Costly Information Acquisition in Decentralized Markets: Experimental Evidence

Elena Asparouhova* Peter Bossaerts Wenhao Yang

Abstract

This study tests the rationality of the decisions to purchase information, the informational efficiency of prices, and the optimality of the resulting allocations in decentralized markets with a series of laboratory experiments. The theory predicts that markets with dispersed information and natural buyers and sellers converge to a fully revealing equilibrium. It is profitable to pay for information and as such, the Grossman-Stiglitz paradox does not emerge. Statistically significant improvements in both price efficiency and allocative efficiency are documented across trading periods. In contrast to the theory, which predicts that the higher the number of informed agents, the stronger the individual incentives to invest in information, the data present the opposite result. The price of information is higher when there are fewer informed agents. The result can be explained by boundedly rational decision making.

JEL Code: D82; C92 G14

Keywords: Experimental Finance; Information Percolation; Decentralized Markets

---

Elena Asparouhova (e.asparouhova@utah.edu) is at the University of Utah, Peter Bossaerts (plb32@cam.ac.uk) is at the University of Cambridge, and Wenhao Yang (wenhao.yang@uncc.edu) is at the University of North Carolina at Charlotte and The Chinese University of Hong Kong, Shenzhen. The authors thank Darrel Duffie, Mark Desantis, Bruce Grundy, Ann Gillette (discussant), Michael Hertzel (discussant), Charles Holt, Matthew Jackson, Shimon Kogan (discussant), Sophie Moinas (discussant), and Gustavo Manso for useful discussions. We also thank participants of the Society of Experimental Economics (2016 North America), FMA (2017), EFA (2018), CICF (2018), WFA (2019), Experimental Economics Workshop organized by the Center for Growth and Opportunity at Utah State University and seminar participants at Brigham Young University, Hebrew University, CKGSB, Tsinghua University, and Chinese University of Hong Kong, Shenzhen. Financial support from the National Science Foundation (Asparouhova, Bossaerts: SES-1426428), the Banque de France Foundation, a Moore Foundation grant to Caltech (2006-13) in support of Experimentation with Large, Diverse and Interconnected Socio-Economic Systems, and the Development Fund of the David Eccles School of Business at the University of Utah is gratefully acknowledged.

* Elena Asparouhova is the corresponding author. She can be contacted at e.asparouhova@utah.edu
University of Utah, David Eccless School of Business, Salt Lake City, Utah 84112.

Costly Information Acquisition in Decentralized Markets: Experimental Evidence

Abstract

This study tests the rationality of the decisions to purchase information, the informational efficiency of prices, and the optimality of the resulting allocations in decentralized markets with a series of laboratory experiments. The theory predicts that markets with dispersed information and natural buyers and sellers converge to a fully revealing equilibrium. It is profitable to pay for information and as such, the Grossman-Stiglitz paradox does not emerge. Statistically significant improvements in both price efficiency and allocative efficiency are documented across trading periods. In contrast to the theory, which predicts that the higher the number of informed agents, the stronger the individual incentives to invest in information, the data present the opposite result. The price of information is higher when there are fewer informed agents. The result can be explained by boundedly rational decision making.

JEL Code: D82; C92 G14

Keywords: Experimental Finance; Information Percolation; Decentralized Markets

1 Introduction

In decentralized markets, parties make opaque bilateral trade agreements. The opaqueness is a common criticism of those markets and, yet, the agreements provide valuable flexibility over standardized exchange agreements in centralized markets. Despite the lack of consensus on the costs and benefits of decentralized markets, the EU recently enacted new rules "...to force trading across all asset classes into open and transparent markets..." (The Economist, A Bigger Bang, Apr 26th 2014).

The transparency that centralized markets provide may, however, give way to the Grossman-Stiglitz paradox; Grossman (1976) and Grossman and Stiglitz (1980) argue that perfectly informationally efficient markets are impossible because when "information is costly, prices cannot perfectly reflect the information which is available, since if it did, those who spent resources to obtain it would receive no compensation."[1] A limited number of papers present alternative setups where information acquisition incentives are restored in the face of revealing prices, see Grundy and McNichols (1989), Manzano and Xavier (2011), and Breon-Drish (2011). Decentralized markets provide an alternative path for relaxing the Grossman-Stiglitz paradox by fully aggregating information when information is dispersed among participants.

To test the theory of costly information acquisition and aggregation in decentralized markets, this paper uses a set of controlled laboratory experiments constructed after theoretically modeled markets. The methodology allows for modulating the distance between the mathematical models and real markets by enabling or disabling of certain market features. While experiments allow for the creation of environments that can be very close to the benchmark theoretical model, the goal of experiments is to provide a bridge between theory and field data. In doing so, the experimental researcher has to perform a cost benefit analysis when enabling/disabling features that are brought to the model for mathematical ease only.

[1] The main argument underlying the paradox is that information acquisition is a strategic substitute. The incentives to collect information decrease as others become more informed, since the price becomes a more informative signal.

Departures from the mathematical model are embedded into the experimental design when those stringent assumptions of the model are substituted with empirically more relevant ones.

The theoretical literature that motivates our study can be divided into two categories based on the underlying information conditions: (i) information amplification, and (ii) aggregation of dispersed information. The information amplification literature, initiated by the seminal paper of Wolinsky (1990), concerns an environment where insiders share identical information. The non-fully revealing prices incentivize information acquisition but at the cost of allocative inefficiency. The aggregation of dispersed information literature, initiated by Duffie, Gârleanu, and Pedersen (2005), models decentralized markets as pairwise random matching of traders. In later papers Duffie and Manso (2007), Duffie, Giroux, and Manso (2010), Duffie, Malamud, and Manso (2009), and Duffie, Malamud, and Manso (2014) gradually allow for more classes of traders based on their knowledge of information and connectivity with one another, while still being able to characterize the dynamics of posterior beliefs in closed-form solutions. Under certain conditions, the above papers show that decentralized exchanges can arrive at an equilibrium where dispersed information is costly *and* markets eventually fully aggregate the information.

Following Duffie, Malamud, and Manso (2014), there are two mechanisms that affect the decision to purchase information in a decentralized exchange. First, information percolates only gradually, providing opportunities to profit from superior information. Rational uninformed traders anticipate this and only trade at prices that reflect their informational disadvantage, incorporating the willingness to incur a small loss in order to become better informed and profit from the information in future encounters. The ability to learn from trading hinders the *ex ante* incentives to acquire information, similar to the way it does in centralized exchanges. This first motive alone makes information acquisition a strategic substitute–the fewer the informed traders, the more valuable the information. Decentralized markets, however, provide a second channel into the valuation of information. After an encounter, traders learn from one another and proceed to their next encounter equipped with the joint information set, theirs and that of their counter-party. As information percolates among traders, the information that traders possess will almost surely converge to the truth. However, a small set of traders will be holding "wrong" information. In a simple example where the "truth" is binary, say 0 or 1, this set of traders will have posteriors that overwhelmingly point towards 0 when the truth is 1 and vice versa.[2] This is known as the martingale convergence property of likelihood ratios (Doob, 1953). The property tells us that the strength of wrong signals increases exponentially as the proportion of ill-informed traders goes to zero. It is the risk of encountering a counter-party in the wrong tail of the information distribution that provides *ex ante* incentives to traders to gather independent information. The comparative statics of the risk-hedging channel, when taken alone, are opposite to the learn-through-trade channel. Namely, the more the informed traders, the more valuable the information, i.e., information acquisition is a strategic complement. When the hedging incentive outweighs the learn-through-trade incentive, a condition guaranteed to be satisfied when the trading horizon is long enough, information gathering becomes a strategic complement. This breaks the spell of the Grossman-Stiglitz paradox in decentralized markets. The experiments are designed in a way that aims to facilitate the dominance of the risk-hedging channel.

The relevant experimental literature that serves as a backdrop for our design starts with early evidence from market experiments that ruled against the ability of decentralized markets to deliver competitive equilibrium (Chamberlin, 1948), even in the absence of asymmetric information. Since the seminal papers of Smith (1962) and Plott and Smith (1978), the view in experimental economics has been that centralized markets are needed to produce competitive equilibrium and its welfare merits. Early experiments with centralized financial markets have also confirmed their capability to amplify information (Plott and Sunder, 1988). The most relevant of past information amplification studies is that of Sunder (1992). Under asymmetric information, and when traders were provided with the opportunity

[2]The measure of the set of such traders converges to zero.

to acquire costly signals, the laboratory markets displayed the properties of the Grossman-Stiglitz paradox. When information was auctioned to participants, information prices quickly plummeted to zero.

Experiments with centralized markets and information aggregation have shown mixed results. The ability of prices to aggregate information depends on the complexity of the signals and the common knowledge of the structure of private values (Corgnet, Desantis, and Porter, 2018; Corgnet, DeSantis, and Porter, 2020; Biais et al., 2005). In a recent large scale meta-study, Page and Siemroth (2020) summarizes the information aggregation literature in centralized markets, and shows that surprisingly little information is incorporated in the asset prices. Page and Siemroth (2017) study costly information acquisition and find that those who acquire information do not recuperate the costs in trading. In related work, Halim, Riyanto, and Roy (2019) studies the effect of social networks on costly information acquisition in centralized markets. They find that higher connectivity in the social network structure increases the incentives to free ride, thus reducing the costly information acquisition. Kendall (2020) studies individual decision making against automated market makers where the cost of information is expressed in time. This study also points to the undesirable conclusion that traders are not willing to pay for information, whether rational or not.[3]

With the early historical record stacked against the decentralized markets setup, it is not surprising that there is very little experimental research within the setup. Asparouhova and Bossaerts (2017) studies information aggregation in decentralized laboratory markets where private information is free and exogenously distributed among traders. In the single asset trading environment, the "no trade theorem" is avoided as trade must occur for diversification motives. The diversification motive is induced by having a non-tradeable portion in the traders' endowments that can be hedged with the traded asset. The information about the value of the traded asset is asymmetric and dispersed among traders. The data reveal that prices aggregate the available information but not in the strict sense of the theory. Instead of

[3]We thank an anonymous referee for the phrasing.

prices converging to the asset's expected payoff under the aggregate of all signals, they fluctuate in a no-arbitrage range formed by the average of the asset's expected payoff under each of the private signals.

The experimental design in this paper differs significantly from Asparouhova and Bossaerts (2017). Our experimental design follows closely the theoretical setup in Duffie, Malamud, and Manso (2014) in terms of the traders' utility, information structure, and the trading mechanism (see details in section 3) due to the different focus of this paper, namely costly information acquisition. Trading in this single asset environment occurs due to buyers and sellers having differential valuation of the asset.

This study is the first to use experiments to address the question of costly information acquisition in decentralized markets, a necessary condition for informationally efficient prices to exist. While the issue is investigated through the lenses of the information percolation model developed in Duffie, Malamud, and Manso (2014), it is not a forgone conclusion that the traders in real markets understand the subtleties behind the reasoning, and are confident that others do as well, so that the predicted costly information acquisition emerges. The study investigates if the informativeness of prices increases as the information percolates, and if allocations exhibit the Pareto optimality properties usually attributed to centralized markets only.

Our first finding is that the participants with no initial information remain willing to pay for information across all replications, in contrast to the centralized markets in Sunder (1992). Price efficiency improves significantly from one trading period to the other. Moreover, trading volume significantly increases as information is exchanged through trades. With increased trading activity, more gains from trade are necessarily realized, and thus allocative efficiency improves. The theory predicts that the desire to invest in information gathering exhibits strategic complementarity; this result, however, does not emerge in the experiments. Instead of bidding more aggressively for information when there are many informed traders, the uninformed traders submit higher bids for signals when fewer traders are informed.

Comparing the gross payoffs between informed and uninformed traders, in treatments with both low and high number of informed traders, the ones who successfully acquire information outperform those who do not. Once the payoffs are adjusted for the cost of information, there is no statistically detectable difference in performance. This result suggests that taking prices of information as given, traders make rational decisions when buying information. Two additional findings might explain why information prices do not conform to the model, however. First, we find that traders make large losses after trading with someone who holds "wrong" signals, an event more likely to occur when more traders are informed. Given that most of the value of acquiring private signals relates to arming oneself against counter-parties with "wrong" signals, this might explains why, overall, bidding low (in comparison to the theoretical prediction) in the information auction in the treatment with many informed traders does not lead to higher average payoffs: participants appear to be unable to make the best use of the information they obtain through the auction, or assess when exactly this information is most valuable.

Second, the results demonstrate that the uninformed traders send conservative offers to the asset market that are far from their valuation, while informed traders send aggressive orders. In the Low information treatment, this results in fewer trades than predicted. Thus, in this treatment, contrary to the theory, trading is unlikely to happen between two uninformed traders. As such, unless informed, one would not realize gains from trade, and this provides incentives for higher bidding for information than what the theory warrants.

The findings have implications for real-world markets. Many markets are organized as loose networks where trading flow is opaque. It has been well documented that successful financial traders rely not only on communication through markets but also through a synchronous verbal communication system, which used to be phone-based, but recently has become made more effective through chat and instant messaging. Such verbal communication can only have as a goal to clarify the intentions of the parties, for the benefit of all involved in communication through markets, see Saavedra, Hagerty, and Uzzi (2011) for a detailed study. Analyses of the nature of (instant) messages which reveals the role in clarification and verification can be found in Cetina, Karin, and Bruegger (2002) and Preda (2012).

Confirming the theory, the experiments show that decentralized institutions are not necessarily detrimental to information aggregation and allocation efficiency. Prices in decentralized markets could ultimately be more informative and, hence, provide better signals to the economy for optimal resource allocation. However, decentralized markets appear to provide incentives for information acquisition that are not fully aligned with those from the theory and more research is needed to understand the precise channels for those deviations.

The rest of this paper is organized as follows: Section 2 describes the theoretical framework. Section 3 concerns the experimental design. Section 4 discusses the results, and Section 5 concludes.

2 Theoretical Framework

The economy in the theoretical model is populated by an infinite number of risk neutral agents who trade a single asset for T periods, one unit at a time. The asset's only payout is through its terminal value at time T. The agents do not know the realization of the terminal value but can either purchase a signal at time 0 or learn about it through trading. Agents meet randomly as pairs of one buyer and one seller at a time (without repetition) and are given the opportunity to trade in a private double auction.[4]

[4]For simplicity, the authors assume that the auction format is a seller's auction, where upon buyer's and seller's orders crossing, the trade execution price equals the seller's offer.

If trading follows the above mechanism, under certain conditions[5], the prices eventually become fully revealing. The decision to purchase information exhibits strategic complementarity. This means that there are stronger incentives for traders to acquire information if they perceive that others are doing the same. The result is an elegant implementation of perfect Bayesian equilibrium in monotone undominated strategies.

2.1 The Traded Asset and Payoffs

There is one risky asset with a payoff that depends on the realization of a binary variable $Y \in \{0,1\}$, and the value is revealed to traders after period T. Whether an agent i is a buyer or a seller is determined by their utility function

$$U_i = v_i 1_{\mathrm{Y}=1} + v_i^H 1_{\mathrm{Y}=0} = v_i Y + v_i^H (1-Y), \tag{1}$$

where $v_i = v_s$ for sellers, and $v_i = v_b$ for buyers, with $v_b > v_s$. Similarly, $v_i^H = v_s^H$ for sellers, and $v_i^H = v_b^H$ for buyers, with $v_b^H > v_s^H > v_b$.

2.2 Information Structure

Agents in this economy are either endowed with or have the option to acquire information in the form of a noisy signal about the state of the world.[6] In each period agents are randomly matched (without repetition[7]) and given the opportunity to trade. When a seller and a buyer submit their offers, based on the strictly increasing bidding strategies, they can then infer the beliefs of their counter-party. Thus, after an encounter, a buyer and a seller have identical posterior beliefs. Because a trader can learn the information held by her counter-parties from the offers they submit, their incentive to buy information

[5] The number of periods T must be large enough.

[6] In the most general model the agents are divided into different classes based on their private values, initial information and likelihood of matching other classes.

[7] Not only do agents not meet again once they have met, but at each point in time, for any particular agent, all prior encounters including that agent are with agents whose extended encounter sets were disjoint with that of the given agent.

ex-ante is reduced. Because posterior beliefs are identical for trading partners, the cross-sectional distribution of these beliefs can be summarized by an evolution equation. As the agents become more informed on average, the cross-sectional distribution becomes fat-tailed, i.e. some agents of a shrinking mass get extreme posteriors. A trader with an extreme and wrong (pointing to $Y = 0$ when $Y = 1$ and vice versa) posterior belief is referred to as a "long shot." The desire to hedge against an encounter with a long shot increases the incentives of traders to acquire information ex-ante.[8] It is the presence of long shots that brings about the information complementarity: if information acquisition is costly, as more agents in the economy are informed and thus the tails of the cross-sectional posterior distribution are heavy, agents are more eager to purchase information to mitigate the effect of counter-parties of the long shot type.

Technically, the long shots emerge as information propagates through trade and each of the trading parties' posterior density is replaced by the convolution of the traders' pre-trade densities. The cross sectional distribution of beliefs becomes fat-tailed due to the martingale convergence property of likelihood ratios. The detailed proof can be found in Doob (1953); below is a recap. If $(X_n : n \geq 1)$ is a sequence of i.i.d. random variables with common density g, and f is another density with the property that whenever $g(x) = 0$, then $f(x) = 0$ then Doob (1953, Sect. VII.9) shows that under g, $L_n = \prod \frac{f(X_i)}{g(X_i)}$, $n \geq 1$, is a martingale with mean 1 that converges almost surely to 0. This can only happen if events that are rare (with probabilities converging almost surely to zero) are also extreme (for the mean to be 1). We are going to refer to this result as the *Martingale Property*, MP. Having rational expectations, agents understand that they can counteract the adverse event of encountering a long shot by purchasing an independent signal: "as agents acquire more information, the average quality of information is improved, but there is also an increased risk of receiving an unusually misleading bid. That is, the tails of the cross-sectional type densities become fatter at both ends, for some period of time."

[8]Note, that the encounter with a long shot itself is profitable. It is the future encounters, after inheriting the signal of the long shot that make the long shot counterparties lose money.

(see p.15 of Duffie, Malamud, and Manso (2014)). How the MP is implemented in the experiment is detailed in Section 3.

2.3 Testable Hypotheses

The first hypothesis takes the broad-brush conclusion of Duffie, Malamud, and Manso (2014) that traders purchase information in equilibrium and serves as a comparison to other information acquisition studies:

Hypothesis 1 *Traders acquire costly information.*

In centralized markets, Sunder (1992) finds that due to markets transmitting the private information, when the cost of information is fixed, the demand quickly dwindles down to zero. On the other hand, Page and Siemroth (2017) in (centralized) prediction market setup find over-investment in information. Similar over-investment results are found by Chen and He (2021) in an environment with centralized matching markets, where investment in the laboratory also positively correlates with the belief that others are investing in information. The latter is a mistake according to the corresponding model, which differentiates it from the conclusion of Duffie, Malamud, and Manso (2014), where the behavior is justified in equilibrium.

The hypothesis is acknowledgement to the issue of information investment that has been central to the literature. The precise prediction that follows from the theoretical development is listed below.

Hypothesis 2 *The higher the number of informed traders in the market, the higher the incentive of each uninformed trader to acquire costly information, i.e., there are strategic complementarities in information acquisition.*

Hypothesis 3 *As trading progresses, prices become more informationally efficient. Allocation efficiency also improves with trading.*

This hypothesis is weaker than the precise prediction of the theory that states that (i) prices converge to fully revealing, and (ii) allocations are efficient. The reason for a weaker hypothesis is the limitations of experimental testing (the limited number of trading rounds) that we discuss in the next section.

3 Experimental Design

The experiments are designed around the goal of testing Hypotheses 1, 2, and 3. While experiments provide control and precision in hypothesis testing, they are subject to some unique challenges. To ensure reliable data collection on information acquisition choices, the design needs a sufficient number of replications of the same situation. In addition, the implementation of the theoretical setup asks for appropriate incentives for buying information to be present in the laboratory. The direct interpretation of the theory would require that the number of periods in each trading replication, T, be large. Large T allows for extreme wrong signals to emerge. The constraint that the experimental methodology imposes on the possible tests of the theory is that a typical trading session lasts about 3 hours, and as a result one cannot have both many replications, and many periods within replication. To shorten the number of periods while still preserving the underlying reason for information acquisition, "long shots" are introduced into the information acquisition stage. If a session has a given number I of participants, then I binomial random variables are drawn until there is an extreme "opposite to the truth" signal, while the likelihood ratio is close to 1 even in the small sample.[9] Notice, that if only a small number of participants are informed there might not be a long shot within every given replication.[10] The introduction of long

[9]The insistence that the likelihood ratio is in the interval [0.95 1.05] is consistent with the martingale property of the likelihood ratio result and it precludes the sample from having too extreme of signals. For example, with an urn of 9 balls, the opposite signal can only have up to 6 balls of the "wrong" color.

[10]If instead of insisting on the occurrence of a long shot within I signals, one insisted that a long shot be distributed as a signal each replication, such choice would change the distribution of long shots between the High and Low information treatments. The latter would not only add to the complexity but it would correspond to two different parameterizations of the model. The code for how the signals for the experiments are drawn (the signals are always drawn during the experiment and not pre-drawn) is included in the Online Appendix A.

shots allows for the number of trading periods to be reduced to only two. Subsection 3.3.3 describes in detail how long shots enter the experimental design.

Another simplification in the experimental design concerns Bayesian updating. After a transaction, each of the two parties is provided direct access to the information of the other. This is an important methodological innovation — if investment in information is observed in the experiment, it cannot be attributed to traders' inability to infer the counter-parties' signals from offers. This design feature ensures the perfectly revealing nature of the encounter, but still leaves it to the traders to decide how to use this information in the formation of their future bid/ask offers. There is no exchange of information between a buyer and a seller who submitted offers but did not trade. This is a deliberate design choice that departs from the theory. This feature prevents traders in the experiment from submitting meaningless bid/ask offers with the sole purpose of collecting information.

The study consists of five experimental sessions, each lasting approximately 3 hours, with the number of participants ranging from 14 to 18, for a total of 80 participants across the sessions. The total participant compensation is a show-up reward plus the cumulative earnings from the each session's replications.[11] Participants receive their compensation in cash at the end of each experiment session. The final payments ranged between \$20 and \$84 with an average of \$52.3 (median \$52).

3.1 The Data-collection Portion of Each Session

There are 10 replications per session.[12] A replication, depicted in the diagram below, consists of (i) an information acquisition period, (ii) trading Period 1, and (iii) trading Period 2. During the replications

[11]The payment protocols for the vast majority of experiments in Economics and Finance involve the payment for only one or two of the replications, chosen at random at the end of each experimental session. In our study, this creates highly variable payoffs among the participants. While a variable payoff is generally a good incentive feature, in this study not all variability is incentives-related. The main concern is the variability coming from the payoffs of those who happen to be randomly allocated the long shot signal. To combat this variability one needs to pay more than one or two replications. The decision to pay for all replications was driven by the desire to diminish the complexity in an already complex experiment.

[12]Only Session 3 has 11 replications.

all interactions such as bidding, signal distribution, and trading, happen through the experimental software. In the very first replication of each session, every participant receives information (thus, no auction for information). This treatment, called “Full” (Information), is not meant to serve as a proper treatment, but a calibration. It ensures that participants experience the receipt of information and the act of trading on it, before initiating replications with asymmetric information. As such, this replication is not counterbalanced against the others. Additionally, the first two sessions included full information treatment in replication 6.[13]

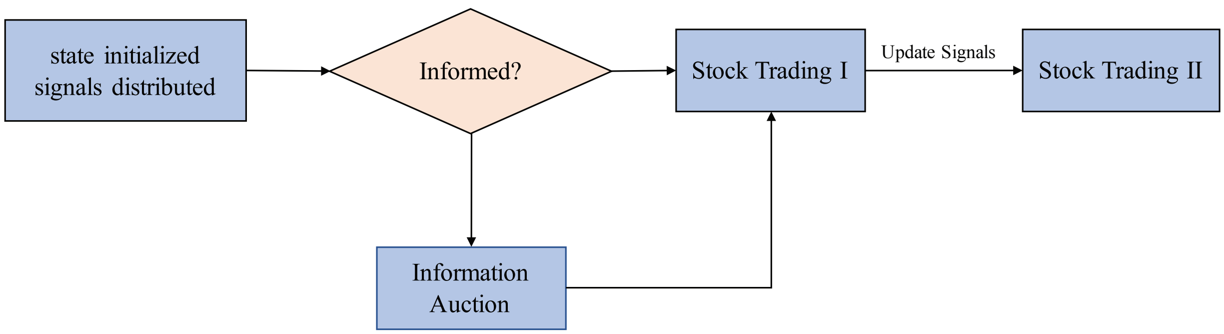

For the replication involving asymmetric information, before the information acquisition period, participants learn how many other participants are informed and how many will have the option to purchase information. Each session implements two treatments, “High” information or “Low” information. For example in the session with 16 participants, in the High (Low) information treatment, 12 (4) are initially informed, all of the remaining 4 (12) traders have the opportunity to bid for information and 1 (3) information signals are awarded (via the auction described below). Thus, the total number of informed traders is 13 (7) in the High (Low) information treatment. The exact details of the experimental sessions and treatments are presented in Table 2.

[13]The first two sessions also have replication 6 being a Full treatment, as part of an initial design that had 5 replications that repeated twice, and each contained 1 Full, 2 High and 2 Low information treatments. With the initial results pointing against complementarity, that design had a low power to reject complementarity as the number of High Treatment observations were too low. This lead to a change in the design to fit 1 Full, 3 Low, and as many High Information treatments as the time of the experiment session allowed. Online Appendix A includes analysis of the possible order effects that this change might have introduced and we do not find such effects.

The identity of the initially informed traders is randomly determined. Those who are left uninformed participate in an information auction, as in Sunder (1992). During the auction, the traders do not know if their trading role will be that of a buyer or a seller (thus preserving the *ex ante* nature of the information purchasing decision). After the auction and before the first of the two trading periods, traders are assigned their roles of buyers and sellers, and those roles remain unchanged for the duration of the replication.

In each of the two trading periods, buyers (sellers) can post at most one buy (sell) offer. After a random matching between buyers and sellers, crossing offers are executed at the midpoint and each trader with a successful transaction "inherits" the information of their counter-party. The second trading period is identical to the first, save for the information content available to those who successfully traded in the first trading period. After the conclusion of the second trading period, the binary state is revealed and the traders' earnings for the replication are computed and presented to them. The next replication repeats the scenario all over, with a different information treatment.

3.2 The Traded Asset

The single security traded in the experiment is called the "Stock." In each of the two trading periods, buyers are endowed with cash and can buy at most one unit of Stock per period, for a total of up to two units per replication. Similarly, the sellers can only sell at most one unit of the Stock per period, and up to two units per replication. The payoff of the Stock to buyers and sellers depends on which of the two states, $\{Orange, Blue\}$, is realized. In the former case $Y = 0$ and in the latter $Y = 1$. The payoff of a unit of the Stock is:

$$U_i = v_i 1_{\mathrm{Y}=1} + v_i^H 1_{\mathrm{Y}=0} = v_i Y + v_i^H (1 - Y),$$

with $v_b^H > v_s^H > v_b > v_s$ and . The latter generates gains from trade both when $Y = 1$ and when $Y = 0$. Table 1 illustrates the payoff structure of the Stock.

3.3 Information

3.3.1 Noisy Signals

Before the start of a replication, participants know that they have equal chances to be a buyer or a seller. The state that determines payoffs in each replication is determined by a coin toss with the outcome revealed only at the end of the replication.[14] If the state is $Y = 0$, then the signals are a sample of 9 balls drawn with replacement from an urn consisting of 2/3 blue balls and 1/3 orange balls (the state is called Blue when $Y = 0$). If $Y = 1$, then the signals are draws of 9 balls from an urn with 2/3 orange balls and 1/3 blue balls. The signal an agent receives is called "the sample" of that agent. After determining the state but before the buyer/seller assignment, some traders receive a sample; these traders are called "initially informed" traders, the rest are called "initially uninformed." The individual sample size is set to 9 to accommodate the long shot definition.[15] A signal that is a long shot has to be rare and extreme. Simulations with 9 balls demonstrate that long shots that consist of 6 balls of the "wrong" color are possible within the requirement that likelihood ratios do not deviate from 1.[16]

An agent's *type* is determined by the number, B, of blue balls in a sample of S balls. The sample size S can be 0, 9 and 18. $S = 0$ for the uninformed traders in either Period 1 or Period 2 of the trading. $S = 18$ for traders who are informed during Period 1 of trading and trade with another

[14] The coin is tossed inside a box in front of all participants, ensuring that the mechanism is apparent and while outcome remains concealed it is later announced and verified by the participants. If the outcome of the coin toss is "heads" ("tails"), the state of the world is $Y = 1$ ($Y = 0$).

[15] The definition is that a long shot is a signal that has probability of being drawn that goes a.s. to zero, but as this probability goes down, its "extremeness" increases, so that to keep the likelihood ratio equal to 1.

[16] For example, it is not possible to satisfy the likelihood ratio restriction with 7 balls of the wrong color. An alternative design could involve signals of 7 balls and long shot being 5 of 7 balls. We do not believe that the results would depend on the above variations.

informed counterparty, and as a result inherit their signal. Thus, $S = 18$ is only possible in Period 2 of the stock trading. For all other information types and trading periods, $S = 9$. Thus, the type of a trader is defined as $\theta(B,S) = \log \frac{1+2^{2B-S}}{1+2^{S-2B}} = \log \frac{2^{2B}}{2^S} = \log(2^{2B-S})$. Hence, $P(Y = 0|\theta) = \frac{e^\theta}{1+e^\theta}$. When a type $\theta(B_1,S_1)$ trades with a type $\theta(B_2,S_2)$, they both emerge out of the encounter with the new type $\theta(B_1+B_2,S_1+S_2) = \theta(B_1,S_1)+\theta(B_2,S_2)$. The probability that the state is $Y = 1$ after observing B blue signals equals $P(Y = 1|B) = \frac{1}{1+2^{2B-S}}$. The unconditional probabilities are $P(Y = 0) = P(Y = 1) = 0.5$.

3.3.2 Information Auction

The initially uninformed traders have the opportunity to participate in an information auction for samples. The auction, conducted as a first price auction, takes place before the buyer/seller assignments. The precise auction design is as follows. If u samples are auctioned off, then u separate first-price auctions are conducted. Every uninformed trader provides a bid, and once submitted, it is randomly assigned to one of the auctions (the number of initially uninformed agents is chosen so that it is divisible by u). The highest bid in each auction wins the sample and pays the highest price.

As in the example provided in subsection 3.1, assume there are 16 participants, then in a High (Low) treatment, there are 12 (4) initially informed and 4 (12) initially uninformed traders. Then 1 (3) samples are auctioned off. The 4 (12) initially uninformed traders submit bids[17], and the bids are then randomly assigned across the auctions with 4 participants each.[18] As in this example, all auctions have the same number of participants, independent of the treatment.[19]

[17]The bid price is allowed to be 0.

[18]In this example the High treatment has 13 participants who eventually have information, while the Low treatment has 7 after the conclusion of the information auction.

[19]Why the choice of First Price Auction (FPA) vs. a multi-unit auction or a BDM mechanism? Because of the common value nature of the "item" auctioned, none of the above mechanisms are incentive compatible. When dealing with the k-th price auction, in Chapter 4 of "The Economics of Risk," Kagel (see (Meyer, 2003)) summarizes the existing experimental literature and concludes that the level of the winner's curse strongly depends on the number of bidders. The design feature of having the same number of bidders for all sessions addresses the dependence of the deviation from the theory (winner's curse) on the number of bidders. One favorable point in the research inquiry is that the variable of interest is the difference in valuation of information between the High and Low Information treatments. Assuming all else equal, in particular keeping the number of bidders constant, taking first differences should cancel out noise/irrationality in bidding. In that sense, any of the mechanisms

3.3.3 Long Shots

Statistically, the cross-sectional distribution of types, as discussed in Subsection 3.3.1, develops fat tails as traders meet and exchange information. The traders on the wrong tail are called "long shots." In the experiment with only a limited number of participants, long shots are introduced by having one (or at most two) of the traders receive information that is wrong, and the distribution of drawn balls is 6:3 in favor of the wrong-colored balls. The possibility of long shots is presented to the participants in the instructions (see the Information Auction part). Then they learn hands on about the long shots in the practice sessions. As traders finish each practice replication, the experimenter reveals the information distribution and asks those who have a sample pointing opposite to the truth to raise their hands.[20] Some participants receive opposite signals by pure chance, but these signals are less extreme, e.g., 5 out of 9 balls are of the wrong color.[21]

In unreported results (available in Online Appendix A), we conduct extensive simulations with 14 to 18 virtual participants to test the complementarity result in the $T = 2$ scenario. The complementarity result holds, and the value of information triples from about \$0.34 to almost \$1.10 in the simulations in the "Low" vs. "High" informational treatments. The simulations also serve as a guidance for setting the auction parameters, in particular, endowing the participants with \$2 to participate in the the auction for information.

would potentially work. The main reason for the choice of FPA then becomes its relative simplicity and also the participants' familiarity with it over other mechanisms.

[20] The initial design included the detailed explanation of long shots in the written instructions, which had the effect of creating a lot of questions and confusion that would only clear at the practice replications. For this reason in the final design the formal introduction of long shots was postponed to the practice replications.

[21] Long shot outcomes are obtained by drawing outcomes in a counter-balanced way so that the martingale property of the likelihood ratio holds in each sample. Specifically, before each replication, we execute a program that draws $I-1$ or $I-2$ 9-ball samples, with replacement, from the urn corresponding to the state of the world. The last or last two 9-ball samples are then chosen to ensure that the average of the likelihood ratios (across the I 9-ball samples) is as close as possible to 1. Whether the long shots will actually affect trading depends on whether they end up in the hands of the informed traders, or are bought in the auction by the uninformed. As a result, not every replication has a realized long shot in it.

3.4 Market Structure

Immediately following the information auction is trading Period 1. Sellers submit at most one ask offer to sell one unit of the Stock, and buyers submit at most one bid offer to buy one unit of the Stock. Then buyers are randomly matched with sellers, and for every pair where the bid exceeds the ask offer, a trade is executed at the midpoint between the bid and the ask. In addition, all traders who successfully participate in a transaction inherit the signal of their counter-parties. Equipped with the new information, traders enter Period 2 of trading. Period 2 is identical to Period 1, except that the exchange does not proceed to a Period 3. Each trader's earnings from trading equal the gains from trade for that participant. For buyers, this is the difference between the value of the Stock and the purchasing price; for sellers, it is the difference between the selling price and the seller's value. [22]

3.5 Parameters for the Laboratory Implementation

We use simulation analysis to calibrate experimental parameters such as the payoff values and auction prices.[23] The simulations provide an upper bound on the value of information and thus we gauge this value by the simulation outcomes. The upper bound is needed because the uninformed traders have to be provided with sufficient, but reasonable amounts of cash to purchase information. The upper bound ranges from \$0.3 to \$1.37 as number of informed traders ranges from 1 to 15. Based on the simulation results, the cash made available to the initially uninformed traders is set to \$2. The traders can use the cash to bid in the information auction. Only the winning bid (or bids in the case of multiple first price auctions) acquires information, pays the bid price, and keeps the rest of the cash to be counted towards

[22] The laboratory where the experiments are conducted has a regulation where subjects are paid \$15 on average. When the average payoffs are below that, everyone's earnings is multiplied by a factor (on for the entire session and all subjects) to bring the earnings to \$15. When the payoff is above \$15, on average, there is no adjustment. Individual researchers cannot influence this rule as it is per all users of the experimental subjects database.

[23] The simulation analysis is available from the authors upon request.

the replication's payoff. The rest of the traders' cash of $2 also goes towards their total replication payoff. None of this initial cash enters the trading periods or counts towards trading performance.

3.6 The Instructional Portion of Each Session

The first half of the session consists of an instruction block followed by three practice replications. During the instruction period the experimenter reads the instructions out loud, in addition to presenting them on a large screen. The screen instructions follow those given to the participants, and where appropriate, they are aided by graphic/picture presentations. The practice replications allow the participants to familiarize themselves with the software, Flex-e-markets, and the incentives that the experimental design imposes on them. At the same time, participants have the opportunity to, and do, ask questions. The experimenters also demonstrate the statistical properties of the signals during this block of time.

4 Results

4.1 Value of Information

The goal of this paper is to study the incentives to acquire costly information in a decentralized market setting. The theory predicts that information will be acquired in equilibrium, even if it is costly. Moreover, the acquisition of information exhibits complementarities: the utility gain from becoming informed increases with the number of informed traders. To test the strategic complementarity, two main treatments are introduced. In the High (Low) information treatment, the majority of traders are initially informed (uninformed). According to the theory, the value of information is predicted to be higher in the High treatment. In both treatments, the uninformed traders have the opportunity to participate in information acquisition auctions, acquire a signal, and become informed. Thus, the first measure that assesses the value of information is a function of the bids in the auction stage of the game. The bids provide an *ex ante* measure of the value of information.

The second measure, an *ex post* one, is the difference between the gains made by informed and uninformed traders in the Stock trading stage of the game. With rational expectations, the cost of information measured through the bids should equal the benefits of information measured by the excess gains from trade to the informed traders. The following subsection presents the results of the treatment effect on the value of information based on these two measures.

We start with a descriptive analysis of the signal-purchasing behavior. We reiterate each of the hypotheses before the presentation of the relevant analysis.

Hypothesis 1 *Traders acquire costly information.*

Every signal that was available for purchase was purchased, in both treatments, and for a positive price. Table 3 Panel A reports the summary statistics of the bidding patterns in the information auction across treatments (High and Low). Traders are willing to pay a strictly positive price for information. The average bid from all information auctions is \$1.24 (with a median of \$1.36). When conditioned on the treatment, participants bid \$1.32 on average (median of \$1.49) in the Low treatment and \$1.08 (median of \$1.17) in the High treatment. The data provides strong support for Hypothesis 1.

The remainder of subsection 4.1 is devoted to testing of the second hypothesis:

Hypothesis 2 *The higher the number of informed traders in the market, the higher the incentive of each uninformed trader to acquire costly information, i.e., there are strategic complementarities in information acquisition.*

4.1.1 The *ex ante* Value of Information from Information Auction

Figure 1 shows that the entire bid distribution is shifted towards higher prices in the Low treatment. To analyze the effect of the number of initially informed traders on information bidding, we conduct pairwise comparisons at the trader level. Specifically, for each trader who participates in the information auction in both treatments, we compute the average bidding price per treatment and conduct a pairwise t-test. The results are included in Panel B of Table 3. As suggested by the t-test results, the average bid in the High treatment is significantly lower than in the Low treatment (with a t-statistic -3.62 and p-value 6.10^{-4}). Figure 2 plots the pairwise comparison of auction bids in a "violin graph." Panel C reports that all information that can be purchases is indeed purchased. In particular, it reports the statistics of positive bids across the two treatments. As shown in the table, about 88% and 86% of participants bid positively in the information auction in the Low and High treatments respectively.

To further examine the treatment effect on participants' information acquisition decisions, three regression models are estimated. The first model is a simple OLS regression where the bidding prices are regressed on the treatment indicator and the replication number. The model is specified as follows

$$b_{i\tau} = \alpha + \beta T_{i\tau} + \gamma\tau + \varepsilon_{i\tau}, \tag{2}$$

where $b_{i\tau}$ is the bid of participant i in replication τ. $T_{i\tau}$ is an indicator variable for the High treatment in replication τ. The coefficient β captures the treatment effect and is the main coefficient of interest. $\gamma\tau$ captures the "time" (replication) trend. The results are reported in column (1) of Table 4. The coefficient β is -0.23 and statistically significant. This implies that the traders are willing to pay \$0.23 less for the information in the High treatment than in the Low treatment, consistent with the univariate results from Table 3.

The second model augments the OLS regression with individual fixed effects:

$$b_{i\tau} = \alpha_i + \beta T_{i\tau} + \gamma \tau + \varepsilon_{i\tau}, \tag{3}$$

where α_i is the individual fixed effect, which absorbs any individual level unobserved time-invariant characteristics that may affect the bidding choice of trader i. The results, reported in column (2) of Table 4, are similar to those of the OLS regression. β is about -0.2 and statistically significant. In addition to the fixed effects model, we estimate the above equation using a random effects model. The results are reported in column 3 of Table 4. The overall results are almost the same as those reported in column (2). The treatment effect is negative and statistically significant. A Hausman test indicates the the fixed effects and random effects models are not statistically different in our context.[24] The time trend coefficient γ is insignificant in the OLS and FE specifications and marginally significant in the RE specification. The latter coefficient is positive, meaning that bids increase in time. This is in contrast to Sunder (1992) where the prices of information goes to zero over time.

The above results lead to the conclusion that traders are willing to pay for information and bid higher for information in the Low treatment. While the fact the traders are willing to purchase information is in line with the conjectured strictly positive value of information, these results go against the strategic complementarity result that emerges from the theory. Further discussion follows in Section 4.1.3.

4.1.2 The *ex post* Value of Information from Stock Trading

The *ex post* performance difference between informed and uninformed traders is a natural measure for the value of information. We use stock-trading profits as a measure of trader performance.[25] In addition

[24]The p-value for the Null (difference in the coefficients are not systematic) is 0.476.

[25]For example, if the state is Orange and a buyer acquires one unit of the Stock in trading Period 1 for \$5 and one unit in trading Period 2 for \$3, and taking into account that the buyer valuation is equal to \$6 in that state, their stock-trading profit is equal to \$4=(\$6-\$5)+(\$6-\$3), etc.

to the stock-trading profit, the initially uninformed traders keep the remainder of the \$2 auction budget they did not spend. Thus, traders with non-winning bids have \$2 added to their payoff. The traders with winning bids have the difference between \$2 and the winning bid added to their accounts. The sum of the auction cash leftover and the stock-trading profit is referred to as the total payoff in the analysis.

The traders can be split into three groups depending on their information. The first group consists of the traders initially endowed with information, termed "Endowed." The second group of traders consists of those who are initially uninformed but successfully acquired information in the auction, termed "Acquired." The first two groups together are the informed traders, termed "Informed." The last group of traders consists of those who are initially uninformed and fail to acquire information. The *ex post* gross and net value of information can be captured by regressing the performance measures (stock-trading profit and total payoff) on the indicators of the information status.[26]

Table 5 Panel A reports the results of regressions with traders' performance as a dependent variable. The first three columns (1) to (3) use the stock-trading profit as performance and columns (4) to (6) use the total payoff (i.e., including the cash remaining from the auction stage). In the first column, the coefficient of "Informed" is significantly positive. This is consistent with the positive value of information found in the auction bids, reported in Section 4.1.1. The traders with information earn on average \$1.40 more than the uninformed. This value is similar in magnitude to the average auction bidding price of \$1.24. In column (2), we split the "Informed" into "Endowed" and "Acquired," and find that both groups generate higher stock-trading profit than the uninformed. Interestingly, those who acquire information in the auction seem to make better use of it (around \$0.50) than those who are endowed but the result is not statistically significant. In the third column, we test the treatment effect on the value of information using a difference-in-difference specification. Specifically, the indicator for Low

[26] we refer the difference in the stock-trading profit as the gross value of information because it does not consider the cost of acquiring it.

treatment T_L is interacted with the variable "Informed." The coefficient of the interaction term captures the difference of relative profits between informed and uninformed traders across the two treatments: $(\Pi_{informed} - \Pi_{uninformed}) \times T_L - (\Pi_{informed} - \Pi_{uninformed}) \times T_H$, where Π_i stands for the stock-trading profit for group i. The coefficient, 1.47, while positive, is not significant. In sum, the *ex ante* value of information is significantly higher in the Low treatment, while the *ex post* value is not statistically different across the treatments. The latter result suggests a possible lack of power, mainly because the noise in the posted offers during trading is a lot higher than the noise in the auction bids. Strategic complementarity in information acquisition requires that the coefficient be negative.

Columns (4) through (6) of Table 5 present the results of the same analysis as above but using the total payoff, equal to the sum of stock-trading profit and the cash remaining from the information auction, as the dependent variable. The coefficient of "Informed" in column (4) is small and not significant. This means that after accounting for the price of information, the final profits of informed and uninformed traders are not statistically different. The coefficient includes both the endowed and those who acquired information. Zooming in, column (5) splits the "Informed" into the previously defined two groups, "Endowed" and "Acquired," with the coefficient of the latter being essentially zero. Note that the only non-trivially affected coefficient in column (5) vs. column (2) is that of "Acquired." The constant must mechanically increase by the $2 that is added to the payments of all those who failed to acquire information. The "Endowed" do not have that amount added, so their coefficient mechanically decreases by $2. Concentrating on the coefficient for those who acquired information, it shows that the winning bids exactly offset the benefit from the information. This finding is important because it validates both the experimental design and the rationality of the participants. Column (6) reports the results of the difference-in-difference test. The coefficient of 1.5 on the interaction term is positive and not significant. Strategic complementarity would require that this coefficient be negative. It may seem puzzling that the lower (higher) cost of information in the High (Low) treatment does not translate into more (less) profit for the informed traders. However, the cost of information and the trading profits are jointly determined. We cannot establish the causal relationship in the current design.

One concern about the results reported in Table 5 Panel A is that the traders who successfully acquired information might be systematically different from those who failed to acquire. Thus, the performance difference might be driven by the unobserved heterogeneity between these two groups of traders. In unreported results, this concern is addressed by conducting the same analysis from above on a sub-sample of observations consisting of only the subjects who bid the highest and the second highest (i.e., barely lost) in the information auction, with all results remaining the same.[27]

We conduct several robustness checks on the results above. Since traders learn from their past trading experience and this can lead to adjustments in their usage of information, we perform additional testing where later replications are differentiated from the early ones. The results remain unchanged in that strategic complementarity fails to emerge and the coefficients of Table 5 remain qualitatively unchanged. The corresponding table that takes into account only the last five replications is presented in Online Appendix A. Since each unique individual in each experimental session sits through all treatments in the sequence provided by the experimenters, effectively this is a within-subject arrangement, and as such we investigate whether ordering effects influence the results. The indicator variables for the sequences High-High and High-Low (leaving out Low-High) are insignificant in all specifications. The results are presented in Online Appendix A.

To summarize, the *ex post* value of information is strictly positive based on the relative performance between informed and uninformed traders. Moreover, when the cost of information is taken into account, the performance difference is no longer statistically different. This suggests that the cost of

[27]The analysis, in the spirit of a regression discontinuity design, is based on the assumption that the highest and the second highest bidders have similar unobserved characteristics. In this sub-sample, the results confirm the main findings. For example, the traders who successfully acquired information earn about $2 more than the second highest bidder. This suggests that the results in Table 5 are not driven by the potential unobservable heterogeneity between winners and losers in the information auction.

information equals the benefit from it. Using the difference-in-difference specification, we do not find strategic complementarity, i.e., the value of information is not higher in the High treatment. Strategic complementarity crucially depends on rational behavior with respect to the long shots. The High treatment has more initially informed traders and therefore higher likelihood to encounter long shots. If traders under-estimate the losses due to encounters with long shots, they will bid too little in the auction in the High treatment.

In sum, the presented analyses in subsections 4.1.1 and 4.1.2 do not provide evidence to support Hypothesis 2.

4.1.3 Insights behind the Deviation from Strategic Complementarity

Below we provide some insights as to why strategic complementarity might be failing in the studied setup. The result that profits net of auction winner costs remain the same across Treatments implies that traders ought not to bid more when more participants are initially endowed with information, contrary to the theoretical claim of "complementarity" in information acquisition. This raises the issue: why is it that those who acquire information at lower prices than theory predicts (those in the High Information treatment) do not make more money? The analysis below aims at bringing some light to this question assuming that traders perhaps do not account for the impact of long shots initially but might learn and adjust their bidding with experience.

We examine if the experience in the previous trading replications affects the subsequent propensity to acquire information. If a trader gets a wrong signal through a trading encounter, the trader might learn from the experience and be more eager to acquire information in the subsequent replications. The driving force of strategic complementarity is to hedge against the downside from trading with a trader with an extreme wrong signal. With multiple replications, the existence of such learning mechanisms can be tested. The results are reported in Table 6 Panel A. In these regressions, the dependent variable is the bidding price submitted by the traders in the information auction. In column (1), the independent variable is "Wrong Signal", which is equal to 1 if a trader had a wrong signal in the previous replication. The coefficient is 0.04 and insignificant. This suggests the wrong signal does not lead to stronger incentive to acquire information in the subsequent auctions. In column (2) we run the similar regression with a dummy, "Traded with Longshot," which is equal to 1 if the trader transacted with a long shot in the previous replication. The coefficient is -0.28 and insignificant, implying that the traders do not seem to take into consideration not only the possibility but the actual encounter of a long shot when bidding for information. In column (3), a variable based on the difference between the conditional expected value and the true value of the asset is constructed. The idea is that traders who were affected by poor information would learn and seek information in subsequent replications. The coefficient for this variable is -0.003 and insignificant. In column (4) to (6), we interact the three variables from column (1) to (3) with the treatment dummy T_L to test for any differential effects between the treatments. All interaction terms are insignificant. To summarize, none of the three measures of "misinformation" have impact on the subsequent bidding prices during the information auction. This might explain why strategic complementarity is not detected, which requires traders to hedge ex ante the possibility of meeting a long shot.

Maybe traders do not account for their own history of being poorly informed because it does not affect their bottom line.[28] To this end, the trading profits of those who were eventually poorly informed due to an encounter with a long shot are analyzed. Specifically, Panel B of Table 6 focuses on the uninformed traders and tests whether encounters with long shots cause trading losses in the next period, and if those losses are higher in the High information treatment. The dependent variable is the Period 2 gains from

[28] In the High treatment, the chance of encountering a long shot is higher than in the Low treatment. In the extreme case when no one is informed, the probability of encountering a long shot is zero. The presence of long shots, therefore, provides for stronger incentives to acquire information in the High treatment.

trade of uninformed traders. In column (1) the dependent variable is regressed on a dummy that equals one for those who happened to trade with a long shot in Period 1. The coefficient suggests those who traded with a long shot lose about $1.9 more than other uninformed traders.[29] Thus, encountering a long shot does pose a significant negative impact on the future trading profits. When the "Traded with Longshot" is interacted with the indicator for Low treatment, T_L, the coefficient is 2.9, but not statistically significant with a t-statistic of 1.5 (based on the 240 observations). If this were significant, it would mean that the traders make more in the Low information treatment, i.e., that the loss due to trading with a long shot is smaller in the Low treatment. The downside risk from trading with a long shot, therefore, is higher in the High treatment as dictated by strategic complementarity. In contrast with the uninformed, when informed traders encounter a long shot, they do not lose money, as column (3) of Panel B displays. In short, the consequences of running into a long shot are severe and statistically significant. However, traders do not seem to be adjusting their bidding strategy based on their own experiences with long shots. While the experimental design eases the Bayesian task of information swapping after trading, the task of Bayesian updating and learning after unfavorable encounters is left to the traders and they seem to not properly account for it. In summary, the conjecture that participants do not know how to make the best use of the information they obtain through the auction, or when exactly this information is most valuable is not rejected by the data. A proper testing of the conjecture would require further study which can only be accomplished by a different experiment that focuses on misunderstanding of long shots.

Secondly, the other side of the coin about information being potentially mispriced is the high price for it in the Low information treatment. Why are traders willing to pay such high prices in the Low Treatment? Simulation analysis and consequent comparison with experimental results shine some light on this dual question. If no one is informed, the trading volume should be maximal as there are ex-ante gains from trade. In simulation analysis provided in Online Appendix A, it is demonstrated that

[29]The trade with a longshot does not result in losses, it is the subsequent trade of the long shot counterparty that causes losses.

the trading volume should be higher in the Low Information treatment that in the High Information treatment. The comparison of the experimental results with the simulations is as follows: if the "base volume" is 100% (everyone trades), the simulations under High treatment have the trading volume at 78%, while under Low treatment, it is 91%. The empirical results in our experiment are 55% and 46% for High and Low respectively. This suggest that when subjects acquire information they bid more aggressively (closer to their valuation) in the stock trading stage in comparison to when they do not have information. Table 7 shows the results of the analysis where the distance of traders' posteriors from their offer prices (known as shading) in the stock market is regressed against their informational status (informed or not) and on the treatment they were in (plus other control variables). The shading of $1.7 for the Uninformed over the Informed traders does not change as the Uninformed traders move between the High and Low Information Environments. This means that when the majority of traders are uninformed, traders miss out on gains from trading. This behavior provides additional explanation for the non-complementarity result: in the Low treatment unless one gets informed, trades are unlikely to happen. It therefore pays off to get informed more than the theory would otherwise predict, driving up the price in the Low treatment.

We would also like to point to the following that our experiment was not designed to test, but future research can address. There is experimental evidence, as summarized in Benjamin (2019), that participants form posteriors in non-Bayesian ways depending on the timing of the provision of information. In our setting, we cannot distinguish updating mistakes due to (i) timing of information delivery or (ii) processing of own information vs. information acquired through trading, as (i) and (ii) fully overlap. But both types of mistakes could explain why subjects update less to subsequent information after having been initially informed, which in turn could explain why they hedged less against longshots.[30]

[30] An analysis that addresses mistakes of type (ii) is included in Online Appendix B.

4.2 Efficiency

This subsection is devoted to the testing of the last hypothesis, done in two parts–testing separately informational and allocative efficiencies:

Hypothesis 3 *As trading progresses, prices become more informationally efficient. Allocation efficiency also improves with trading.*

4.2.1 Informational Efficiency

The conclusion from the series of papers on information percolation[31] is that prices eventually become fully revealing, a desirable outcome that was deemed one of the most important advantages of centralized markets over decentralized ones. The gradual convergence together with the opportunity to re-trade in decentralized markets make costly information acquisition worthwhile (with the feedback between prices and the decision to acquire information being incorporated into the traders' equilibrium strategies), and hence, free the decentralized market from the classic Grossman-Stiglitz paradox.

With $T=2$ trading periods, the decision to purchase information is analyzed with two measures of price efficiency across the periods.[32] The first measure is the deviation of the individual offer prices from the Stock value (different for buyers and sellers). The second measure is the difference of transaction prices from the average (across buyers and sellers) stock Value, which we call the Stock's "fundamental value." Based on the payoff structure of the asset, the fundamental value is \$15 when the state is Blue (\$17 for the buyers and \$13 for the sellers). When the state is Orange, the fundamental value is \$3 (\$0 for the sellers and \$6 for the buyers).

[31] Duffie and Manso (2007); Duffie, Malamud, and Manso (2009, 2014)

[32] The experimental design prevents testing whether prices eventually become fully revealing. While important, this question cannot be studied without first asking if traders do actually acquire information. The issue of price convergence is left to future research.

We perform a simple regression to test whether prices move closer to the respective fundamental values in the second trading period. The regression results are reported in Table 8. Column (1) reports the deviation of the order prices, while column (2) reports the deviation of transaction prices. The intercepts in the regression results in Table 8, at about \$4.20 for order prices and \$2.80 for transaction prices, indicate that order prices and transaction prices are far from fully revealing on average. At the same time, the absolute deviation of transaction prices from fundamental values decreases by about \$0.74 in the second period (\$0.65 for deviations of order prices), confirming the prediction that informational efficiency improves over time. This effect is statistically significant. As expected, the deviations in the Low information treatment compared to the High information treatment are much larger, by \$1.30 for order prices and \$1.60 for transaction prices. The prices for participants who receive counterfactual signals (long shots) deviate more from the fundamental value that the rest of the traders. The replication number does not have significant coefficient, suggesting that learning across replications is not the main driving force behind the efficiency improvements.

There can be a concern that the informational efficiency improvements documented here are somewhat mechanical since transactions in trading Period 1 lead to automatic information transmission.[33] Indeed, it is a desired property of any market mechanism that more information in the population should translate into more informative prices. While the result is trivial theoretically, it is not a foregone conclusion empirically (as demonstrated by the meta study of Page and Siemroth (2020)). A stronger test of the theory would have traders infer the information from the bid of their counter-party. The simplification of the Bayesian inference problem, however, is a major methodological feature, deliberately put in place to give the percolation theory the best chance. As mentioned before, price convergence should be investigated in a separate study.

[33]We thank Shimon Kogan for pointing this out.

4.2.2 Allocative Efficiency

An economy has achieved allocative (Pareto) efficiency when all gains from trade have been exploited. In the setup of the experiment, as well as in the theory (Duffie, Malamud, and Manso, 2014), gains from trade increase with the number of transactions, because buyers always have higher valuation for the Stock than sellers.

Examining gains from trade across the two trading periods and across treatments, the average number of transactions in trading Period 1 is 3.7, while it is 4.1 in trading Period 2, an increase of 15%. There is a difference in the number of transactions across treatments as well. In the Low treatment the average number of Period 1 transactions is 2.8 while in Period 2 it is 3.7. In the High treatment those numbers are 4.1 and 4.4, respectively.

To formally confirm the finding, Column (1) of Table 9 reports the results from a regression analysis. The dependent variable is the number of transactions, the independent variables are an indicator for Period 2, an indicator for the Low treatment, and finally an indicator for the Orange state. The results show a significant increase in the number of transactions in Period 2 in comparison to Period 1, as well as a significant difference across treatments, with fewer (marginally significant) transactions occurring in the Low treatment.

Since the number of transactions depends on the samples of information distributed to the participants, a second measure of allocative efficiency, called "propensity to trade," is implemented. The numerator is equal to the total number of successful transactions when the participants' actual bid (p_b) and offer(p_s) prices are considered and all possible hypothetical buyer-seller matches are formed.[34] The denominator is computed in the same way as the numerator, only that the traders' bids and offers are substituted with the traders' conditional expected Stock values based on their private signals, denoted p_b^* and p_s^* for

[34]In contrast, in the experiment only one such match is formed in each replication.

buyers and sellers respectively. The formula below describes this measure. For B buyers and S sellers, using b and s to index each buyer and seller, we have

$$Pr[Trade] = \frac{\sum_b \sum_s 1_{p_b > p_s}}{\sum_b \sum_s 1_{p_b^* > p_s^*}} \tag{4}$$

On average, the propensity to trade equals 0.57 in Period 1 and 0.75 in Period 2. The analysis with regards to this measure is presented in column (2) of Table 9. The propensity to trade is about 18% higher in the second period. The result is highly significant.

Both analyses imply that there is information transmission across periods, resulting in more transactions in Period 2. As noted above, more transactions immediately imply improved allocative efficiency in Period 2.

4.2.3 Correlation Between Informational and Allocative Efficiency

To investigate the empirical validity of the claim that informational efficiency and allocative efficiency are related, their correlation is analyzed in Table 10 using the efficiency measures constructed in section 4.2.1 and section 4.2.2. The analysis is performed at the replication level. Specifically, the left panel shows the results from the regression of the number of successful transactions per replication on the deviations of transaction prices (column(1)) or bid/ask mid-points (column (2)) from fundamental values. The fundamental value is \$3 (the midpoint between the seller's value of \$0 and the buyer's value of \$6) in the Orange state, and \$15 (the midpoint between the sellers' value of \$13 and the buyers' value of \$17) in the Blue state. The right panel (columns (2) and (3)) repeats the same regressions but there the dependent variable is the "Success Ratio."

The results show that lower deviations (i.e., higher price efficiency) lead to more transactions. For example, one standard deviation decrease in mis-pricing in the order price leads to 0.4 more transactions and a 5.36 percent increase in the propensity to trade. The results robustly show that higher informational efficiency in the experimental economy leads to better allocative efficiency.

4.2.4 Additional Discussion on Informational and Allocative Efficiency

Future work should compare the data obtained from decentralized markets like the ones in this study directly to a centralized market holding the other features of the design fixed.[35] While the present study establishes sufficiency of a decentralized market to induce information demand, the causal role of decentralization is only indirectly demonstrated, and such causality is extremely important from a market design perspective.

The contrast between the results in the Low and High treatment shows that the amount of information in the market matters. Ideally, one would like to have a breakdown of the Grossman-Stiglitz effect and the positive demand for information. Taking once again a market design perspective, such an answer would help assess whether the information purchases induced by decentralization is worth the costs to society in terms of decentralization fractions. This paper cannot provide a complete answer to this question. Theoretically, the experimental economy would be better off under no information (due to existing ex-ante gains from trade) than under moderate amounts of information.[36] That the release and purchasing of information can have overall negative effects has been addressed by Hirshleifer (1971).[37] The results in this paper show, however, that the release of information is less detrimental than envisioned in Hirshleifer (1971), due to the conservative bidding and the resulting low trading volume in the Low Information treatment in comparison to the theoretical benchmark. The deviation from theory is such

[35] We thank an anonymous referee for raising the points presented in this subsection.

[36] In the setup for this project, the allocational efficiency takes a U-shaped form (see Online Appendix A).

[37] Hirshleifer (1971) argues that if information is released before agents are able to trade, it has an adverse impact on risk sharing and as a result reduces welfare

that the economy with more information has not lower but higher welfare. Ideally, one would also have a treatment where no information is distributed. While the current design does have a "Full" treatment where everyone receives information, it is not a treatment in the usual sense of the word, since it was not counterbalanced against the High and Low treatments; it was only meant to calibrate participants' beliefs.

Additional partial answers to the issues raised above are provided in Asparouhova and Bossaerts (2017). There, information is free, however, and while the exchange mechanism is decentralized, it is organized as private continuous double auctions, and therefore very different from the mechanism employed here. However, that design allows for easy comparison between centralized and decentralized versions of the same market. The results from that study show that in decentralized markets efficiency ranges from as low as 0 and as high as 49% (measured as the reduction in risk and, in particular, as the reduction in the imbalance of a two-stock portfolio). Significantly, there appears to be no relationship between increases in allocational efficiency and the extent of information aggregation in pricing. Evidently, informational efficiency appears to be neither necessary nor sufficient for allocational improvements. These findings are contrasted with those obtained in an experiment where everything is the same, except that participants trade in a centralised market. The centralised market generates even less allocational efficiency: it ranges from a low of -28% to a high of only 30%.

The issues surrounding informational and allocational efficiency, and the relation to market design are extremely important topics, that have not received the attention they deserve, mainly because they require the study of counterfactuals. The advent of modern experimental techniques and software is rapidly removing the barriers to researching the important issues, of which the current and past research studies have only scratched the surface.

5 CONCLUSIONS

This paper presents the results from a series of controlled experiments designed to study the incentives to acquire costly information and the dissemination of this information in decentralized markets. The experimental design is inspired by the theory advanced in Duffie, Malamud, and Manso (2014). While the focus is on the particular model, the design of the experiment encompasses a broader set of trading structures. One important feature of the setup in Duffie, Malamud, and Manso (2014) that we adhere to, and one that facilitates the percolation of information, is that market participants are natural buyers and sellers, and each knows the role (buyer or seller) of the counter-party. It is reasonable to believe that such knowledge is present in functioning decentralized markets, e.g., a bank knows that Toyota is the natural seller of US dollars and buyer of Japanese Yen.

In order to address the costly acquisition of information in this study, the experimental design asks for many replications of the trading scenario under different informational treatments. This necessarily pushes the second question about prices "eventually" transmitting all of the dispersed information into the background. One aspect of markets that this study abstracts away from is the costly acquisition of connectivity, a crucial feature of the decentralized market dealer network (e.g., see Glode and Opp (2018) for related theoretical work on dealer network and Alfarano et al. (2020) for related experimental work). The model in Duffie, Malamud, and Manso (2014) allows for both information and connectivity acquisition. The latter will be the subject of future work; here, the focus is on information acquisition.

The first result is that as predicted by the theory, virtually all traders are willing to pay to acquire information. Traders not only bid aggressively for information but also manage to use the acquired information in subsequent trading to recuperate the costs.

With two treatments, we test a crucial feature of information acquisition, namely that of strategic complementarity. Strategic complementarity requires that the value of information be greater when more traders are informed. The second result is that when an *ex ante* measure of value of information is considered, this value is greater when fewer agents are informed, opposite to the theoretical prediction. When an *ex post* measure is considered, the value of information is not statistically different between the High and Low information treatments. Taken together, those results suggest that strategic complementarity does not emerge in this setting. This finding is likely due to the inability of participants to fully comprehend the subtle martingale property of the likelihood ratio. The profits of traders who encounter long shots are lower than those who do not, but they seem to fail to account for this in the *ex ante* stage when bidding for information. The traders exhibit overweighting of own informationa and underweighting the information provided to them after a transaction. This behavior can also serve as a partial explanation for why traders do not seem eager to hedge against encounters with traders who have the wrong signal.

The third result is about the link between informational and allocative efficiency. We find that price quality improves in Period 2. Allocative efficiency also improves significantly over the two periods, and is higher when informational efficiency increases. In sum, the predictions of the theory are largely upheld in the experiment except for the strategic complementarity.

While directly related to the theory of information percolation, our results have some policy relevance. In particular, this paper provides the first evidence-based research on the incentives to acquire costly information in dark markets. We show that the decentralized environment is conducive to costly information acquisition. In turn, the prices in such markets can eventually be more informative. Thus, the often heard pessimism about decentralized markets may be unfounded.

References

Alfarano, S., A. Banal-Estanol, E. Camacho-Cuena, G. Iori, and B. Kapar. 2020. Centralized vs decentralized markets in the laboratory: The role of connectivity. *Working Paper* .

Asparouhova, E., and P. Bossaerts. 2017. Experiments on percolation of information in dark markets. *The Economic Journal* 127:F518–44.

Benjamin, D. J. 2019. Errors in probabilistic reasoning and judgment biases. *In B. Douglas Bernheim and Stefano DellaVigna and David Laibson, (eds), Handbook of Behavioral Economics - Foundations and Applications 2 North-Holland: Elsevier* (pp. 69–186).

Biais, B., D. Hilton, K. Mazurier, and S. Pouget. 2005. Judgemental overconfidence, self-monitoring, and trading performance in an experimental financial market. *The Review of Economic Studies* 72:287–312.

Breon-Drish, B. 2011. Asymmetric information in financial markets: Anything goes. *Working Paper* .

Cetina, K., Karin, and U. Bruegger. 2002. Global microstructures: The virtual societies of financial markets. *American Journal of Sociology* 107:905–50. doi:10.1086/341045.

Chamberlin, E. H. 1948. An experimental imperfect market. *The Journal of Political Economy* 56:95–108.

Chen, Y., and Y. He. 2021. Information acquisition and provision in school choice: An experimental study. *Games and Economic Behavior* 197:1–32.

Corgnet, B., M. Desantis, and D. Porter. 2018. What makes a good trader? on the role of intuition and reflection on trader performance. *The Journal of Finance* 73:1113–37.

Corgnet, B., M. DeSantis, and D. Porter. 2020. The distribution of information and the price efficiency of markets. *Journal of Economic Dynamics and Control* 110:103671–. ISSN 0165-1889. doi:https://doi.org/10.1016/j.jedc.2019.02.006.

Doob, J. L. 1953. *Stochastic processes*. John Wiley & Sons, New York.

Duffie, D., N. Gârleanu, and L. H. Pedersen. 2005. Over-the-counter markets. *Econometrica* 73:1815–47.

Duffie, D., G. Giroux, and G. Manso. 2010. Information percolation. *American Economic Journal: Microeconomics* 2:100–11.

Duffie, D., S. Malamud, and G. Manso. 2009. Information percolation with equilibrium search dynamics. *Econometrica* 77:1513–74.

———. 2014. Information percolation in segmented markets. *Journal of Economic Theory* 153:1–32.

Duffie, D., and G. Manso. 2007. Information percolation in large markets. *The American Economic Review* 97:203–9.

Glode, V., and C. C. Opp. 2018. Over-the-counter vs. limit-order markets: The role of traders' expertise. *Working Paper* .

Grossman, S. 1976. On the efficiency of competitive stock markets where trades have diverse information. *The Journal of Finance* 31:573–85.

Grossman, S., and J. Stiglitz. 1980. On the impossibility of informationally efficient markets. *The American Economic Review* 70:393–408.

Grundy, B. D., and McNichols. 1989. Trade and the revelation of information through prices and direct disclosure. *The Review of Financial Studies* 2:495–526.

Halim, E., Y. E. Riyanto, and N. Roy. 2019. Costly information acquisition, social networks, and asset prices: Experimental evidence. *The Journal of Finance* 74:1975–2010. doi:10.1111/jofi.12768.

Kendall, C. 2020. Market panics, frenzies, and informational efficiency: Theory and experiment. *American Economic Journal: Microeconomics* 12:76–115.

Manzano, C., and V. Xavier. 2011. Public and private learning from prices, strategic substitutability and complemen-tarity, and equilibrium multiplicity. *J. Math. Econ.* 47:346–69.

Meyer, D. J. 2003. *The economics of risk*. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

Page, L., and C. Siemroth. 2017. An experimental analysis of information acquisition in prediction markets. *Games and Economic Behavior* 101:354–78.

———. 2020. How much information is incorporated in financial asset prices? experimental evidence. *Review of Financial Studies* 34:4412–49.

Plott, C., and S. Sunder. 1988. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica* 56:1085–118.

Plott, C. R., and V. L. Smith. 1978. An experimental examination of two exchange institutions. *The Review of Economic Studies* 45:133–53.

Preda, A. 2012. Tags, transaction types and communication in online anonymous markets. *Socio-Economic Review* 11:31–56. doi:10.1093/ser/mws011.

Saavedra, S., K. Hagerty, and B. Uzzi. 2011. Synchronicity, instant messaging, and performance among financial traders. *Proceedings of the National Academy of Sciences* 108:5296–301.

Smith, V. L. 1962. An experimental study of competitive market behavior. *Journal of Political Economy* 70:111–37. ISSN 00223808, 1537534X.

Sunder, S. 1992. Market for information: Experimental evidence. *Econometrica: Journal of the Econometric Society* 60:667–95.

Wolinsky, A. 1990. Information revelation in a market with pairwise meetings. *Econometrica* 58:1–23.

Tables

Table 1. Payoff Structure of the Stock

	Buyers	Sellers
Orange	$6	$0
Blue	$17	$13

Table 2. Session Information

Panel A. Summary Session Information

session	N. Subjects	N. Replications	N. Initially Informed Traders
2017-02-22	16	10	(4,12)
2017-03-30	18	10	(3,13)
2017-04-06	14	11	(3,12)
2017-05-30	16	10	(4,12)
2017-06-05	16	10	(4,12)

Panel B. Detailed Session Information

Session	Number of Subjects	Round Number	Treatment	Number of Initially Informed	Number of Signals Auctioned
2017-02-22	16	1	Full	16	0
2017-02-22	16	2	High	12	3
2017-02-22	16	3	Low	4	1
2017-02-22	16	4	High	12	3
2017-02-22	16	5	Low	4	1
2017-02-22	16	6	Full	16	0
2017-02-22	16	7	High	12	3
2017-02-22	16	8	Low	4	1
2017-02-22	16	9	High	12	3
2017-02-22	16	10	Low	4	1
2017-03-30	18	1	Full	18	0
2017-03-30	18	2	Low	3	1
2017-03-30	18	3	High	13	3
2017-03-30	18	4	Low	3	1
2017-03-30	18	5	High	13	3
2017-03-30	18	6	Full	18	0
2017-03-30	18	7	Low	3	1
2017-03-30	18	8	High	13	3
2017-03-30	18	9	Low	3	1
2017-03-30	18	10	High	13	3
2017-04-06	14	1	Full	14	0
2017-04-06	14	2	High	11	4
2017-04-06	14	3	Low	2	1
2017-04-06	14	4	High	11	4
2017-04-06	14	5	High	11	4
2017-04-06	14	6	Low	2	1
2017-04-06	14	7	High	11	4
2017-04-06	14	8	High	11	4
2017-04-06	14	9	Low	2	1
2017-04-06	14	10	High	11	4
2017-04-06	14	11	High	11	4
2017-05-30	16	1	Full	16	0
2017-05-30	16	2	High	12	4
2017-05-30	16	3	High	12	4
2017-05-30	16	4	Low	4	1
2017-05-30	16	5	High	12	3
2017-05-30	16	6	High	12	3
2017-05-30	16	7	Low	4	1
2017-05-30	16	8	High	12	3
2017-05-30	16	9	High	12	3
2017-05-30	16	10	Low	4	1
2017-06-05	16	1	Full	16	0
2017-06-05	16	2	High	12	3
2017-06-05	16	3	High	12	3
2017-06-05	16	4	Low	4	1
2017-06-05	16	5	High	12	3
2017-06-05	16	6	High	12	3
2017-06-05	16	7	Low	4	1
2017-06-05	16	8	High	12	3
2017-06-05	16	9	High	12	3
2017-06-05	16	10	Low	4	1

Table 3. Summary of Statistics for Auction Bidding Prices

Panel A reports the summary statistics of bidding prices in the information auction. High treatment is when we have a majority of the participants initially endowed with information, while in the Low treatment the majority of the participants start uninformed. In both cases the uninformed traders are given the opportunity to acquire information in the auction. Panel B reports the paired t-test of auction bidding prices. For each participant who experiences both treatments, we compute the average bidding price per treatment and compare through a paired t-test. Panel C reports the frequency and percentage of the positive bids across treatments

Panel A. All Auction Bids

Treatment	mean	std	median	p10	p25	p75	p90
High	1.08	0.70	1.17	0.01	0.45	1.71	2.00
Low	1.32	0.64	1.49	0.37	0.97	1.95	2.00
Whole Sample	1.24	0.67	1.36	0.10	0.75	1.89	2.00

Panel B. Paired T-Test

	Value
Mean Differences (High minus Low)	-0.19
t-statistic	-3.62
P-value	0.0006
Degrees of Freedom	63

Panel C. Positive Bids

Treatment	Number of Auction Participants	Number of Positive Bids	Percentage
Low	202	190	94.06%
High	99	90	90.91%

Table 4. Treatment Effect on Auction Bidding

This table reports regressions that identify the treatment effect on bidding prices in the information acquisition auction. The independent variables in both regression models are the bidding prices submitted in the information auction. Column (1) reports the standard OLS regression. Column (2) reports the within estimator where the individual fixed effects are included. Column (3) reports the between estimator based on Random Effects. T_H is a dummy for the High treatment. "Replication #" is the replication number. All standard errors are clustered at the session level. *p<0.1; **p<0.05; ***p<0.01

	Dependent variable: bidding price		
	OLS	*Fixed Effects*	*Random Effects*
	(1)	(2)	(3)
T_H	-0.2310^{***}	-0.2039^{***}	-0.2062^{***}
	(0.0726)	(0.0482)	(0.0476)
Replication #	0.0089	0.0144	0.0133^{*}
	(0.0158)	(0.0105)	(0.0111)
Intercept	1.2599^{***}		1.1874^{***}
	(0.1471)		(0.1066)
Num. obs.	301	301	301
R^2 (overall)	0.0218	0.0276	0.0278

Table 5. Traders' Performance

We regress the payoffs to the participants on a set of variables to examine the effect of information acquisition on profitability. The dependent variables in columns (1)-(3) are the profit extracted purely from stock trading, thus excluding any cash remaining from the information auction. The dependent variables in columns (4)-(6) are the total payoff that includes the cash remaining from information auction (i.e., total earnings minus cost of information acquisition). "Informed" is an indicator for the participants who either initially receive information for free or acquire information in the information auction. "Endowed" indicates participants who are initially endowed with information for free. "Acquired" is an indicator for the (initially uninformed) participants who successfully acquire information. The standard errors are clustered at the trader level. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Dep. Var.	Stock Trading Profit (1) - (3)			Total Payoff (4) - (6)		
	(1)	(2)	(3)	(4)	(5)	(6)
Informed	1.3973***		0.6736	−0.5668		−1.3002*
	(0.4561)		(0.7493)	(0.4561)		(0.7496)
Endowed		1.3108***			−0.6892	
		(0.4992)			(0.4992)	
Acquired		1.8022**			0.0062	
		(0.7465)			(0.7548)	
T_L			−0.8459			−0.8459
			(0.8892)			(0.8892)
Informed $\times T_L$			1.4732			1.5131
			(1.1234)			(1.1237)
Constant	1.6018***	1.6018***	2.1728***	3.6018***	3.6018***	4.1728***
	(0.3786)	(0.3788)	(0.6531)	(0.3786)	(0.3788)	(0.6531)
Observations	700	700	700	700	700	700
Adjusted R^2	0.0091	0.0082	0.0087	0.0003	−0.00001	0.00004

Table 6. Mechanisms on Information Acquisition Incentives

We explore the effect of meeting a long shot on the trading profit. We also test the effect of receiving a wrong signal on subsequent information acquisition decisions. The dependent variable in Panel A is the bidding price submitted during the information auction. The dependent variables in Panel B are the stock trading profit during the second period for uninformed traders (columns 1-2) and informed traders (columns 3-4). “Wrong Signal” is a dummy variable indicating whether the trader had the signal indicating the opposite state of the world in the previous replication. “Traded with Longshot” indicates if a trader traded with a long shot in the previous round (Panel A) or in the first period (Panel B). $|EV - FV|$ is the absolute difference between the conditional expected value of the true fundamental value, where EV is computed based on the trader’s posterior probabilities according to Equation **??**. T_L is a dummy for Low treatment. The standard errors are clustered at the trader level. *p<0.1; **p<0.05; ***p<0.01

Panel A. Subsequent Information Acquisition Incentives

	Auction Bid Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Wrong Signal	0.0440			0.1789		
	(0.0843)			(0.1870)		
Traded with Longshot		−0.2858			−0.3436	
		(0.2033)			(0.3590)	
$\|EV - FV\|$			−0.0033			−0.0240
			(0.0111)			(0.0213)
Wrong Signal $\times T_L$				−0.1985		
				(0.2180)		
Traded with Longshot $\times T_L$					0.0867	
					(0.4322)	
$\|EV - FV\| \times T_L$						0.0293
						(0.0264)
T_L	0.2327***	0.2330***	0.2304***	0.2735***	0.2295***	0.1180
	(0.0742)	(0.0738)	(0.0734)	(0.0924)	(0.0748)	(0.1354)
Constant	1.0733***	1.0938***	1.0958***	1.0461***	1.0961***	1.1803***
	(0.0947)	(0.0920)	(0.0990)	(0.1048)	(0.0928)	(0.1261)
Observations	301	301	301	301	301	301
Adjusted R^2	0.0212	0.0276	0.0207	0.0212	0.0245	0.0212

Panel B. Second Period Profit After Meeting a Long Shot

	Dependent variable: Second Period Profit			
	Uninformed Traders		Informed Traders	
	(1)	(2)	(3)	(4)
Traded with Longshot	−1.9141*	−3.7059*	−0.2852	−0.1630
	(1.0355)	(1.9593)	(1.3708)	(1.3665)
T_L		−0.3204		0.4898
		(0.5209)		(0.3931)
Traded with Longshot $\times T_L$		2.9454		
		(1.9504)		(0.0000)
Constant	0.8641***	1.0809***	1.6420***	1.5198***
	(0.2659)	(0.3510)	(0.1879)	(0.1814)
Observations	240	240	460	460
Adjusted R^2	0.0004	−0.0045	−0.0020	−0.0010

Table 7. Bidding Aggressiveness

This table reports how traders with no information at the time of trading submit orders. The dependent variable is the absolute difference between the trader's posterior implied expected value and the order price submitted by the trader. *Uninformed* indicates the traders who failed to acquire information in a given replication. Other variable definitions are the same as in previous tables. Standard errors are clustered at the session level. *p<0.1; **p<0.05; ***p<0.01

	Dependent variable: $\lvert EV - OrderPrice \rvert$		
	(1)	(2)	(3)
Uninformed	1.7044***	1.6870***	1.5775***
	(0.1804)	(0.2638)	(0.2988)
T_L		−0.0202	−0.0068
		(0.2162)	(0.2008)
Uninformed $\times T_L$		0.0389	0.0624
		(0.4330)	(0.4219)
Seller			0.8870***
			(0.1988)
Second Period			−0.3430***
			(0.1003)
Replication #			0.0231
			(0.0237)
Longshot			−0.2269
			(0.2150)
O State			0.1361
			(0.1020)
Constant	3.3525***	3.3567***	3.2717***
	(0.1826)	(0.2153)	(0.5112)
Observations	1,592	1,592	1,372
Adjusted R^2	0.0731	0.0719	0.1075

Table 8. Informational Efficiency

In this table we regress the deviations from fundamental value of the stock on a set of variables. The FV (fundamental value) is \$15 in Blue state and \$3 in Orange state, i.e. the average payoff across buyers and sellers (see Table 1). The dependent variable in columns (1) and (3) is the absolute difference between the order price and the fundamental value. The dependent variable in the column (2) is the absolute difference between the transacted price and the fundamental value. "Second Period" indicates the second stock trading period within each replication. "Seller" indicates the sell orders. "Replication #" is the replication number within each session. "Longshot" will take value of 1 if either buyer or seller is a Longshot. "O State" is 1 if the state is orange and 0 if it's blue. T_L is a dummy indicating the Low treatment. Standard errors are clustered at the session level.*p<0.1; **p<0.05; ***p<0.01.

	Dependent variable:	
	|FV − Order Price|	|FV − Transaction Price|
	(1)	(2)
Second Period	-0.6785^{***}	-0.7725^{***}
	(0.1234)	(0.1639)
Seller	0.7716^{**}	
	(0.3551)	
Replication #	0.0291	0.0865
	(0.0353)	(0.0703)
Longshot	5.7498^{***}	4.5944^{***}
	(0.7269)	(0.7496)
O State	1.0489^{***}	1.0546^{**}
	(0.2279)	(0.4929)
T_L	1.3032^{***}	1.5468^{***}
	(0.1653)	(0.2397)
Intercept	4.2380^{***}	2.8718^{***}
	(0.6568)	(0.6415)
Num. of Obs.	1,592	413
Adjusted R^2	0.1092	0.1839

Table 9. Allocative Efficiency

In this table we examine whether allocative efficiency is improved. The dependent variable in column (1) is number of transactions in each trading period. The dependent variable in column (2) is the ratio between number of simulated successful transactions based on actual submitted order prices relative to the number of simulated successful transactions based on values (prices) implied by the correct posterior for each trader. To derive this measure, we exhaust all possible matches between buyers and sellers in each trading period. Then we count, out of all possible matches, the number of successful transactions according to the actual submitted order prices, and the number of successful transactions according to values implied by the correct posterior (i.e., when traders submit orders that are equal to the expected value of the stock according to their private signals). Then we take the ratio between the former and the latter."O state" indicates Orange state. Seller indicates the order is a selling order. "Second Period" indicates the second period of stock trading. "Replication #" is the replication number within each experiment session. T_L indicates the Low treatment. All standard errors are clustered at session level.

	Dependent variable:	
	N. of Transactions	Propensity to Trade
	(1)	(2)
T_L	-0.6968	-0.1091^{*}
	(0.4770)	(0.0582)
Second Period	0.6471^{***}	0.1792^{***}
	(0.1433)	(0.0406)
O State	-0.1928	-0.0241
	(0.1792)	(0.0343)
Intercept	4.0523^{***}	0.6198^{***}
	(0.1636)	(0.0376)
Observations	102	102
Adjusted R^2	0.0813	0.2407

Table 10. Correlation between Informational and Allocative Efficiency

In this table we examine the correlation between informational and allocative efficiency. The dependent variable in column (1) and (2) is the number of transactions in each trading period. The dependent variable in column (3) and (4) is the propensity to trade, computed as the ratio between number of simulated successful transactions based on actual submitted order prices relative to the number of simulated successful transactions based on posterior implied prices for each trader (see details in table 9). |FV − Order Price| is the average value (per replication) of the absolute difference between the fundamental value and the order prices submitted by traders. |FV − Transaction Price| is the average value (per replication) of the absolute difference between the fundamental value and the actual transaction prices. The fundamental value is the average payoff on the stock to buyers and sellers depending on the states: $15 in the Blue state and $3 in the Orange state. Standard errors are clustered at the session level. *p<0.1; **p<0.05; ***p<0.01.

	Dependent variable:			
	N. of Transactions		Propensity to Trade	
	(1)	(2)	(3)	(4)
|FV − Order Price|	-0.4149^{***}		-0.0573^{**}	
	(0.1333)		(0.0249)	
|FV − Transaction Price|		-0.1787^{*}		-0.0330^{*}
		(0.0926)		(0.0178)
Constant	6.0894^{***}	4.6803^{***}	0.9453^{***}	0.7810^{***}
	(0.7336)	(0.3815)	(0.1085)	(0.0553)
Observations	51	51	51	51
Adjusted R^2	0.2314	0.0633	0.2191	0.1380

Figure 1: Histogram of Bids in the Information Auction

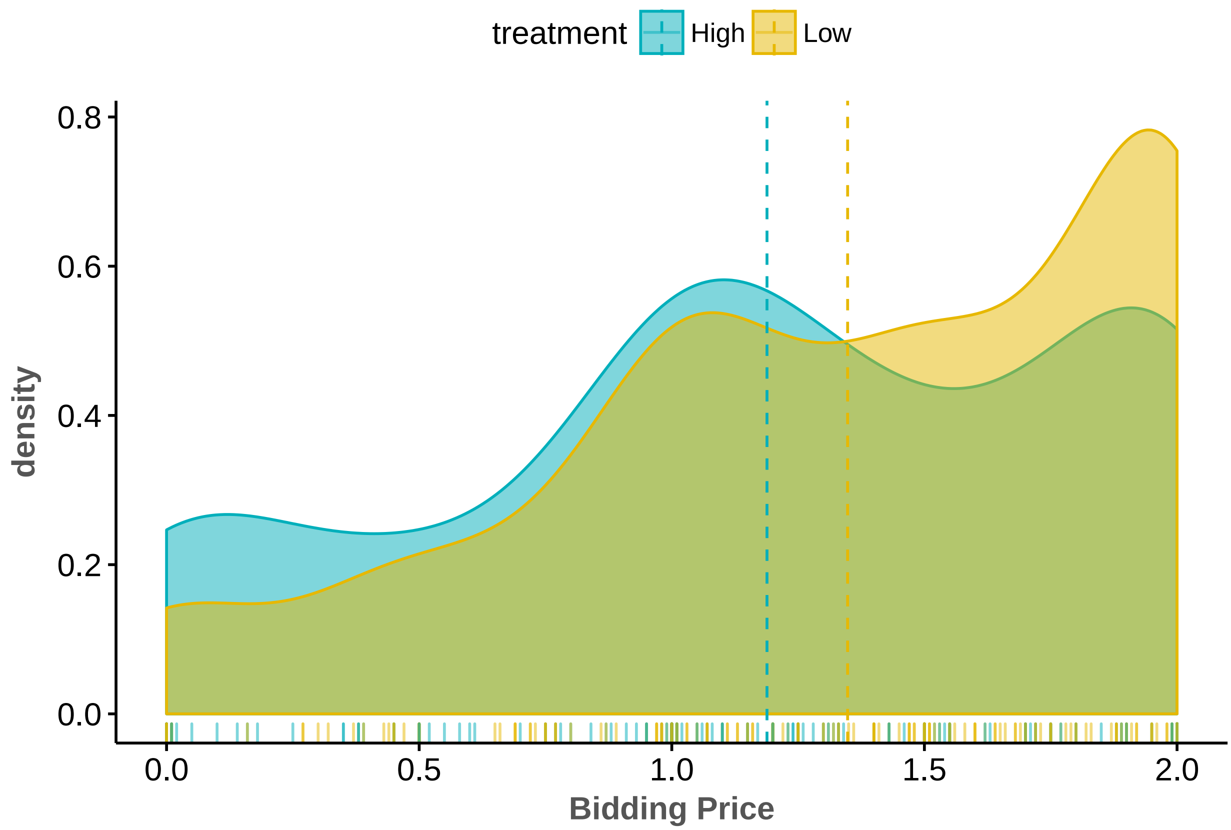

Figure 2: Pairwise Comparison of Auction Bids