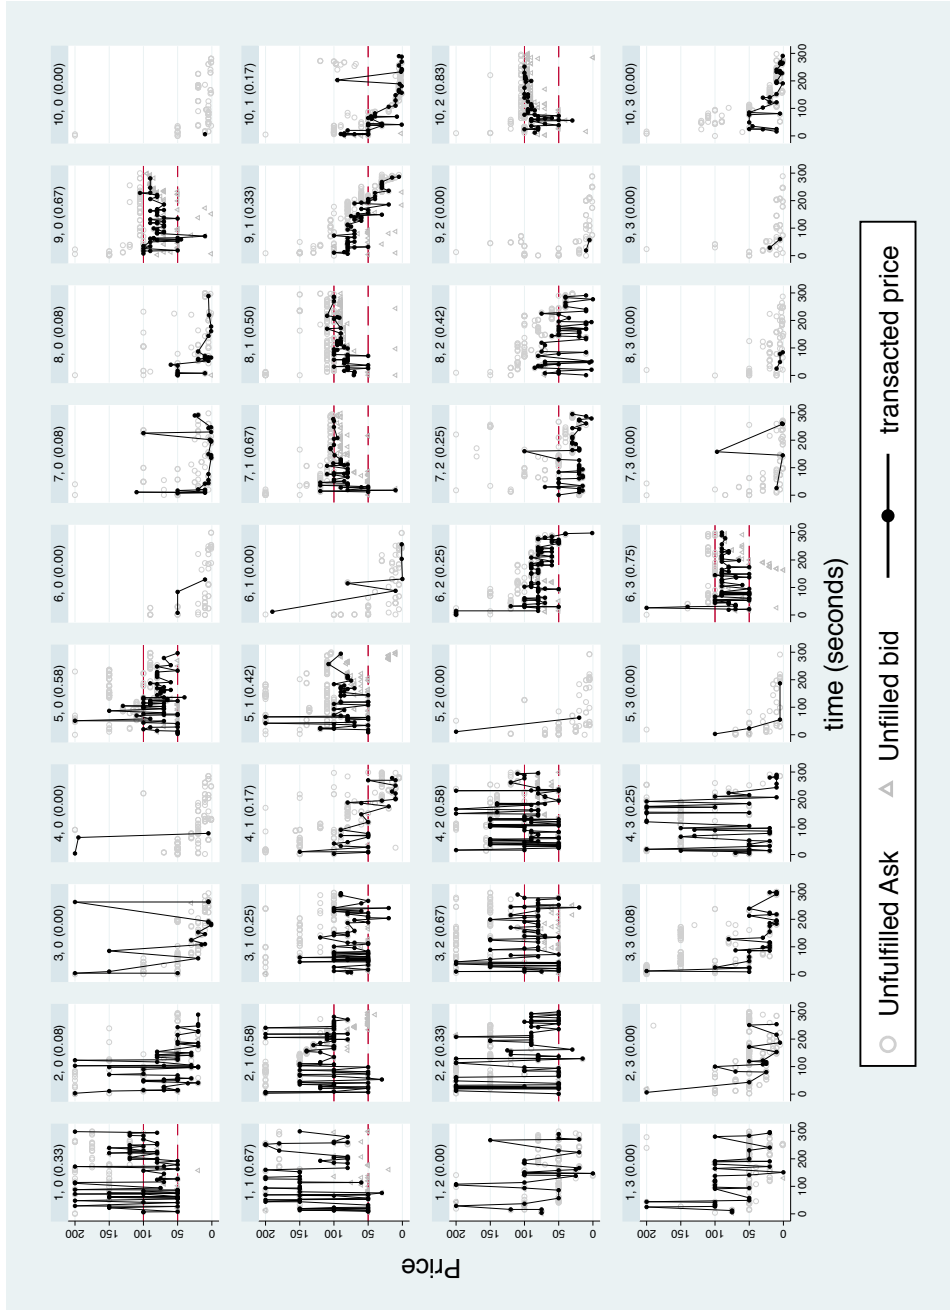
The header of each panel details the round, followed by the asset type and finally, the proportion of subjects who submitted beliefs $b = \alpha$ (in parenthesis).

49

Figure C6: Session 6 Data.

The header of each panel details the round, followed by the asset type and finally, the proportion of subjects who submitted beliefs $b = \alpha$ (in parenthesis).

Figure C7: Session 7 Data.

The header of each panel details the round, followed by the asset type and finally, the proportion of subjects who submitted beliefs $b = \alpha$ (in parenthesis).

51

C.2 Further information about relative prices

In section 4.1.1 of the main paper, we provide a measure of market efficiency based on the relative prices of the highest (labelled as asset A) and second highest (labelled as asset B) assets at time window t. Based on the above, we derive f_t and f_t^*, which can be defined for any $\beta \geq 1$. In the main paper we plot f_t and f_t^* for $\beta = 1, 1.5, 2, 2.5, 3$. On Figure C8, we plot f_t and f_t^* for $\beta = 1, 1.1, 1.2, 1.3, 1.4, 1.5$. As with higher values of β, f_t and f_t^* decrease as β increase. f_t and f_t^* remain rather constant across the duration of the round when subjects are inexperienced, but trend upwards for experienced subjects.
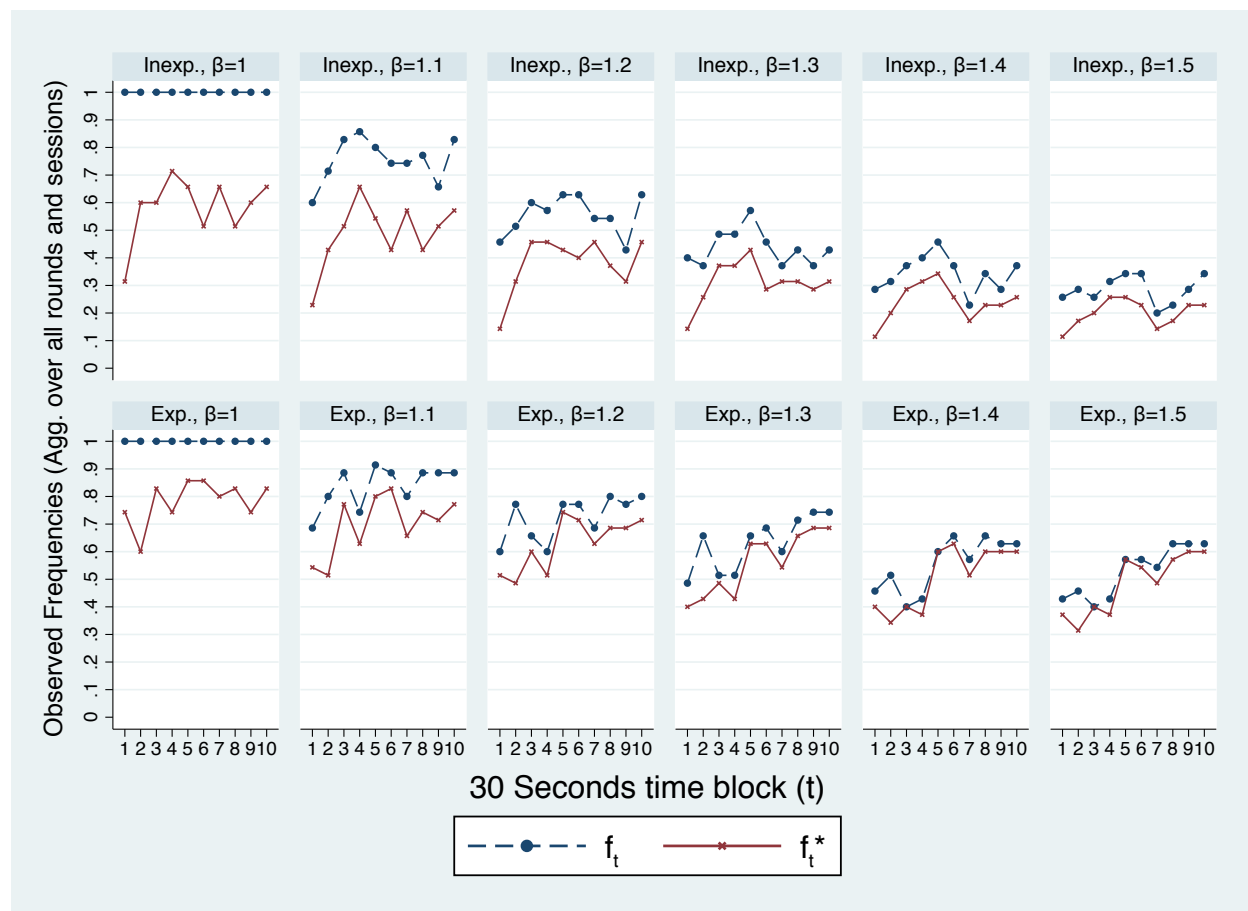
Figure C8: f_t and f_t^* for $\beta = 1, 1.1, 1.2, 1.3, 1.4, 1.5$.

Appendix D: Subject level data

D.1 Seemingly irrational purchases and bankruptcy

Purchases of assets above 100 ECU are "irrational" in the sense that the maximum redemption value of each asset is 100 ECU. It is, however, rational to purchase an asset at a high price in order to manipulate other traders' beliefs and sell more units of the asset at a gain, or if the buyer believes she will be able to resale the asset at a higher price. Such seemingly irrational purchases form 19% and 9% of all purchases for inexperienced and experienced subjects, respectively. Subjects incurred bankruptcies in 75 (17%) and 39 (9%) instances for inexperienced and experienced subjects, respectively.[43] In this section, we elaborate on these two observations and how they might be linked.

We assigned a unique ID number to each of the 84 subjects: 1–12 (session 1), 13–24 (session 2), 14–36 (session 3), 37–48 (session 4), 49–60 (session 5) 61–72 (session 6) and 73–84 (session 8). Figure D1 presents the number of irrational purchases made by inexperienced and experienced subjects. For inexperienced subjects, irrational purchases seem to be evenly spread across all subjects. In contrast, irrational purchases for experienced subjects seem to cluster around a handful of subjects, most notably subject 6, 11 and 37. As we will later show, experienced subject 6 will be of special interest as he was observed to make more than 70 irrational purchases.

Figure D2 presents the number of bankruptcies for inexperienced and experienced subjects. Except for experienced subject 6, there seems to be a correlation between the frequency of irrational purchases and bankruptcies (Spearman correlations are $\rho = 0.63$ and $\rho = 0.35$ for inexperienced and experienced subjects, respectively). This should be expected given that our payoff structure implies that subjects who often purchase assets above 100 ECU or a non-true asset above 50 ECU make a negative payoff in expectations.

It is somewhat peculiar that experienced subject 6 frequently made seemingly irrational purchases, but never went bankrupt. To shed some light on this matter, we focus on the purchasing behaviours of subjects. From Figure D1, we observe that experienced subject 6 had frequently purchased the true asset. Some of such also occurred above 100 ECU at the range between 101–120 ECU. At the same time, the same subject also frequently sold the non-true assets at prices above 0 ECU. As such, experienced subject 6 often compensated the losses from seemingly irrational purchases with the gains from selling off the non-true

---

[43] A subject is said to be bankrupt if his payoffs at the end of the round are negative.

assets. Figure D4 details the average share of true asset held in the subjects' inventory. Here, we see that the true asset form around 80% of the experienced subject 6's inventory. For the above reasons, the subject avoided bankruptcy.
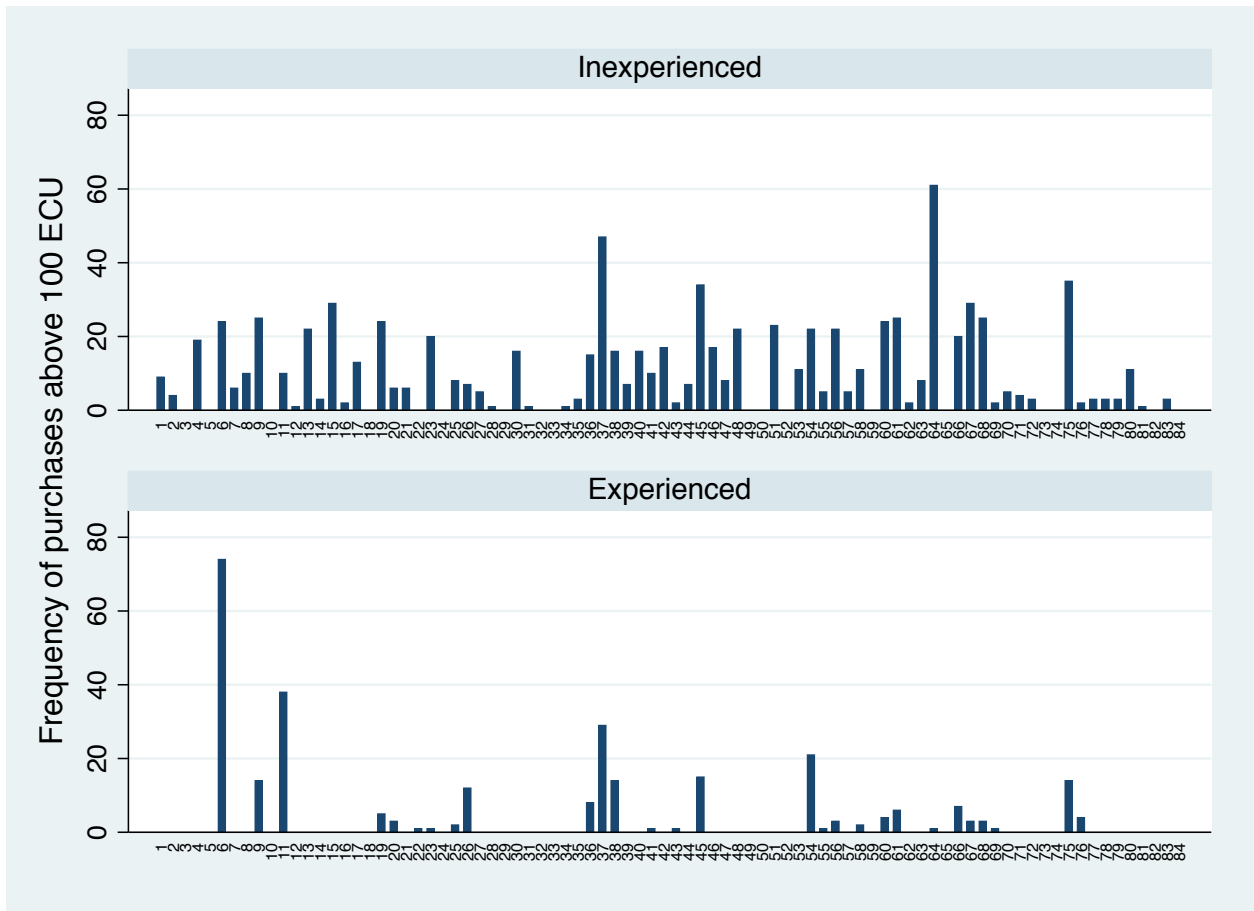

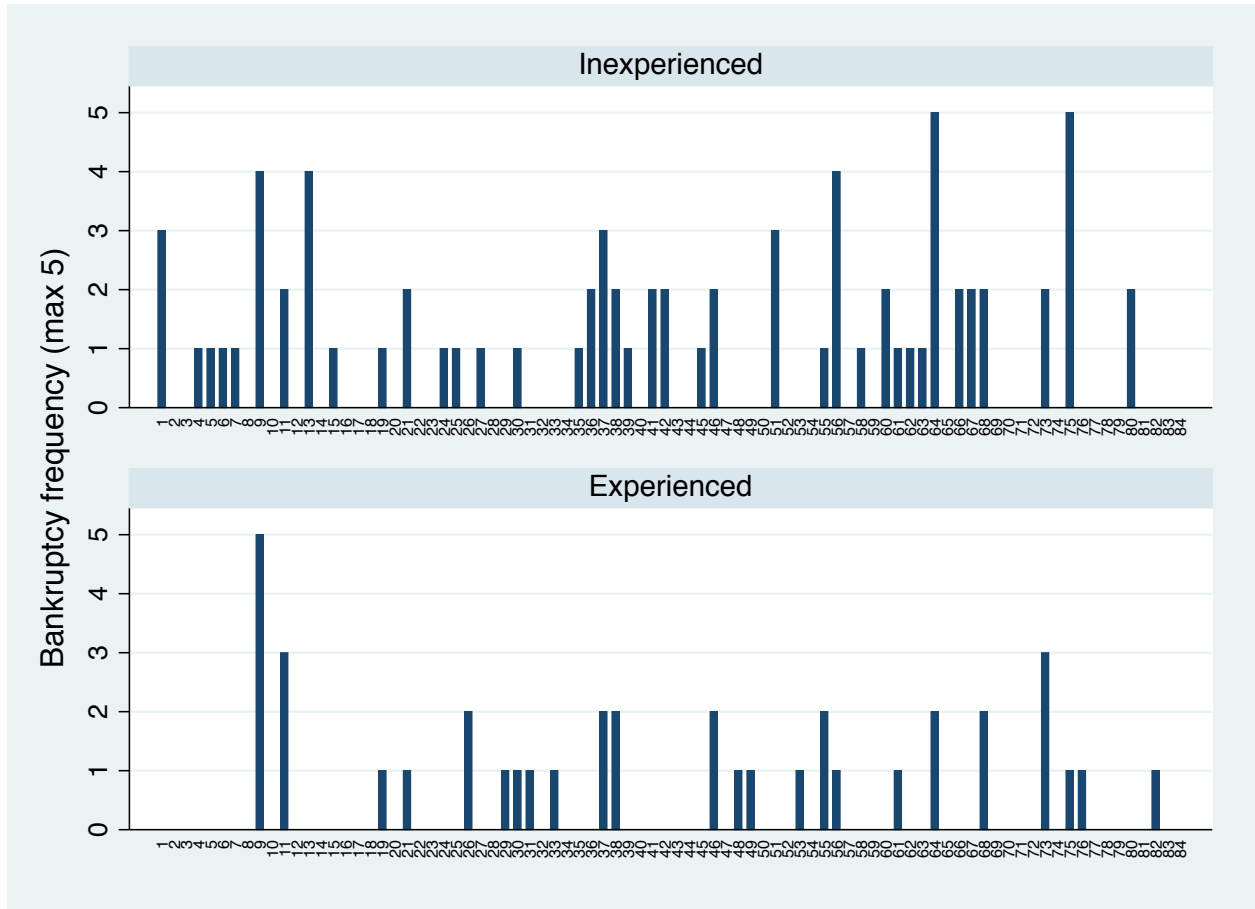Figure D1: Frequency of purchases at prices above 100 by subject ID.

Figure D2: Frequency of bankruptcies by subject ID.

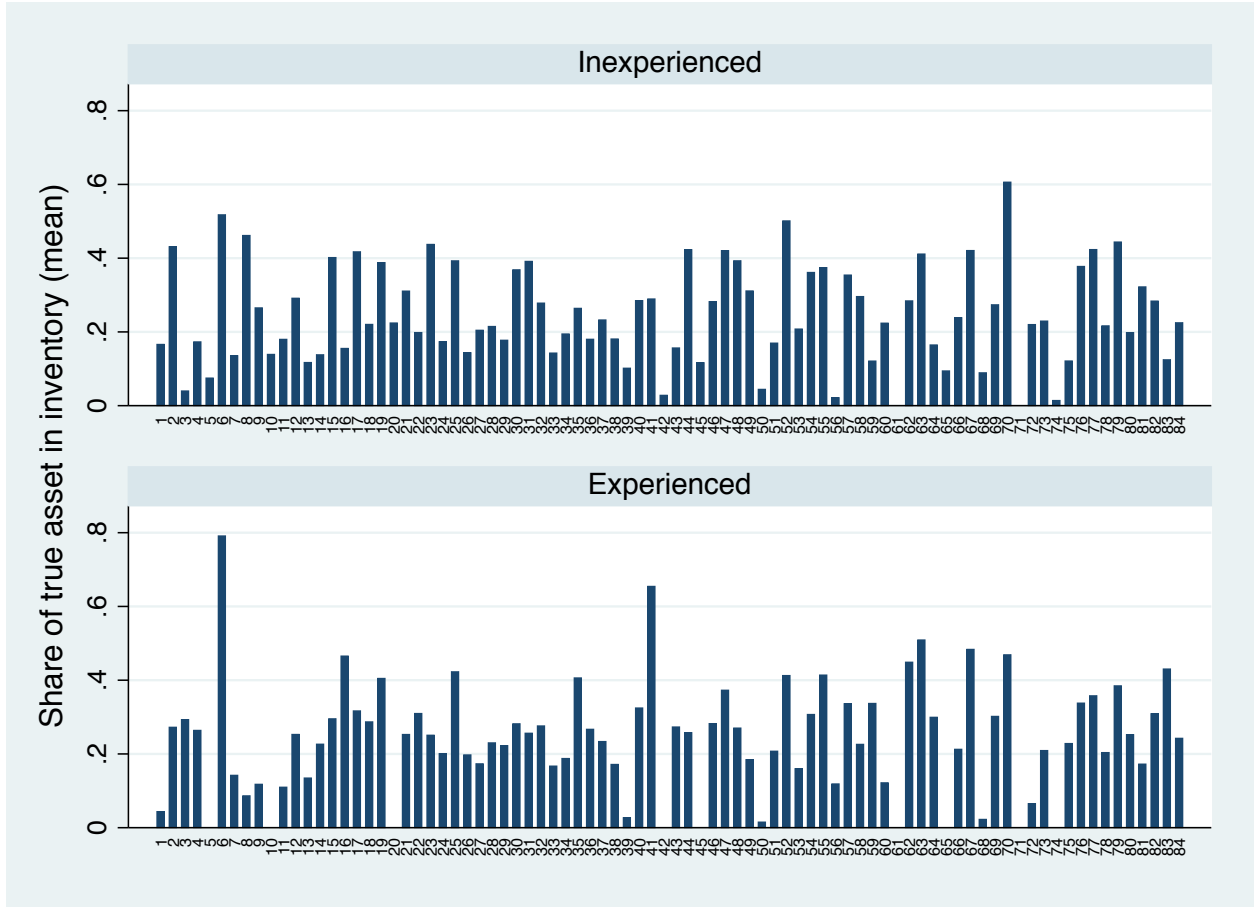Figure D3: Frequency of true asset owned by subject ID.

Figure D4: Proportion of true asset purchases by subject ID.

D.2 Non-Minority and minority subject differences

The following graphs provide details as to how the behaviours of the non-minority and minority subjects differ across the rounds. The graphs are based on data from the heterogeneous information events, where minority traders exist.

Figure D5 presents the proportion of correct minority and non-minority subjects at each round. With the exception of rounds 4 and 5, minority subjects are more frequently observed to be correct.

Figure D6 presents the mean proportion (%) of true assets in the inventory of minority and non-minority subjects at each round. Here, we observe that from round 7 onwards, minority subjects clearly hold a greater proportion of the true asset in their inventory. The differences from the earlier rounds are less obvious.

Figure D7 presents the mean number of convergent purchases by minority and non-minority subjects at each round. Here, we see some evidence that mean convergent purchases are more frequently higher for minority subjects.

Figure D8 presents the mean absolute deviation of purchase prices from FRE. Here, we see that deviations are somewhat higher for non-minority relative to minority subjects, though the differences seem to be decreasing with rounds.

Figure D9 details the mean market stage payoffs for minority and non-minority subjects. There does not seem to be any obvious pattern as to differences in payoffs between both groups of subjects.
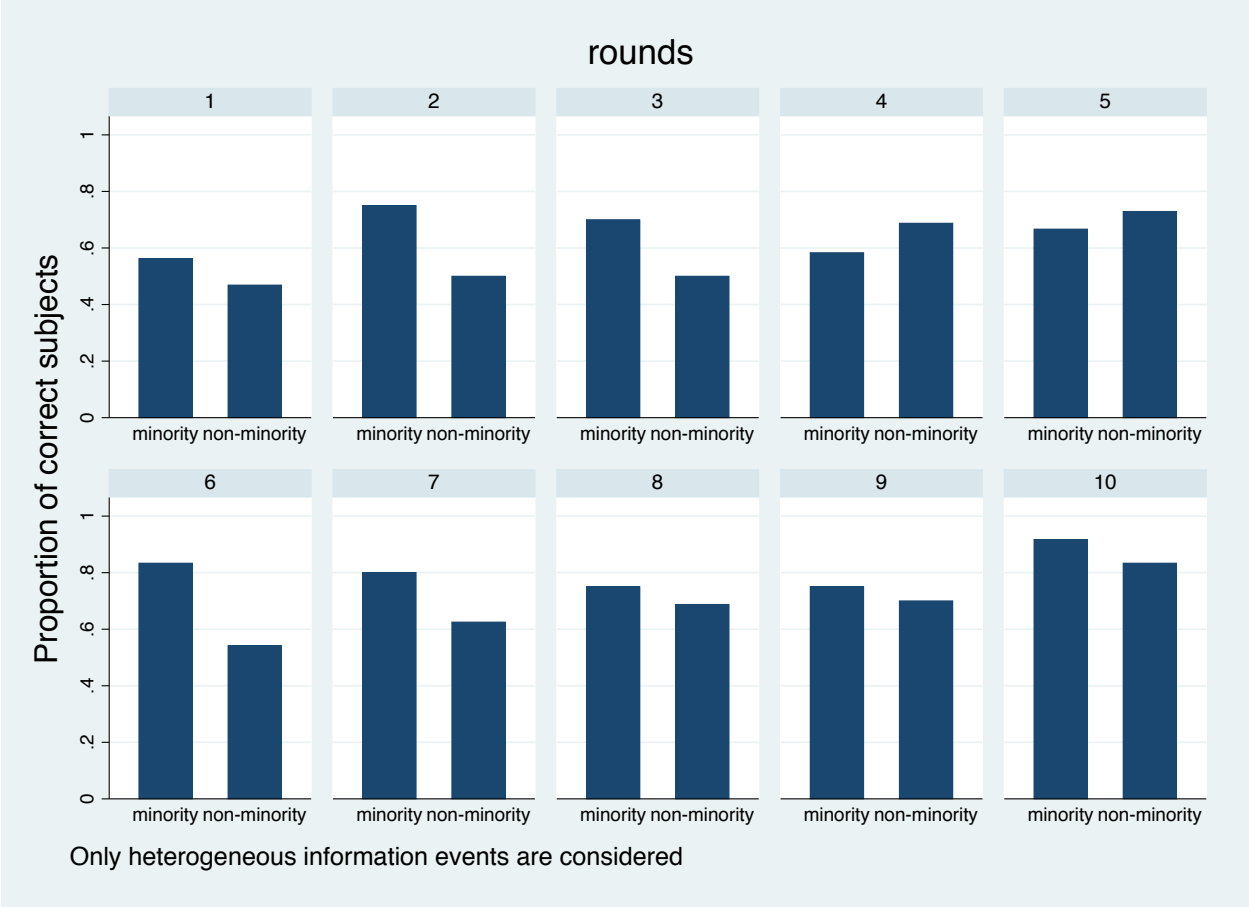
Figure D5: Proportion of correct minority and non-minority subjects.
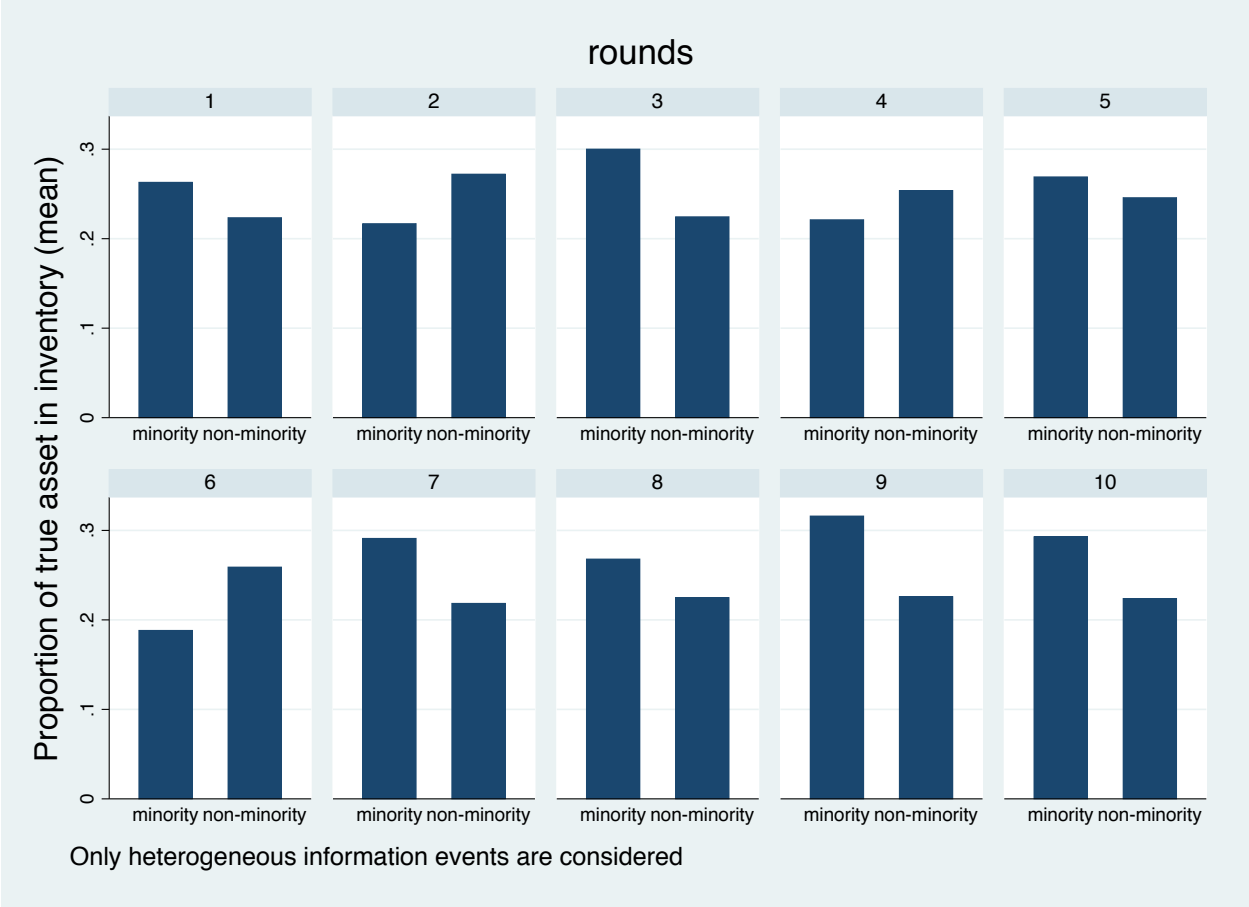
Figure D6: Proportion of True assets in inventory of minority and non-minority subjects.
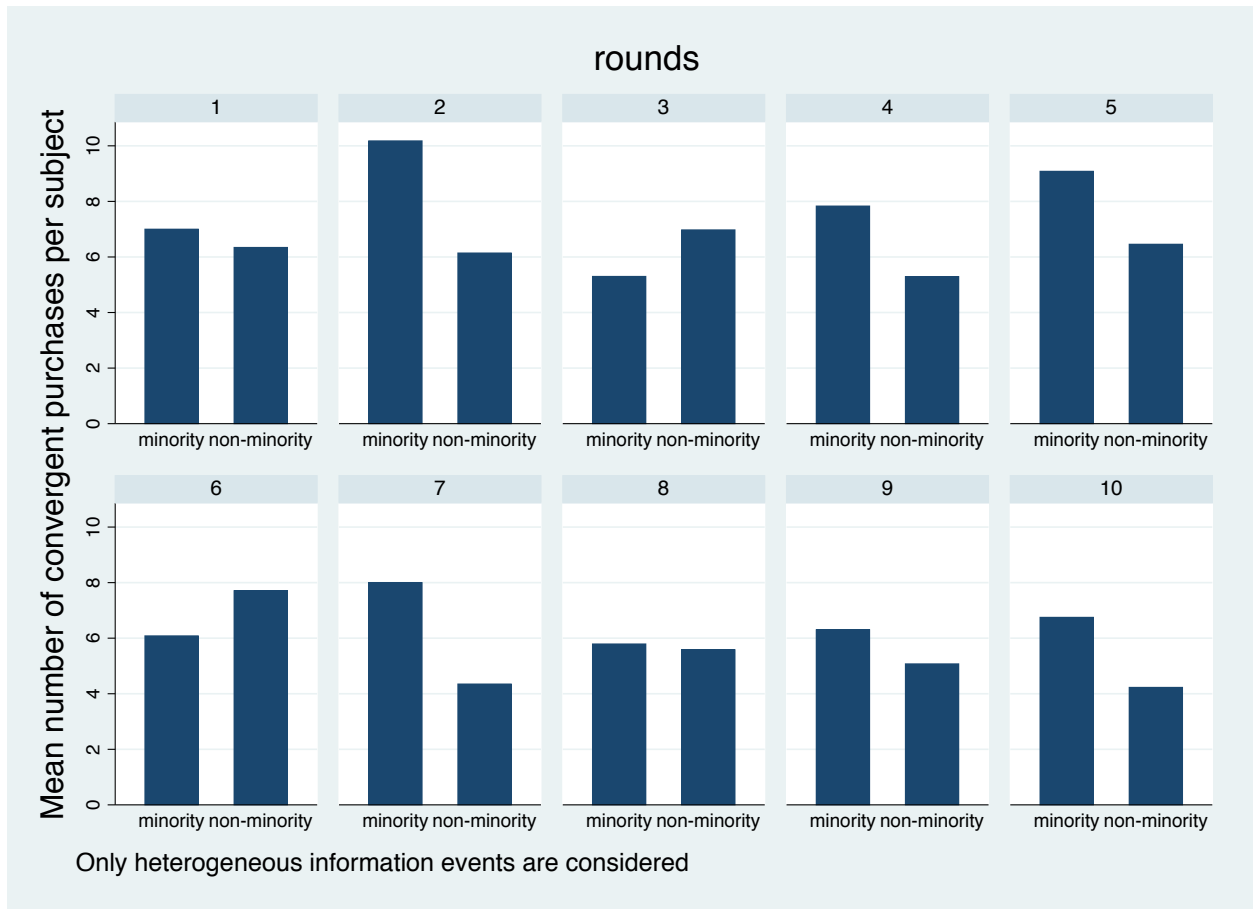
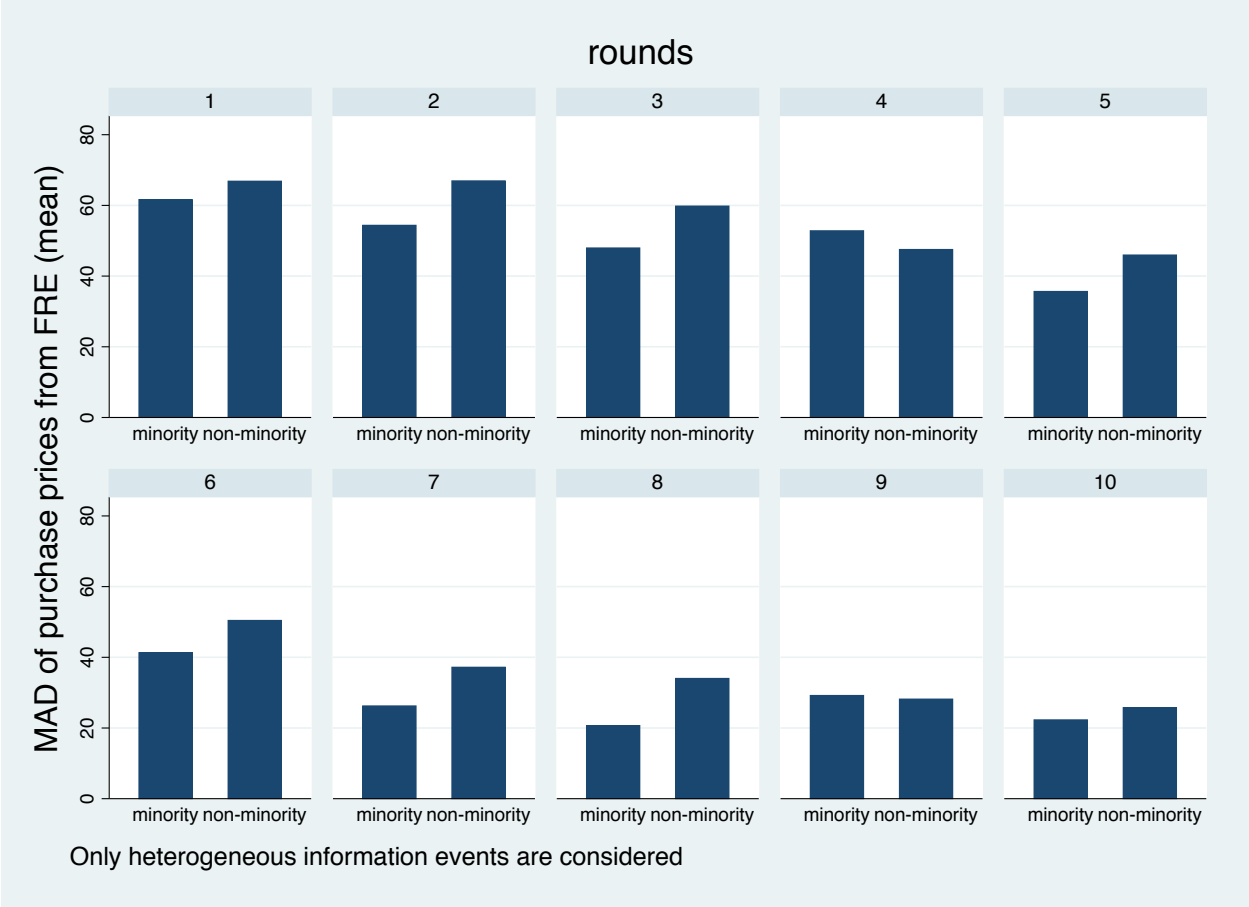Figure D7: Number of convergent purchases per minority and non-minority subjects.

Figure D8: Mean absolute deviation of purchase price from FRE for minority and non-minority subjects.
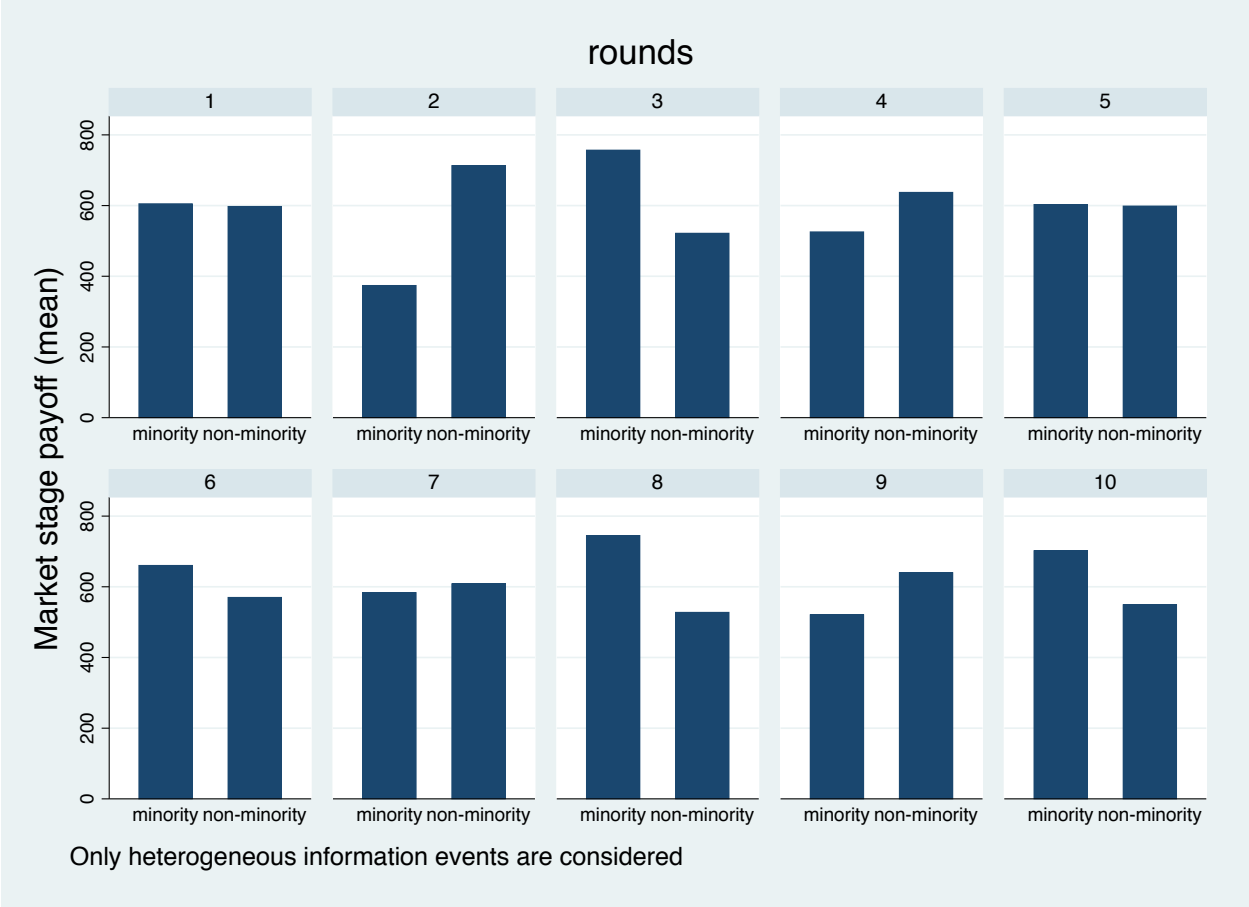
Figure D9: Market stage payoff for minority and non-minority subjects.