10-2018

The Distribution of Information and the Price Efficiency of Markets

Brice Corgnet
Chapman University, corgnet@chapman.edu

Mark DeSantis
Chapman University, desantis@chapman.edu

David Porter
Chapman University, dporter@chapman.edu

The Distribution of Information and the Price Efficiency of Markets

Brice Corgnet*, Mark DeSantis** and David Porter**

Abstract

Apparently contradictory evidence has accumulated regarding the extent to which financial markets are informationally efficient. Shedding new light on this old debate, we show that differences in the distribution of private information may explain why informational efficiency can vary greatly across markets. We find that markets are informationally efficient when complete information is concentrated in the hands of competing insiders whereas they are less efficient when private information is dispersed across traders. A learning model helps to illustrate why inferring others' private information from prices takes more time when information is more dispersed. We discuss the implications of our findings for understanding the potential consequences of lowering the cost of information on the informational efficiency of markets.

Keywords: Information aggregation, information dispersion, market efficiency, experimental asset markets, learning models, cognitive finance.

JEL CODES: C92, G02, G14.

* EMLYON Business School, GATE L-SE 69130 Ecully, France
** Argyros School of Business and Economics & Economic Science Institute, Chapman University, Orange, CA, 92866

1. Introduction

Markets are essential to well-functioning societies. Yet, the debate regarding the efficiency of markets is still rampant. To shed light on this conundrum, we study whether differences in the structure of information between markets can account for the differences in their informational efficiency. In particular, we posit that the concentration of private information across market participants is a key determinant of the informational efficiency of a market.

Because it is impossible to clearly assess the informational efficiency of markets with archival data (Fama, 1991), we make the methodological choice to rely upon experimental markets. A growing number of studies have demonstrated the unique benefits of the experimental method in allowing researchers to control and thus clearly identify the extent to which private information is captured by market prices (e.g., Bossaerts, 2009; Frydman et al. 2014; Noussair and Tucker, 2014).

We extend the seminal work by Plott and Sunder (1988) where they find that with a complete set of Arrow Debreu spanning securities markets are efficient information gatherers. The structure of their environment provides a particular insight into how information is quickly and efficiently transmitted in the market. With spanning markets, one-half of the traders are *fully* informed in all but one of the markets. Thus, these markets are populated by competing insiders who will reveal the states that cannot occur thus making the true asset value transparent to all traders.

We build on this work to assess whether the presence of insiders is indeed the driving force leading markets to aggregate private information. We consider market experiments in which the asset can take one of three possible values. In our markets, traders can either be *uninformed*, *partially* informed or *fully* informed. The *uninformed* trader only knows the prior distribution of the three possible asset values. The *partially* informed trader is given one private *signal* regarding the value (out of the three possible ones) the asset will not take. The *fully* informed trader knows with certainty the value the asset will assume at the end of the market. If markets are informationally efficient then prices should reflect all the available private information regardless of the distribution of private information (Fama, 1970). It follows that the presence or absence of *fully* informed traders will not affect prices as long as the aggregate information of all traders in the market is complete.

To derive conjectures regarding the informational efficiency of our experimental markets, we rely on a learning model following the work of Friedman (1991) and Copeland and Friedman (1987; 1991). We extend their approach to the case in which a proportion of the traders may be

boundedly rational. Following the work of Corgnet, DeSantis and Porter (2015; 2018) on the cognitive underpinnings of trading in experimental markets, we assume traders to be either *reflective*, in which case they properly apply Bayes' rule to infer the true asset value from market orders,[1] or *non-reflective* in which case they exclusively rely on their private information to value the asset. Corgnet, DeSantis and Porter (2015) show that the distinction between *reflective* and *non-reflective* traders is crucial for understanding the informational efficiency of markets. This distinction echoes previous research in Finance regarding people's limited capacity to learn others' private information from prices (e.g., see Eyster, Rabin and Vayanos, 2018).

Using model-based simulations, we establish several conjectures regarding the informational efficiency of markets and learning dynamics. Keeping the number of private *signals* constant across markets, the simulations indicate that informational efficiency should be higher in markets populated by both *fully* informed and *uninformed* traders than in markets solely populated by *partially* informed traders. Our simulations show that *uninformed* or *partially* informed traders learn the true asset value substantially faster when *fully* informed traders are present in the market. Intuitively, orders involving *fully* informed traders carry substantially more information than orders involving *partially* informed traders thus facilitating Bayesian inference and the transmission of private information into prices. *Fully* informed traders provide a clear signal of the asset's price relative to its true value as they follow a riskless strategy of buying (selling) when the price is below (above) the true value.[2] Our simulations also predict that the efficiency of markets will increase and trading volumes will decrease as the number of *fully* informed traders rises. Finally, our simulations also deliver a point prediction according to which a market populated with at least three *fully* informed traders should reach a higher level of informational efficiency than a market populated entirely by *partially* informed traders.

Our experimental findings provide clear-cut support for our conjectures including our point prediction. Our experimental results also support our model by showing that markets populated by a higher number of *reflective* traders (assessed using the cognitive reflection test, Frederick, 2005) reach higher levels of informational efficiency.

---

[1] We define the term "market order" to include bids, asks, and contract prices rather than to distinguish between limit orders and immediately executable orders. As such, throughout the manuscript this term includes not only transaction prices but also "limit orders".
[2] Unlike the works of Hanson, Oprea and Porter (2006) and Veiga and Vorsatz (2010), our simulated traders do not engage in strategic manipulations of the market.

2. Contribution to previous literature

Efficient Market Hypothesis

Our findings show that the distribution of private information has a significant effect on the informational efficiency of the market. This may explain why some markets have been found to be informationally efficient whereas others have not (e.g., Fama, 2008).[3] Future research should assess whether there is, at least, some correlational evidence between the presence of insiders in stock markets and their informational efficiency. Our study can thus be seen as a starting point for reconciling the seemingly contradictory pieces of evidence fueling the debate between classical and behavioral finance.

Because our findings are based on experimental markets, we are able to identify the causal effect of the distribution of information on the informational efficiency of markets thus making our findings immune to the usual data-mining critique (e.g., Fama, 2008). We have isolated the distribution of private information as a systematic factor influencing the informational efficiency of markets.

Models of Informational Efficiency

Our experimental data are consistent with the predictions of our learning model thus providing further support for the work of Friedman (1991) which stressed the need for a simplification of the modeling of continuous double auctions settings to abstract away from game-theoretic dimensions.[4]

At the same time, our findings are also compatible with the implications of the noisy rational expectation equilibrium (*NREE*) models (e.g., Grossman and Stiglitz, 1980; Hellwig, 1980; Diamond and Verrecchia 1981; Brennan and Cao, 1996) regarding 'signal amplification' and the fact that prices do not perfectly reflect the true asset value (see Bossaerts, Frydman and Ledyard, 2013). It is worth noting that our learning model shares an important feature with *NREE* models as it considers two types of traders who differ in their level of sophistication. In both models, one type of trader fails to learn from asset prices whereas the other type of trader uses Bayesian updating. Our model differs, however, from *NREE* as *non-reflective* traders use their private information and thus do not randomly submit orders. More importantly, our approach differs from

[3] See also the dialogue between Eugene Fama and Richard Thaler: http://review.chicagobooth.edu/economics/2016/video/are-markets-efficient

[4] In continuous double auctions, traders can submit, at any time, offers to buy or sell the asset. Traders can also accept current offers to buy or sell an asset.

NREE in its aim of modeling a continuous double auction market that mimics the experimental markets which we subsequently use to test our main conjectures. We thus trade off the strategic concerns inherent to the equilibrium approach with the applicability of our model to the complex environment of continuous double auctions.[5] Because our ultimate goal is to test our conjectures in a laboratory setting, we identify our two different *types* of traders using simple cognitive tests. This is achieved by establishing a direct correspondence between the two types of traders and a specific cognitive skill (in our case we consider cognitive reflection as measured by the Cognitive Reflection Test, Frederick, 2005). This link between traders' types and cognitive reflection has previously been established by Corgnet, DeSantis and Porter (2015; 2018). Because noise traders in *NREE* are typically introduced for convenience in order to prevent the application of the no-trade theorem, they are treated as a 'black box'. Our research can thus be seen as an attempt to open the 'black box' and provide a cognitive profile for these traders so as to assess the extent to which they populate actual market settings.

Experimental Finance

Lab and field experiments on information aggregation have mainly focused on institutional aspects of informational efficiency (e.g., Plott, Wit and Yang, 2003). For example, the work of Plott and Sunder (1988) led to the idea that information aggregation was more easily achieved in markets in which a complete set of Arrow-Debreu securities was available. In Plott and Sunder (1988) Arrow-Debreu securities markets, the traded asset had a positive value only if one of three possible states occurred. It follows that traders receiving a signal that a given state cannot occur become de facto *fully* informed that the Arrow-Debreu security associated with that state is worthless. They can thus sell that security to secure a profit at no risk. Thus, in the Plott and Sunder (1988) setting, Arrow-Debreu securities transform a market in which traders are originally *partially* informed into a market in which half of the traders are *fully* informed and half are *uninformed*. Our findings shed light on Plott and Sunder's (1982, 1988) results that Arrow-Debreu securities facilitate the aggregation of information because these securities create markets in which some traders are *fully* informed and private information is less dispersed.[6]

---

[5] Trying to do both at the same time has had limited success (see e.g., Cason and Friedman, 1996 for a discussion).

[6] Our work also relates to a strand of the literature that focuses on other institutional features of markets such as the existence of, as well as the level of competition between, market makers (Cason, 2000; Krahnen and Weber, 1999, 2001), the existence of dark markets (Asparouhova and Bossaerts, 2017) and the presence of futures markets (Friedman, Harrison and Salmon, 1984).

Our work directly relates to a number of studies which have assessed the effect of the timing of information release on the informational efficiency of markets. Copeland and Friedman (1987), as well as Barner, Feri and Plott, (2005), have extended the original experiments of Plott and Sunder (1982) to different structures of information including sequential information release. Their findings generally suggest that the sequential arrival of information tends to hamper informational efficiency. Also, asset markets in which traders make sequential decisions and in which traders' decision times are random have been found to generate information cascades in which some private information fails to be transmitted to prices (Nöth and Weber, 2003).[7]

The complexity of the informational structure has also been found to affect the informational efficiency of markets. Plott, Wit and Yang (2003) find that information aggregation is less pronounced in settings in which pooling all private traders' information does not unveil the true asset value with certainty. In addition to aggregate uncertainty, Camerer and Weigelt (1991) show that the nature of the private signals plays an important role in the informational efficiency of markets. The authors show that private information may fail to aggregate in cases in which traders do not know with certainty whether half the market is composed of insiders. O'Brien and Srivastava (1991) find that asset prices do not converge to fundamentals in an environment with multiple assets in which asset valuation can depend upon a previous realization of dividends.

In line with previous research, our model and our experimental data show that market complexity, whether it arises from the complexity of the information structure or other institutional features, is a crucial ingredient to understanding the informational efficiency of markets.

Finally, our paper largely builds on the work of Bossaerts, Frydman and Ledyard (2013) who compare the informational efficiency of markets populated with different numbers of *fully* informed traders. They find evidence for 'signal amplification' as prices tended to reflect the true value more closely as the number of *fully* informed traders in the market increased. In addition to replicating their findings for markets populated with *fully* informed traders, we extend their work by comparing the informational efficiency of market environments that only differ in the dispersion of information.

[7] Other works have focused on the presence of manipulators who could distort prices and hamper the informational efficiency of the market (Hanson, Oprea and Porter, 2006; Veiga and Vorsatz, 2009, 2010). Hanson, Oprea and Porter (2006) incentivized manipulators to distort market prices but find that they were ultimately unable to affect informational efficiency. Veiga and Vorsatz (2009, 2010) used computerized market manipulators and showed that these operators could distort asset prices.

In doing so, we show, using both a learning model and experimental markets, that concentrating private information tends to promote informational efficiency.

3. Model and conjectures

Following Friedman (1991), our model assumes that traders do not strategically place orders. We thus posit an intermediate level of sophistication in the modeling of agents interacting in a continuous double auction setting. At the lower end of the sophistication spectrum is the zero-intelligence model of Gode and Sunder (1993) which only constrains traders' behavior to avoid trading at a loss. In this model, traders do not update their beliefs regarding the true asset value in contrast to Friedman (1991). At the upper end of the sophistication spectrum are the game-theoretic double auction models in which traders are assumed to behave strategically (e.g., Wilson, 1987). However, these models have been found to underperform less sophisticated models in explaining traders' behavior in double auction market experiments (e.g., Cason and Friedman, 1996). Based on the current state of the literature we thus employ a modeling approach similar to Friedman (1991) as such models have been found to be the most successful in explaining experimental data. In this setting, traders establish their beliefs regarding the true asset value based on observed market orders. Traders then follow a *reservation-price* strategy that consists of buying (selling) the asset whenever the price is below (above) their belief of the true asset value.

3.1. Model

3.1.1. Trading and information

Even though we present our model for the specific parameters used in our experiments (see Section 4 for the experimental design), it is important to note that our findings apply more generally. In Section 5.3.1, we stress test the model by assessing the robustness of our conjectures to different distributions of the asset value. We also consider the robustness of our findings to different cash and shares endowments as well as to the case of markets populated by a larger number of traders.

Our model allows for a market populated with traders possessing different levels of private information. Depending on the market design, traders can be either *uninformed*, *partially* informed or *fully* informed. In particular, we take the parameters of Market 9 from Plott and Sunder (1988) as our baseline. This market is populated by 12 *partially* informed traders who receive the same endowment of cash (1,200) and shares (4). The asset value can only take one of three possible

values: 50, 240 or 490 with probabilities 35%, 45% and 20%, respectively.[8] Half of the *partially* informed traders receive one private signal regarding which possible value the asset could not take (e.g., "Not 50") whereas the other half of the *partially* informed traders receive the other possible signal (e.g., "Not 240") so that the aggregate information available to traders in the market was complete. Building on this market design, we then consider the case in which the market is populated by both *uninformed* and *fully* informed traders (see e.g., Plott and Sunder, 1982, or Bossaerts, Frydman and Ledyard, 2013). An *uninformed* trader does not receive a private signal, while a *fully* informed trader receives two *different* signals (e.g., "Not 50"; "Not 240") thus making the true asset value (e.g., 490) clear.

One caveat of the market design we employ is that the asset value is the same for all traders implying no gains from trade and thus no trade in a rational expectation equilibrium. Our approach thus focuses on the informational efficiency of markets rather than on the allocative efficiency of markets. We made this choice purposefully as we intend to first consider a setting in which information aggregation had been evidenced (Treatment C in Plott and Sunder, 1988) rather than considering the case in which gains from trade exist but information aggregation failed to be observed (Treatment A in Plott and Sunder, 1988). Notwithstanding, the conjectures derived from our model simulations appear to be robust to all treatments and all specifications used in Plott and Sunder (1988) including markets in which gains from exchange exist (see Section 5.3.1). Finally, while the no trade theorem may technically apply to our market design, the extensive literature in experimental markets shows that trade does indeed occur even in the absence of gains from trade (e.g., Palan, 2013; Noussair and Tucker, 2014). One possible reason for this is the existence of noise (e.g., Grossman and Stiglitz, 1980) or boundedly rational traders (e.g., Friedman, 1991; Gjerstad and Dickhaut, 1998; Peng and Xiong, 2006; Esponda, 2008; Eyster and Rabin, 2005; Mondria, 2010; Corgnet, DeSantis and Porter, 2015; Kacperczyk, Van Nieuwerburgh and Veldkamp, 2016; Vives and Yang, 2017; Eyster, Rabin and Vayanos, 2018) who fail to immediately infer others' private information from prices.[9]

[8] The other two specifications (Markets 7 and 8) by Plott and Sunder (1988) used in this treatment are almost identical. These markets are such that the probabilities of occurrence of each of the three possible values are the same. Market 7 uses the same three values as Market 9 (50, 240 and 490) whereas Market 8 uses values 125, 375 and 525. The authors did not report any differences across these markets. Importantly, the conjectures we derive from our simulations are robust to alternate specifications (see Section 5.3.1).

[9] Risk-sharing motives may also be a reason to trade in our markets if all private information is not captured in the price in which case no risk remains.

Our model reproduces the decentralized continuous double auction mechanism used in the experimental markets.[10] At the beginning of each market, a trader (selected at random) posts a bid-ask spread. A second trader is then selected and given the option to either accept the current bid or ask or improve the bid-ask spread.[11,12] Traders (selected at random) continue to improve the bid-ask spread until a trade occurs.[13] If a trade occurs, then the price, p, is set at the current best bid (or ask) and the trading book is updated.[14]

The decision to improve the bid-ask spread or trade depends on a trader's belief regarding the true asset value. Specifically, a trader will improve the current bid-ask spread if her belief lies within the current spread. In that case, the updated bid (ask) will be drawn from a uniform distribution between the current bid (ask) and the trader's belief. If traders' beliefs are below (above) the current best bid (ask), then they would sell (buy) the asset. This trading strategy can be seen as an example of a *reservation-price* strategy (see Friedman, 1991). In our model, the reservation-price strategy consists of trading according to a trader's current belief of the true asset value. While this strategy can only be justified if we assume risk neutrality for the traders, this assumption is widely used in double auction models (see Copeland and Friedman, 1987; Friedman, 1991) as well as in game-theoretic microstructure models (e.g., Glosten and Milgrom, 1985; Kyle, 1985).[15]

Similar to our experimental design, trading follows a bid-ask improvement rule. In our market environment, traders observe market prices as well as bids and asks. We thus assume that traders can use all market orders (bids, asks and contract prices) to update their beliefs.

3.1.2. Reflective and non-reflective traders

---

[10] A market maker (or a specialist) is commonly utilized in the market microstructure literature (e.g., Glosten and Milgrom, 1985; Easley and O'Hara, 1987). In order to mimic our experimental design as closely as possible we do not include a market maker in our model.

[11] A bid-ask spread is only updated when the newly selected trader has enough cash (shares) to cover the bid (ask) position. In line with our experimental markets, traders are not permitted to sell short or borrow funds. Moreover, when a trader is selected to act, the trader first removes any bids or asks from the order book that are inconsistent with the trader's current belief of the asset value. For example, if the trader previously submitted a bit that is now greater than her current belief, then this bid is automatically canceled. Note that traders in our experimental markets also had the ability to cancel orders.

[12] Our model differs from Friedman (1991) as the randomly selected trader is forced to act, if possible. Thus, the trader cannot choose to not act in an attempt to gain better terms of trade.

[13] The sequence of trades is not random, however. Traders either will or will not trade at a given bid-ask spread depending upon their current belief regarding the true asset value.

[14] The accepted bid/ask is removed from the book, and the next best bid/ask, if one exists, becomes the current best bid/ask. We thus consider the case of an open book similar to our experimental design (see Section 4).

[15] The zero-intelligence model of Gode and Sunder (1993) also disregards issues related to risk attitudes.

Even though current learning models in which all traders use Bayes' rule have had some success in explaining data in experimental markets with double auctions, they still 'leave much to be desired' (Cason and Friedman, 1996, page 1332). The authors point to a promising direction (page 1333) by stating: "Given that the most 'rational' of available models did not perform especially well on its 'home turf', perhaps the most appropriate next step is to explore new models that incorporate significant bounded rationality." Following this suggestion, we populate our market with traders who fail to apply Bayes' rule to learn from market orders.[16] We consider a model in which heterogeneous agents interact in the market. In particular, we assume our traders to be either *reflective* or *non-reflective*. *Reflective* traders apply Bayes' rule to infer the true asset value as they observe market orders. We use the term *reflective* for these traders as the failure to use Bayes' rule to update beliefs on the face of new information has been closely connected to ones' cognitive reflection capacity (e.g., Oechssler, Roider and Schmitz, 2009; Campitelli and Labollita, 2010; Hoppe and Kusterer, 2011; Lesage, Navarrete and De Neys, 2013; Toplak, West and Stanovich, 2011, 2014; Sirota, Juanchich and Hagmayer, 2014; Corgnet, DeSantis and Porter, 2015, 2018) as measured by Frederick's (2005) Cognitive Reflection Test (henceforth, CRT).

Non-reflective traders only update their beliefs once, on the basis of their private information, and do not use market data to infer other traders' private information. These identifying characteristics are consistent with several important models in the literature: the Walrasian model (Lintner, 1969), the no revelation of expectations model (Copeland and Friedman, 1987; 1991), models stressing investors' inattention and information processing costs (e.g., Peng and Xiong, 2006; Mondria, 2010; Kacperczyk, Van Nieuwerburgh and Veldkamp, 2016; Vives and Yang, 2017), the behavioral equilibrium model (Esponda, 2008) and the cursed equilibrium model (Eyster and Rabin, 2005; Eyster, Rabin and Vayanos, 2018).

[16] This can be seen as being inspired from quasi-Bayesian learning models such as Daniel, Hirshleifer and Subrahmanyam (1998) and Rabin (2002). Daniel, Hirshleifer and Subrahmanyam (1998) account for stock market under- and over-reaction to news by developing a model in which traders overvalue the precision of their private information and update their beliefs self-servingly by downplaying information which may not be consistent with their prior beliefs. Rabin (2002) considers the case in which investors assign too much weight to small samples to account for stock market under- and over-reaction to news. Both Daniel, Hirshleifer and Subrahmanyam (1998) and Rabin (2002) build their models on extensive cognitive psychology literature showing that people largely fail to apply Bayes' rule suffering instead from a number of common and long-lasting biases (e.g., Tversky and Kahneman, 1974; Stanovich, 2009; Kahneman, 2011).

According to the Walrasian and the no revelation of expectations models, traders make decisions based on their private information and fail to infer other traders' information from market prices.[17] The behavior of traders in these models as well as in models of investors' inattention closely relates to the widely-documented phenomenon of the winner's curse (Bazerman and Samuelson, 1983; Thaler, 1988, 1991). The winner's curse occurs when auction participants fail to anticipate the informational content of other individuals' bids and end up paying too much for the auctioned item. Applied to financial markets, Eyster, Rabin and Vayanos (2018) refer to those who "neglect the informational content of prices" as *cursed* traders.

3.1.3. Learning

In our model, learning occurs after each market event. Market events either correspond to an improvement of the bid-ask spread or to a transaction. Based on these market events, *reflective* traders can update their beliefs regarding the true asset value by applying Bayes' rule to infer other traders' private information. Each event corresponds to one of the following actions: (1) best bid is accepted; (2) best ask is accepted; or (3) the bid-ask spread is improved. Let μ_{n,s_j}^{η} represent the belief of trader j of type η after event n, which is denoted by $\{e(n)\}$. Trader types, $\eta = \tau \times s$, identify whether the trader is reflective or not ($\tau = R$ or NR) as well as the signal the trader received ($s = Not\ 50, Not\ 240$ or $Not\ 490$ for *partially* informed traders and $s = I$ or U for *informed* or *uninformed* traders).

Reflective traders update their beliefs using Bayes' rule with $i \in \{50,240,490\}$ as follows:

$$\mu_n^{R,s} = E\left[v|\mu_{n-1}^{R,s}, \{e(n)\}\right] = \sum_i P\left[v = i|\mu_{n-1}^{R,s}, \{e(n)\}\right] \times i$$

where $\mu_{n-1}^{R,s}$ represents the prior belief of the trader with the signal "s" and v is the true asset value.[18]

In Appendix B, we detail the Bayesian learning procedure to calculate $\mu_n^{R,s}$ for each of the three possible types of events. We then use these detailed procedures to simulate our model.

3.1.4. Model example

[17] It is important to note that the two models differ because traders are myopic in the Walrasian model whereas they trade based on equilibrium prices in the no revelation of expectations model.
[18] Note that *informed reflective* traders do not update their belief as they already know the true asset value.

To illustrate the intuition behind the model's learning mechanism, consider the following simple example. Suppose the true asset value is 490. Also, assume all traders are *reflective* and that this assumption is commonly known by all traders. The *fully* informed traders know the true asset value, i.e. $\mu_0^{R,I} = 490$, while the *uninformed* traders' belief of the asset value is $\mu_0^{R,U} = 223.5$ ($= 0.35 \times 50 + 0.45 \times 240 + 0.20 \times 490$). Consider the following market events.

First market event: bid-ask spread improvement: 200-360

The initial event is always an improvement of the bid-ask spread because the book is empty. This is equivalent to assuming that the initial bid-ask spread is 50-490, i.e. the minimum-maximum possible value of the asset. The *fully* informed traders do not update their beliefs; however, the *uninformed reflective* traders must update their belief regarding the likelihood of occurrence of each of the three values of the asset by determining the probability of occurrence of the first market event given the asset value is either 50, 240 or 490. If the true value was 50 or 490, then this new bid-ask pair (200-360) could only have been submitted by another *uninformed* trader. Indeed, since traders follow a reservation price strategy based upon their belief, if the true value is 50 (490), then *fully* informed traders could not submit a bid (ask) above (below) 50 (490). If the true value was 240, then any trader could have generated this first event. Thus, while the *uninformed reflective* traders revise their initial belief of 223.5, they do not significantly alter it based upon this event, i.e. $\mu_1^{R,U} = 218.7$.

Second market event: trade at 360

A trader can only accept the current best ask if their belief is greater than the ask value. This implies that this ask must have been accepted by an *informed* trader with the belief of 490 because *uninformed reflective* traders hold a belief of 218.7. Thus, the *uninformed reflective* trader learns from the second market event that the true asset value is 490.

This example illustrates how quickly *uninformed reflective* traders can learn in the presence of *fully* informed traders. In the absence of *fully* informed traders in the market, *uninformed reflective* traders could not have determined the true asset value as early as the second market event. Inclusion of *non-reflective* traders will tend to slow this process down as *non-reflective uninformed* traders will stick to their initial belief of 223.5 thus preventing *reflective* traders to learn any private information from the market orders submitted by these *non-reflective* traders.

3.2. Conjectures

We conduct simulations of the learning model to derive a series of conjectures that can be tested experimentally. Each simulation is conducted with the same number of traders (12), and the same endowment of cash (1,200 francs) and shares (4) used in the original setting of Plott and Sunder (1988) which will be the basis for our experiments (see Section 5.3.1 for robustness checks with different parameter values). A total of 25,000 simulations were run for (each of the) various combinations of two exogenous model parameters. The first parameter corresponds to the asset value $v \in \{50, 240, 490\}$. The second parameter is the proportion of *reflective* traders ($\alpha \in \{0, \frac{1}{12}, \frac{2}{12}, \frac{3}{12}, \dots, 1\}$) in the market which we assume to be common information known by all traders. Each simulation runs until 30 trades have been executed so as to achieve a trading volume comparable to experimental markets using a similar design (see Corgnet, DeSantis and Porter, 2018).[19]

Our first conjecture compares informational efficiency in markets that differ in the relative number of *partially* informed, *fully* informed and *uninformed* traders. However, we hold the number of private signals in each market constant. Thus, we are able to compare, for example, a market with 12 *partially* informed traders (each endowed with one signal) to a market with 6 *fully* informed (each endowed with two signals) and 6 *uninformed* (not endowed with a signal) traders. Our simulations show that markets populated by both *fully* informed and *uninformed* traders are more informationally efficient than markets solely populated by *partially* informed traders. To show this, we measure the informational efficiency of a market by calculating the absolute distance between a market price and the true asset value for each transaction in a given simulation. We then take the mean of this measure over a given simulation. Averaging this mean value across all 25,000 simulations yields the mean absolute deviation (MAD) for a given distribution of private information. In Figure 1, we compare MAD values between a market populated by 100% (50%) *partially* informed traders with a market populated by 50% (25%) *fully* informed traders and the rest *uninformed* traders.[20] In both comparisons, the number of private signals is identical between the market with *partially* informed traders and the market with *fully* informed traders. We observe that informational efficiency, which is inversely related to MAD values, is highest in markets with *fully* informed traders compared to markets with *partially* informed traders.

[19] Our results are, however, robust to different stopping rules including stopping a simulation after either 5 or 60 transactions (results available upon request from the authors).

[20] We will follow this naming convention for the remainder of the manuscript. That is, if the title of a specific simulation or experimental treatment does not explicitly identify all traders, then it is assumed that the unnamed traders are *uninformed*.

Figure 1 also shows, in line with our model, that the informational efficiency of markets increases (MAD values decrease) as the proportion of *reflective* traders in the market increases.

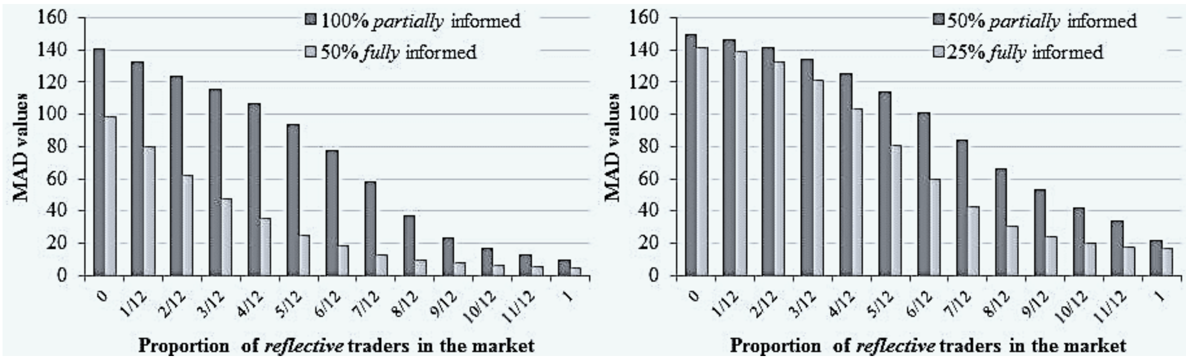

Figure 1. These figures represent the mean absolute price deviation from the true asset value in simulated markets. In the left panel, we compare a market populated by 100% *partially* informed traders and a market populated by 50% *fully* informed traders. In both markets, the number of private signals is identical and equal to 12. In the right panel, we compare a market populated by 50% *partially* informed traders to a market populated by 25% *fully* informed traders. In both markets, the number of private signals is identical and equal to 6. MAD values are averaged across all three asset values, 50, 240 and 490.

In addition, the informational efficiency is similar in the market populated by 100% *partially* informed traders (MAD = 72.8) and the market populated by only 25% *fully* informed traders (MAD = 71.4). This holds even though there are twice as many private signals in the market with *partially* informed traders (12 signals) than in the market with 25% *fully* informed traders (6 signals). We investigate this further by comparing the informational efficiency of the market in which all traders are *partially* informed with markets populated by different proportions of *fully* informed traders (with all other traders uninformed). In Figure 2, we show that a market populated by 100% *partially* informed traders is less informationally efficient than a market populated by only 25% (3 out of 12) *fully* informed traders whereas it is more efficient than a market populated by 17% (2 out of 12) *fully* informed traders. Even though the MAD values reported in Figure 2 correspond to averages across all possible proportions of *reflective* traders in the market, the ordering of the different markets holds for each of the possible proportions of *reflective* traders. It follows that the qualitative nature of Conjecture 1b, 1c and 1d does not hinge upon the proportion of *non-reflective* traders in the market.

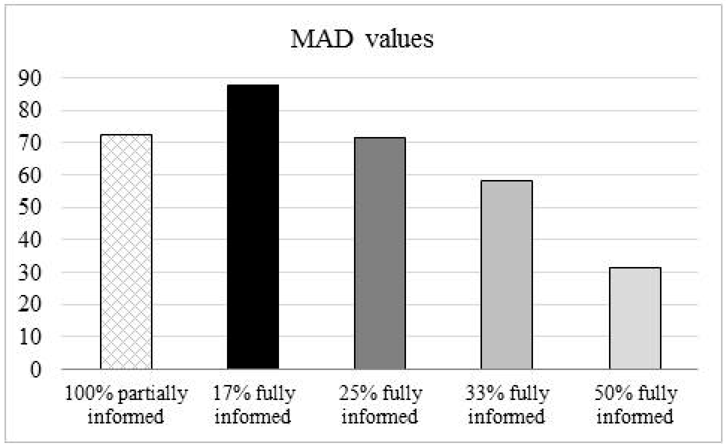

Figure 2. These bar charts represent the mean absolute price deviation from the true asset value (averaged across all possible proportions of *reflective* traders in the markets), for different numbers of *partially* and *fully* informed traders. MAD values are averaged across all three asset values, 50, 240 and 490.

Figure 2 also highlights, in line with *NREE* models (see Bossaerts, Frydman and Ledyard, 2013, Hypothesis 3), that the increase in the number of *fully* informed traders will lead to an increase in the informational efficiency of market prices, which is in line with the idea of 'signal amplification'.

Based on these simulations, we derive the following conjecture.

Conjecture 1. (Informational Efficiency of Markets)

a) *The higher the proportion of reflective traders in the market the higher the informational efficiency of the market.*

b) *Markets with the same number of private signals are more informationally efficient when signals are concentrated in the hands of fully informed traders rather than dispersed across partially informed traders.*

c) *The higher the number of fully informed traders in a market the higher the informational efficiency of the market.*

d) *Markets with two (three) fully informed traders should be less (more) informationally efficient than markets in which all traders are partially informed.*

Markets populated by different proportions of *fully* informed traders not only vary in the informational efficiency of prices but also in trading volumes. By construction, simulations were stopped after 30 transactions.[21] Thus, we assess trading intensity (rather than volume) by

[21] While the primary stopping criteria for the simulations was 30 transactions, simulations also halted if (i) 10,000 iterations of the simulation occurred or (ii) 1,000 iterations passed without an event. One iteration of the simulation corresponds to the random selection of a trader and that trader's action: improve the bid/ask spread, accept the current best bid or ask, or nothing. So, if the simulation randomly selected a trader who did not update the bid-ask spread or accept the current best bid or ask 1,000 consecutive times, then the simulation halted in order to avoid an infinite loop. The assumption is that (i) traders were unable to trade at this point due to a lack of cash/shares or (ii) traders beliefs'

calculating the average number of transactions that occur within a block of five iterations in each simulation. In Figure 3, we show that trading intensity decreases as the number of *fully* informed traders in the market increases (see last four bar charts). This likely holds because, as the number of informed traders increases, prices move closer to the true asset value. This implies that a trader's belief regarding the asset value is more likely to be aligned with other traders' beliefs. Because the discrepancy in traders' beliefs drives trading, markets in which private information is more likely to be aggregated generate fewer transactions.

We also show that trading intensity is lower for markets in which all traders are *partially* informed compared to markets populated by less than 50% *fully* informed traders. Finally, trading intensity is essentially identical for the 100% *partially* informed market and the market populated by 50% *fully* informed traders. These findings form the basis for Conjecture 2.

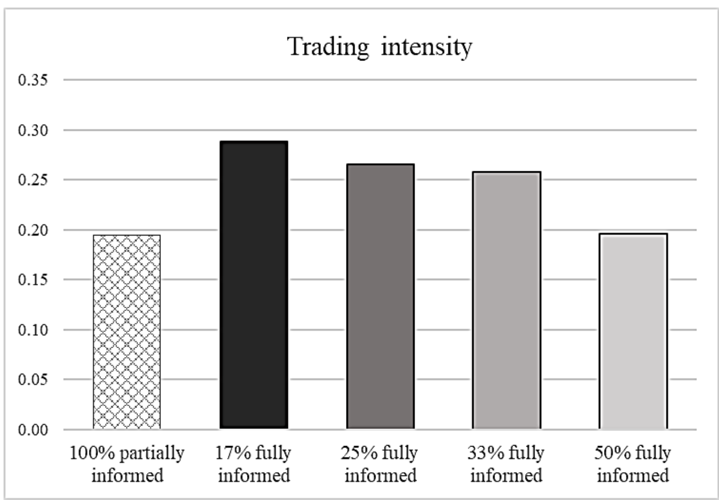

Figure 3. These bar charts represent trading intensity (averaged across all possible proportions of *reflective* traders in the markets, across all 25,000 simulations and across all three asset values) for different types of markets which differ in the number of *partially* and *fully* informed traders. Trading intensity is calculated as the average number of transactions completed in a market for blocks of five iterations of the algorithm.

Conjecture 2. (Trading volumes)

a) *Trading intensity is lower for markets in which all traders are partially informed compared to markets populated by less than 50% fully informed traders.*

b) *Trading intensity is the same for the 100% partially informed market and the market populated by 50% fully informed traders.*

c) *Trading intensity decreases as the number of fully informed traders in the market increases.*

4. Experimental design

were close enough to the true value that they were unable to update the spread. Note that a bid-ask spread of was not allowed to be an empty interval.

To test our conjectures, we use experimental markets that have the same characteristics as those described in our model simulations. We consider a market setting in which subjects could trade an experimental asset that can only take three possible values: 50, 240 or 490 francs (each franc was worth $0.001) with probabilities 35%, 45% and 20%, respectively. Trading takes place using a computerized continuous double auction mechanism.[22] We conduct a total of 5 treatments using a between-subject design. Treatments only vary in the distribution of private signals across traders. In the '*100% partially informed*' treatment, each of the 12 traders in the market is privately informed of a value the asset could not take. As half of the traders were given one signal (e.g., "Not 50") and the other half were given the other possible signal (e.g., "Not 240"), the aggregate information available to traders in the market was complete.

We contrast the '*100% partially informed*' treatment with four *'fully informed'* treatments in which information is concentrated in the hands of a few *fully* informed traders with all other traders *uninformed*.[23] From the point of view of the amount of aggregate information, the '*100% partially informed'* treatment and the *'fully informed'* treatments do not differ. In all cases, the aggregate information reveals the true asset value. To test our conjectures, we consider *'fully informed'* treatments in which the number of *fully* informed traders is either 2, 3, 4 or 6.

We recruited a total of 418 individuals from a subject pool of more than 1,500 individuals at a major Western US University. We conducted a total of 35 sessions. Ten sessions were conducted for both the '*100% partially informed'* and the *'50% fully informed'* treatments. Five sessions were conducted for each of the other three treatments. We conducted twice as many sessions for the first two treatments as they are the key treatments for the testing of our conjectures given that the total number of private signals (12) is the same in both markets.

Each market was composed of 12 traders.[24] In line with our model, traders were endowed with 1,200 francs in cash and 4 shares of the asset. Each session consisted of 17 markets with independent draws for the asset value.[25]

[22] The experiment was conducted using Zocalo which is an open-source software used for experimental markets.
[23] In the '*fully informed*' treatments the *informed* traders receive two signals (e.g., if the true value is 240, then these traders receive the signal "Not 50" as well as the signal "Not 490").
[24] Two sessions in the *'50% fully informed'* treatment were conducted with 11 traders due to an elevated number of last-minute cancellations.
[25] We used the same draw for all sessions to facilitate comparisons across sessions. This draw corresponds to the asset values used in Market 9 of Plott and Sunder (1988).

Before the trading phase of each session started, subjects completed a training exercise regarding a random device (a spinning wheel) that represented the probabilistic distribution of the true asset value (50, 240 or 490 francs). During the training, subjects had to predict the outcome of the spinning wheel over 10 trials (see online Appendix O1, Instructions Part 1). Each correct prediction was rewarded 25 cents, and each incorrect answer incurred a 10-cent penalty. The instructions concluded with a comprehension quiz and took approximately 20 minutes for subjects to complete. Average earnings (including a $7 show-up payment) for the two and one-half hours experiments were equal to $48.

At the end of all experimental sessions, we elicited demographic information and evaluated subjects' level of cognitive reflection using a 7-item version of the test developed by Frederick (2005) (see Toplak, West and Stanovich, 2014). We also collected theory of mind scores (see Baron-Cohen et al. 1997) (see Appendix D for a description of the survey). As is standard in the literature, these tests were not incentivized; however, subjects received $3 for completion of the survey.

5. Experimental Results

We test our conjectures sequentially, starting with an analysis of informational efficiency and then volumes.

5.1. Informational efficiency of markets

The top panel of Figure 4 illustrates the sharp difference in asset prices between our *'100% partially informed'* treatment and our *'50% fully informed'* treatment thus providing support for Conjecture 1b.[26] The bottom panel of Figure 4 illustrates that, in line with Conjecture 1c, prices are closer to the true asset values as the number of *fully* informed traders increases in the market.

[26] In Appendix E, we show the graphs of market prices for each experimental session.

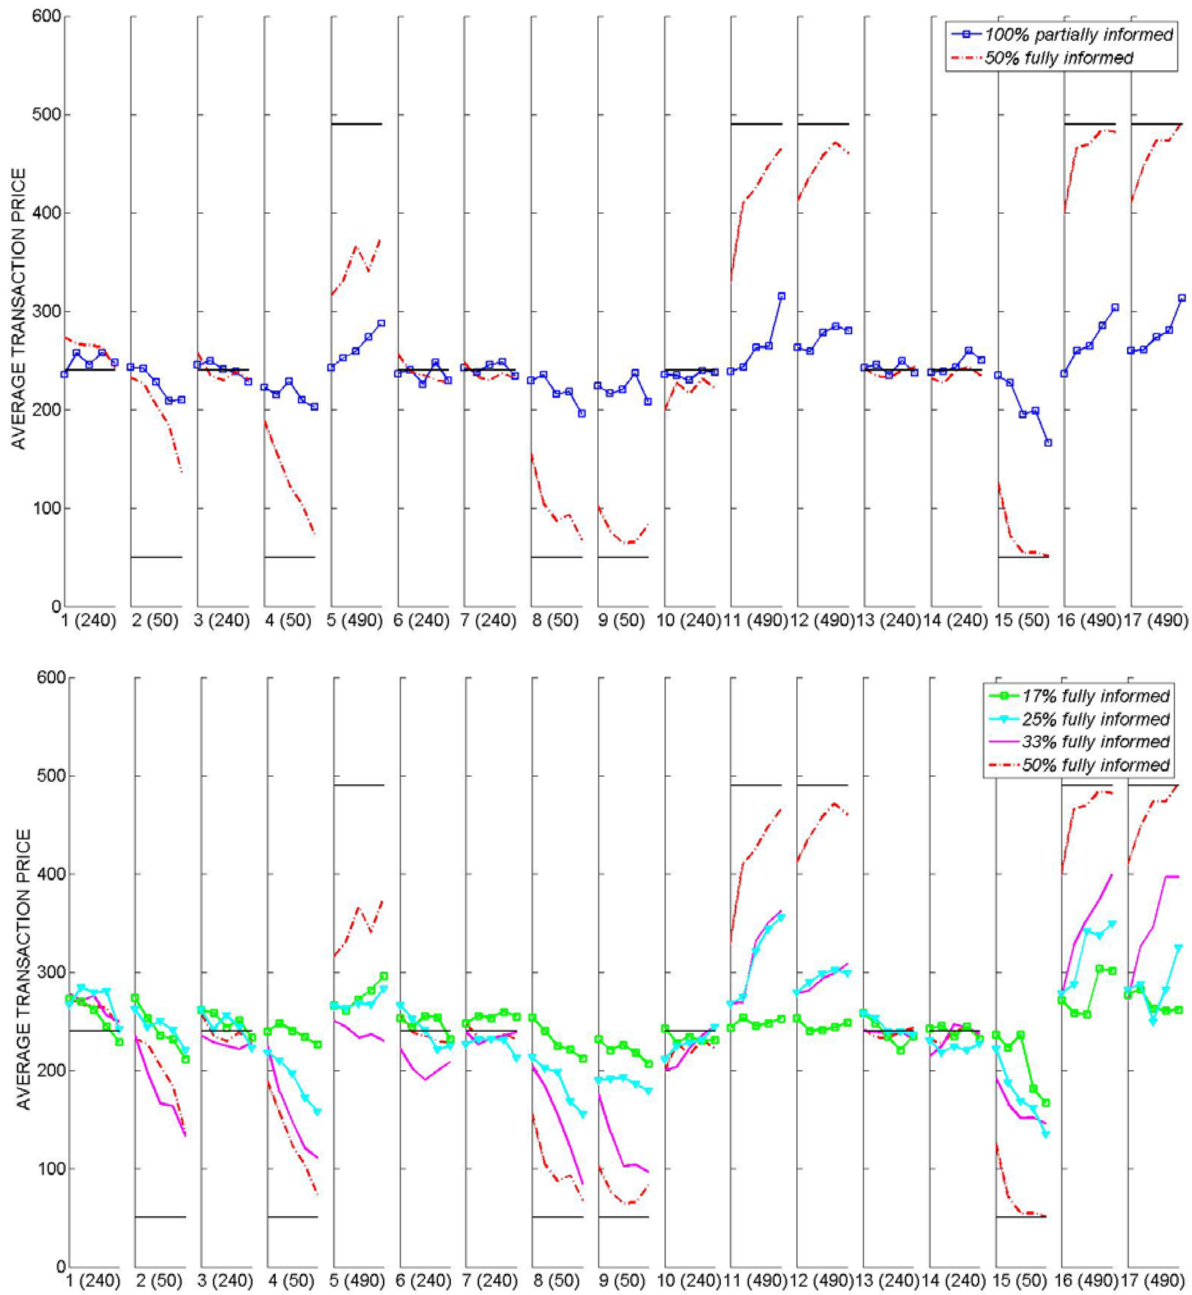

Figure 4. Top Panel. Average price per minute over the 10 sessions for each of the seventeen markets in the '*100% partially informed'* (solid blue line with square markers) and the '*50% fully informed'* treatments (dash-dot red line). Bottom Panel. Average price per minute for the '*50% fully informed*' treatment is repeated along with the average price per minute over the five sessions for each of the seventeen markets in the '*17% fully informed'* (solid green line with square markers), '*25% fully informed'* (solid cyan line with triangle markers) and '*33% fully informed'* (solid magenta line) treatments. The true asset value is denoted in parentheses at the bottom of each subfigure and is represented by a solid (black) horizontal line.

We assess the statistical significance of our findings in the regression reported in Table 1. Coefficients of the treatment dummy variables should be interpreted as positive or negative effects

with respect to the '*100% partially informed*' treatment for which we omitted the dummy variable. In these regressions, we use treatment dummies which take value one when a market was conducted under the corresponding treatment and value zero otherwise. To assess Conjecture 1a, we also control for the proportion of *reflective* traders in the market. To that end, we make use of the cognitive reflection test as a way to assess how *reflective* a trader is. An individual's score on the CRT, as a measure of cognitive reflection, is a key determinant of her capacity to properly use Bayes' rule (see Corgnet, DeSantis and Porter, 2015; 2018).

We consider a trader to be *reflective* (*non-reflective*) when her score is above (below) the median CRT scores of all 418 subjects who participated in this study. Alternately, in regression [2] we control for the average CRT score of all traders in a given session.[27] We also control for the true asset value in a given market and for the number of market periods in which the trader has participated (*Market number* variable). The negative and significant coefficient of the proportion of *reflective* traders (see regression [1]), as well as the negative and significant coefficient of the average trader's CRT score in a session (see regression [2]), confirm Conjecture 1a. Indeed, both variables are negatively related to our measure of informational efficiency (MAD). We confirm Conjecture 1b because the coefficient for the Dummy '*50% fully informed'* is negative and significant (see regressions [1] and [2]) showing that informational efficiency is higher when half of the traders receive two signals than when all traders receive only one signal. This finding is actually consistent with the recent results of Barreda et al. (2017) who reported, in an experimental setup in which private signals provide probabilistic clues for the true value of the asset, that prices are closer to the fundamental value when the signals are concentrated in the hands of 'quasi-insiders' rather than dispersed across traders. We also find support for Conjecture 1c through regression [3] as the coefficient on the number of *fully* informed traders is negative and significant. Finally, we support Conjecture 1d because the coefficient for the Dummy *'17% fully informed'* is positive and significant whereas the coefficient for Dummy *'25% fully informed'* is negative and significant in regressions [1] and [2]. This implies that markets in which all traders are *partially* informed are more efficient than markets with two *fully* informed traders but less efficient than those with three *fully* informed traders.

Table 1. Random effects panel regression of MAD values per market as a function of cognitive reflection and treatment dummies

[27] In line with Corgnet, DeSantis and Porter (2018), we also find that traders earnings are positively related to individual CRT scores (see Table A2 in Appendix A).

MAD	All treatments		Only treatments with *fully* informed traders
Specification:	[1]	[2]	[3]
Dummy *'17% fully informed'*	13.592***	14.222***	-
	(4.963)	(4.902)	
Dummy *'25% fully informed'*	-14.477*	-16.158**	-
	(8.128)	(8.141)	
Dummy *'33% fully informed'*	-26.636**	-26.482**	-
	(10.636)	(10.627)	
Dummy *'50% fully informed'*	-66.824***	-69.518***	-
	(7.200)	(6.194)	
Number of *fully* informed	-	-	-19.237***
			(2.223)
Proportion of *reflective* traders	-35.560*	-	-
	(19.802)		
Average CRT score of traders	-	-14.910***	-
in a session		(4.443)	
Market number	-0.196	-0.196	-1.138
	(0.589)	(0.589)	(0.733)
True asset value	0.131***	0.131***	0.128***
	(0.027)	(0.028)	(0.0328)
Intercept	110.824***	137.784***	158.024***
	(11.762)	(16.400)	(14.869)
Observations (sessions)	595 (35)	595 (35)	425 (25)
Prob > χ^2	0.000	0.000	0.000
R^2	0.186	0.190	0.198

*p-value<0.10, **p-value<0.05 and ***p-value<0.01. Robust standard errors clustered at the session level in parentheses.

5.2. Trading volumes (Conjecture 2)

In Figure 5, we report, in line with Conjecture 2a, that trading volumes are lower for markets populated by 100% *partially* informed traders compared to all markets in which less than 50%

fully informed traders were present. In line with Conjecture 2b, trading volumes appear to be similar between markets populated by 100% *partially* informed traders and markets populated by 50% *fully* informed traders. By contrast with Conjecture 2c, trading volumes do not linearly decrease in the number of *fully* informed traders because markets with 33% *fully* informed traders exhibit unexpectedly high volumes. However, trading volumes are lower for markets with 50% *fully* informed traders compared to markets with only 17%, 25% or 33% *fully* informed traders.

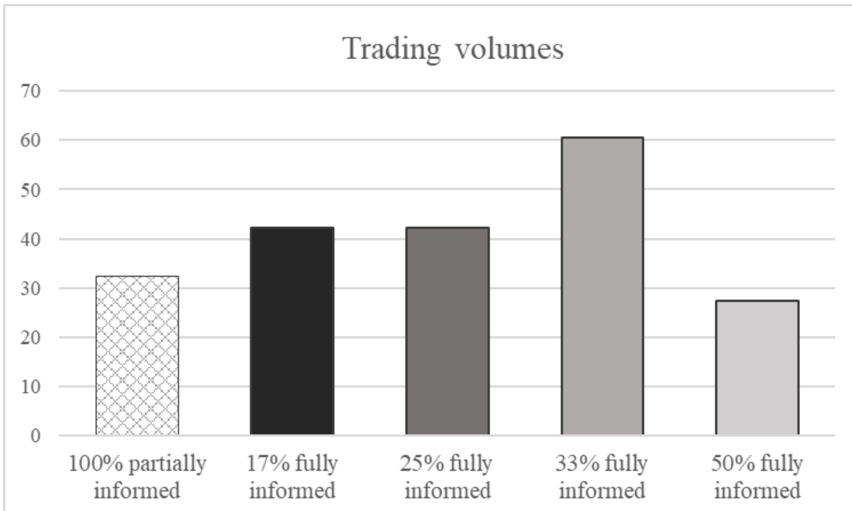

Figure 5. These bar charts represent trading volumes for different types of markets which differ in the number of *partially* and *fully* informed traders.

In Table 2, we replicate the analysis of Table 1 with trading volumes as the dependent variable. In regressions [1] and [2], coefficients of the treatment dummy variables should be interpreted as positive or negative effects with respect to the '*100% partially informed*' treatment for which we omitted the dummy variable. In regression [4], all treatment dummy variables are omitted with the exception of the one corresponding to the *'33% fully informed'* treatment. We find support for Conjecture 2a, because the coefficient for the Dummies *'17% fully informed'*, *'25% fully informed'* and *'33% fully informed'* are all positive (see regressions [1] and [2]). These coefficients are also statistically significant with the exception of the Dummy *'25% fully informed'*. In line with Conjecture 2b, the coefficient for the Dummy *'50% fully informed'* is either non-significant (see regression [1]) or negative and marginally significant (see regression [2], *p*-value = 0.088). We find mixed support for Conjecture 2c. In regression [3] the coefficient for the number of *fully* informed traders in a market is negatively and significantly associated with trading volumes, which supports Conjecture 2c. However, given Figure 5, we further explored this issue with regression

[4]. Here we see that trading volumes are higher in markets with *'33% fully informed'* traders compared to all the other markets which is at odds with Conjecture 2c.

Table 2. Random effects panel regression of trading volumes per market as a function of cognitive reflection and treatment dummies

Trading volumes	All treatments		Only treatments with *fully* informed traders	
Specification:	[1]	[2]	[3]	[4]
Dummy *'17% fully informed'*	9.380** (3.993)	11.173*** (4.228)	-	-
Dummy *'25% fully informed'*	9.286 (6.809)	7.114 (6.427)	-	-
Dummy *'33% fully informed'*	28.514*** (7.884)	26.112*** (7.614)	-	25.644*** (7.841)
Dummy *'50% fully informed'*	-6.324 (4.623)	-7.099* (4.153)	-	-
Number of *fully* informed	-	-	-4.482*** (1.589)	-
Proportion of *reflective* traders	-29.792** (13.996)	-	-	-
Average CRT score of traders in a session	-	-9.150*** (3.066)	-	-
Market number	-1.418*** (0.187)	-1.418*** (0.187)	-1.589*** (0.216)	-1.589*** (0.216)
True asset value	-0.006** (0.003)	-0.006** (0.003)	-0.007** (0.003)	-0.007** (0.003)
Intercept	52.134*** (4.208)	70.537*** (9.160)	74.984 (7.186)	51.031*** (3.977)
Observations (sessions)	595 (35)	595 (35)	425 (25)	425 (25)
Prob > χ^2	0.000	0.000	0.000	0.000
R^2	0.432	0.466	0.244	0.353

**p*-value<0.10, ***p*-value<0.05 and ****p*-value<0.01. Robust standard errors clustered at the session level in parentheses.

5.3. Robustness checks

5.3.1. Simulation results for alternative specifications

Although we presented our model for the specific parameters used in our experiments, our findings apply more generally. In Appendix C, we stress test the model by assessing the robustness of our main conjecture (1b) according to which markets with the same number of private signals are more informationally efficient when signals are concentrated in the hands of *fully* informed traders rather than dispersed across *partially* informed traders. We test whether this conjecture holds for different distributions of the asset value, different cash and share endowments, and for markets populated by a larger number of traders.

One limitation of our current market setup is that it does not allow traders to mutually gain from exchange. To address this, we conducted simulations for markets in which traders were assigned different values for the asset. Thus, mutual gains from exchange can be realized when high-value traders buy shares from low-value traders. In Appendix C, we report simulation results showing that allocative efficiency in markets in which traders can mutually gain from exchange is higher when private information is concentrated in the hands of *fully* informed traders rather than dispersed across *partially* informed traders (see Figure C1). These simulation results thus extend Conjecture 1b.

5.3.2. Mixing partially and fully informed traders

We have shown thus far that our experimental findings are largely consistent with the predictions of our model. As a robustness test, we ran an additional experiment in which we mixed *partially* and *fully* informed traders. In particular, we consider the case in which the market is populated by 33% (4 out of 12) *fully* informed, 25% (3 out of 12) *partially* informed and 42% (5 out of 12) *uninformed* traders (*'33% fully informed / 25% partially informed'*). For such a market, our model simulations predict an average MAD value of 50.78, which is less than the MAD of 59.67 for the case of 33% *fully* informed traders and greater than the MAD of 32.55 for the case of 50% *fully* informed traders. That is, our model predicts that the level of informational efficiency will only be slightly improved by adding three *partially* informed traders to a market populated by four *fully* informed traders. In addition, the informational efficiency of the *'33% fully informed / 25% partially informed'* treatment is expected to be lower than for the *'50% fully informed'* treatment. Our experimental results confirm our model predictions as the informational efficiency of the *'33% fully informed / 25% partially informed'* treatment is found to be lower than the informational

efficiency of the *'50% fully informed'* treatment, while being only slightly (but not significantly) higher than the *'33% fully informed'* treatment (see Table A1 in Appendix A).

As a final consistency check, we confirm a basic prediction of our model according to which *reflective* traders systematically earn more than *non-reflective* traders (see Table A2 in Appendix A and Appendix B2). These findings are in line with the previous results of Noussair, Tucker and Xu (2014), Corgnet et al. (2015) and Corgnet, DeSantis and Porter (2018).

6. Conclusion

What elements of the market institution have the ability to enhance its informational efficiency? We examine one aspect of the question by building on prior experimental designs to develop an asset market environment in which the distribution of private information is analyzed.

We derived conjectures based on a learning model with heterogenous agents. In our model, traders differ in their ability to apply Bayes' rule to infer private information from market orders. We define traders as being either *reflective* or *non-reflective* depending on whether they use Bayesian updating or not. Our model shows that, for the same level of aggregate information in a market, informational efficiency is substantially higher when private information is concentrated in the hands of a few *fully* informed traders than when private information is widely dispersed. Our model also allowed us to test a point prediction according to which markets populated by three *fully* informed traders would be more informationally efficient than markets populated solely by *partially* informed traders. We confirm these conjectures by comparing the informational efficiency of experimental asset markets in a series of treatments that only differed in the distribution of private information. Because we collected data on traders' *cognitive reflection* scores, we were also able to show that, in line with our learning model, markets with a higher proportion of *reflective* traders achieved a higher level of informational efficiency. Finally, we largely confirmed our conjecture that trading volumes should be lower in markets in which the informational efficiency is predicted to be higher.

Our findings help reconcile different views regarding the informational efficiency of markets by showing that the extent of the transmission of private information to asset prices crucially hinges on the distribution of information in the market. Our results suggest that markets in which information is concentrated in the hands of insiders are more likely to achieve strong-form efficiency than markets in which this is not the case. This implies the level of informational efficiency might substantially differ across markets and industries. Future research should

investigate further the inability of markets to aggregate highly fragmented information because, as Hayek (1945, page 520) puts it, the ultimate economic problem considers the “utilization of knowledge which is not given to anyone in its totality”.

7. References

Asparouhova, E. and P. Bossaerts (2017) “Experiments on Percolation of Information in Dark Markets,” *The Economic Journal* 127 (605), F518-F544.

Barner, M., Feri, F. and C. Plott (2005) “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.

Baron-Cohen, S., Jolliffe, T., Mortimore, C. and M. Robertson (1997) “Another Advanced test of Theory of Mind: Evidence From Very High Functioning Adults With Autism or Asperger Syndrome,” *Journal of Child Psychology and Psychiatry* 38, 813-822.

Barreda Tarrazona, I. Grimalda, G; Morone, A., Nuzzo, S. and A. Teglio (2017) “Centralizing Information Improves Market Efficiency More Than Increasing Information: Results From Experimental Asset Markets,” Kiel Working Paper, 2072.

Bazerman, M. and W. Samuelson (1983) “I Won the Auction But Don’t Want the Prize,” *Journal of Conflict Resolution* 27(4), 618-634.

Bossaerts, P. (2009) “What Decision Neuroscience Teaches us About Financial Decision Making,” *Annual Review of Financial Economics* 1, 383-404.

Bossaerts, P., Frydman, C. and J. Ledyard (2013) “The Speed of Information Revelation and Eventual Price Quality in Markets with Insiders: Comparing two Theories,” *Review of Finance* 18 (1), 1-22

Brennan, M. and H. Cao, H (1996) “Information, Trade, and Derivative Securities,” *The Review of Financial Studies* 9(1), 163-208.

Bruguier, A., Quartz, S. and P. Bossaerts (2010) “Exploring the Nature of ‘Trader Intuition’,” *Journal of Finance* 65, 1703-1723.

Camerer, C. and K. Weigelt (1991) “Information Mirages in Experimental Asset Markets”, *Journal of Business* 64, 463-93.

Campitelli, G., and M. Labollita (2010) “Correlations of Cognitive Reflection with Judgments and Choices,” *Judgment and Decision Making* 5, 182-191.

Cason, T. (2000) “The Opportunity for Conspiracy in Asset Markets Organized with Dealer Intermediaries,” *Review of Financial Studies* 13, 385-416.

Cason, T. and D. Friedman (1996) "Price formation in double auction markets," *Journal of Economic Dynamics and Control* 20(8), 1307-1337.

Copeland, T. and D. Friedman (1987) "The Effect of Sequential Information Arrival on Asset Prices: An Experimental Study," *Journal of Finance* 42, 763-797.

Copeland, T. and D. Friedman (1991) "Partial Revelation of Information in Experimental Asset Markets," *Journal of Finance* 46, 265-295.

Corgnet, B., DeSantis, M. and D. Porter (2015) "On the Cognitive Foundations of Information Aggregation in Asset Markets: Reflective Learning & Market Efficiency," Economic Science Institute working papers, 15-15.

Corgnet, B., M. DeSantis and D. Porter (2018) "What Makes a Good Trader? On the Role of Intuition and Reflection on Trader Performance," *Journal of Finance* 73(3), 1113-1137.

Corgnet, B., R. Hernan-Gonzalez, P. Kujal and D. Porter (2015) "The Effect of Earned Versus House Money on Price Bubble Formation in Experimental Asset Markets," *Review of Finance* 19 (4), 1455-1488.

Daniel, K., Hirshleifer, D. and A. Subrahmanyam (1998) "Investor Psychology and Security Market Under- and Overreactions," *Journal of Finance* 53, 1839-1885.

De Martino, B., O'Doherty, J., Debajyoti, R., Bossaerts, P. and C. Camerer (2013) "In the Mind of The Market: Theory of Mind Biases Value Computation During Financial Bubbles," *Neuron* 79, 1222-1231.

Diamond, D. and R. Verrecchia (1981) "Information Aggregation in a Noisy Rational Expectations Economy, *Journal of Financial Economics*," 9 (3), 221-235.

Easley D. and M. O'Hara (1987) "Price, Trade Size, and Information in Securities Markets," *Journal of Financial Economics* 19(1), 69-90.

Esponda, I. (2008), "Behavioral Equilibrium in Economies with Adverse Selection," *American Economic Review* 98, 1269-1291.

Eyster, E. and M. Rabin (2005) "Cursed Equilibrium," *Econometrica* 73(5), 1623-1672.

Eyster, E., Rabin M. and D. Vayanos (2018) "Financial Markets where Traders Neglect the Information Content in Prices," *Journal of Finance*, In Press.

Fama, E. (1970) "Efficient Capital Markets: A Review of Theory and Empirical Work," *The Journal of Finance* 25(2), 383-417.

Fama, E. (1991) "Efficient Capital Markets: II," *The Journal of Finance* 46(5), 1575-1617.

Fama, E. (2008) "A Brief History of the Efficient Market Hypothesis," *American Finance Association* speech.

Frederick, S. (2005) “Cognitive Reflection and Decision Making,” *Journal of Economic Perspectives* 19(4), 25-42.

Friedman, D. (1991) “A Simple Testable Model of Double Auction Markets,” *Journal of Economic Behavior & Organization* 15(1), 47-70.

Friedman, D., G. Harrison, and J. Salmon (1984) “The Informational Efficiency of Experimental Asset Markets,” *Journal of Political Economy* 92, 349-408.

Frydman, C., Barberis, N., Camerer, C., Bossaerts, P. and A. Rangel (2014) “Using Neural Data to Test a Theory of Investor Behavior: An Application to Realization Utility,” *Journal of Finance* 69 (2), 907-946.

Gjerstad, S. and J. Dickhaut (1998) “Price Formation in Double Auctions,” *Games and Economic Behavior* 22(1), 1-29.

Glosten, L. and P. Milgrom (1985) “Bid, Ask and Transaction Prices In A Specialist Market With Heterogeneously Informed Traders,” *Journal of Financial Economics* 4(1), 71-100.

Gode, D. and S. Sunder (1993) “Allocative Efficiency of Markets With Zero-Intelligence Traders: Market as a Partial Substitute For Individual Rationality,” *Journal of political economy* 101 (1), 119-137.

Grossman, S. and J. Stiglitz (1980) “On the Impossibility of Informationally Efficient Markets,” *American Economic Review* 70(3), 393- 407.

Hanson, R., Oprea, R. and D. Porter (2006) “Information Aggregation and Manipulation in an Experimental Market,” *Journal of Economic Behavior and Organization* 60, 449-459.

Hayek, F. (1945) “The Use of Knowledge in Society,” *American Economic Review* 35(4), 519-530.

Hefti, A., Heinke, S. and F. Schneider (2016) “Mental Capabilities, Trading Styles and Asset Market Bubbles: Theory and Experiment,” Working Paper, Research Gate.

Hellwig, M. (1980) “On the Aggregation of Information in Competitive Markets,” *Journal of Economic Theory* 22(3), 477-498.

Hoppe, E., and D. Kusterer (2011) “Behavioral Biases and Cognitive Reflection,” *Economics Letters* 110, 97-100.

Kacperczyk, M., Van Nieuwerburgh, S. and L. Veldkamp (2016), “A Rational Theory of Mutual Funds' Attention Allocation,” *Econometrica* 84, 571-626.

Kahneman, D. (2011) “Thinking Fast and Slow,” *New York: Farrar, Straus and Giroux*.

Krahnen, J. and M. Weber (1999) “Does Information Aggregation Depend on Market Structure?, *Zeitschrift für Wirtschafts- und Sozialwissenschaften* 119, 1-22.

Krahnen, J. and M. Weber (2001) “Market making in the Laboratory: Does Competition Matter?” Experimental Economics 4(1) 55-85.

Lesage, E., G. Navarrete, and W. De Neys (2013) “Evolutionary Modules and Bayesian Facilitation: The Role of General Cognitive Resources,” *Thinking & Reasoning* 19, 27-53.

Lintner, J. (1969) “The Aggregation of Investor’s Diverse Judgments and Preferences in Purely Competitive Security Markets,” *Journal of Financial and Quantitative Analysis* 4(4), 347-400.

Mondria, J. “Portfolio Choice, Attention Allocation, and Price Comovement,” *Journal of Economic Theory* 145, 5, 1837-1864.

Nöth, M. and M. Weber (2003) “Information Aggregation with Random Ordering: Cascades and Overconfidence,” *Economic Journal* 113, 166-189.

Noussair, C. and S. Tucker (2014) “A Collection of Surveys on Market Experiments,” *Journal of Economic Surveys* 27(3), 395-397.

Noussair, C., S. Tucker and Y. Xu (2014) “A Futures Market Reduces Bubbles but Allows Greater Profit for More Sophisticated Traders,” Working Papers in Economics 14/12, University of Waikato, Department of Economics.

O’Brien, J and S. Srivastava (1991): “Dynamic Stock Markets with Multiple Assets: An Experimental Analysis,” *Journal of Finance* 46(5), 1811-1838.

Oechssler, J., Roider, A. and P. Schmitz (2009) “Cognitive Abilities and Behavioral Biases,” *Journal of Economic Behavior & Organization* 72(1), 147-152.

Palan, S. (2013) “A Review of Bubbles and Crashes In Experimental Asset Markets,” *Journal of Economic Surveys* 27 (3), 570-588.

Peng, L. and W. Xiong (2006) “Investor Attention, Overconfidence and Category Learning,” *Journal of Financial Economics* 80, 563-602.

Plott, C. and S. Sunder (1982) “Efficiency of Experimental Security Markets with Insider Information: An Application to Rational Expectations Models,” *Journal of Political Economy* 90, 663-698.

Plott, C. and S. Sunder (1988) “Rational Expectations and the Aggregation of Diverse Information In Laboratory Security Markets,” *Econometrica* 56, 1085-1118.

Plott, C, Wit, J. and W. Yan (2003) “Parimutuel Betting Markets as Information Aggregation Devices: Experimental Results,” *Economic Theory* 22, 311-351.

Rabin, M. (2002) “Inference by Believers in the Law of Small Numbers,” *The Quarterly Journal of Economics* 117(3), 775-816.

Sirota, M., M. Juanchich, and Y. Hagmayer (2014) “Ecological rationality or nested sets? Individual differences in cognitive processing predict Bayesian reasoning,” *Psychonomic Bulletin & Review* 21, 198–204.

Stanovich, K. (2009) “What Intelligence Tests Miss: The Psychology of Rational Thought,” *Yale University Press*.

Thaler, R. (1988) “Anomalies: The Winner's Curse,” *Journal of Economic Perspectives* 2(1), 191-202.

Thaler, R. (1991) “The Winner's Curse: Paradoxes and Anomalies of Economic Life,” *New York: Free Press*.

Toplak, M., West, R., and K. Stanovich (2011) “The Cognitive Reflection Test as a Predictor of Performance on Heuristics and Biases Tasks,” *Memory & Cognition* 39, 1275-1289.

Toplak, M., West, R., and K. Stanovich (2014) “Assessing Miserly Information Processing: An Expansion of the Cognitive Reflection Test,” *Thinking & Reasoning* 20(2), 147-168.

Tversky, A. and D. Kahneman (1974) “Judgment Under Uncertainty: Heuristics and Biases,” *Science* 185, 1124-1131.

Veiga, H., & Vorsatz, M. (2009) “Price Manipulation in an Experimental Asset Market,” *European Economic Review* 53 (3), 327-342.

Veiga, H., & Vorsatz, M. (2010) “Information Aggregation in Experimental Asset Markets in The Presence of a Manipulator,” *Experimental Economics* 13(4), 379-398.

Vives, X. and L. Yang (2017) “Costly Interpretation of Asset Prices”, Working Paper.

Wilson, R. (1987) “On equilibria of Bid-Ask Markets,” In Arrow and the Ascent of Modern Economic Theory, *Macmillan Press*, London, G. Feiwell (Ed.), 375-414.

Appendix A. Additional analyses

We present two robustness checks in this appendix. First, Table A1 shows that adding three *partially* informed traders to a market populated by four *fully* informed traders only has a slight impact on the informational efficiency of the market (see the coefficient of the '*33% fully informed'* dummy variable). Moreover, these markets with seven (*partially* and *fully*) informed traders do not perform as well as (with respect to informational efficiency) as markets with six *fully* informed traders (see the coefficient of the *'50% fully informed'*). This result holds even when controlling for participants' CRT scores via the Proportion of reflective traders variable or the Average CRT score variable.

Table A1. Random effects panel regression of MAD values per market as a function of cognitive reflection and treatment dummies

MAD	Treatments '*33% fully informed*', '*50% fully informed*' and '*33% fully informed / 25% partially informed*'		
Specification:	[1]	[2]	[3]
Dummy '*33% fully informed*'	3.541 (12.279)	0.785 (12.618)	4.696 (12.139)
Dummy '*50% fully informed*'	-37.472*** (8.997)	-43.038*** (8.102)	-38.910*** (9.354)
Proportion of *reflective* traders	-34.650 (38.228)	-	-
Average CRT score in a session	-	-19.559** (8.115)	-
Market number	-2.545*** (0.732)	-2.541*** (0.732)	-2.544*** (0.731)
True asset value	0.136*** (0.032)	0.136*** (0.032)	0.136*** (0.032)
Intercept	103.166*** (18.016)	139.563*** (16.400)	88.725*** (11.909)
Observations (sessions)	340 (20)	340 (20)	340 (20)
Prob > χ^2	0.000	0.000	0.000
R^2	0.173	0.183	0.169

*p-value<0.10, **p-value<0.05 and ***p-value<0.01. Robust standard errors clustered at the session level in parentheses.

In Table A2, we present our second robustness check which demonstrates that *reflective* traders earn more than *non-reflective* traders when pooling all treatments together (see the positive coefficient for CRT score). The CRT score is calculated as the number of correct answers of a

given trader on the CRT. We also control for theory of mind scores, using the eye gaze test (Baron-Cohen et al. 1997), to assess subjects' capacity to infer other's intentions. In this task, participants looked at images of people's eyes and had to choose one of four feelings that best described the mental state of the person whose eyes were shown. Our theory of mind score is defined as the number of correct answers to the 36 question, 10-minute test. In line with Bruguier, Quartz and Bossaerts (2010) and De Martino et al. (2013), Hefti et al. (2016), Corgnet et al. (2018) we find a positive relationship between trader's performance and theory of mind scores. We refer to Signal "Not *x*" Dummies as dummy variables set to 1 for subjects who received the signal "Not *x*". The "Two signals Dummy" is set to 1 for subjects who received two signals and are thus *fully* informed. The Top 25% Reflective Dummy and Top 25% Theory of Mind Dummy variables are set to one if the participant scored in the top 25% of all participants in this study (across all 35 experimental sessions) on the CRT or the Theory of Mind Eye Gaze test, respectively, and zero otherwise.

Table A2. Random effects panel regression of trader earnings as a function of cognitive reflection and information held by traders

Market earnings (in francs)	All treatments	
	[1]	[2]
Top 25% Reflective Dummy	100.004***	-
	(27.009)	
CRT score	-	32.774***
		(6.337)
Top 25% Theory of Mind Dummy	50.917*	-
	(26.543)	
Theory of Mind score	-	4.414*
		(2.720)
Gender Dummy	69.424***	60.507**
(1 if male)	(25.477)	(25.078)
Market number	0.609	0.628
	(1.438)	(1.438)
True asset value	3.977***	3.977***
	(0.075)	(0.075)
Signal "Not 50" Dummy	142.048***	146.514***
	(41.064)	(40.635)
Signal "Not 240" Dummy	-96.761**	-93.721**
	(41.620)	(41.058)
Signal "Not 490" Dummy	151.921***	155.379***
	(28.964)	(28.443)
Two signals Dummy	246.405***	246.742***
	(14.683)	(14.693)
Intercept	1,039.539***	890.062***
	(31.723)	(81.832)
Observations	6,511	6,511
Prob > χ^2	0.000	0.000
R^2	0.660	0.662

*p-value<0.10, **p-value<0.05 and ***p-value<0.01. Robust standard errors clustered at the trader level in parentheses. Note that the number of observations is lower than the total number of participants times the number of markets (7,106) because we did not collect theory of mind scores for three sessions in the '*100% partially informed*' treatment.

Appendix B. Model conjectures

B1. The model

B1.1. Model description

This appendix provides a detailed description of the model presented in Section 3. It is based upon Appendix G in Corgnet et al. (2015) and has been modified to account for *informed* and *uninformed* traders rather than *partially* informed traders. We assume there are two types, τ, of traders. *Reflective* traders ($\tau = R$) utilize market orders (bids, asks and contract prices) to update their beliefs of the asset's value, while *non-reflective* traders ($\tau = NR$) do not update their initial beliefs. A market order implies an update of the bid-ask spread or a transaction. Each order corresponds to one of the following events: (1) best bid is accepted; (2) best ask is accepted; or (3) the bid-ask spread is updated.

Traders will either be *informed* (i.e., know the true asset value with certainty), $s = I$, or *uninformed* (i.e., only know the possible asset values along with their probabilities), $s = U$.[28] We thus have four possible trader types: $\eta = (\tau, s)$ where $\tau \in \{R, NR\}$ and $s \in \{U, I\}$. Let μ_n^{η} represent the belief of a trader of type η after event n, which is denoted by $\{e(n)\}$. *Uninformed reflective* traders update their beliefs using Bayes' rule with $i \in \{50,240,490\}$ as follows:

$$\mu_n^{R,U} = E\left[v|\mu_{n-1}^{R,U}, \{e(n)\}\right] = \sum_i P\left[v = i|\mu_{n-1}^{R,U}, \{e(n)\}\right] \times i$$

where $\mu_{n-1}^{R,U}$ represents their prior belief and v represents the asset value.

This trader updates her belief according to the following formula:

$$\begin{aligned}
\mu_n^{R,U} &= \\
&= E\left[v|\mu_{n-1}^{R,U}, \{e(n)\}\right] \\
&= P\left[v = 50|\mu_{n-1}^{R,U}, \{e(n)\}\right] \times 50 + P\left[v = 240|\mu_{n-1}^{R,U}, \{e(n)\}\right] \times 240 + P\left[v = 490|\mu_{n-1}^{R,U}, \{e(n)\}\right] \\
&\qquad \times 490 \\
&= \frac{P[\{e(n)\}|v = 50]Prior^{50}(n-1)}{P[\{e(n)\}]} \times 50 + \frac{P[\{e(n)\}|v = 240]Prior^{240}(n-1)}{P[\{e(n)\}]} \times 240 \\
&\qquad + \frac{P[\{e(n)\}|v = 490]Prior^{490}(n-1)}{P[\{e(n)\}]} \times 490
\end{aligned}$$

where

[28] The variable s represents the informational type of the trader. In the case of *partially* informed traders, s identifies the signal received by the trader, i.e. "Not 50", "Not 240" or "Not 490".

$$P[\{e(n)\}] = P[\{e(n)\}|v = 50] \times Prior^{50}(n-1) + P[\{e(n)\}|v = 240] \times Prior^{240}(n-1) + P[\{e(n)\}|v = 490] \times Prior^{490}(n-1).$$

After the $(n-1)^{st}$ event, the *uninformed reflective* trader believes that the probability the true asset value is v is given by $Prior^{v}(n-1)$.

The term, $P[\{e(n)\}|v]$, i.e. the probability that event $e(n)$ occurred given the asset's true value is v, must be calculated. The derivation of this term is dependent upon the type of event as well as the region in which the event occurred. The set of possible prices is the interval 50 through 490 on the real number line, as we restrict bids and asks to this interval which corresponds to lowest (highest) possible asset value. We subdivide this interval into several subintervals, called regions, which are delineated by the traders' beliefs along with the overall interval's endpoints, 50 and 490. Suppose the *uninformed reflective* trader believes the asset value is $\mu_{n-1}^{R,U}$ = 130. As part of the learning process, this trader must determine the likelihood of event $e(n)$ occurring assuming the true asset value is 50, 240 or 490. Suppose this trader assumes the true asset value is 240. Then this trader assumes that all *informed* traders know the asset value is 240, while any *uninformed non-reflective* traders hold their initial belief of 223.5.[29] In this scenario the interval [50,490] would be divided into 4 regions: [50,130], [130,223.5], [223.5,240], and [240,490].

Now this trader will update her belief based upon the probability of a particular event occurring with a particular region. And, this probability is based upon the likelihood of the different traders acting in each region. Recall that traders submit offers to buy below their belief and offers to sell above their belief. So, if a *reflective* trader observes a new offer to buy at 225, then she knows this could only have been submitted by a *fully* informed trader with a belief of 240. Let b_{n-1} (a_{n-1}) represent the current best bid (ask) when $e(n)$ occurs, while b_n (a_n) represents the improved bid (ask) corresponding to $e(n)$.[30] We denote the region in which an event occurred by X_r, where $r \in \{b_{n-1}, a_{n-1}, b_n, a_n\}$ If $r = 50$ (240) [490] and the trader assumes the asset value is 50 (240) [490], then region, X_r, will consist of the single point, 50 (240) [490] because the *informed* traders' belief coincides with r.

We consider the following three cases characterized by event type.

[29] Note that the initial beliefs of *reflective* and *non-reflective* traders with the same clue are identical.
[30] Note that b_{n-1} (a_{n-1}) may not correspond to the $(n-1)^{st}$ bid (ask). Rather, the subscript corresponds to the event number. Suppose, for example, that the order book contains bids of 75 and 100. The current best bid, b_{n-1}, would be 100. Further, suppose $e(n)$ corresponds to a trader accepting the $b_{n-1} = 100$. Thus, the new current best bid would be $b_n = 75$.

Case 1: *the current best bid is traded,* $\{e(n)\} := \{T(b_{n-1})\}$

Suppose event n corresponds to the acceptance of the current best bid, b_{n-1}. Then the conditional probability that event $\{T(b_{n-1})\}$ occurred, i.e. the current best bid b_{n-1} was traded, given the true asset value is v, is

$$P[\{e(n)\}|v] = P[\{T(b_{n-1})\}|v] = \rho_{sell}^{X_{b_{n-1}}}$$

where $\rho_{sell}^{X_{b_{n-1}}}$ represents the probability that a trader is willing to sell in the region $X_{b_{n-1}}$, i.e. the region in which the $(n-1)^{st}$ bid, b_{n-1}, resides.

To calculate the probability of selling in region $X_{b_{n-1}}$ we consider the beliefs of four possible trader types: *informed reflective* traders, *informed non-reflective* traders, *uninformed reflective* traders, and *uninformed non-reflective* traders.[31] Suppose $v = 50$ in the above equation. This implies that the beliefs of the *informed reflective* and *informed non-reflective* traders are 50. The belief of the *uninformed non-reflective* traders is 223.5. The probability that a trader is willing to sell in $X_{b_{n-1}}$ is $\chi^{R,U} \times \alpha/2 + \chi^{\tau,I} \times 1/2 + \chi^{NR,U} \times (1-\alpha)/2$, where $\chi^{\tau,s}$ is one if the trader's belief is less than or equal to the lower bound of the corresponding region ($X_{b_{n-1}}$) and zero otherwise. That is, the *uninformed reflective* trader calculates the proportion of traders holding a belief less than or equal to the current best bid. The exogenous parameter α represents the actual proportion of reflective traders in the market.[32] The proportion of *(un)informed non-reflective* traders is represented by $(1-\alpha)$.

Case 2: *the current best ask is traded,* $\{e(n)\} := \{T(a_{n-1})\}$

Suppose the current best ask, a_{n-1}, is accepted. Then the conditional probability that event $\{T(a_{n-1})\}$ occurred, i.e. the current best ask a_{n-1} was traded, given the true value of the asset is v, is:

$$P[\{e(n)\}|v] = P[\{T(a_{n-1})\}|v] = \rho_{buy}^{X_{a_{n-1}}}$$

where $\rho_{buy}^{X_{a_{n-1}}}$ represents the probability that a trader is willing to buy in region $X_{a_{n-1}}$.

[31] The *informed reflective* and *informed non-reflective* types could be consolidated into a single group as neither type will update its belief given it knows the actual value.

[32] *Reflective* traders assume that other *reflective* traders share their belief of the proportion of *reflective* traders in the market. It is possible that *reflective* traders' belief of the proportion of *reflective* traders in the market is not accurate. If this proportion is not common information known by all traders, then Corgnet, DeSantis and Porter (2015) consider the situation in which this proportion is not common information known by all traders. They show this assumption impacts both the simulation, as well as the experimental, results.

In this case we determine whether the beliefs of the four different types of traders (η) is greater than or equal to the upper bound of the region and add $\alpha/2$ ($[1-\alpha]/2$) to $\rho_{buy}^{X_{a_{n-1}}}$, i.e. we check if their belief is greater than or equal to the current best ask.

Case 3: *both the current best bid and ask are improved,* $\{e(n)\} := \{\mathrm{NT}(b_{n-1}) \cap \mathrm{NT}(a_{n-1})\}$

Suppose neither the current best bid nor the current best ask is traded. Then both the bid and the ask are updated from the current best bid, b_{n-1}, and ask, a_{n-1}, to the new (improved) best bid, b_n, and ask, b_n. Then the conditional probability that event $\{e(n)\} = \{\mathrm{NT}(b_{n-1}) \cap \mathrm{NT}(a_{n-1})\}$ occurred, i.e. both the current best bid and the best ask were "not traded" (NT), given the true value of the asset is v, is:

$$
\begin{aligned}
P[\{e(n)\}|v] &= P[\{\mathrm{NT}(b_{n-1}) \cap \mathrm{NT}(a_{n-1})\}\,|v, b_{n-1}, a_{n-1}] \\
&= P[\mathrm{NT}(b_{n-1})\,|v, b_{n-1}, a_{n-1}] && [i] \\
&\times P[\mathrm{NT}(a_{n-1})|v, b_{n-1}, a_{n-1}, \mathrm{NT}(b_{n-1})] && [ii] \\
&\times P\big[b_n \in X_{b_n} \,\cap\, a_n \in X_{a_n}\big|v, b_{n-1}, a_{n-1}\big] && [iii]
\end{aligned}
$$

Consider each term in this product separately.

[i] The first term is equivalent to one minus the probability that the current best bid, b_{n-1}, was not traded. This is equal to one minus the probability that traders are willing to sell in the region $X_{b_{n-1}}$. Thus, we have:

$$P[\mathrm{NT}(b_{n-1})\,|v, b_{n-1}, a_{n-1}] = \left(1 - \rho_{sell}^{X_{b_{n-1}}}\right).$$

[ii] The second term may be rewritten as:

$$
\begin{aligned}
P[\mathrm{NT}(a_{n-1})|v, b_{n-1}, a_{n-1}, \mathrm{NT}(b_{n-1})] &= 1 - P[\mathrm{T}(a_{n-1})|v, \mathrm{NT}(b_{n-1})] \\
&= 1 - P\big[a\ trader\ buys\ in\ X_{a_{n-1}}|v, a\ trader\ buys\ in\ X_{b_{n-1}}\big] \\
&= 1 - P\big[a\ trader\ buys\ in\ X_{b_{n-1}}|v, a\ trader\ buys\ in\ X_{a_{n-1}}\big] \\
&\times \frac{P\big[a\ trader\ buys\ in\ X_{a_{n-1}}\big]}{P\big[a\ trader\ buys\ in\ X_{b_{n-1}}\big]} = 1 - \rho_{buy}^{X_{a_{n-1}}}/\rho_{buy}^{X_{b_{n-1}}}
\end{aligned}
$$

where we make use of Bayes' rule as well as the fact that since $X_{a_{n-1}}$ must lie to the right of (or possibly coincide with) $X_{b_{n-1}}$ on the number line, $P\big[a\ trader\ buys\ in\ X_{b_{n-1}}|v, a\ trader\ buys\ in\ X_{a_{n-1}}\big] = 1$.

[iii] Finally, the third term may be expressed as:

$$P\left[b_n \in X_{b_n} \cap a_n \in X_{a_n} \middle| v, b_{n-1}, a_{n-1}\right] = P\left[b_n \in X_{b_n} \middle| v, b_{n-1}, a_{n-1}\right] \times P\left[a_n \in X_{a_n} \middle| v, b_{n-1}, a_{n-1}, b_n\right]$$

$$= \left[\sum_{\eta} \frac{\gamma}{2} \left\{ \frac{min\left(u_{X_{b_n}}, a_{n-1}\right) - max\left(l_{X_{b_n}}, b_{n-1}\right)}{min\left(\mu_{n-1}^{\eta}, a_{n-1}\right) - max(50, b_{n-1})} \right\}\right]$$

$$\times \left[\sum_{\eta} \frac{\gamma}{2} \left\{ \frac{min\left(u_{X_{a_n}}, a_{n-1}\right) - max\left(l_{X_{a_n}}, b_n\right)}{min(490, a_{n-1}) - max\left(\mu_{n-1}^{\eta}, b_n\right)} \right\}\right]$$

where the sums are taken over the four trader types, $\eta = (\tau, s)$: *informed reflective* traders, *informed non-reflective* traders, *uninformed reflective* traders, and *uninformed non-reflective* traders. As the true value is assumed to be v, it is also assumed that the *informed* traders believe the true asset value is v in these calculations. In addition, u_{X_r} (l_{X_r}) represents the upper (lower) bound of the region X_r, $r \in \{b_{n-1}, a_{n-1}, b_n, a_n\}$. The variable γ is set to α for *reflective* trader *types* and $(1 - \alpha)$ for *non-reflective* trader *types*.[33]

Thus, we have:

$$P[\{e(n)\}|v] = \left(1 - \rho_{sell}^{X_{b_{n-1}}}\right) \times \left(1 - \frac{\rho_{buy}^{X_{a_{n-1}}}}{\rho_{buy}^{X_{b_{n-1}}}}\right) \times \left[\sum_{\eta} \frac{\gamma}{2} \left\{ \frac{min\left(u_{X_{b_n}}, a_{n-1}\right) - max\left(l_{X_{b_n}}, b_{n-1}\right)}{min\left(\mu_{n-1}^{\eta}, a_{n-1}\right) - max(50, b_{n-1})} \right\}\right]$$

$$\times \left[\sum_{\eta} \frac{\gamma}{2} \left\{ \frac{min\left(u_{X_{a_n}}, a_{n-1}\right) - max\left(l_{X_{a_n}}, b_n\right)}{min(490, a_{n-1}) - max\left(\mu_{n-1}^{\eta}, b_n\right)} \right\}\right]$$

when the n^{th} event is an improvement to the bid-ask spread.[34]

B1.2. Model example

We demonstrate the learning model by providing an illustrative example. Suppose there are six *informed* traders and four *reflective* traders, i.e. $\alpha = 4/12$. For simplicity, assume two of the *reflective* traders are *informed* and two are *uninformed*. Let the true value of the asset be 50.

Based on their prior information *uninformed* traders' initial belief of the true asset value (prior to the occurrence of any event) is $\mu_n^{R,U} = \mu_n^{NR,U} = 0.35 \times 50 + 0.45 \times 240 + 0.20 \times 490 =$

[33] When running the simulations described in Section B2, traders are first designated as either *informed* or *uninformed*. Next, the *reflective* traders are randomly identified. For example, if $\alpha = 5/12$, then five traders would randomly be identified as *reflective* with the remaining seven identified as *non-reflective*. Thus, it is possible that the number of *reflective* traders is not evenly distributed across the *informed* and *uninformed* types. However, this more closely resembles the laboratory environment.

[34] Strictly speaking, there exist cases in which only the bid or the ask is improved. However, we assume traders do not update their beliefs on these unlikely events. These cases occur, for example, when a trader wants to update both the bid and the ask, but does not have the requisite amount of cash/shares to cover the updated bid/ask as well as any outstanding bids/asks the trader might have.

223.5. *Informed* traders' initial belief is $\mu_n^{R,I} = \mu_n^{NR,I} = 490$. The *informed* traders will not update their belief of the asset value, while the *uninformed reflective* traders will. And, while the *uninformed reflective* traders know that *informed* traders exist, they do not know their belief as they do not know the true asset value. Thus, throughout the updating process, when these traders calculate the probability of an event occurring given the value (50, 240, 490), they must assume that the informed traders' belief coincides with the assumed asset value.

First market event: bid-ask spread {180, 360}

We consider the *uninformed reflective* traders. Suppose the first trader submits an initial bid and ask spread of $\{b_1 = 180, a_1 = 360\}$. The *uninformed reflective* traders will update their belief to 227.3 by applying the following Bayesian formula:

$$\mu_{n,U}^R = \frac{P[\{e(n)\}|v = 50]Prior^{50}(n-1)}{P[\{e(n)\}]} 50 + \frac{P[\{e(n)\}|v = 240]Prior^{240}(n-1)}{P[\{e(n)\}]} 240 + \frac{P[\{e(n)\}|v = 490]Prior^{490}(n-1)}{P[\{e(n)\}]} 490$$

where n would be one and $P[\{e(n)\}] = P[\{e(n)\}|v = 50]Prior^{50}(n-1) \times P[\{e(n)\}|v = 240] \times Prior^{240}(n-1) + P[\{e(n)\}|v = 490] \times Prior^{490}(n-1)$. The initial prior probabilities are $Prior^{50}(0) = 0.35$, $Prior^{240}(0) = 0.45$ and $Prior^{490}(0) = 0.20$, respectively.

It remains to calculate $P[\{e(1)\}|v]$ for $v \in \{50,240,490\}$. This probability is equal to the probability that neither the current best bid nor the current best ask was traded given the value v.[35] In this case the probability of the $\{180, 360\}$ bid-ask pair is equal to the following:

$$P[\{e(1)\}|v] = P[\{\mathrm{NT}(b_0) \cap \mathrm{NT}(a_0)\} | v, b_0, a_0]$$
$$= P[\mathrm{NT}(b_0) | v, b_0, a_0] \times P[\mathrm{NT}(a_0) | v, b_0, a_0, \mathrm{NT}(b_0)] \times P\left[b_1 \in X_{b_1} \cap a_1 \in X_{a_1} \middle| v, b_0, a_0\right]$$
$$= P[b_1 > b_0 | v, b_0, a_0] \times P[a_1 < a_0 | v, b_0, a_0, b_1 > b_0] \times P\left[b_1 \in X_{b_1} \cap a_1 \in X_{a_1} \middle| v, b_0, a_0\right] \quad (1)$$

where the default initial bid, b_0, is set to 50 and the default ask, a_0, is 490. X_{b_1} and X_{a_1} represent the regions in which the new bid (180) and new ask (360) reside. We assume all bids and asks, and hence prices, reside in the interval $[50,490]$. This interval is then divided into regions determined by the traders' beliefs as depicted in Figure B1.

[35] As previously noted in Section B1.1, this probability is equal to the product of three terms: (1) probability that the bid was not accepted given the value v, (2) probability that the ask was not accepted given the value v and the event that the bid was not traded and (3) probability that the new bid and ask lie in their respective regions.

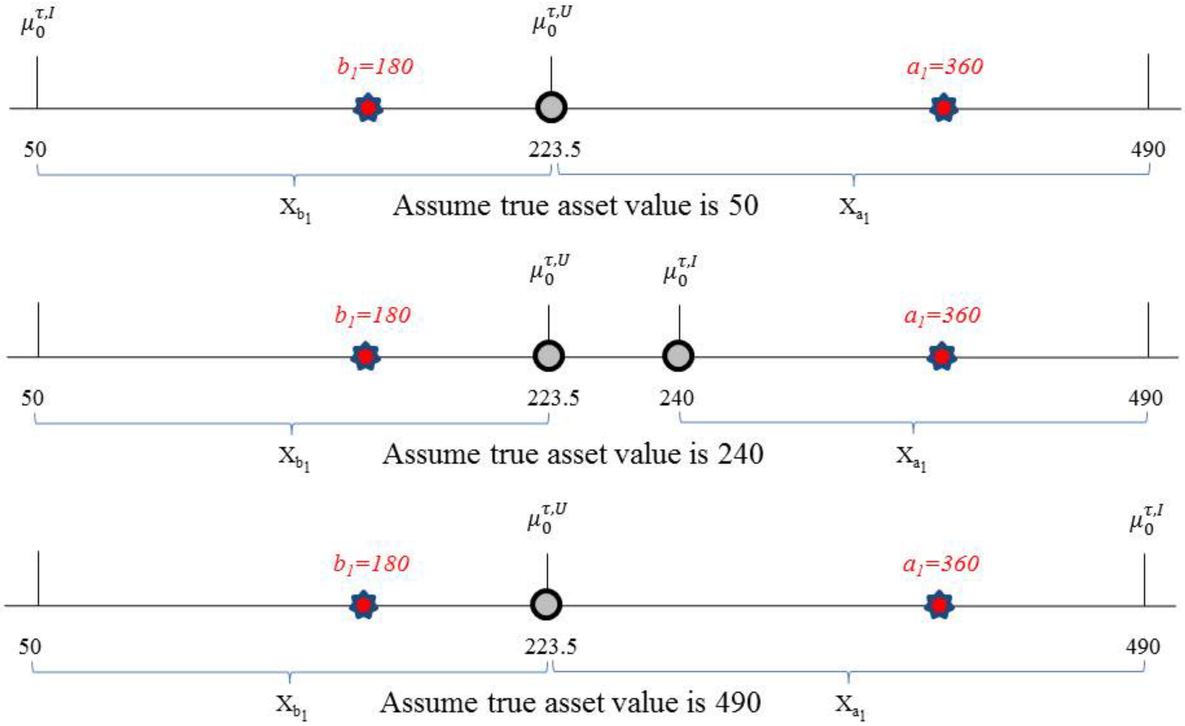

Figure B1. Representation of traders' initial beliefs, the first bid-ask improvement event, and the regions in which the new bid and ask exist. Note that reflective and non-reflective traders hold the same initial beliefs. Moreover, the regions, X_{a_1} and X_{b_1}, depend upon the assumed true asset value.

In words, the probability of the bid-ask improvement given the true value of the asset is v is equal to the product of the following terms:

[i] The probability that the new bid improves the old bid given the value of the asset is v, and given the old bid-ask pair.

[ii] The probability that the new ask improves the old ask given the value, v, the old bid-ask pair (b_0, a_0), and that the new bid improves the old bid.

[iii] The probability that the new bid resides in the region $X_{b_1} = [50{,}223.5]$ and the new ask is in $X_{a_1} = [223.5{,}490]$ given the value of the asset is $v = 50$ or 490, and given the old bid-ask pair. If the asset value is assumed to be 240, then these regions would be $X_{b_1} = [50{,}223.5]$ and $X_{a_1} = [240{,}490]$.

We address each term in the product in order. Assume the true value is 50 (240) {490}. The term [i] is equivalent to one minus the probability that a bid in the region [50,50] ([50,223.5]) {[50,223.5]} was accepted/traded, i.e., one minus the probability of a trader selling in this region, $\left(1 - \rho_{sell}^{X_{b_{n-1}}}\right)$. If the value is assumed to be 50, then this region consists of a single point in which

only informed traders would transact. If the value is assumed to be 240 or 490, then no traders would trade in this region as their beliefs would be greater than the upper bound of the region. Thus, *reflective* traders assess this probability to be 0.5 (0) {0}, which makes this first term equal to 0.5 (1) {1}.

The term [*ii*] is equivalent to one minus the ratio of the probability that a trader buys in the region X_{a_0} to the probability that a trader buys in the region X_{b_0}, i.e. $1-\rho_{buy}^{X_{a_0}}/\rho_{buy}^{X_{b_0}}$. If the asset value is assumed to be 50, then $X_{b_0}=[50{,}50]$ and $X_{a_0}=[223.5{,}490]$ and the probability of buying in each region is 1 and 0, respectively. Indeed, all *informed* traders would buy in the first region, while no traders would buy in the second region as all traders' beliefs would be less than the lower bound of the region. If the asset value is assumed to be 240, then $X_{b_0}=[50{,}223.5]$ and $X_{a_0}=[240{,}490]$ and the probability of buying in each region is 1 and 0, respectively. If the asset value is assumed to be 490, then $X_{b_0}=[50{,}223.5]$ and $X_{a_0}=[490{,}490]$ and the probability of buying in each region is 1 and 0.5 (again, all *informed* traders), respectively. Thus, this second term is equal to 1 (1) {0.5} given the trader assumes the true value to be 50 (240) {490}.

Finally, the term [*iii*] consists of a product of two sums. The first sum is the probability of a trader buying in the region containing the new bid, which is equivalent to the probability that a trader would submit a bid in that region (i.e., her belief is greater than the upper bound of the region). The second sum, the probability of a trader selling in the region containing the new ask given the new bid, is calculated in the same manner (i.e., her belief is less than the lower bound of the region). The term [*iii*] is calculated as follows:

$$\left[\sum_{\eta}\frac{\gamma}{2}\left\{\frac{min\left(u_{X_{b_1}},a_0\right)-max\left(l_{X_{b_1}},b_0\right)}{min\left(\mu_0^{\eta},a_0\right)-max(50,b_0)}\right\}\right]\times\left[\sum_{\eta}\frac{\gamma}{2}\left\{\frac{min\left(u_{X_{a_1}},a_0\right)-max\left(l_{X_{a_1}},b_1\right)}{min(490,a_0)-max\left(\mu_0^{\eta},b_1\right)}\right\}\right]$$

where $\eta=(\tau,s)$ with $\tau\in\{R,NR\}$ and $s\in\{I,U\}$. If the trader *type* being considered involves *reflective* (*non-reflective*) traders, then the variable γ is set to α $(1-\alpha)$ (i.e., $\gamma=\frac{1}{3}$ for this example).

To calculate this third term, the trader must account for the beliefs of all four trader *types*: *uninformed reflective*, *uninformed non-reflective*, *informed reflective* and *informed non-reflective*. Suppose the trader assumes the true asset value is 50. Then $X_{b_1}=[50{,}223.5]$ and $X_{a_1}=[223.5{,}490]$. This trader then assumes all *informed* traders hold the belief of 50, while all *uninformed* traders hold the belief of 223.5. Note that in this case only *uninformed* traders could

have submitted a bid of 180. Thus, the two types of *informed* traders would be excluded from the first sum in this third term. Moreover, as the initial beliefs of *reflective* and *non-reflective* traders are identical, we may combine the terms for these types ($\frac{\gamma}{2}+\frac{(1-\gamma)}{2}=\frac{4/12}{2}+\frac{8/12}{2}=\frac{1}{2}$) to yield:

$$\left[\frac{1}{2}\left\{\frac{223.5-50}{223.5-50}\right\}\right]\times\left[\frac{1}{2}\left\{\frac{490-223.5}{490-223.5}\right\}+\frac{1}{2}\left\{\frac{490-223.5}{490-180}\right\}\right].$$

The probability that this first event occurred given the true asset value is 50 is therefore given by

$$P[\{e(1)\}|v=50]=0.5\times 1\times\left[\frac{1}{2}\left\{\frac{223.5-50}{223.5-50}\right\}\right]\times\left[\frac{1}{2}\left\{\frac{490-223.5}{490-223.5}\right\}+\frac{1}{2}\left\{\frac{490-223.5}{490-180}\right\}\right]=0.23.$$

Next, suppose the *uninformed reflective* trader assumes the true asset value is 240. Then $X_{b_1}=[50{,}223.5]$ and $X_{a_1}=[240{,}490]$. The X_{a_1} region differs from the case in which the trader assumes the value is 50 because the trader now believes that *informed* traders hold the belief of 240. Note that any trader could have submitted this bid-ask combination. Thus, this third term is given by (where we again combine the *reflective* and *non-reflective* types given their common initial beliefs):

$$\left[\frac{1}{2}\left\{\frac{223.5-50}{223.5-50}\right\}+\frac{1}{2}\left\{\frac{223.5-50}{240-50}\right\}\right]\times\left[\frac{1}{2}\left\{\frac{490-240}{490-223.5}\right\}+\frac{1}{2}\left\{\frac{490-240}{490-240}\right\}\right].$$

The probability that this first event occurred given the true asset value is 240 is therefore given by

$$\begin{aligned}&P[\{e(1)\}|v=240]\\&\quad=1\times 1\times\left[\frac{1}{2}\left\{\frac{223.5-50}{223.5-50}\right\}+\frac{1}{2}\left\{\frac{223.5-50}{240-50}\right\}\right]\times\left[\frac{1}{2}\left\{\frac{490-240}{490-223.5}\right\}+\frac{1}{2}\left\{\frac{490-240}{490-240}\right\}\right]\\&\quad=0.93.\end{aligned}$$

Finally, suppose the *uninformed* trader assumes the true asset value is 490. Then $X_{b_1}=[50{,}223.5]$ and $X_{a_1}=[223.5{,}490]$. Note that in this case only *uninformed* traders holding the belief of 223.5 could submit the ask of 360 as *informed* traders hold the belief of 490. Thus, this third term is given by:

$$\left[\frac{1}{2}\left\{\frac{223.5-50}{223.5-50}\right\}+\frac{1}{2}\left\{\frac{223.5-50}{490-50}\right\}\right]\times\left[\frac{1}{2}\left\{\frac{490-223.5}{490-223.5}\right\}\right].$$

The probability that this first event occurred given the true asset value is 490 is therefore given by

$$P[\{e(1)\}|v=490]=1\times 0.5\times\left[\frac{1}{2}\left\{\frac{223.5-50}{223.5-50}\right\}+\frac{1}{2}\left\{\frac{223.5-50}{490-50}\right\}\right]\times\left[\frac{1}{2}\left\{\frac{490-223.5}{490-223.5}\right\}\right]=0.17.$$

Note that $0.23\times 0.35+0.93\times 0.45+0.17\times 0.20=0.533$. Then, after observing this first event, the *uninformed reflective* traders update their belief of the true asset value to:

$$\mu_1^{R,U}=\frac{0.23\times 0.35}{0.533}\times 50+\frac{0.93\times 0.45}{0.533}\times 240+\frac{0.17\times 0.20}{0.533}\times 490=227.3$$

Recall that the *uninformed non-reflective* traders, as well as the *informed* traders do, not update their beliefs. Figure B2 shows the traders' beliefs after this first event.

Second market event: sale at 180

To conclude this descriptive example, suppose the bid of 180 is accepted by the second randomly selected trader. The *uninformed non-reflective* trader updates her belief based upon this event by first calculating the probability of this event occurring given the asset value is 50, 240 or 490, respectively. That is, this trader calculates the probability that a given trader would be willing to sell in region $X_{b_1} = [50{,}223.5]$ given the true asset value is 50, 240 or 490, respectively. Note that this is region is the same for each assumed asset value (see Figure B2). Suppose the asset value is assumed to be 50. Then, only the *informed* traders would be willing to sell in this region. Thus, we have

$$P[\{e(2)\}|v = 50] = P[\{T(b_1)\}|v = 50] = \rho_{sell}^{X_{b_1}} = 1/2$$

as one-half (six) of the traders are informed (two are *reflective* and four are *non-reflective*).

If the *uninformed non-reflective* trader assumes the true asset value is 240, then no trader would be willing to sell in the region $X_{b_1} = [50{,}223.5]$ as each trader would hold a belief greater than the upper bound of this region. Indeed, given the assumption that the true asset value is 240, *informed* traders would believe the asset value is 240. And, *uninformed reflective* traders believe the asset value is 227.3, while *uninformed non-reflective* traders believe the asset value is 223.5. Thus,

$$P[\{e(2)\}|v = 240] = P[\{T(b_1)\}|v = 240] = \rho_{sell}^{X_{b_1}} = 0.$$

Similarly, if the *uninformed non-reflective* trader assumes the true asset value is 490, then no trader would be willing to sell in the region $X_{b_1} = [50{,}223.5]$ as, again, each trader would hold a belief greater than the upper bound of this region. Thus,

$$P[\{e(2)\}|v = 490] = P[\{T(b_1)\}|v = 490] = \rho_{sell}^{X_{b_1}} = 0.$$

The *uninformed non-reflective* trader therefore updates her belief according to the formula

$$\mu_{2,U}^{R} = \frac{P[\{e(2)\}|v = 50]Prior^{50}(1)}{P[\{e(2)\}]}\,50 + \frac{P[\{e(2)\}|v = 240]Prior^{240}(1)}{P[\{e(2)\}]}\,240 + \frac{P[\{e(2)\}|v = 490]Prior^{490}(1)}{P[\{e(2)\}]}\,490$$

$$\mu_2^{R,U} = \frac{1/2 \times 0.23}{1/2 \times 0.23} 50 + \frac{0 \times 0.93}{1/2 \times 0.23} 240 + \frac{0 \times 0.17}{1/2 \times 0.23} 490 = 50.$$

Thus, the uninformed non-reflective trader learns that the true asset value is 50. The new current best bid would be set to 50 (by default), while the current best ask remains 360. This is the case because, as in the experiment, we do not clear the order book after a transaction. The simulation would continue with the random selection of a third trader.

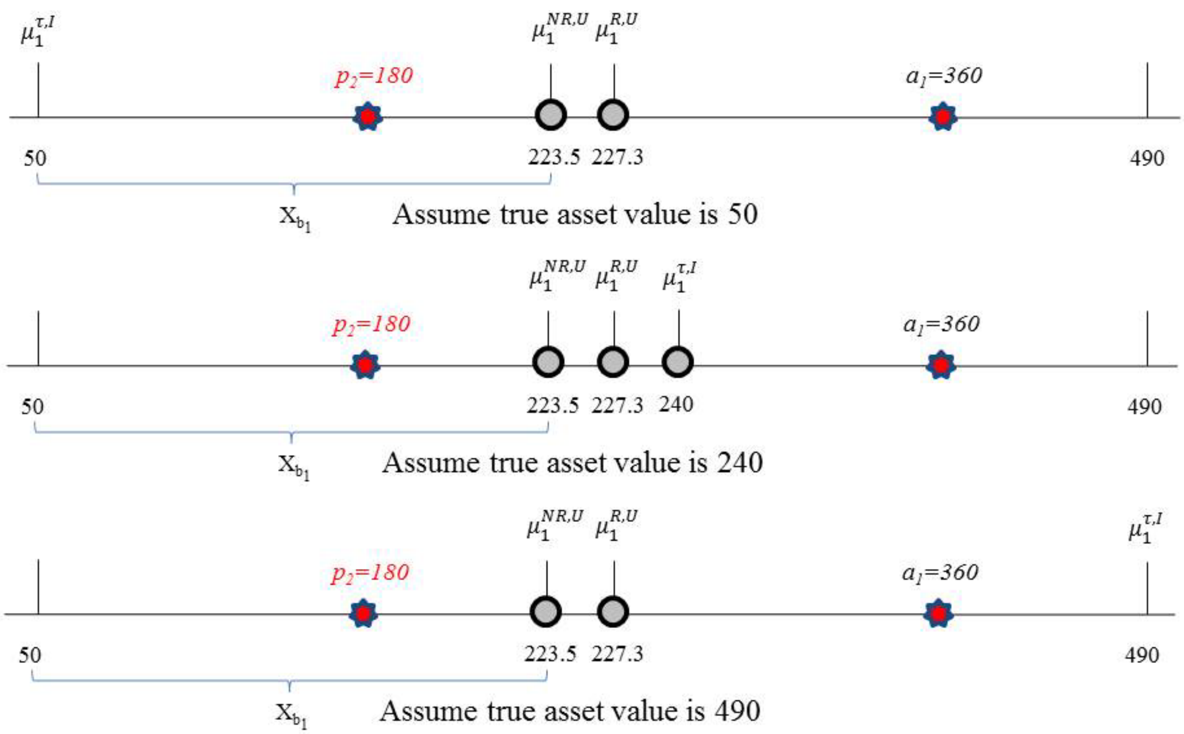

Figure B2. Representation of traders' updated beliefs after the first event, the trade at a price of 180 (the second event), and the region in which the accepted bid exists. Note that *reflective informed* and *non-reflective informed* traders hold the same belief, while *reflective uninformed* and *non-reflective uninformed* traders hold different beliefs. While the region X_{b_1} will typically depend upon the assumed true asset value, in this example the region is identical for each assumed value.

B2. Earnings comparison between *reflective* and *non-reflective* traders

As described in Section 3.2 we ran simulations of the model to develop testable conjectures. In this section we show that the earnings of *reflective* traders exceed those of *non-reflective* traders in markets populated by 50% *fully* informed traders as well as in markets populated by 100% *partially* informed traders.

For each combination (υ, α) with $\alpha \in \{1/12, 2/12, \ldots, 1\}$ we execute 25,000 simulations. Each simulation is conducted as follows: a randomly selected trader either accepts the current best bid (ask) or strictly improves the bid-ask spread.[36] Note that any of this trader's outstanding bids or asks that do not agree with her current belief (i.e., bids above or asks below her belief) are canceled before she acts. This is consistent with our experimental setup in which traders can cancel their own orders. Traders then update their beliefs based upon the event, and the program moves to the next iteration and restarts the process by randomly selecting a trader. The simulation finishes once 30 transactions have been made or 10,000 iterations have been completed.

The following table provides support for the model's prediction that *reflective* traders will earn more than *non-reflective* traders. We utilize a stopping criterion of 30 transactions though the results are robust to alternate criteria (e.g., 5 transactions).

Table B2. Ending wealth for *reflective* and *non-reflective* traders in both treatments.

	50% fully informed		*100% partially informed*	
α	*Reflective*	*Non-reflective*	*Reflective*	*Non-reflective*
1/12	2,555.69	2,211.30	2,676.36	2,200.33
2/12	2,479.25	2,192.15	2,627.51	2,162.50
3/12	2,414.81	2,181.73	2,582.44	2,125.85
4/12	2,361.38	2,179.31	2,538.15	2,090.93
5/12	2,320.05	2,182.82	2,489.64	2,061.68
6/12	2,290.45	2,189.55	2,437.04	2,042.96
7/12	2,269.84	2,198.23	2,376.95	2,048.27
8/12	2,256.99	2,206.02	2,314.99	2,090.02
9/12	2,249.09	2,212.72	2,273.92	2,138.23
10/12	2,244.53	2,217.36	2,254.20	2,168.98
11/12	2,241.73	2,220.93	2,244.79	2,187.26
1	2,240.00	na	2,240.00	na

Appendix C. Robustness checks simulations

C1. Stress tests of the model

To test the robustness of Conjecture 1, we ran additional simulations of our model varying (1) the number of traders, (2) the traders' initial endowments, and (3) the asset values and probabilities.

[36] Any action by the trader (either acceptance of the current best bid/ask or submission of a new bid and ask) is conditional upon the trader's finances. That is, a trader may only submit a new bid (accept the best ask) if she has enough cash to cover the new bid (accepted ask) plus all of her outstanding bids. Similarly, the trader may only submit a new ask (accept the best bid) if she has enough shares to cover the new ask (accepted bid) plus all of her outstanding asks.

Analogous to Figure C1, the mean absolute price deviation from the true asset value averaged across 25,000 simulations for each asset value is reported in Table C1. Consistent with Conjecture 1b, we report that markets in which private information is concentrated (50% *fully* informed) leads to lower MAD values than markets in which information is dispersed (100% *partially* informed).

Table C1.- This table reports the mean absolute price deviation from the true asset value. The reported value corresponds to the average across 25,000 simulations for each asset value (50, 240 and 490) and value of α.

	Market with twice the number of traders [1]		Market with three times more cash and shares [2]		Market with uniform distribution of market value [3]	
Asset Values	*50-240-490*		*50-240-490*		***100-200-300***	
Probabilities	*0.35-0.45-0.20*		*0.35-0.45-0.20*		***1/3-1/3-1/3***	
Endowment	*1,200 francs, 4 shares*		***3,600 francs, 12 shares***		*1,200 francs, 4 shares*	
Number of traders	***24 traders***		*12 traders*		*12 traders*	
	MAD					
α	100% *partially* informed	50% *fully* informed	100% *partially* informed	50% *fully* informed	100% *partially* informed	50% *fully* informed
0	148.81	94.26	137.25	88.67	62.84	35.67
1/24	147.24	85.37				
2/24 (1/12)	145.27	76.32	131.59	70.96	58.75	28.47
3/24	142.86	67.50				
4/24 (2/12)	139.85	59.56	125.27	55.26	53.92	22.32
5/24	136.21	51.77				
6/24 (3/12)	132.06	44.31	116.55	42.19	48.80	17.08
7/24	127.21	38.12				
8/24 (4/12)	121.79	32.41	103.91	31.54	43.49	12.94
9/24	115.75	27.29				
10/24 (5/12)	109.30	23.10	86.91	23.14	36.61	9.58
11/24	102.43	19.31				
12/24 (6/12)	95.30	16.23	67.43	16.96	28.40	6.98
13/24	87.92	13.73				
14/24 (7/12)	79.85	11.78	48.44	12.47	20.71	5.27
15/24	72.22	10.10				
16/24 (8/12)	64.47	8.88	33.07	9.35	15.42	3.93
17/24	57.40	7.85				
18/24 (9/12)	51.90	7.09	22.62	7.21	10.48	3.10
19/24	46.11	6.49				
20/24 (10/12)	41.54	5.89	16.33	5.94	7.95	2.53
21/24	37.40	5.42				
22/24 (11/12)	33.35	5.03	12.35	5.04	5.07	2.12
23/24	28.98	4.73				
1	22.18	4.44	8.99	4.40	6.33	1.95

C2. Allocative efficiency

We use simulations to assess the effect of concentrating information in the hands of insiders in lieu of distributing evenly across traders on the allocative efficiency of markets (see Conjecture 1b). To that end, we simulate markets in which the value of the asset differs across traders (see Plott and Sunder 1988, Series A). In particular, we consider three possible states of the world denoted X, Y and Z. In state X (Y) [Z], the asset value is equal to 100 (240) [300] francs for one-half of the traders and 290 (190) [160] for the other half (see Corgnet et al. 2018b). In these private values simulations, we used the same trading rules and number of traders as in the main treatments with the exception that traders were endowed with 1,500 francs and 3 shares of the asset.

To measure the allocative efficiency of a market, we follow Plott and Sunder (1988) and Corgnet et al. (2018b) by comparing the sum of asset payouts received by all traders in a given market (*Actual Payouts*) with the payouts which would have been made if traders knew the state of the world (*Max Payouts*). More specifically, we assess the extent to which a market allocation of shares improves upon the no-trade allocation. We then calculate our efficiency measure as: *Allocative Efficiency* $\coloneqq \frac{Actual\ Payouts - No\text{-}trade\ Payouts}{Max\ Payouts - No\text{-}trade\ Payouts}$, where the *No-trade Payouts* are those obtained by traders when they keep their initial portfolio (without making a trade). Below, we represent the average allocative efficiency (across all possible proportions of *reflective* traders in the market) for markets with 100% *partially* informed traders and for markets with 50% *fully* informed traders. We show that allocative efficiency is higher when the market is populated by *fully* informed traders rather when the market is only populated by *partially* informed traders.

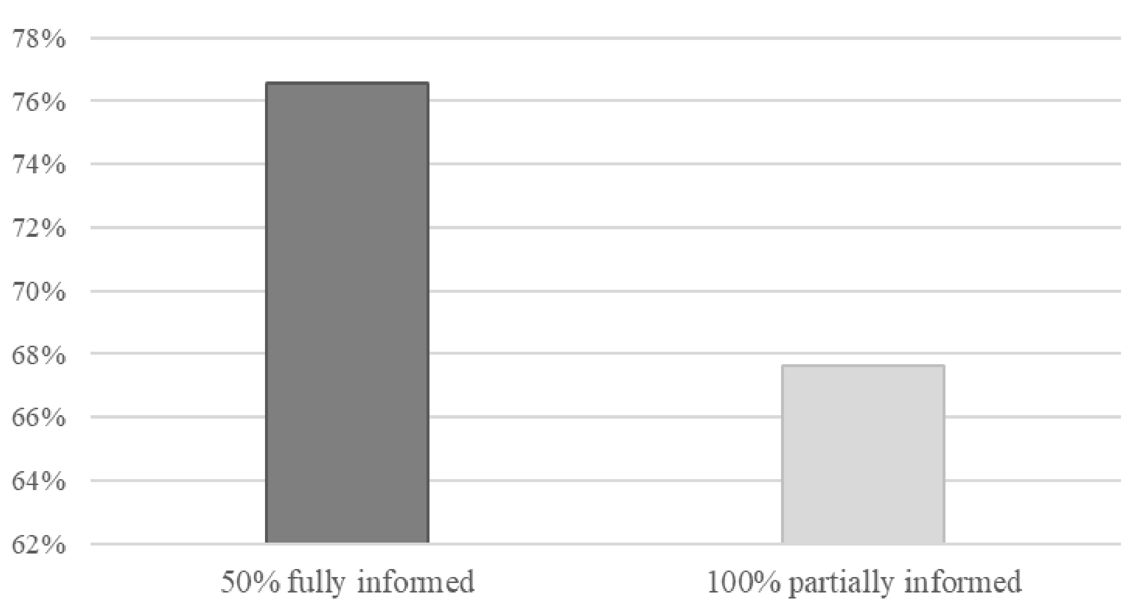

Figure C1. Allocative efficiency simulations for markets with 100% *partially* informed traders and for markets with 50% *fully* informed.

Appendix D. Survey

Cognitive reflection test (CRT) (5 minutes)

Taken from Frederick (2005):

(1) A bat and a ball cost $1.10 in total. The bat costs a dollar more than the ball. How much does the ball cost? ____ cents
[Correct answer: 5 cents; intuitive answer: 10 cents]

(2) If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets? ____ minutes
[Correct answer: 5 minutes; intuitive answer: 100 minutes]

(3) In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake? ____ days
[Correct answer: 47 days; intuitive answer: 24 days]

Taken from Toplak et al. (2014):

(4) If John can drink one barrel of water in 6 days, and Mary can drink one barrel of water in 12 days, how long would it take them to drink one barrel of water together? _____ days
[correct answer: 4 days; intuitive answer: 9]

(5) Jerry received both the 15th highest and the 15th lowest mark in the class. How many students are in the class? ______ students
[correct answer: 29 students; intuitive answer: 30]

(6) A man buys a pig for $60, sells it for $70, buys it back for $80, and sells it finally for $90. How much has he made? _____ dollars
[correct answer: $20; intuitive answer: $10]

(7) Simon decided to invest $8,000 in the stock market one day early in 2008. Six months after he invested, on July 17, the stocks he had purchased were down 50%. Fortunately for Simon, from July 17 to October 17, the stocks he had purchased went up 75%. At this point, Simon has: a. broken even in the stock market, b. is ahead of where he began, c. has lost money
[correct answer: c; intuitive response: b]

Theory of the mind Test (10 minutes)

This is an example of the 36 questions in the test of Baron-Cohen et al. (1997):

Figure D1. Example of an eye gaze test question

Appendix E. Market figures

The average price per market period is listed at the top of each subfigure, and the true asset value is denoted at the bottom of each subfigure. Each transaction is denoted by a red dot. The rational expectations value is indicated by a horizontal line, and the Walrasian model value is indicated by a dashed line.

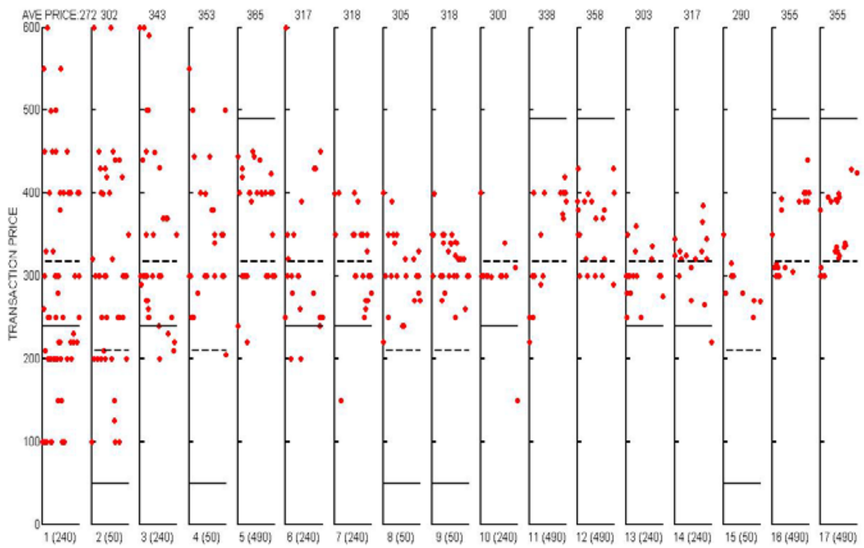

Figure 1.1. *'100% partially informed'* Session 1.

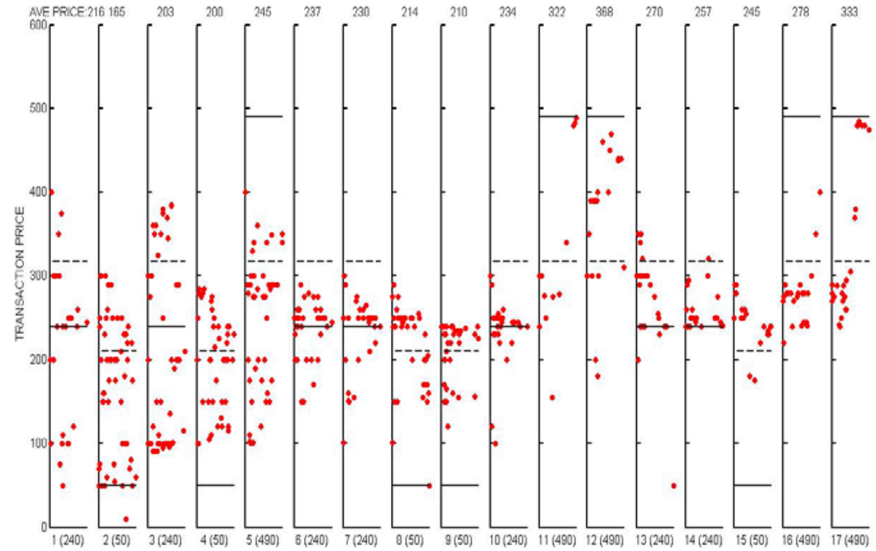

Figure 1.3. *'100% partially informed'* Session 3.

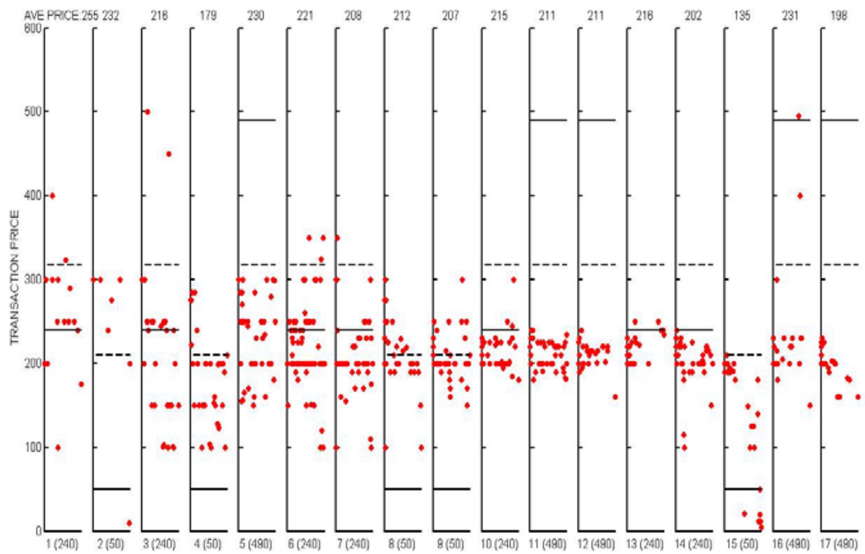

Figure 1.2. *'100% partially informed'* Session 2.

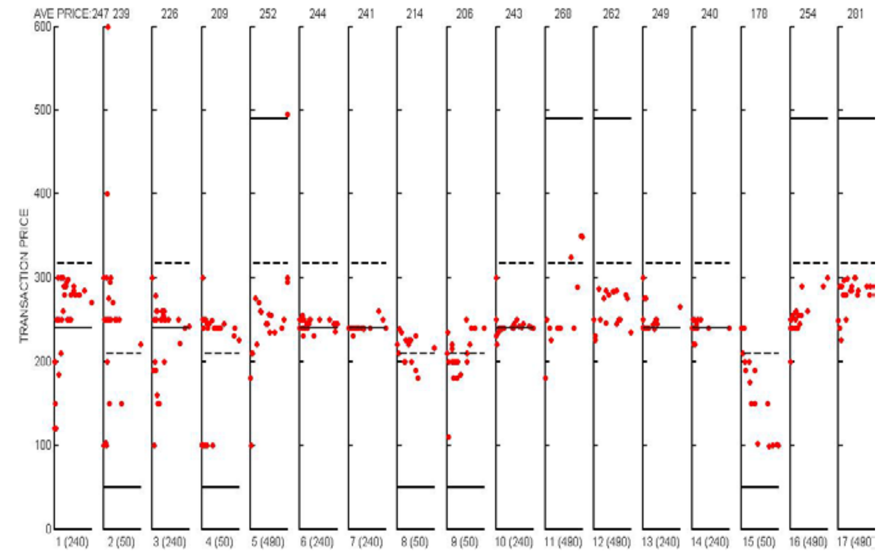

Figure 1.4. *'100% partially informed'* Session 4.

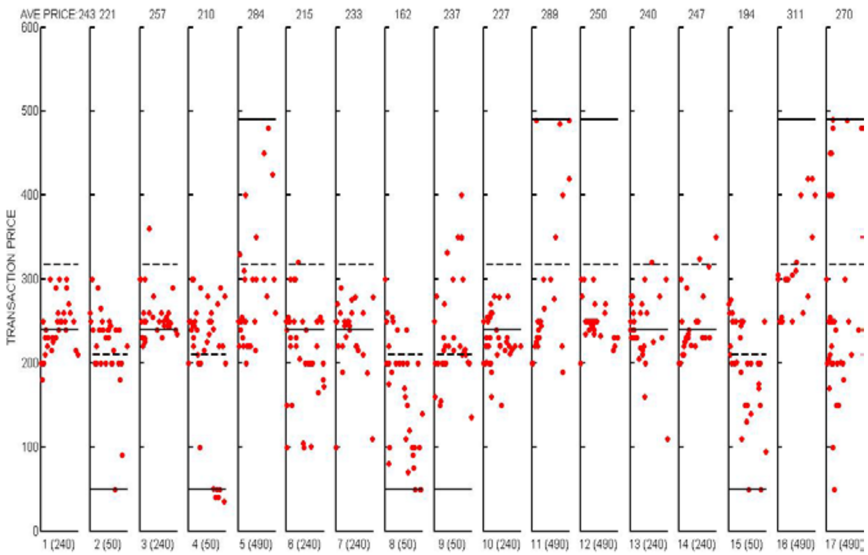

Figure 1.5. *'100% partially informed'* Session 5.

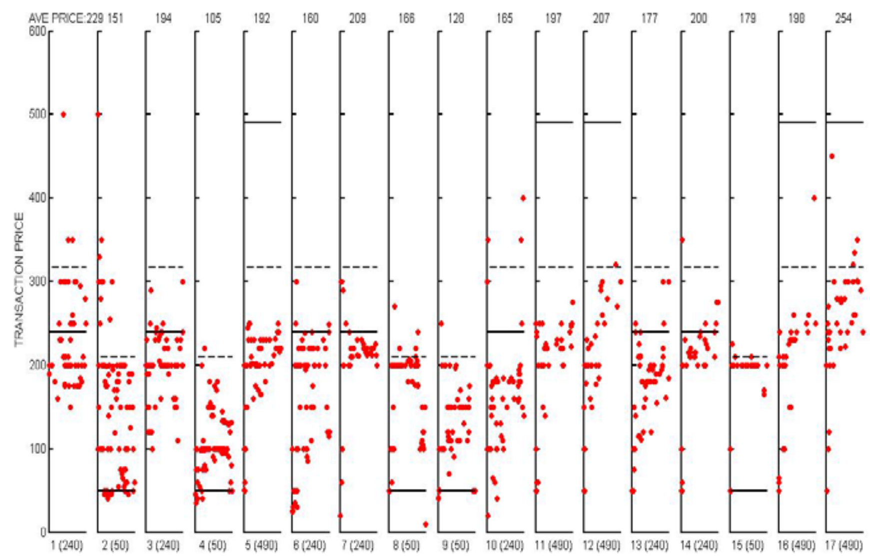

Figure 1.7. *'100% partially informed'* Session 7.

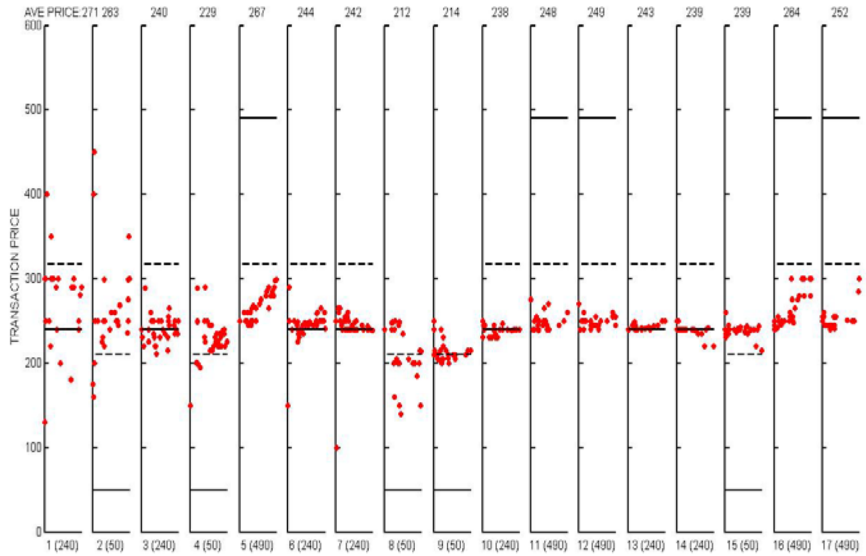

Figure 1.6. ‘*100% partially informed*’ Session 6.

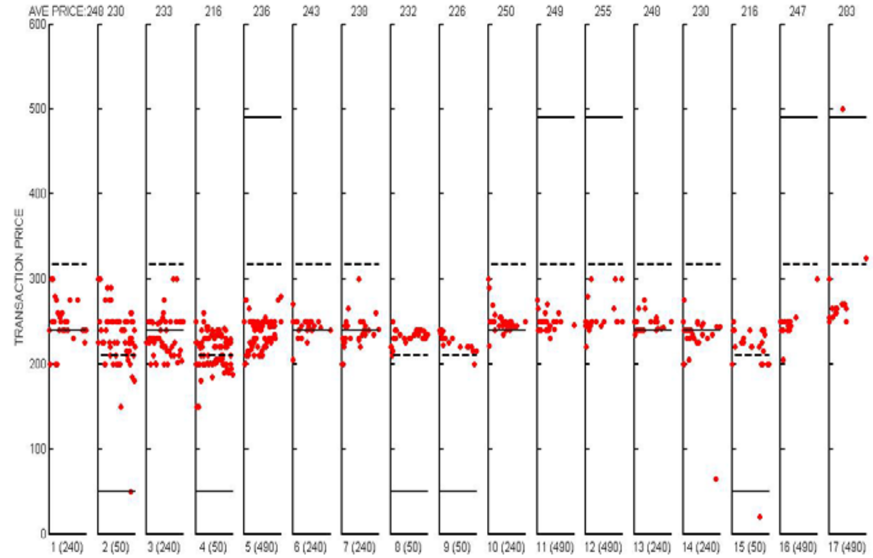

Figure 1.8. ‘*100% partially informed*’ Session 8.

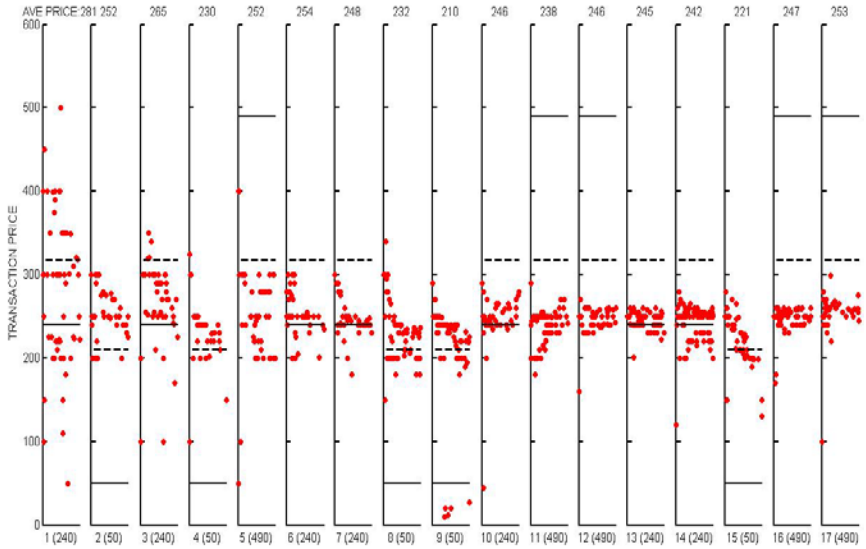

Figure 1.9. *'100% partially informed'* Session 9.

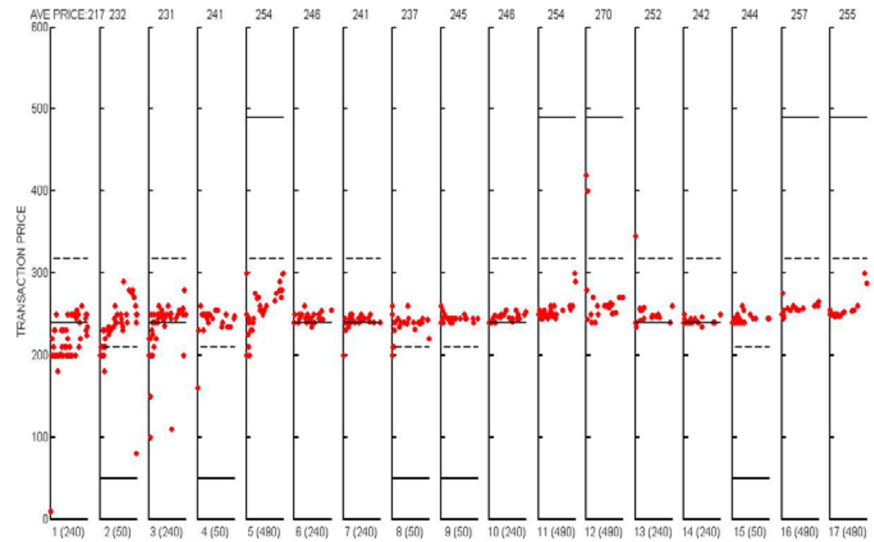

Figure 1.10. *'100% partially informed'* Session 10.

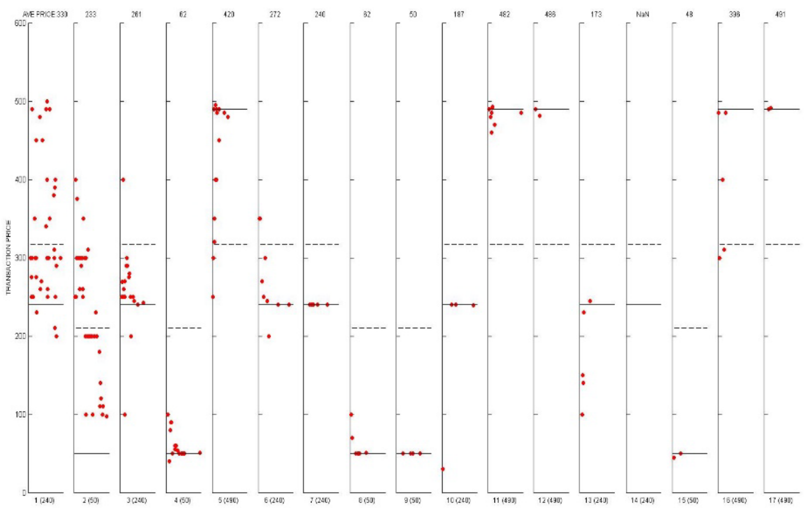

Figure 1.11. *'50% fully informed'* Session 1.

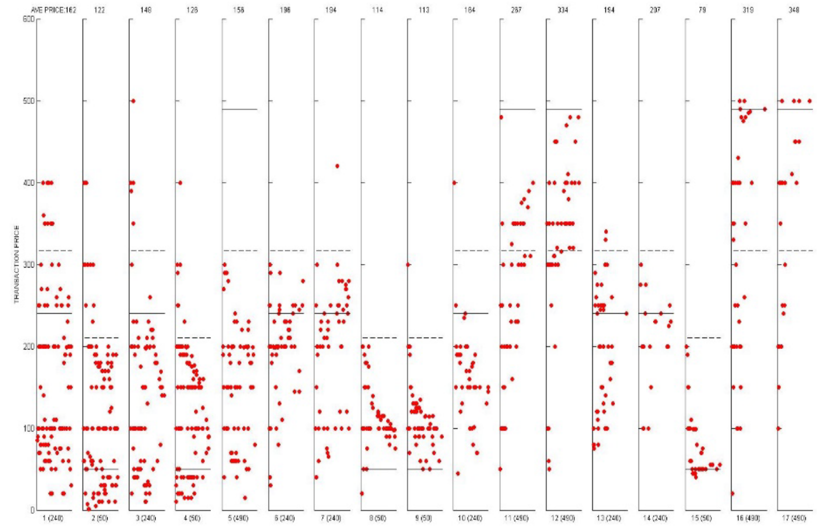

Figure 1.12. *'50% fully informed'* Session 2.

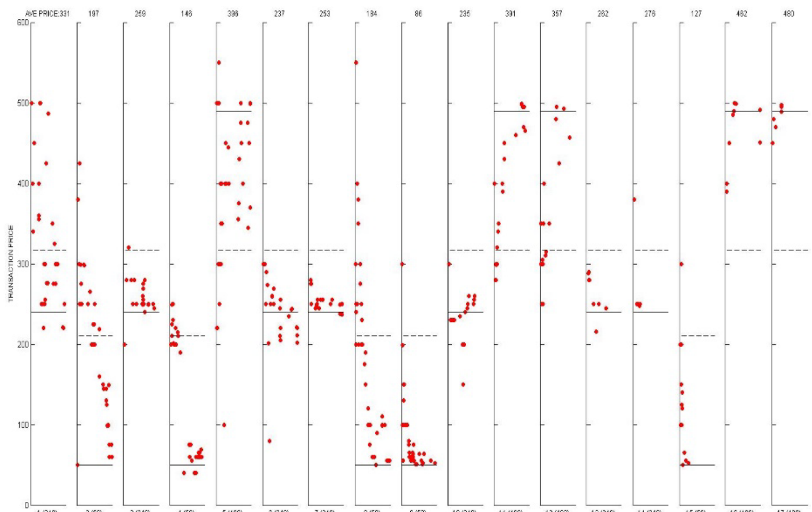

Figure 1.13. *'50% fully informed'* Session 3.

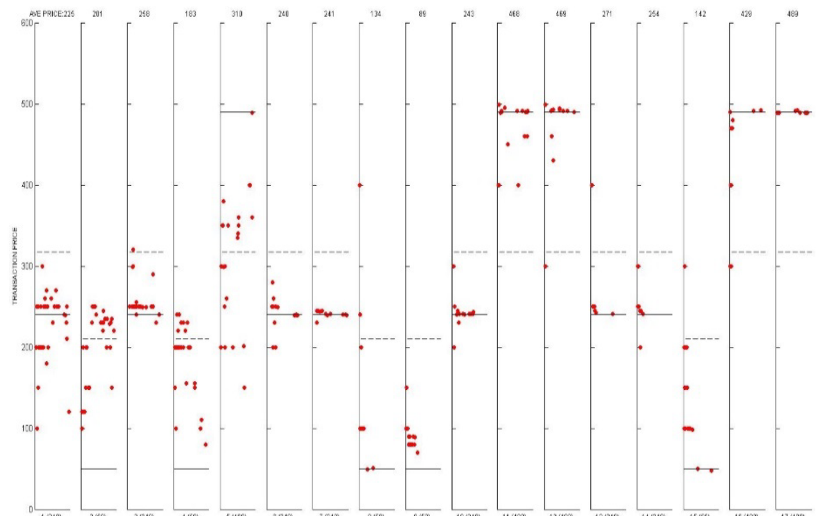

Figure 1.14. *'50% fully informed'* Session 4.

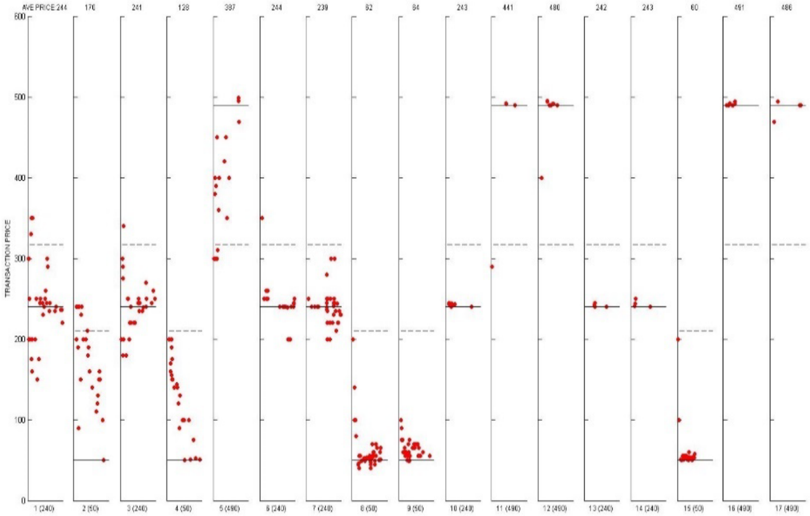

Figure 1.15. *'50% fully informed'* Session 5.

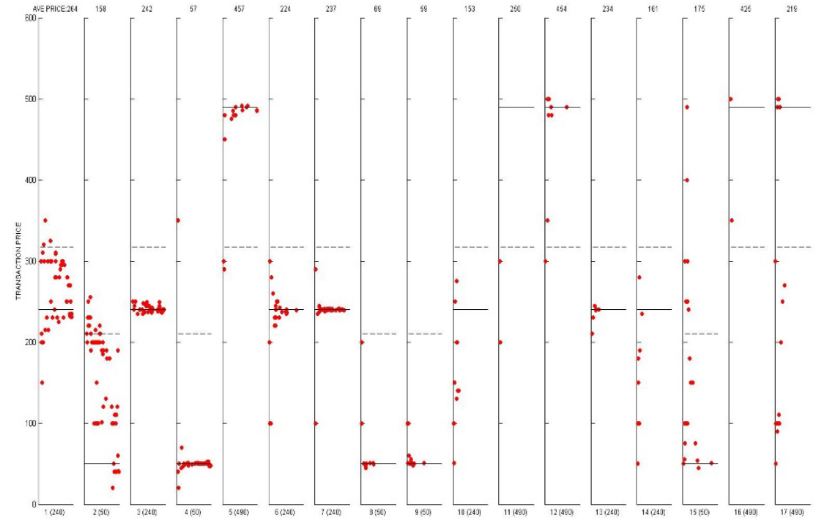

Figure 1.16. *'50% fully informed'* Session 6.

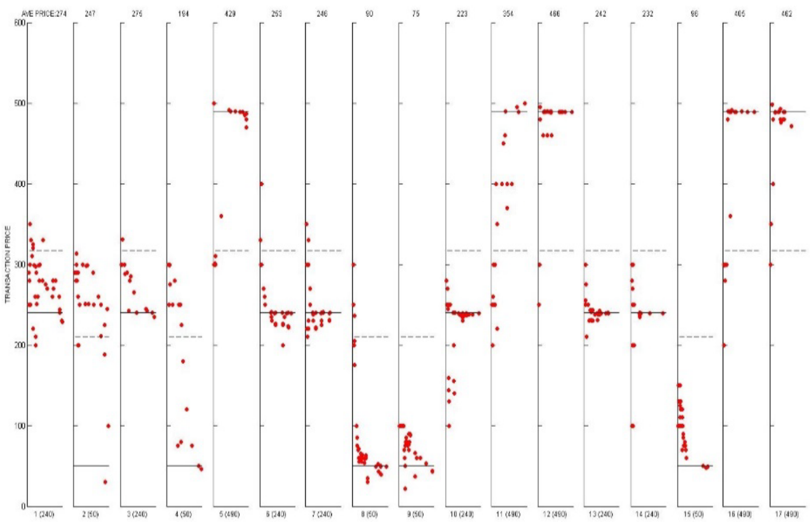

Figure 1.17. *'50% fully informed'* Session 7.

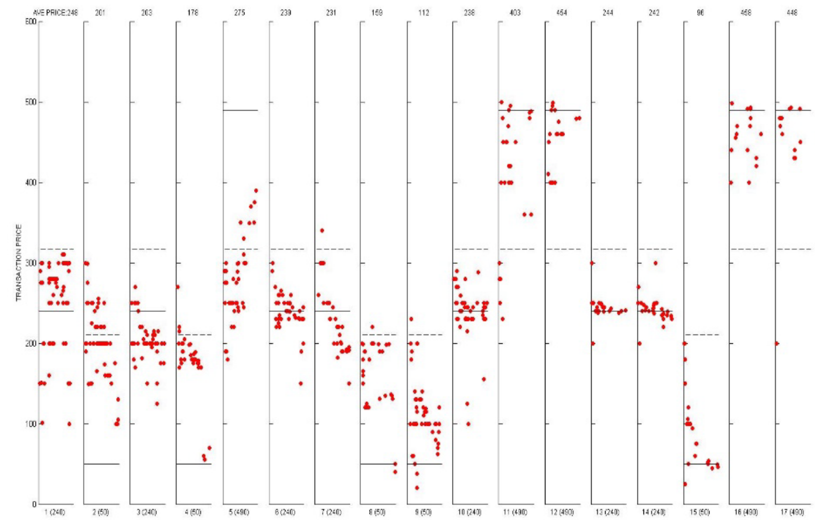

Figure 1.18. *'50% fully informed'* Session 8.

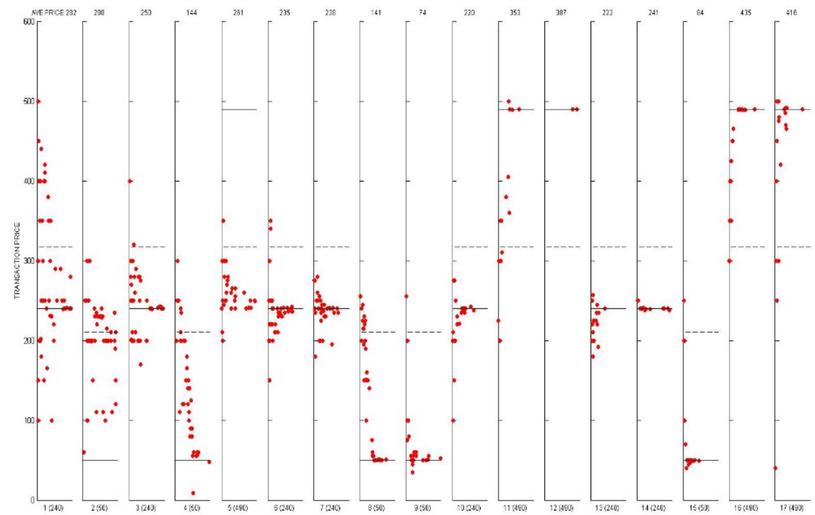

Figure 1.19. *'50% fully informed'* Session 9.

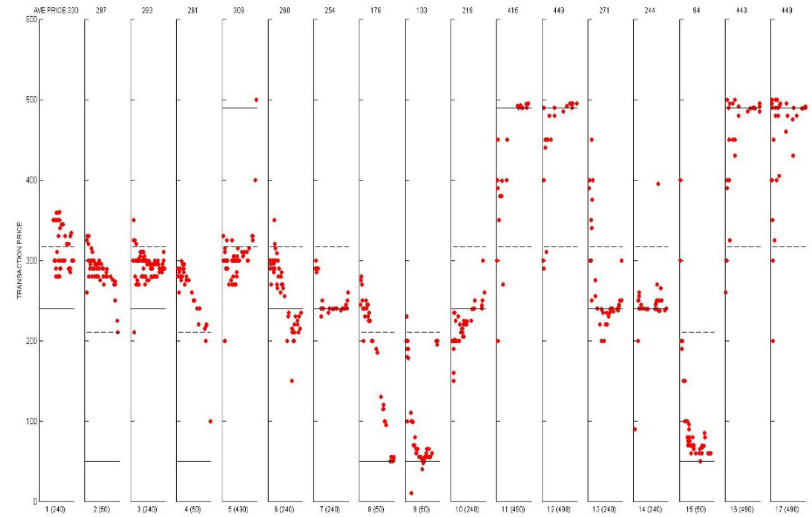

Figure 1.20. *'50% fully informed'* Session 10.

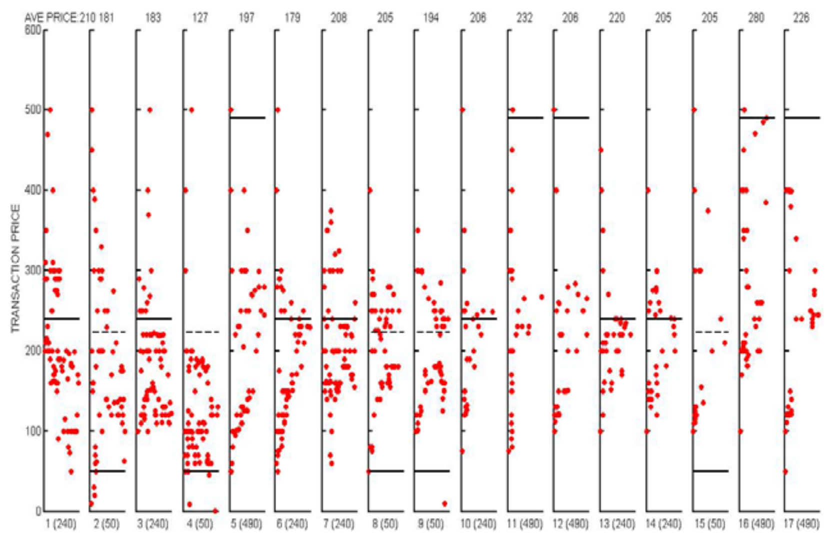

Figure 1.21. *'17% fully informed'* Session 1.

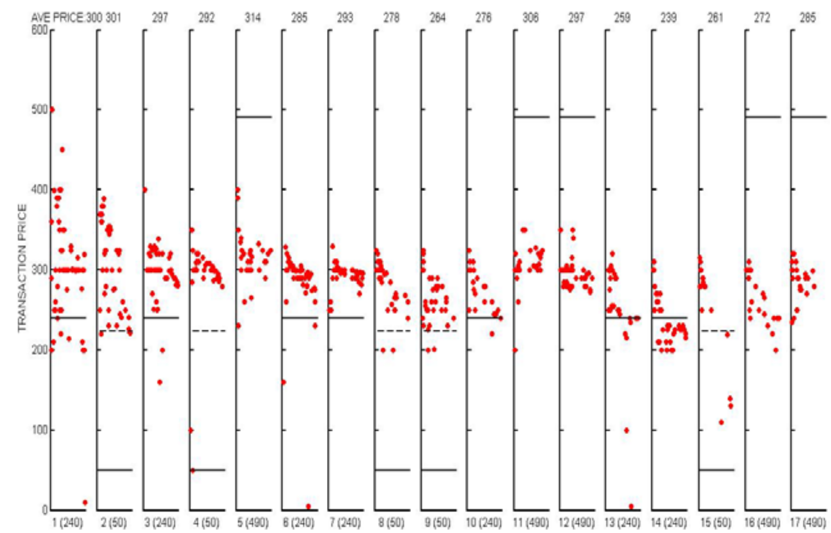

Figure 1.22. *'17% fully informed'* Session 2.

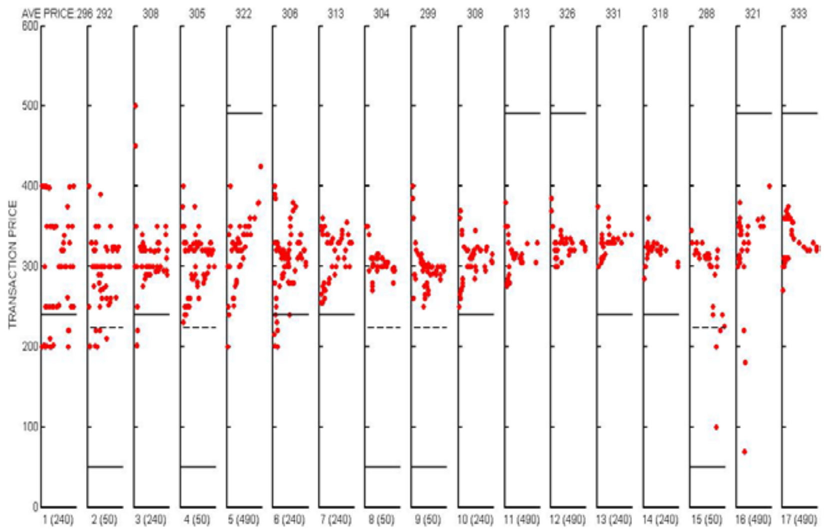

Figure 1.23. *'17% fully informed'* Session 3.

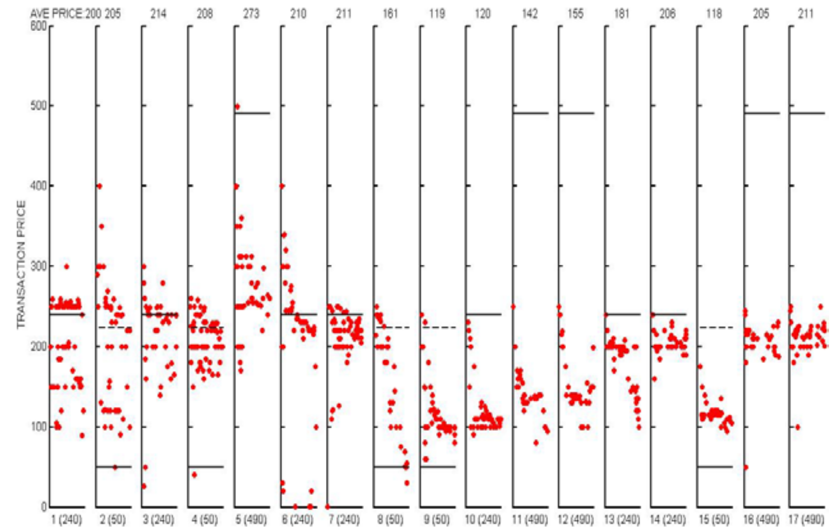

Figure 1.24. *'17% fully informed'* Session 4

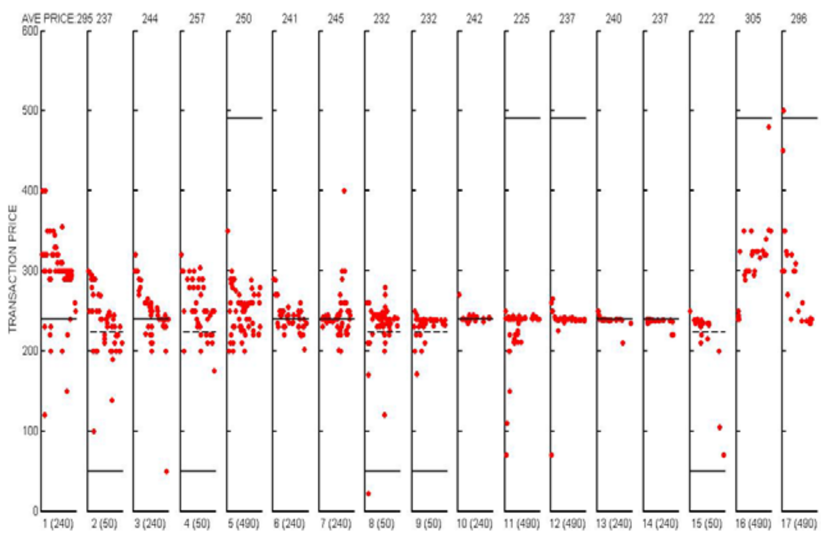

Figure 1.25. *'17% fully informed'* Session 5.

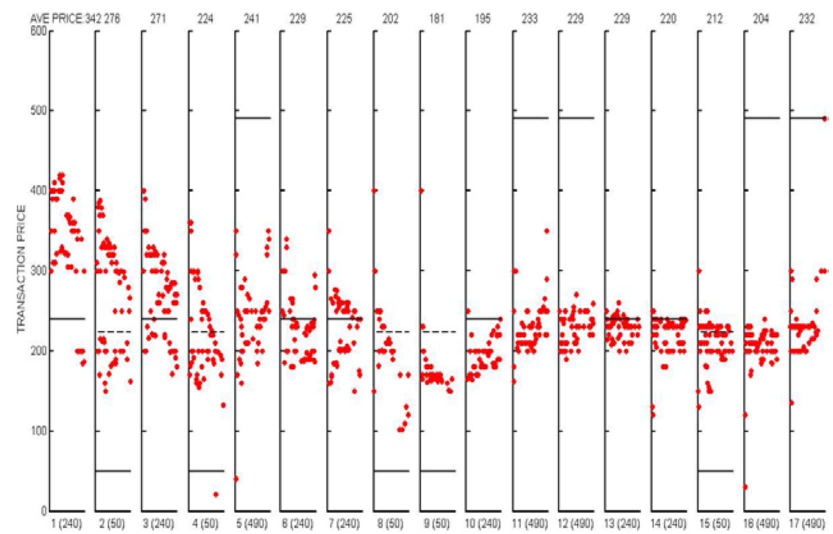

Figure 1.26. *'25% fully informed'* Session 1.

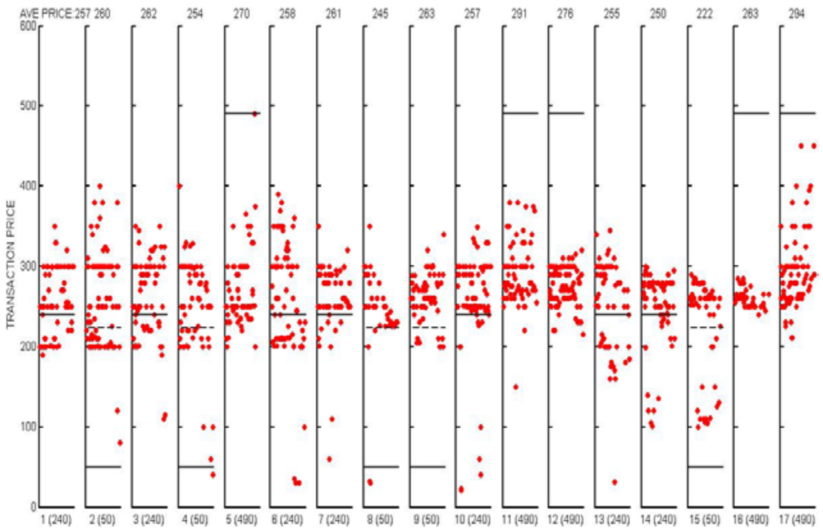

Figure 1.27. *'25% fully informed'* Session 2.

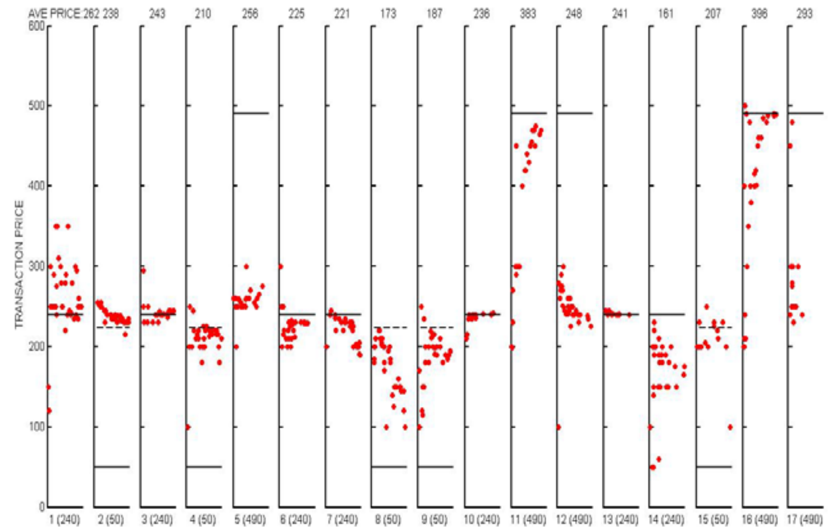

Figure 1.28. *'25% fully informed'* Session 3.

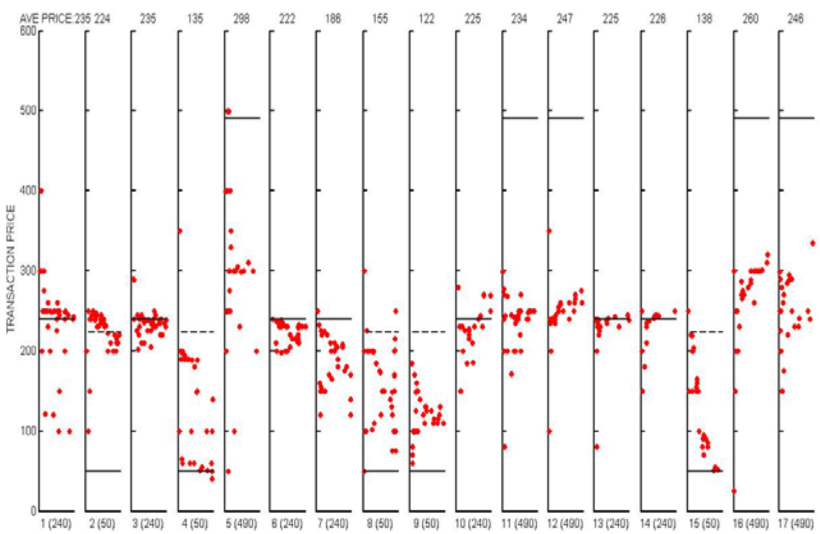

Figure 1.29. *'25% fully informed'* Session 4.

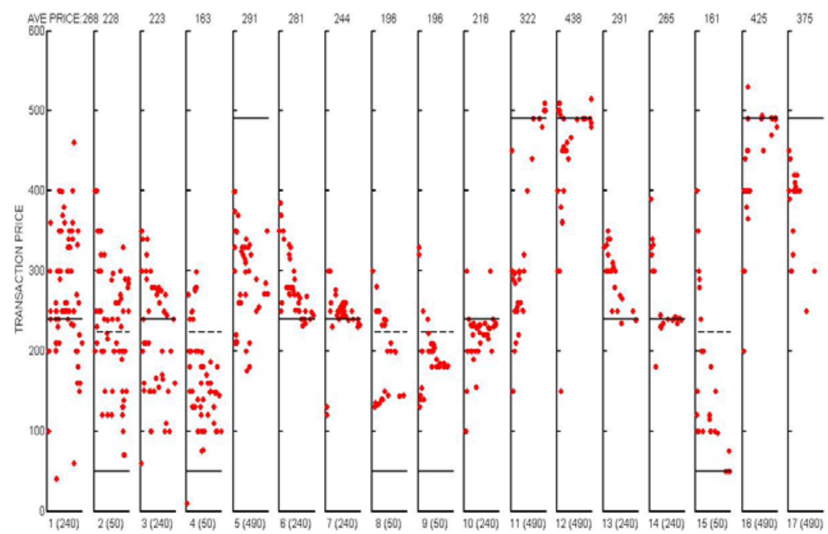

Figure 1.30. *'25% fully informed'* Session 5.

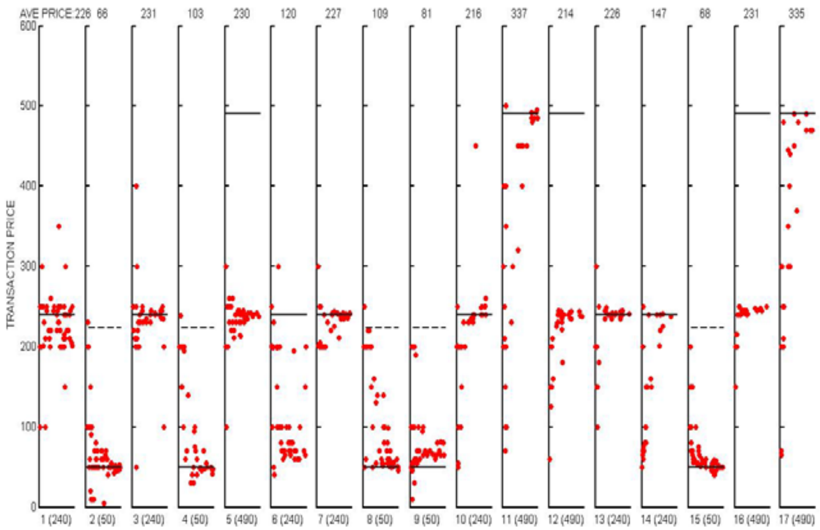

Figure 1.31. *'33% fully informed'* Session 1.

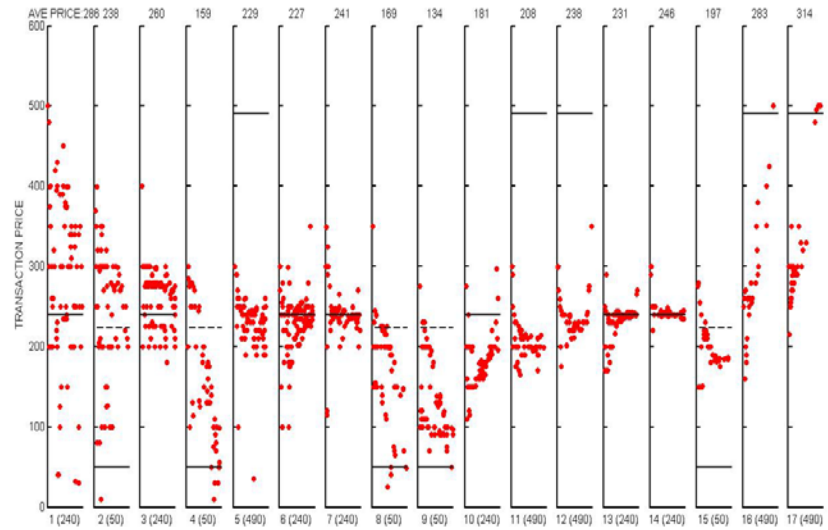

Figure 1.32. *'33% fully informed'* Session 2.

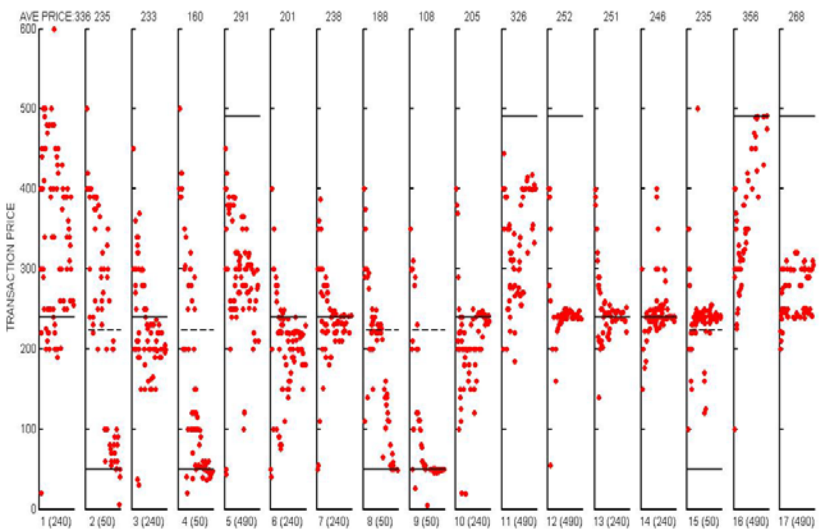

Figure 1.33. *'33% fully informed'* Session 4.

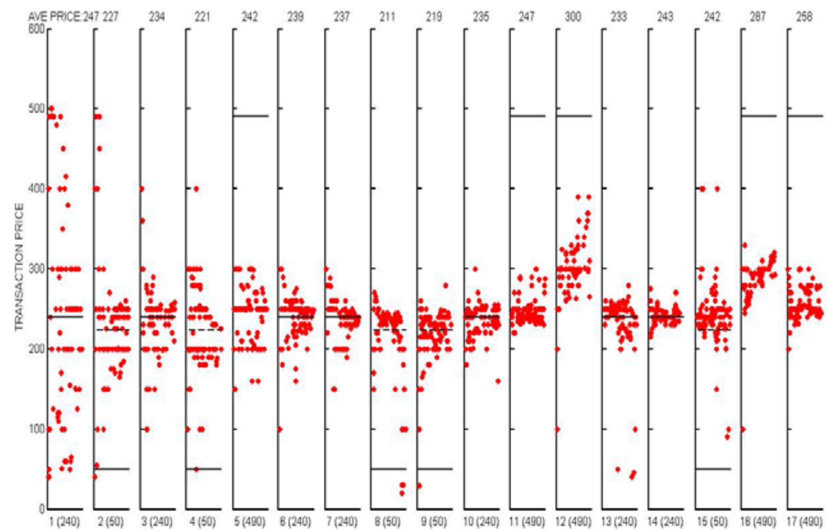

Figure 1.34. *'33% fully informed'* Session 4.

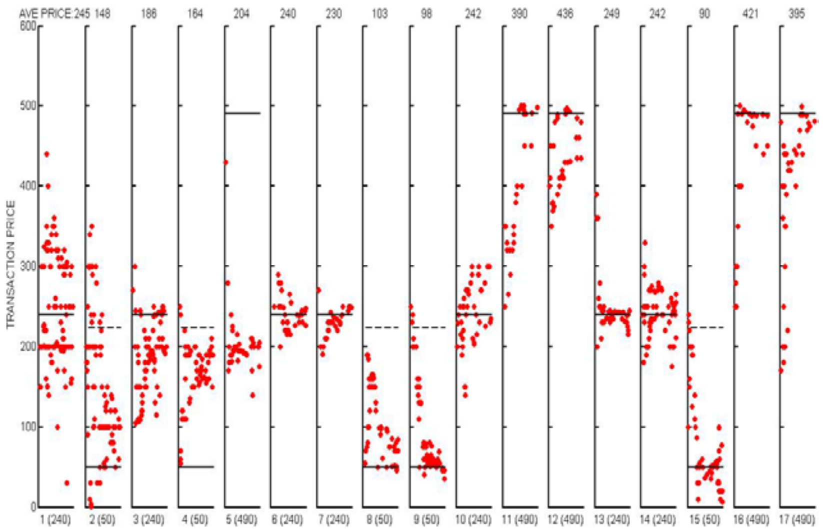

Figure 1.35. *'33% fully informed'* Session 5.

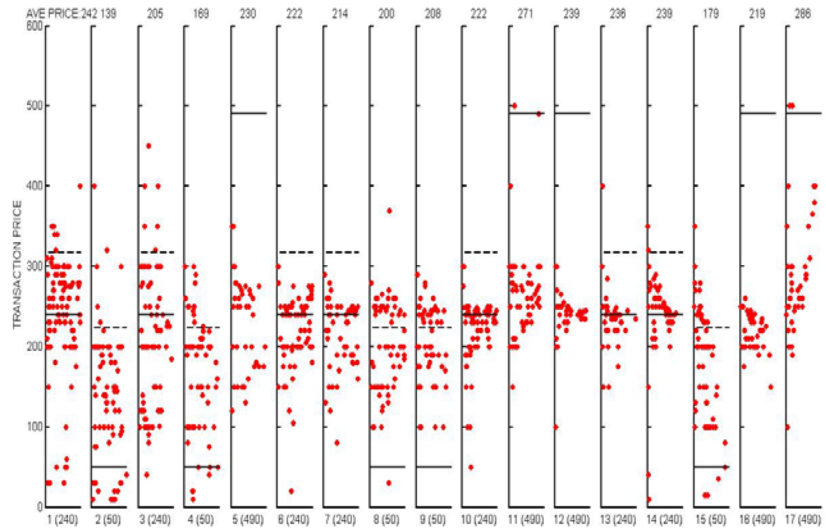

Figure 1.36. *'33% fully informed / 25% partially informed'* Session 1.

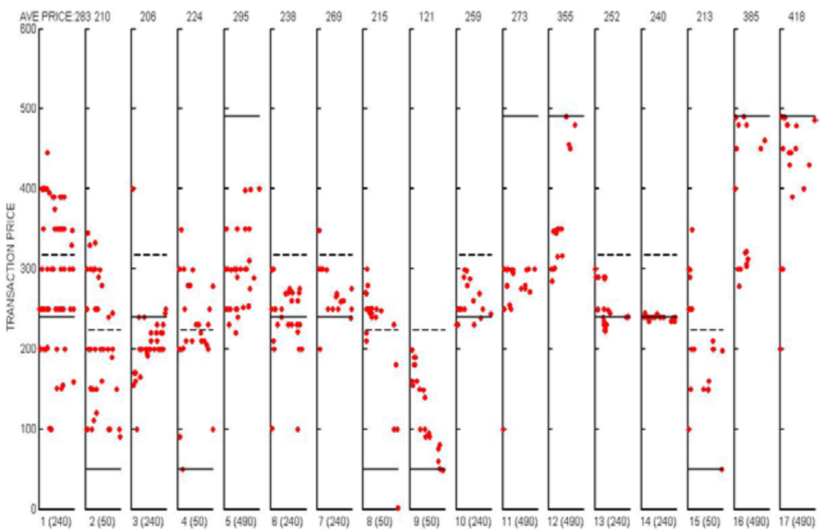

Figure 1.37. *'33% fully informed / 25% partially informed'* Session 2.

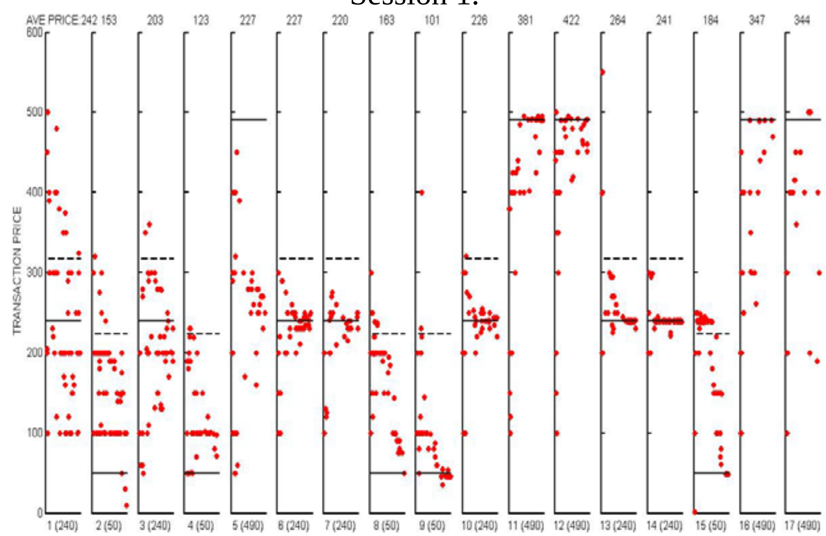

Figure 1.38. *'33% fully informed / 25% partially informed'* Session 3.

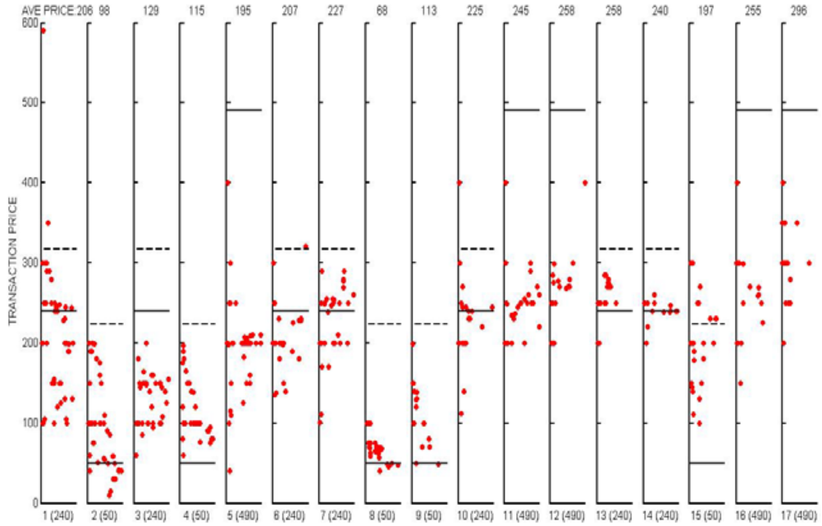

Figure 1.39. *'33% fully informed / 25% partially informed'* Session 4.

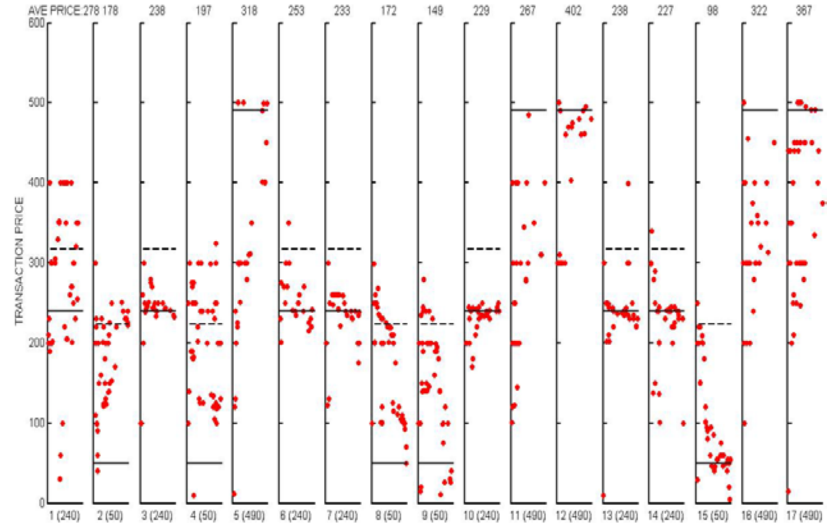

Figure 1.40. *'33% fully informed / 25% partially informed'* Session 5.

Appendix O1. Instructions (online appendix)

We display the screenshots for the *'50% fully informed'* treatment.

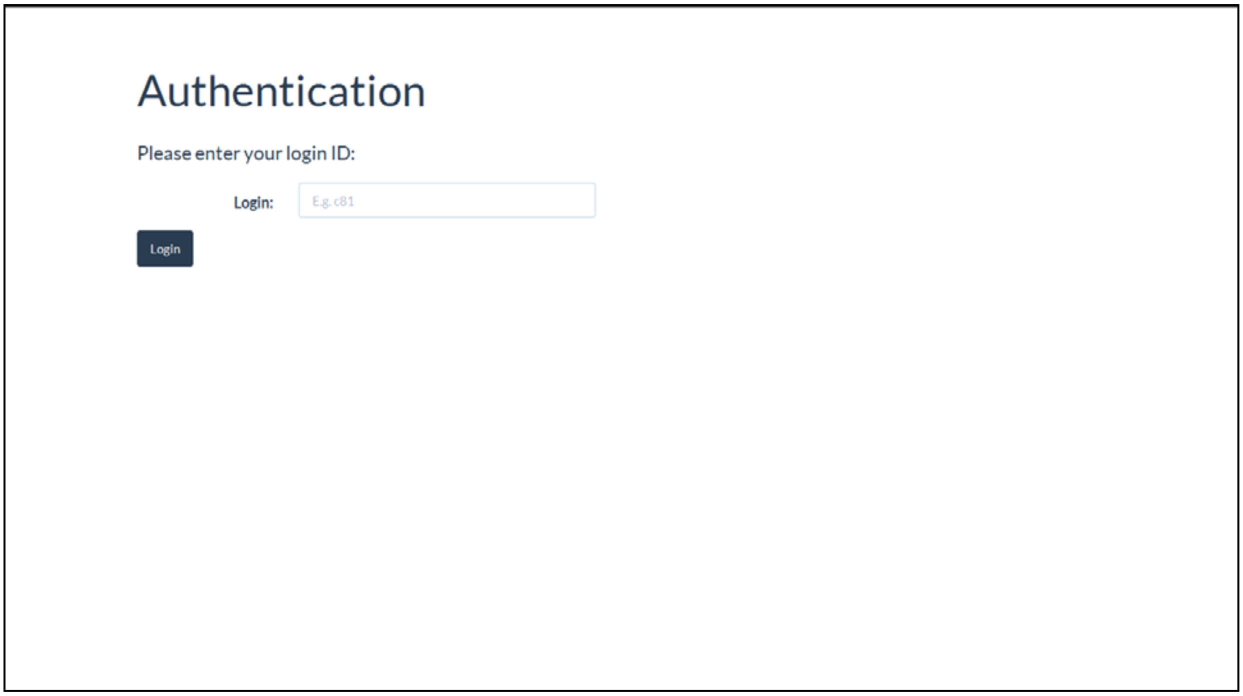

Instructions Part 1

This is an experiment in the economics of market decision making. The instructions are simple, and if you follow them carefully and make good decisions, you might earn a considerable amount of money which will be paid to you in cash.

In this experiment we are going to simulate a market in which you will buy and sell certificates in a sequence of 17 market years.

At the end of each year, each certificate you hold will be worth either 50, 240, or 490. The specific amount will be determined by the spin of a wheel. The features of this spinning wheel are described below.

NEXT PAGE

Instructions Part 1

Consider a wheel subdivided into 100 equal-sized pieces labeled with the numbers 1 through 100. Each year we spin this wheel, and the spinner (arrow) is equally likely to land on any of the numbered pieces. If the spinner lands on a number between 1 and 35 (including 1 and 35), X-dividend of 50 is paid; if the number landed on is between 36 and 80 (inclusive), Y-dividend of 240 is paid; and if the number landed on is between 81 and 100 (inclusive), Z-dividend of 490 is paid.

On the following screen you will be asked to predict (one at a time) the outcome of ten spins of the wheel. That is, you will predict whether the X-, Y-, or Z-dividend is paid. If your individual prediction is correct, you win \$0.25; if wrong, you lose \$0.10. Once you make your prediction, click the "Spin the Wheel!" button. After the spin has occurred you will be prompted to make your next prediction and the software will update your earnings.

Note to simplify the appearance of the wheel, the wheel is subdivided into 3 wedges representing the range of values for which each dividend (X, Y, or Z) is paid.

PREVIOUS PAGE

NEXT PAGE

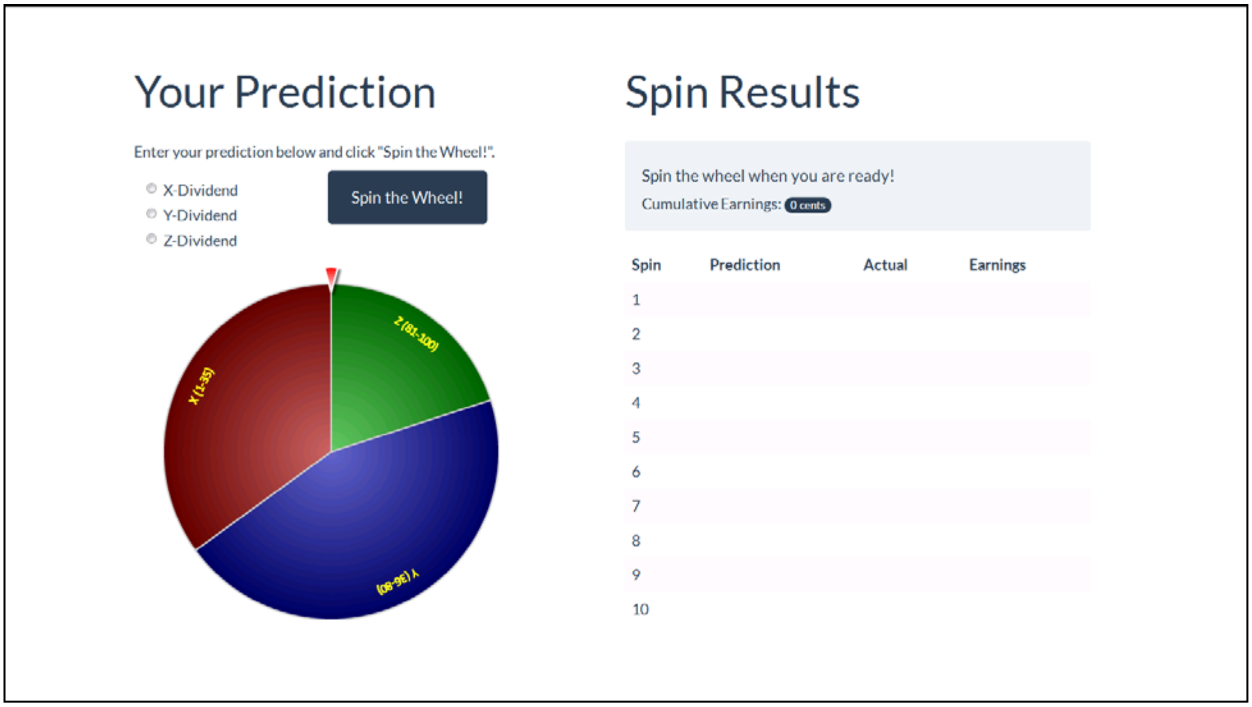
Your Prediction
Enter your prediction below and click "Spin the Wheel!".
X-Dividend
Y-Dividend
Z-Dividend
Spin the Wheel!
Z (81-100)
X (1-35)
Y (36-80)
Spin Results
Spin the wheel when you are ready!
Cumulative Earnings: 0 cents
Spin
Prediction
Actual
Earnings
1
2
3
4
5
6
7
8
9
10

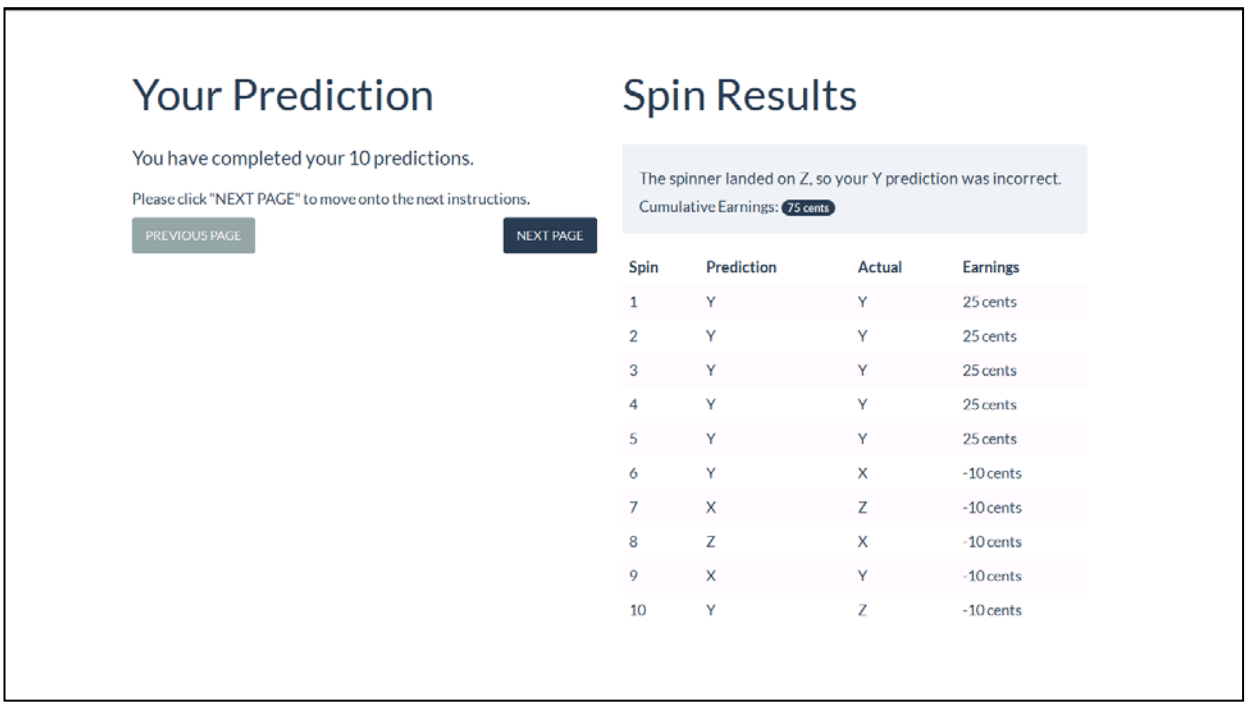
Your Prediction
You have completed your 10 predictions.
Please click "NEXT PAGE" to move onto the next instructions.
PREVIOUS PAGE
NEXT PAGE
Spin Results
The spinner landed on Z, so your Y prediction was incorrect.
Cumulative Earnings: 75 cents
Spin Prediction Actual Earnings
1 Y Y 25 cents
2 Y Y 25 cents
3 Y Y 25 cents
4 Y Y 25 cents
5 Y Y 25 cents
6 Y X -10 cents
7 X Z -10 cents
8 Z X -10 cents
9 X Y -10 cents
10 Y Z -10 cents

Instructions Part 2

The type of currency used in this market is francs. All trading and earnings will be in terms of francs. Each franc is worth $0.001 to you. At the end of the experiment your francs will be converted to dollars at this rate, and you will be paid in dollars. Notice that the more francs you earn the more dollars you earn.

You will spend the next few minutes learning how to use your computer to trade. Talking between participants is not allowed, all communication will take place by using the computer screen in front of you. If you are found to violate this rule or disturb the experiment by inappropriate remarks or otherwise, we WILL ask you to leave.

PREVIOUS PAGE NEXT PAGE

Instructions Part 2 Continued

There are 12 traders in this experiment.

Your profits come from two sources - from collecting certificate earnings on all certificates you hold at the end of the year and from buying and selling certificates. During each year you are free to purchase or sell as many certificates as you wish, provided you follow the rules below.

For each certificate you hold at the end of the year you will be given one of three numbers of francs that will be provided in the Information section of your screen. An example is provided below. One of these three numbers is selected each year using the spinning wheel described in Part 1 of the instructions.

Information

X-Dividend = 50 (Spinner lands on number in range 1-35)

Y-Dividend = 240 (Spinner lands on number in range 36-80)

Z-Dividend = 490 (Spinner lands on number in range 81-100)

Your total certificate earnings for a period are computed for you (and displayed on your screen) by multiplying the earnings per certificate by the number of certificates held. That is,

> (Number of Certificates Held) x (Earnings per Certificate) = Total Certificate Earnings

Suppose, for example, that you hold 5 certificates at the end of year 1. If for that period your earnings are 240 francs per certificate (i.e. it is the Y-Dividend), then your total certificate earnings in the year would be 5x240 = 1,200 francs.

PREVIOUS PAGE NEXT PAGE

Instructions Part 2 Continued

Sales from your certificate holdings increase your francs on hand by the amount of the sale price. Similarly, purchases reduce your francs on hand by the amount of the purchase price. Thus you gain or lose money on the purchase and resale of certificates.

At the beginning of each year you are provided with an initial holding of certificates that is listed on your computer screen as below. You may sell these if you wish or you may hold them. For each certificate you hold at the end of the year, you receive the "earnings per certificate."

In addition, at the beginning of each year you are provided with an initial amount of francs on hand, which is also listed on your computer screen as below. You may keep this if you wish or you may use it to purchase certificates.

Thus at the beginning of each year you are endowed with holdings of certificates and francs on hand. You are free to buy and sell certificates as you wish according to the rules below. Your francs on hand at the end of a year are determined by your initial amount of francs, earnings on certificates owned at the end of the year, and by gains and losses from purchases and sales of certificates. These are your profits for the year. Your holdings of cash and shares are not carried over from one year to the next.

Your Holdings

Cash	1200
Certificates	4

Cumulative Profits	0

PREVIOUS PAGE NEXT PAGE

Information about Dividends

Whether the dividend you receive from the certificates you hold is the X-dividend, the Y-dividend, or the Z-dividend is determined by the experimenter at the beginning of the year by spinning a wheel with 100 equally-spaced numbers on it. Each number is equally likely to be selected (i.e., have the spinner land on it). If the number landed on is between 1 and 35 (inclusive), X-dividend is paid; if the number landed on is between 36 and 80 (inclusive), Y-dividend is paid; and if the number landed on is between 81 and 100, Z-dividend is paid. For example, if the number 37 is landed on, then the Y-dividend will be paid. That is, each participant would receive 240 francs for each certificate held at the end of the year. The wheel will only be spun once at the beginning of each period. You will not be shown the spin of the wheel, but you will be given a hint as to its outcome (described on next page). Moreover, the outcome of the spin, that is the actual dividend for the year, will be displayed on the computer screen at the end of the year.

Your Holdings

Cash	1010
Certificates	4

Cumulative Profits	1970

Certificate Value this Year	Total Franc Balance (Francs + Certificate value)
240	1970

Information

PREVIOUS PAGE NEXT PAGE

Instructions Part 2 Continued

At the beginning of each year, before trading starts, the experimenter will provide half of you with two clues about whether the outcome is X, Y, or Z in each period. After the experimenter has spun the wheel and determined the outcome for the period, your clues (if any) will appear on your computer screen. It will be one of the following:

Clue Options

Option 1	Option 2	Option 3
The certificate is not worth 50 Francs. The certificate is not worth 490 Francs. (Not X-Dividend / Not Z-Dividend)	The certificate is not worth 50 Francs. The certificate is not worth 240 Francs. (Not X-Dividend / Not Y-Dividend)	The certificate is not worth 240 Francs. The certificate is not worth 490 Francs. (Not Y-Dividend / Not Z-Dividend)

Information

Alert:

Your clue

The certificate is not worth 240 Francs.
The certificate is not worth 490 Francs.
(Not Y-Dividend / Not Z-Dividend))

Shared Message

Each certificate will be worth 50 (X-Dividend), 240 (Y-Dividend) or 490 (Z-Dividend). If the spinner lands on 1-35, it is X-Dividend; if 36-80, it is Y-Dividend; if 81-100, it is Z-Dividend.

If your clue is "The certificate is not worth 240 Francs. The certificate is not worth 490 Francs (Not Y-Dividend / Not Z-Dividend)" then the outcome for that period will be the X-dividend. Similarly, the clues "The certificate is not worth 50 Francs. The certificate is not worth 490 Francs" and "The certificate is not worth 50 Francs. The certificate is not worth 240 Francs" inform you that the outcome for the period will be the Y-dividend or will be the Z-dividend, respectively.

Each period, the computer randomly divides the traders into two groups of six each. One group will receive the clues, while the other group will not receive any clues. For example, when the outcome is the Z-dividend, six traders will receive the clue "The certificate is not worth 50 Francs. The certificate is not worth 240 Francs." and six will not receive a clue. Which group receives the clues is determined randomly via the computer. A similar procedure is followed for each year.

PREVIOUS PAGE NEXT PAGE

Instructions Part 2 Continued

The market for these certificates is organized as follows. The market will be conducted in a series of years. Each year lasts for 5 minutes. Anyone may make an offer to buy one certificate at a specified price, and anyone with certificates to sell is free to accept or not accept the offer. Likewise, anyone wishing to sell a certificate is free to make an offer to sell one certificate at a specified price.

The process by which you can make offers to buy and/or offers to sell is described on the following pages.

PREVIOUS PAGE NEXT PAGE

Offer Book

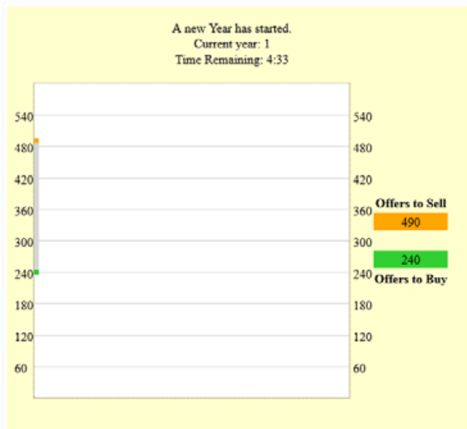

Instructions

During every year, participants can buy or sell certificates from one another by making offers to buy or to sell. The existing offers are shown on the Market Graph to the left. On top of the graph, the current year is shown. Below that, the Time Remaining for the trading year is shown. Each year lasts five minutes. The vertical axis lists the Price for the offers.

Every time someone makes an offer to buy a certificate, a GREEN dot will appear on the graph to the left. Every time someone makes an offer to sell, an ORANGE dot will appear on the graph to the left. Once a trade is actually made, the trade will be shown as a BLACK dot in the graph.

Offers are also listed on the Offer Book to the right of the graph.

PREVIOUS PAGE NEXT PAGE

Making Offers

To accept an existing offer from another participant, click the Buy or Sell button in the Immediate Offer section below. The Immediate Offer section shows you the best prices to buy, or sell, that are currently available in the market.

Suppose you want to place an offer to buy, it must be **higher** than the current best offer to buy, which is now 240. Say you want to buy at 250. You will have to type in 250 and click the Buy button. Please do exactly that on your screen by typing 250 in the Submit New Offer (buy) box and clicking on the Buy button.

PREVIOUS PAGE

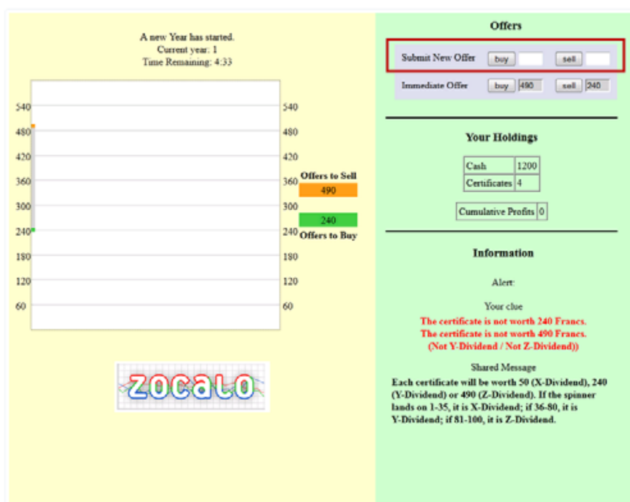

Making Offers

Notice your 250 Offer to Buy now appears in both the Market Graph and the Offer Book.

Next, suppose you want to place an offer to sell, it must be **lower** than the current best offer to sell, which is now 490. Say you want to sell at 470. You will have to type in 470 and click the Sell button. Please do exactly that on your screen by typing 470 in the Submit New Offer (sell) box and clicking on the Sell button.

PREVIOUS PAGE

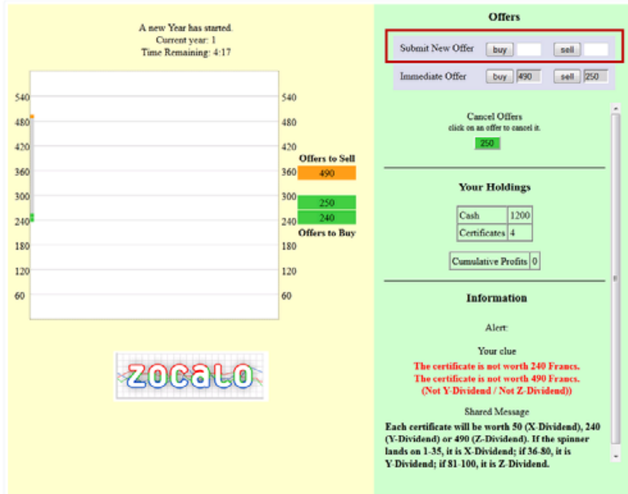

Making Offers

Notice your 470 Offer to Sell now appears in both the Market Graph and the Offer Book.

Now that you have submitted new offers, click the NEXT PAGE button to practice accepting existing offers.

PREVIOUS PAGE NEXT PAGE

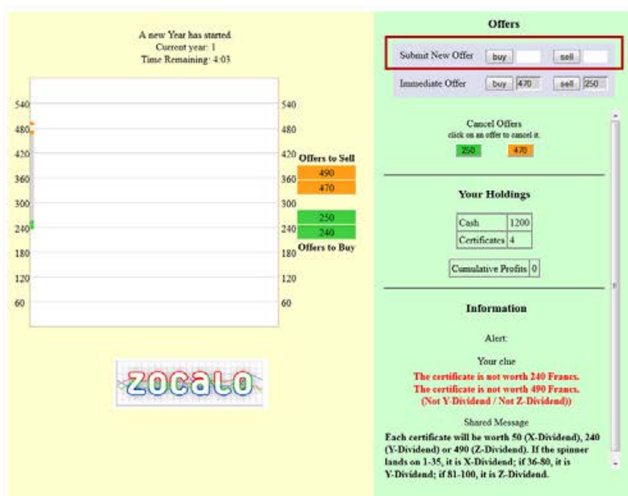

Accepting Offers

To accept an existing offer from another participant, click the Buy or Sell button in the Immediate Offer section below. This section shows you the best prices to buy, or sell, that are currently available on the market.

By clicking on the Sell button, you sell at the listed price. The current best offer to buy is 260, if you click Sell, you sell a certificate at the price of 260 immediately. Your certificates decrease by 1, and your cash increases by 260 francs. This is reflected in the "Your Holdings" box. **Click the Sell button below to continue.**

Submit New Offer buy sell

Immediate Offer buy 450 sell 260

PREVIOUS PAGE

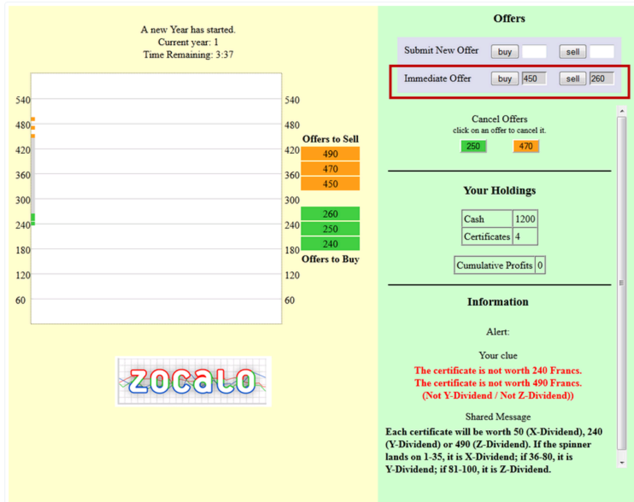

Accepting Offers

Notice your certificate sale of 260 is shown as a black dot in the Market Graph. Also, as noted in the "Your Holdings" box, your cash increased by 260 francs and your certificates decreased by 1.

Next, by clicking on the Buy button, you buy at the listed price. The current best offer to sell is 450, if you click Buy, you buy a certificate at the price of 450 immediately. Your certificates will increase by 1, and your francs will decrease by 450. This is reflected in the "Your Holdings" box. Your francs will decrease by 450. **Click the Buy button below to continue.**

Submit New Offer buy sell

Immediate Offer buy 450 sell 250

PREVIOUS PAGE

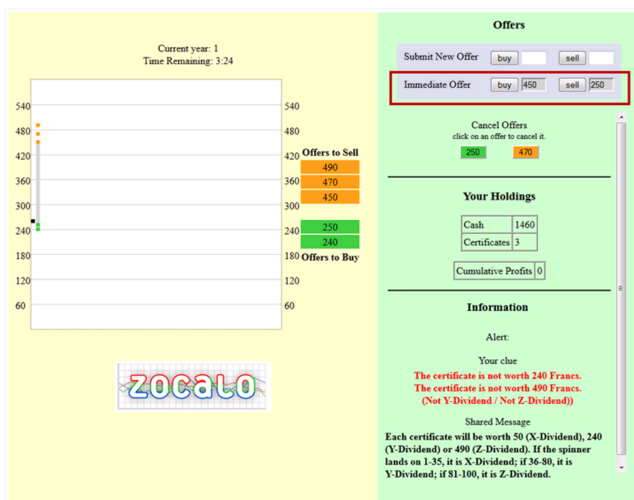

Accepting Offers

Notice your certificate purchase of 450 is shown as a black dot on the Market Graph. Also, as noted in the "Your Holdings" box, your cash decreased by 450 francs and your certificates increased by 1.

Now that you have accepted immediate offers, click the NEXT PAGE button to continue.

Submit New Offer buy sell

Immediate Offer buy 470 sell 250

PREVIOUS PAGE NEXT PAGE

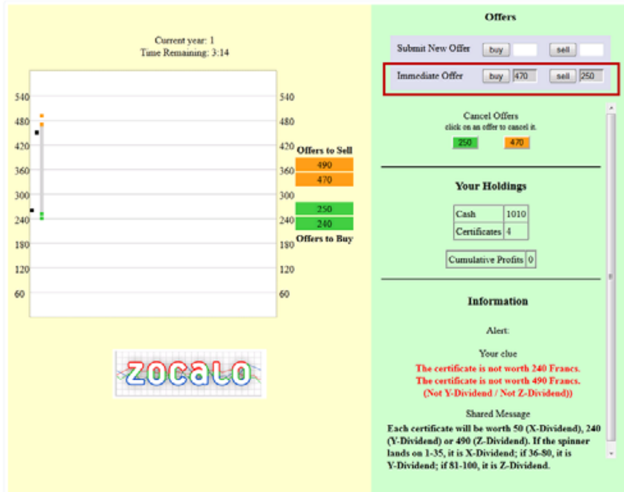

Your Holdings

As demonstrated on the previous pages, as you BUY certificates, your holdings of certificates will automatically increase and your cash will decrease. As you SELL certificates, your cash will automatically increase and your holdings of certificates will decrease. Your holdings of certificates and shares as well as your Cumulative Profits (sum of the profits for each of the market years) are displayed in the "Your Holdings" box.

PREVIOUS PAGE NEXT PAGE

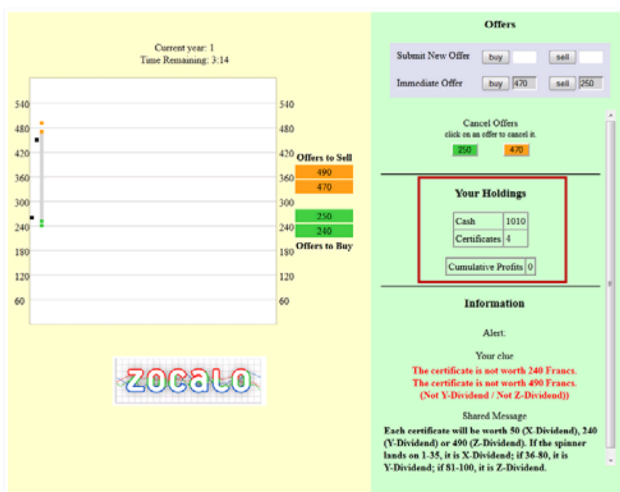

Cancelling Offers

Whenever you enter new offers to buy, or sell, you will have those offers appear as Buttons below the offer box. Your outstanding offers to buy cannot exceed your cash holding; your outstanding offers to sell cannot exceed the number of certificates you currently own. Therefore, you may have to delete your current offers to complete additional transactions. To do this, look for "Cancel Orders" and click on the corresponding button to take your offer out of the market. By clicking on these buttons, you can take them out of the market.

For example, please remove your Offer to Buy of 250 from the market by clicking the green Offer to Buy button labeled "250."

PREVIOUS PAGE

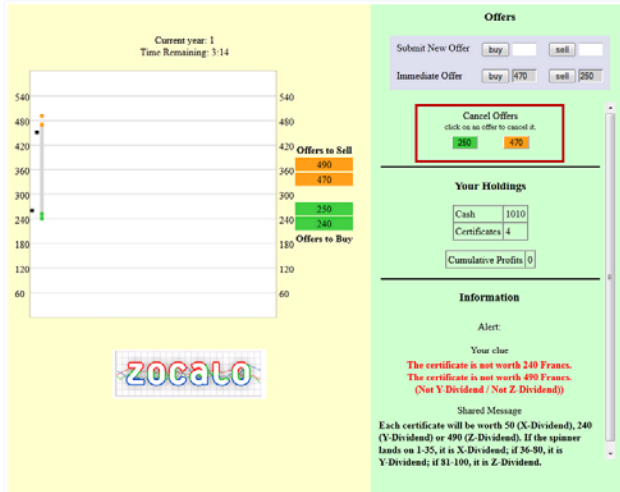

Cancelling Offers

Notice your Offer to Buy of 250 is no longer represented by a green dot in the Market Graph or listed in the Offer Book.

Click NEXT PAGE to continue.

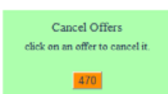

PREVIOUS PAGE NEXT PAGE

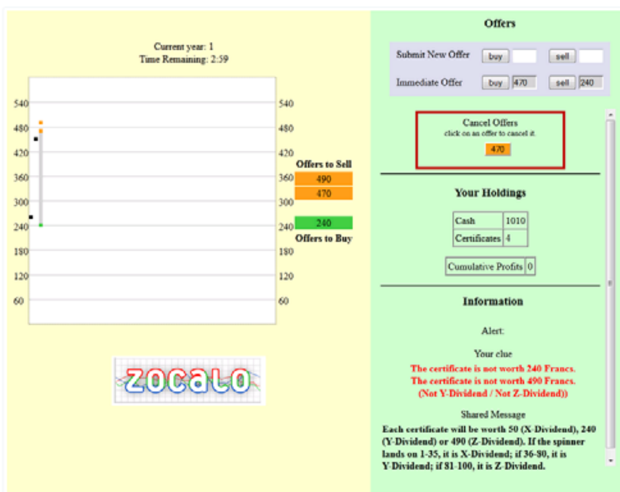

Instructions Part 3

Please answer the following questions. After all participants have correctly answered them, as part of the instructions there will be a practice trading year to ensure you understand the software.

Question 1: At the end of each period, each certificate earns a dividend of?

- A. 0.
- B. 50.
- C. Either 50, 240 or 490.
- D. 490.

Question 2: Which of the following statements is correct?

- A. All dividend values (50, 240 or 490) are equally likely.
- B. 50 is more likely to occur than the other dividend values.
- C. 490 will occur on average 50% of the time.
- D. 240 is more likely to occur than the other dividend values.

Question 3: You can put a new offer to buy in the market by:

- A. Clicking on the *buy* button when submitting a new offer to the market
- B. Clicking on the *sell* button when submitting a new offer to the market
- C. Clicking on the *sell* button to accept an immediate offer available in the market
- D. Clicking on the *buy* button to accept an immediate offer available in the market

Question 4: You can accept an existing offer to sell in the market by:

- A. Clicking on the *buy* button when submitting a new offer to the market
- B. Clicking on the *sell* button when submitting a new offer to the market
- C. Clicking on the *buy* button to accept an immediate offer available in the market
- D. Clicking on the *sell* button to accept an immediate offer available in the market

Question 5: The experiment will consist of a total of:

- A. 12 market years
- B. 13 market years
- C. 15 market years
- D. 17 market years

Question 6: If I receive the clue "The certificate is not worth 50 Francs. The certificate is not worth 490 Francs." at the beginning of a given year, this means:

- A. Each certificate will earn a dividend of 490 at the end of the period.
- B. Each certificate will NOT earn a dividend of 240 at the end of the period.
- C. Each certificate will earn a dividend of 50 at the end of the period.
- D. Each certificate will earn a dividend of 240 at the end of the period.

Question 7: Which of the following statements is correct?

- A. All the traders will receive clues.
- B. No trader will receive a clue.
- C. Each trader will receive a different clue.
- D. Half of the traders will receive clues and the other half will not receive a clue.

PREVIOUS PAGE SUBMIT


Quiz Completed
You have successfully completed the quiz.
The practice round will begin when everyone is ready.
Please wait quietly.