Costly Information Acquisition in Decentralized Markets: An Experiment

Elena Asparouhova

University of Utah

Peter Bossaerts

University of Melbourne

Wenhao Yang

University of South Carolina

September 14, 2017

The authors thank Darrel Duffie, Matthew Jackson for useful discussion. Financial support from the National Science Foundation (Asparouhova, Bossaerts: SES-1426428), 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.

Abstract

In a series of controlled laboratory experiments, we test the rationality of the decisions to purchase information, the informational efficiency of prices and the optimality properties of the resulting allocations in a decentralized markets. The theory predicts that markets with dispersed information and identifiable 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. We find statistically significant improvements for both price efficiency and allocational efficiency in the course of the experiment. The participants remain willing to pay for information, in contrast with centralized markets, where information effectively aggregates very fast (Plott and Sunder (1988)), and where the information auction price drops to zero (Sunder (1992)).

An over-the-counter (OTC) marketplace allows two parties to contract for an immediate or a future date trade and to leave the details of the contract known only to them, at least for a period of time. The most frequent criticism of the OTC market has been its opaqueness and for that reason it has been labelled as a "dark market."[1] The proponents of the OTC markets have put forward the importance of allowing flexibility in the contract specifications as to best suit the counter-parties' risk exposure, a feature that standardized exchange contracts do not allow for.[2] While there is no unanimous agreement on the costs and benefits of the OTC, the EU will be operating under new rules agreed in the revision to the Markets in Financial Instruments Directive and new Regulation (MiFID II/MiFIR), set to take effect from 3 January 2018. "Its main thrust will be to force trading across all asset classes into open and transparent markets–not just equities, the focus of MiFID 1, or derivatives, the focus of EMIR's clearing rules" (The Economist, A Bigger Bang, Apr 26th 2014).

However, the movement to centralized markets is not without its drawbacks as those are under the spell of the Grossman-Stiglitz paradox. Grossman (1976) and Grossman and Stiglitz (1980) argue that perfectly informationally efficient markets are impossible. The argument boils down to the following: "because 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.

Dark markets, on the other hand, most often modeled as a game of decentralized bilateral trading, can loosen the grip of the paradox. The properties of the equilibrium outcome as well as the convergence dynamics and the incentives of the traders to reveal their information along the way highly depend on the informational structure and market design of the trading game. In all cases, however, the speed of dissemination of information is slower than that in centralized markets and thus incentives for information collection

[1] McCrank, John (6 April 2014), Dark markets may be more harmful than high-frequency trading, New York: Reuters, retrieved 12 April 2014

[2] Chapter 3 in Understanding Derivatives: Markets and Infrastructure, by Richard Heckinger, vice president and senior policy advisor, Ivana Ruffini, senior policy specialist, financial markets, and Kirstin Wells, vice president and risk officer.

are restored.[3] This sometimes comes at the cost of the equilibrium not being perfectly revealing (in a similar way that centralized markets restore incentives for information collection–through noisy prices) or not having all gains from trade realized.

The theoretical literature, concerning both centralized and decentralized markets, can be further divided into two categories based on the underlying information conditions. To distinguish between the two, we refer to the case of dispersed information as "information aggregation" (or "information percolation" when specifically addressing the decentralized setting), and to the case where insiders share identical information as "information amplification."

The first strand of literature, initiated by with the seminal paper of Wolinsky (1990), concerns an environment of information amplification, whereby all informed traders initially possess an identical signal and have the option to trade a single indivisible good at given prices, an assumption later relaxed by other models.[4] The steady state equilibria are shown to exhibit inefficiencies as not all gains of trade are realized.[5] The non-fully revealing prices help the markets escape the Grossman-Stiglitz paradox, but the cost of that is the resulting allocational inefficiencies.

The second strand of literature concerns the aggregation of dispersed information. This literature, initiated by Duffie, Gârleanu, and Pedersen (2005), models OTC 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 informational knowledge and connectivity with one another, while still being able to characterize the dynamics of posterior beliefs in closed form solutions. The authors argue that the negative

[3] See, for example, Wolinsky (1990), Blouin and Serrano (2001), Duffie and Manso (2007), Duffie, Giroux, and Manso (2010), Duffie, Malamud, and Manso (2009), and Duffie, Malamud, and Manso (2014), Ostrovsky (2012), Golosov, Lorenzoni, and Tsyvinski (2014), and Asriyan, Fuchs, and Green (2017).

[4] Golosov, Lorenzoni, and Tsyvinski (2014) relax the indivisibility of the traded good and the pricing rule assumptions to obtain the result that, "in the long run, the equilibrium converges to an ex post Pareto efficient allocation and the value of information goes to zero."

[5] The ability of centralized markets to amplify information and obtain prices that are fully revealing is a well-established prediction of the rational expectations equilibrium (Radner (1979)). Wolinsky's propositions were tested (indirectly) in Bossaerts, Frydman, and Ledyard (2014).

attitude towards decentralized exchange may have been premature. Specifically, when information is dispersed among participants, decentralized exchange can eventually fully aggregate the information. A critical driving force behind the full revelation result is that information gathering is subject to strategic complementarity. This means that if a trader knows that others pay for information, it is in this trader's own interest to pay for information as well.

In sum, the small but rapidly growing theoretical literature on decentralized markets has made significant advances in the understanding of the informational and allocational properties of the resulting equilibria. In sharp contrast with centralized markets, decentralized dark markets may escape the Grossman-Stiglitz paradox, and for that reason, ruling them out might be suboptimal. However, all of the equilibrium properties of the models within this literature depend on carefully chosen assumptions and it remains an empirical question whether these assumptions are upheld in the markets populated by humans, whether the desired equilibria are reached, and whether there is a robust set of results that holds across setups.

We argue that experiments are well, if not the best, suited to address many of those outstanding questions. The power of the methodology is in providing the researchers with the ability to closely mimic theoretically modeled markets while having the freedom to switch on and off certain market features that modulate the distance between the mathematical models and real markets. As with any other methodology, and similar to experiments in other fields of science, the purpose is to provide a bridge between theory and field analysis, in large part through the lenses of the robust findings that emerge from the experimental analyses.

Market experiments have long denied support to the argument that decentralized trading mechanisms can deliver the competitive equilibrium outcome (Chamberlin (1948)), even in the absence of asymmetric information. More generally, since the seminal papers of Smith (1962) and Plott and Smith (1978), the view in experimental economics has been that a (particular form of) centralized market is needed to produce the competitive equilibrium outcome and its welfare merits.

Early experiments with centralized financial markets have confirmed their capability to amplify information (Plott and Sunder (1988)). Under asymmetric information, those markets also display the quintessential properties of the Grossman-Stiglitz paradox (Sunder (1992)). Namely, when information is auctioned to participants, information prices quickly plummet to zero.

Experiments with centralized markets and information aggregation have shown mixed results. It appears that the ability of prices to aggregate information depends both on the complexity of the signals and the common knowledge of the structure of the traders' private values. For example, Plott and Sunder (1988) study an environment where the traders are informed about the structure of the signal but not about the structure of the private values of the market participants. In this scenario prices do not reveal the underlying information. In cases where the structure of payoffs is given to all traders but the signal is complex[6], like in a treatment of Plott and Sunder, replicated by Corgnet et al. (2015) and Biais et al. (2005), the prices once again fail to aggregate the underlying information.

Here we revisit the long standing claim about the inability of decentralized market to achieve efficiency. We also address the question about the individual decision to invest in information gathering. To the best of our knowledge, this is the first experimental study of costly information acquisition in a decentralized setting. The conclusions of past experimental studies have likely changed because of the drastic increase in speed and connectivity in modern markets that could not be envisioned even a decade ago. Researchers have only recently been able to adapt their methodologies in order to reflect these aspects of decentralized markets. For this and other series of studies, we have developed and made publicly available a software that mimics a lot of those crucial market features. Traders connect via bilateral computer connections and each connection maintains an order book between the two counter-parties, visible and accessible only by them. Order submissions and cancelations are continuous, while trading within the software can be either continuous or via batch call auctions.[7]

[6]The signal is complex because under aggregation the economy is one without aggregate risk while out of equilibrium there can a substantial amount of risk in the economy.

[7]See www.flexemarkets.com.

Because of the technological challenges experimental papers on decentralized markets are all but missing. In a recent study, Asparouhova and Bossaerts (2016) use controlled laboratory experiments to test the information percolation theory in a decentralized trading setting where information can be freely obtained. They show that prices aggregate the available information in a decentralized market but not in the strict sense of the theory - the latter would require that transaction prices converge exponentially fast to expected payoffs conditional on the aggregate information. Prices instead fluctuate within narrow no-arbitrage bands centered around the average of the private signals. The speed with which the average trade price aggregated information may have been slowed down by the following features of the environment: (i) ability to submit orders to multiple counter-parties at once, (ii) ability to re-engage with a counter-party at a later time in the absence of an immediate response (counter-bid).

In line with the categories in the theoretical literature, we divide the experimental studies of decentralized markets into two categories: information amplification and information aggregation. After a further bifurcation, the studies can be separated into free and costly information acquisition. Asparouhova and Bossaerts (2016) concerns information aggregation with free information. Faced with a field of open questions, we take the information aggregation setup and ask if investors are willing to purchase information. We argue that this is the most pressing of all questions. The traders' willingness to acquire costly information is a necessary condition for informationally efficient prices to exist. We also ask if prices are informationally efficient, and if allocations exhibit the desired optimality properties of centralized markets.

While we focus on the model developed in Duffie, Malamud, and Manso (2014), the design of the experiment necessarily encompasses a broader set of trading structures than those envisioned in the theory. One important feature of this setup is that market participants are natural buyers and sellers, and each knows the role of the counter-party. We choose this setup as we think such knowledge is present in the functioning OTC markets.[8]

[8]This requirement to know the role of one's counter-party corresponds to the requirement that a trader is aware of the structure of the payoffs in the economy, discussed above in the context of the ability of centralized markets to achive informational efficiency.

For example, when a company like Nissan contacts the bank, the bank knows that they are the natural sellers of US dollars and buyers of Japanese Yen.

In order to address the costly acquisition of information, the experimental design asks for many replications of the trading scenario under different informational treatments. This necessarily brings the second question about prices "eventually" transmitting all of the dispersed information into the background. This study will only be able to provide partial answer to this question due to the limited number of trading rounds that each trading replication can accommodate.[9]

We find that price efficiency, while initially low, improves significantly with trading. Moreover, we observe significant increase in the number of trades as information is exchanged through trades. With increased number of trades, we necessarily observe higher gains from trades, and thus an increased allocational efficiency. Participants with no initial information remain willing to pay for information, in contrast with the centralized markets in Sunder (1992). The theory predicts that the desire to invest in information gathering exhibits strategic complementarity. However, this result does not seem to emerge in our experiments. Instead of bidding more aggressively for information when there are many informed traders, the uninformed traders submit higher bids when only a few traders are informed.

Comparing the gross payoffs between informed and uninformed traders, the ones who successfully acquire information outperform the ones who do not. Once the payoffs are adjusted for the cost of information, there is no statistically detectable difference in performance. Thus, the break with complementarity does not appear to be due to irrationality given the experimental market outcomes. Further details on how it is possible that the complementarity that drives the equilibrium predictions does not emerge but no one seems to be "hurt" is provided in the Results section.

[9] In an usual experimental session that lasts about 3 hours the design can include many replications with a few trading rounds or a few replications with many trading rounds. The main research question calls for the former.

Overall, the results suggest that dark markets support information acquisition albeit not in the strategic complementarity way the theory prescribes. The rest of the predictions of the theory however, are upheld in the data.

The findings have important implications for real-world markets. Many markets are organized as loose networks where trading flow is opaque. Confirming the theory, our experiments show that such organization is not detrimental to information aggregation. Most importantly, the decentralized organization induces acquisition of costly information, unlike centralized, transparent markets. As such, prices in decentralized markets could ultimately be more informative, and hence, provide better signals to the economy for optimal resource allocation. However, decentralized markets successfully incentivize information acquisition for reasons that are different from those in the theory.

The rest of the paper is organized as follows. Section I presents a short summary of the model and the corresponding experimental design. Section II concerns the recap of the original percolation theory, the parametrization of the experimental design and the simulation results. Section III describes the experimental design in detail. Section IV discusses the results, and Section V concludes.

I. An Abridged Description

In brief, the economy in the theoretical model is populated by infinite number of risk neutral agents who trade a single T-period lived asset, a unit at a time. The terminal value of the asset is a binary variable and agents can either purchase information about its realization or learn about the realization through trading. Agents meet randomly one pair of a buyer and a seller at a time (and without repetition) and are given the opportunity to trade in a private double auction.[10]

Under certain conditions, if everyone has a signal about the outcome of the binary state, and trading follows the above mechanism, the prices eventually become fully revealing.

[10] 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.

Most importantly, if information can be purchased, the decision to do so exhibits strategic complementarity. As such, from the moment one trader acquires information, there are strong incentives for others to do so too, as long as T (the time horizon) is large enough.

The result is an elegant implementation of perfect Bayesian equilibrium in monotone undominated strategies. It provides a strategic foundation for a break from the Grossman-Stiglitz paradox. But to achieve this, the authors need to impose restrictive assumptions both on the market mechanism[11] and on the type of security that is traded[12]. The notion of equilibrium comes with high demands on each player's cognitive ability, in particular it requires unrealistic levels of common knowledge of rationality. Mathematically, the strategic complementarity is an application of the martingale property of Bayesian likelihood ratios (see Bossaerts, Yang, and Asparouhova (2017)). Its implementation requires traders to not only use Bayes' rule in updating their information but also to be able to deduce and apply the martingale property in their forecasts of future trading opportunities.

For our experimental design, we relax some of the restrictive modeling assumptions and recast the theory in its simplest possible form, but without losing the richness of the model in order to be able to test its predictions. It is an empirical question about how robust the results are to relaxing the assumptions and this paper attempts to do that.

While more details about the experimental design as it relates to the theory follow in the theory section, here we provide the general description. The study consists of five experimental sessions, with the numbers of participants ranging from 14 to 18 each. Each session lasts approximately 3 hours. The first half of the session consists of an instruction period 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 for the participants to familiarize themselves with the software, Flex-e-markets, and the incentives that the experimental design imposes on them. This is also when the

[11] For example–give details that lead to the application of the law of large numbers–requirements on the filtration of information.

[12] The security is a very specific one and it does not become clear how it can be generalized

participants have the opportunity to and do ask questions. It is the block of time when the experimenter(s) demonstrate the statistical properties of the informational signals.

During the second half of the experimental session there are several replications of similar situations, that differ only by the informational structure provided to the participants. During those replications interactions among participants, and between the experimenter and the participants (including signal distribution) happen through the software.

The replications are referred to as "trading situations." Each replication consists of (i) an information acquisition period, (ii) trading period I, and (iii) trading period II. Before the information acquisition period participants learn how many other participants will be informed and how many will have the option to purchase information. For example in what we call "High" information session, out of 16 participants, 12 are initially informed.

With the exception of one, there are 10 replications per experimental session (Session 3 has 11 replications). The first replication in all sessions is one where all traders receive a signal. The number of uninformed traders as well as the number of signals to be auctioned off is publicly announced before each replication. Each session consists of two treatments, "High" information or "Low" information. The identity of initially informed traders is randomly determined. Those who are left uninformed participate in an informational auction, in analogy with Sunder (1992). Information is in the form of noisy signals about the binary state of the world. When traders acquire information they do not know if their trading role will be that of a buyer or a seller (thus preserving the ex-ante nature of information purchasing decision). For example, in the Low information treatment only 3 or 4 traders possess information initially. Another 3 or 4 signals are auctioned off among the uninformed participants. More concretely, in Session 1, 4 participants are initially informed, 3 signals are available to be auctioned, and the 12 uninformed traders submit bids for those 3 signals.

After the information 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 trading periods buyers (sellers) can submit up to one buy (sell) order. After a random matching between buyers and sellers, crossing orders are executed at the midpoint[13] and each trader with a successful transaction "inherits" the informational piece of their counter-party. In the original theory, information is "inherited" because the common knowledge assumption and its consequence of everyone knowing the equilibrium (strictly monotonic) bidding functions. As a result traders are able to invert bids and get a full revelation of the counter-party's signal. In our setting, we give the each of the two parties direct access to each other's information. Information transmission therefore happens effortlessly (i.e., no common knowledge and inversion of bids into information is asked from traders). This is an important methodological innovation–if we observe investment in information under our setting, it cannot be attributed to traders investing and trying to conceal their current information for longer than a transaction time's frame. With informational exchange after a transaction, we ensure the perfectly revealing nature of the encounter, but we still leave it to the traders to decide how to use this information in the formation of their stock offers.

The second trading period is the same as the first, save for the informational content of those who successfully traded in the first trading period. After the conclusion of second trading session the binary state (coin toss) is revealed and the traders' earnings are computed and presented to them. The next replication repeats the scenario all over, with different information treatment. The payoff from each experimental session equals a show up fee plus the earnings from all replications.

II. Specifics

In what follows we recap the theory using the parametrization applied in the experimental design.

[13]While the most elegant version of the theory is presented under sellers' auction, Duffie, Malamud, and Manso (2014)'s results remain intact unter alternative auciton mechanisms.

A. The Traded Asset

Theory

There is one risky asset whose payoff depends on the realization of a binary variable $Y \in \{0,1\}$, and the value is realized at time T. There are (infinitely) many agents, split between buyers and sellers. Whether an agent i is a buyer or a seller is determined by their utility function

$$U_i = v_i 1_{Y=1} + v^H 1_{Y=0} = v_i Y + v^H(1-Y).$$

If an agent is a seller, then $v_i = v_s$, while $v_i = v_b$ if the agent is a buyer, with $v_b > v_s$.

Experiment

The single security traded in the experiment is called the "Stock". In each of the two trading sessions buyers are endowed with cash only and can buy at most one unit the Stock per round. Similarly, the sellers can only sell at most one unit of the Stock per round. The payoff of the stock to buyers and sellers depends on which of the two states of the world $s = \{O(range), B(lue)\}$ realizes. In the former case $Y = 0$ and in the latter $Y = 1$.

After running pilots with the payoffs specified by the theory, we changed the no-gains from trade condition in the state of $Y = 0$ to

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

with $v_b > v_s$ as before but also $v_b^H > v_s^H$. The change is driven purely by the experimental implementation. An equilibrium with no trade is unusual to implement in the laboratory, as the mistakes of the participants cannot me modeled using standard econometric techniques when the equilibrium strategy is on the boundary of the action space. Table 1 illustrates the payoff structure of the Stock.

B. Information

Theory

Agents are endowed with information regarding the probability of the state of the world[14] Every period they are randomly matched (without repetition[15]) and given the opportunity to trade.

When a seller and a buyer meet, they reveal to one another their posterior probability that has incorporated all of their information up to the point of that encounter. Such revelation occurs through the observation of the bids submitted in an auction, as outlined below.

Imposing some restrictive assumptions for tractability, Duffie, Malamud, and Manso (2014) show that there exists a unique perfectly revealing equilibrium in monotone strategies. Based on the strictly increasing bidding strategies, traders can then infer the posterior beliefs of their counter-parties. Thus, after trading, a buyer and a seller have each a new and identical posterior belief. Based on this result, the cross-sectional distribution of the posterior beliefs can be summarized by an evolution equation. As the agents become more informed on average, the cross-sectional distribution gets fatter tails especially on the right, i.e. some agents of a shrinking mass get extreme posteriors (pointing to $Y = 0$ when $Y = 1$ and vice versa). We refer to an trader with an extreme posterior belief as a "long shot."

The presence of long shot types is a crucial feature. This feature 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. The basic intuition is that the agents want to avoid counter-parties of the long shot type. Such counter-parties become less likely as more samples are acquired, but conditional on obtaining information of the wrong type,

[14] 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.

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

the sample is more extreme. [16] The fact that wrong signals become less likely but more extreme is a consequence of the martingale property of likelihood ratios. While likelihood ratios will converge to zero as sampling increases (the likelihood of signals given the false hypothesis becomes smaller and smaller), with each signal addition, the expected likelihood ratio equals the prior likelihood ratio.[17]

Experiment

If the state is $Y = 0$, or "Blue," then signals are a sample 9 balls drawn from an urn consisting of 2/3 blue balls and 1/3 orange balls. If $Y = 1$, then the signals are draws of 9 balls from an urn with 2/3 orange balls and 1/3 blue balls. We call the signals that an agent receives, "the sample" of that agent.

Let B denote the random variable that counts the number of blue balls from the sample (thus $9 - B$ is the number of orange balls). Then the conditional and unconditional probabilities of observing "Y=1" and Y=0" are as follows.

$P(Y = 1|sample) = \frac{1}{1+2^{2B-9}}$

$P(Y = 0|sample) = \frac{1}{1+2^{9-2B}}$

$P(Y = 0) = P(Y = 1) = 0.5.$

An agent's *type* is then defined as $\theta = \log \frac{1+2^{2B-9}}{1+2^{9-2B}} = \log \frac{2^{2B}}{2^{9}} = \log(2^{2B-9})$. Hence, $P(Y = 0|\theta) = \frac{e^{\theta}}{1+e^{\theta}}$. In any case, when θ meets ϕ, they both emerge out of the encounter

[16] Most important for our experimental investigation, is the part of the theory dealing with endogenous information acquisition. Intuitively, "because trade is private, it takes time for information to become incorporated into prices. Informed investors may therefore profitably invest in gathering information. Furthermore, incentives to gather information may improve if agents anticipate trading with more informed counter-parties." Theorem 5.2 gives the sufficient conditions for information acquisition to be a strategic complement. "In order to understand the general complementary result for large T, we note that the incentives to gather information in an OTC market are determined by two basic effects: a "learning effect" and an opposing "pricing effect." Specifically, for any class i of agents, a change in the information acquisition policy of some other class agents influences the gain from information acquisition via two channels: (i) the profit from trade with a counter-party of any given type changes because the double-auction equilibrium bid and ask prices change (the pricing effect), and (ii) the distribution of types of future trading counter-parties changes (the learning effect). The pricing effect always reduces the information gathering incentives of individual sellers and buyers as their counter-parties become better informed. The learning effect is more subtle. 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." The latter is the effect of the martingale property of the posterior likelihood ratio mentioned earlier and elaborated in Appendix A.

[17] Doob (1949) and Bossaerts (2004) and Bossaerts, Yang, and Asparouhova (2017).

with new type $\theta + \phi$. Duffie and Manso (2014) provide the sufficient conditions for the existence of equilibria in strictly monotone undominated strategies in the above setup.[18]

Before each trading replication starts, and before Trading Session I, participants know that they have equal chance to be a buyer or a seller. A coin is tossed inside a box in front of all participants (thus the mechanisms is apparent and the outcome remains concealed but will be later announced and verified by the participants), and if the outcome is "heads," the state of the world is "Orange." The state is "Blue" in the case of a "tails" coin toss. Some traders are endowed with private information about the payoff relevant state of the world, which will be revealed at the end of each replication. The private information is distributed among traders at the beginning of each replication before the roles of buyers or sellers are assigned to them. The private information is represented by the 9-ball sample from the appropriate urn. Recall that given the probability of getting a certain orange(blue) ball is $p = 2/3$, the likelihood of drawing x orange (blue) balls in a sample of 9 draws is equal to

$$P(\#Orange = x) = \binom{9}{x} p^x (1-p)^{(9-x)} \tag{1}$$

Information Auction

The uninformed traders have the opportunity to participate in an information auction for the 9-ball samples. The auction takes place before the assignment of buyer/seller roles.

The auction design is as follows. If a number u of signals are auctioned off, then u separate first price auctions are conducted. Every uninformed trader provides their bid, and once submitted, their bid randomly goes towards one of the auctions (the number of uninformed agents chosen so that it is divisible by u). The highest bid wins the signal and pays the highest price. For example, in a pool of 16 traders, 4 are initially informed, and 12 are uninformed. 3 signals are auctioned off. The 12 traders submit bids (the bid price can be 0), and the bids are then randomly split across the three auctions. The highest bid within each auction wins the signal.

[18] See Propositions 4.3 to 4.6 and Theorem 4.7 in Duffie, Malamud, and Manso (2014).

Long Shots

We introduce long shots by having one (or at most two) of the traders receive information that is wrong, and the distribution of drawn balls is at least 6:3 in favor of the wrong colored balls. Participants learn about the long shots in the practice sessions. In those, as traders finish each of the three practice sessions the experimenter asks those who have a sample that corresponds to the truth to raise their hand, then from those who have the "opposite signal" the experimenter asks for the 6:3 draw to raise their hand. Of course, there are several participants who receive the opposite signal by chance but the composition of that signal is invariably 5:4 in favor of the ball colored opposite to the true state.

Trading of the Stock

After the information auction is concluded, trading period I of the stock market opens. Sellers submit at most one ask offer to sell a unit of the stock, and buyers submit at most one bid offer to buy a 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. The transaction price is the midpoint between the bid and ask offers. In addition, each trader who successfully participates in a transaction inherits the signal of their counter-party. Thus those who trade in period 1 enter period 2 with either no sample (if both the trader and her counter-party) are uninformed, a sample of 9 balls (if one trader is informed and the other is not), or with 18 balls (if both traders are informed). Equipped with the new information, the traders enter period 2 of trading. Period 2 is identical to period 1, except that the exchange does not proceed to period 3. After the matching is and the trades are

Note that in our replications within a session, trading happens over two trading periods only. Thus, our tests will only speak to price improvement from period 1 to period 2. The experimental logistics have proven challenging to our participants, and in a 3 hour experiment, about 1.5 hours are dedicated to training. In the remaining time, we as experimenters are solving a constrained optimization problem. We can either fit more replications of two periods or fewer replications, of, say 3 periods. We choose to have more replications to ensure robustness and replicability of our experimental results.

The payoff to a market participant per replication equals the sum of the Stock payoffs across the two trading periods. The session earnings equal the sum of the replication earnings.

C. Simulation Results

The preliminary results include simulations of the above economy. A minor deviation in the simulations vs. the theory is that the auction is not a seller's auction as in the original article but instead when orders cross, the trade is executed at the midpoint. A more significant deviation is that agents are not strategic in our simulations, i.e, they bid sincerely their valuations of the traded security.

Also, we directly use the theoretical result of full revelation of private information after a trading encounter and assign both an agent of type θ and her counter-party of type ϕ the new type $\theta + \phi$ immediately following the trading encounter. In the simulated economy, we focus on the payoff difference between informed and uninformed agents. The results robustly show that when the economy become more informative (fewer uninformed agents), the payoff difference between informed and uninformed increases, confirming that the information acquisition is a strategic complement. We find that strategic sophistication is not needed for the complementarity result. It only reinforces it. Complete simulation results will be provided in an online Appendix[19]).

III. Results

A. Costly Information Acquisition

One of the main implications of the theoretical models of decentralized markets is that the frictions in those markets make it worth the while for traders to invest in the acquisition of information. The model of Duffie, Malamud, and Manso (2014) provides an even more

[19] See https://www.dropbox.com/s/rnmtq7ytooicaiv/ElenaResults%20%28version%201%29.xlsx?dl=0

precise prediction, namely that the decision to acquire information is subject to strategic complementarity. To enable testing for the dependence of the value of information on the proportion of initially informed traders, we introduce two main treatments. In the High (Low) information treatment, the majority of traders is initially informed (uninformed).

Table 3 reports the summary statistics of the bidding patterns in the information auction across treatments (High and Low). The first result is that traders are willing to pay a strictly positive price for information. The average bid from all information auctions is \$1.24 (and a median of \$1.36). When conditioned on the treatment, participants bid \$1.32 on average (median of \$1.49) in the Low information treatment and \$1.08 (median of \$1.17) in the High information treatment. Examining the bid distributions across treatments, shown in Figure 1, we find that the entire bid distribution is shifted to right in the Low information 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 pairwise t-test. The results are included in Table 4. As suggested by the t-test, the average bid in the High treatment is significantly lower than in the Low treatment (with a t-statistic -3.62 and p-value 0.0006). Figure 2 plots the pairwise comparison of auction bids in a "violin graph."

To examine the treatment effect on subjects' information acquisition decisions more rigorously, we test the treatment effect within a regression framework. We estimate β using two methods. The first method utilizes individual fixed effects. The model is specified as follows

$$b_{ij} = \alpha_i + \beta H_j + \gamma X_j + \varepsilon_{ij}, \tag{2}$$

where b_{ij} is the bids for participant i in replication j. α_i in equation (2) are dummies (intercepts) for each subject, which absorbs any individual level unobserved time-invariant characteristics that may affect the bidding choice. H_j is an indicator for High treatment. The coefficient β captures the treatment effect and is our main coefficient of interest. X_j

includes other control variables such as trading rounds. The index on those variables is j because their values are collected at the replication level. The results are reported in column 1 of Table 5 . Consistent with the summary statistics of Table 3, β is about -0.2 and statistically significant.

The second method relies on a linear mixed model (or random coefficients model). Since each subject in our experiment participates multiple times in both the Low and High treatments, the linear mixed model is particularly suited in this repeated measurements design. The model is described in the following equation.

$$b_{ij} = (\alpha_0 + \alpha_i) + (\beta_0 + \beta_i)H_j + \gamma X_j + \varepsilon_{ij} \tag{3}$$

α_0 and β_0 are the "fixed effects,"[20] while α_i and β_i are the "random effects," which captures the individual level responses to the treatments. The results are reported in column 2 of Table 5. The overall results are almost the same as those reported in the column 1. The treatment effect is negative and statistically significant. Figure 3 depitcts the distribution of individual level coefficients. There is a substantial variation for the treatment effect across subjects, but the entire distribution is in the negative region.

To summarize, through various statistical tests and different experimental designs, we conclude that traders are willing to pay for information and bidding higher for information in the Low information treatment. This presents evidence in favor of strictly positive value of information (and in line with our simulation estimates) but against the strategic complementarity result that emerges from the theory.

B. The Effect of Costly Information Acquisition on Performance

One possible explanation of the result that uninformed traders place higher bids for information in the Low treatment is that they are overbidding in that treatment. This would mean that the value of information, measured by the difference in performance between

[20]Notice here the term "fixed effects" is distinct from the individual fixed effect in equation 2. Here the fixed effects are meant to direct overall effect from the covariates.

informed and uninformed traders, is lower than what uninformed traders bid. If this is the case, the traders who bid for information but did not acquire any (because their bid was too low) would be performing better than those who did acquire information.

The performance of the traders in the experiment is measured as the gains from trade. For example if a buyer in the Orange state, when her valuation for the stock is \$6 acquires one unit of the stock in trading period 1 for \$5 and one unit in trading period 2 for \$3, then her performance is equal to \$4=(\$6-\$5)+(\$6-\$3). In addition those traders who are initially uninformed are given \$2 that they can use to purchase information. Those traders who do not acquire information have the \$2 added to their replication payoff. Those traders who bid the highest in the auctions have the bid price subtracted from the \$2 and the remainder is added to their replication payoff. We refer to the portion of the payoff that comes from gains from trades as profits. The payoff that includes the payment of the auction leftover is refereed to as the total payoff.

Table 6 shows the average profits and total payoffs of all traders across different treatments and information acquisition status. The traders are split into three groups. The first group is the traders that are endowed with (free) information. The second group of traders is those who are initially uninformed but successfully acquired information. The last group of traders is those who are initially uninformed and fail to acquire information. Panel A of Table 6 reports the average *profits* from stock trading in the full sample. What emerges from this table is that the group of informed traders (initially endowed or having purchased information) generates higher profits from trading than the group of uninformed traders, although the differences are not statistically significant (a formal regression analysis accounting for repeated trader observations follows below). Panel B of Table 6 reports the results based on *total payoff*. These performance numbers suggest that information acquisition happens exactly at the price where traders are indifferent between obtaining the signal or not.

Next, using a regression framework, we examine the treatment effect on subject's performance. The dependent variable is on of the two variables that measure performance–stock trading profit or total payoff. The regression results are reported in Table 7. The dependent variables are stock trading profit/loss for the first three columns and the total payoff for columns 4-6. The coefficient of "With Information" is significantly positive for stock trading profit but insignificant for total payoff. This formally confirms the observation from table Table 6 that uninformed traders earn less than informed but the price of information exactly equals the difference between those payoffs. In particular, informed subjects on average earn \$1.65 more than uninformed subjects during the stock trading, suggesting a strictly positive value of information. However, when the leftover of the proceeds from the information auction are added to the profits, the performance, measured as the total payoff is not longer different between the informed and uninformed traders.

Column 2 of Table 7 presents a comparison between the three groups of traders. Those who fail to acquire information perform significantly worse than those initially endowed with information, at about \$1.5. Interestingly, those who acquire information in the auction seem to make a better use of it (around 50 cents) than those who are endowed but the result is not statistically significant. When comparing the profits of types of traders interacted with the treatment dummies, we find that traders with information, either acquired or endowed, perform significantly better in the Low treatment than those who failed to acquire information. The difference is \$2.2 for group with acquired information and \$2.6 for the group with information endowment. The interaction terms with High treatment suggest insignificant differences. Because the value of information varies by treatment and it is higher in the Low information treatment, rational bidders for information should bid more aggressively in the Low information treatment. That this is exactly what bidders do is suggested by the result presented in columns 4-6. Recall that those columns have as a dependent variable the total payoff to the traders. For traders endowed with information this is the same as their stock trading profits, for those who did not acquire information, the unspent \$2 are added to the profits, and for those who acquire information, it is the profits plus the leftover from the \$2 after paying the price of information. Independent of the specification, the differences between the informed and uninformed traders' payoffs are insignificant. Most informative is the specification in column 6 that now those who acquire information do not perform significantly better than those who do not. The difference of those endowed with information with those who did not acquire is not significant either but that is purely mechanical result, due to the difference of $2 added to the payoff of those who fail to acquire information (as displayed by the difference of $2 between the "Endowed$\times T_L$" coefficients in the specifications of columns 3 and 6). The meaningful comparison is that of those who acquire and those who fail and those differences are insignificant.

Table 7 presents several other statistics speaking to the rationality of traders (carefully addressed in the next section). The first is that traders' performances are influenced by the realization of the state. Recall that in the Orange state the difference between a buyer and a seller's valuation is equal to $6, while it is $4 in the Blue state, thus the gains from trade are $3 and $2 respectively per trader. This difference of $1 is exactly the value of the dummy variable "O State" across all specification. Another important coefficient is that of the performance of those traders who happened to have the wrong signal. Those traders would only trade if the offer they make is unfavorable to them. Thus, conditional on trading in trading period 1, they make relatively large losses, theoretically equal to $13 in the Blue state if a long shot is a seller (a long shot seller would ask to sell at $0 when they should ask to sell at $13). This loss is equal to $11 for buyers in the Orange state. The buyers in the Blue state and sellers in the Orange state fail to trade in the first trading period. In the second period if traders infer the information from the above trading realizations, they should make the average profits of $2 and $3 in the Blue and Orange state in the second trading periods. Thus the average loss from being a long shot should equal to approximately $2 (=(-13-11+2+3+0+0+2+3)/8=-14/8). Our results indicate that our long shot traders do not seem to incorporate the information from their trades (or lack thereof) and account for the informational disadvantage. This finding is consistent with our not finding strategic complementarity in the decision to purchase information. The strategic complementarity crucially depends on the fully rational behavior of the long shots.

C. Traders Rationality

Individuals are known to have difficulties to perform Bayesian updating. In this section we examine whether traders in our experimental markets are making bids and offers to buy and sell the stock according to their signals. This check is important because the theory relies on the basic assumption that traders behave according to their signals. The stock trading consists of two period for each replication. In each period, traders make their bids/asks based on a sample of signals they have (some are uninformed). To check whether they are bidding according to those signals, we compute each trader's expected value of the stock according to this trader's signal. The expected value is computed as

$$EV = \mathbb{P}(O|S_i)V_i^O + \mathbb{P}(B|S_i)V_i^B \tag{4}$$

In equation (4), S_i is the signal of trader i and V_i^O (V_i^B) is the value of the stock in Orange (Blue) state for that trader. The actual values are tabulated in Table 1.

In Figure 4 we plot the actual bids submitted by traders against their expected values. At each level of the expected value we make a box plot of the bids submitted. If the traders are perfectly rational, we should expect the average bids (horizontal bar inside the box) to form a 45 degree line passing through the origin. Although not perfect, the actual plot is displaying the desired property of Bayesian updating. The average bids increase monotonically as the expected values increase. When the offers are separated between those of the buyers and sellers, as the graph displays, the buyers tend to bid more consistently according to their posteriors.

To better understand the underlying bidding functions, for both buyers and sellers, we regress the bids on a set of covariates that can potentially influence the bidding strategy. Table 8 report the results. Column 1 shows that traders are generally trading according to their private signals. Perfect rationality implies a coefficient of 1 for expected value (EV), while we get 0.84 on average. The sellers tend to bid around \$4.3 higher than the buyers. In the Low information treatment, the traders bid slightly less (with a difference of approximately 42 cents) than in the High treatment or in the Baseline treatment (where all the traders are initially endowed with information). Columns 2 to 6 present the results from regression specifications that give more precise understanding of the bidding strategies used by traders. The state of the world shows no impact on the rationality as shown in column 2. Column 3 shows that the sellers are bidding less according to their signals than buyers do. The buyers are responding almost perfectly, with EV loading at 0.95, while sellers respond roughly 18 cents less for a unit of change in EV. Column 4 shows that subjects improve in terms of bidding according to their posterior in the second trading period. Interestingly, traders' Bayesian responses worsen over time as the experiment progresses as indicated by the negative coefficient (although statistically insignificant) on the interaction term between "Rounds" and EV. Lastly, in the Low information treatment, the subjects are less Bayesian. However, the economic magnitude is small.

Overall, the traders, particularly the buyers, seem to respond to their information in a Bayesian way when making stock trading decisions. Therefore the fact that we do not find strategic complementarity cannot be attributed to a bounded rationality argument.

D. Efficiency

D.1. Informational Efficiency

The striking conclusion from the series of papers of Duffie and coauthors (2007, 2010, 2014) is that their decentralized market converges to a fully revealing equilibrium, a desirable outcome that was deemed one of the most important advantages of centralized markets over decentralized markets. At the same time, the strategic complementarity of information acquisition makes the fully revealing equilibrium in decentralized market immune from the classic Grossman-Stiglitz paradox.

To get a better idea about the informational efficiency of prices and whether it improves over time, we examine the deviation of bid/ask and price from the fundamental value. Based on the designed payoff structure of the asset, the (expected) fundamental value of the stock is on average \$15 (or \$17 for buyer and \$13 for seller) when the state is Blue, and \$3 (\$0 for seller and \$6 for buyer) when the state is Orange. Since we expect that the deviations from fundamental value to become smaller in the second trading period, we run a simple regression to test this basic hypothesis. The regression results are reported in Table 9. Column 1 reports the deviation of the bids/asks, while column 2 reports the deviation of transaction prices.

The intercepts in the regression results in Table 9, at about \$4.2 for bids/asks and \$2.8 for prices, indicate that bids/asks and prices are far from fully revealing on average. At the same time, the absolute deviation of price from fundamental value decreases by about 74 cents in the second period (65 cents for deviations of bids/asks), confirming the prediction that information efficiency improves over time. This effect is statistically significant. Intuitively, deviations in the Low treatment is much larger by \$1.3 for bids/asks and \$1.6 for prices. Almost mechanically, subjects who receive the opposite signals (long shot) deviate a lot from the fundamental value. The trading round does not have significant coefficients, suggesting learning is not the main driving force behind the efficiency improvements.

D.2. Allocation Efficiency

Here, we examine whether allocation efficiency is improved from the first to the second trading period. We look at allocation efficiency in two different ways. First, we check the number of transactions per trading period. By design there is guaranteed Pareto improvement for any successful transactions, thus number of transactions suffices to measure the gains from trade. Second, we run simulations based on the bid/ask orders submitted by subjects in each trading period. Since the transactions in stock trading are facilitated by random matching (demanded by theory), the current sample only represents one realization of the possible matchings. In order to assess whether the stock trading is robustly leading to Pareto efficiency improvements, we compute the success rate transactions from all possible matches. Specifically, for each trading period, we compute the total number of possible matches between buyers and sellers. Then out of those matches, we count in how many cases the buyer is bidding higher than the seller. The success rate is computed as the number of crossing matches divided by the number of total matches.

Table 10 reports the results on allocation efficiency. In column 1 the dependent variable is the number of actual transactions happened. The result shows more transactions during the second period, suggesting Pareto improvements. In column 2 we examine the simulated success rate. The success rate is about 10% higher in the second period. The result is highly significant. Overall, the allocation efficiency improves in the second period.

IV. Conclusions

The paper presents the results from a series of controlled experiments designed to study information revelation in decentralized markets and the incentives to acquire information costly. Our experimental setting is inspired by the theory advanced in Duffie, Malamud, and Manso (2014). Consistent with the theoretical predictions, we find that, while initially noisy, transaction prices become significantly more informative in the second round. Allocational efficiency also improves significantly over the time. Second, traders in our environment bid aggressively for information–information samples cost around $1.65 meaning that most of the bids reside between $1 and $2. Unlike in centralized markets, participants pay positive prices for information. Third, we were not able to detect the strategic complementarity in the information acquisition. We attribute this to the inability of subjects to fully comprehend the subtule martingale property for the likelihood ratio, which is necessary for the strategic compelmentarity to emerge in the theory. Overall, we find support for the main predictions of the theory.

Appendix A: Martingale Property

Definition 1 *Let $(M_n : n \geq 0)$ be a sequence of real-valued random variables. Then, $(M_n : n \geq 0)$ is said to be a martingale (with respect to the sequence of random elements $(Z_n : n \geq 0)$ if:*

(i) $E(M_n) < \infty$ for $n \geq 0$.

(ii) for each $n \geq 0$, there exists a deterministic function $g_n()$ such that $M_n = g_n(Z_0, Z_1, ..., Z_n)$;

(iii) $E(M_{n+1}|Z_0, Z_1, ..., Z_n]) = M_n$, for $n \geq 0$.

The critical component of the martingale definition is condition (iii). If we view M_n as the fortune of a gambler at time n, then condition (iii) is asserting that the gambler is involved in playing a "fair game," in which he/she has no propensity (in expectation) to either win or lose on any given gamble. For example, a random walk with independent mean-zero increments is a martingale.

Martingales inherit many of the properties of mean-zero random walks. In view of the analogy with random walks, it is natural to consider the increments $D_i = M_i - M_{i-1}$, $i \geq 1$, namely, the martingale differences. The following proposition is a clear generalization of two of the most important properties of mean-zero random walks.

Proposition 1 *Let M_n be a martingale with respect to Z_n. Then, $E(M_n) = E(M_0)$. In addition, if $E(M_n^2) < \infty$ then $Cov(D_i, D_j) = 0$, $i \neq j$, so that $Var(M_n) = Var(M_0) + \sum Var(D_i)$.*

Proposition 2 *Let $(X_n : n \geq 1)$ be a sequence of iid random variables with common density g. Suppose that f is another density with the property that whenever $g(x) = 0$, then $f(x) = 0$. Set $L_0 = 1$ and*

$$L_n = \prod \frac{f(X_i)}{g(X_i)}$$

Then, $(L_n : n \geq 0)$ is a martingale with respect to $(X_n : n \geq 1)$.

Proof: $E(L_{n+1}|X_1, ..., X_n) = E(L_n \frac{f(X_{n+1})}{g(X_{n+1})}|X_1, ..., X_n) = L_n E(\frac{f(X_{n+1})}{g(X_{n+1})}|X_1, ..., X_n) =$

$L_n E(\frac{f(X_{n+1})}{g(X_{n+1})}) = L_n \int (\frac{f(x)}{g(x)}) g(x) dx = L_n$, since f is a density that integrates to 1.

This is known as a likelihood ratio martingale.

To show why the likelihood ratio martingale arises naturally, suppose that we have observed an iid sample from a population, yielding observations $X_1, X_2, ..., X_n$. Assume that the underlying population is known to be iid, either with common density f or with common density g. To test the hypothesis that the X_i's have common density f (the "f-hypothesis") against the hypothesis that the X_i's have common density g (the "g-hypothesis"), the Neyman Pearson lemma asserts that we should accept the "f-hypothesis" if the relative likelihood

$$\frac{f(X_1)...f(X_n)}{g(X_1)...g(X_n)}$$

is sufficiently large, and reject it otherwise. So, studying L_n in the case where the X_i's have common density g corresponds to studying the test statistic above when the "state of nature" is that the "g-hypothesis" is true. Given this interpretation, it seems natural to expect that L_n converges to zero as the sample size n goes to positive infinity. This is because for a large sample size n, it is extremely unlikely that such a sample will be better explained by the "f-hypothesis" than by the other one. The fact that L_n ought to go to zero as $n \to \infty$ is perhaps a bit surprising, given that $E(L_n) = 1$ for $n \geq 0$.

To prove that $L_n \to 0$ almost surely as $n \to \infty$, note that

$$\log(L_n) = \sum log(\frac{f(X_i)}{g(X_i)}).$$

Then the strong law of large numbers guarantees that

$$\frac{1}{n} \log L_n \to \int \log(\frac{f(x)}{g(x)}) g(x) dx.$$

The right hand side of the above is known as the *relative entropy*. Since log is strictly concave, Jensen's inequality asserts that

$$E \log(\frac{f(x)}{g(x)}) < \log(E(\frac{f(x)}{g(x)})) = 0.$$

As a consequence, not only does L_n converge to zero as $n \to \infty$ a.s, but the rate of convergence is exponentially fast. It is worth noting that this is an example of a sequence of random variables L_n for which $L_n \to 0$ a.s. and yet $EL_n \nrightarrow 0$ as $n \to \infty$ (in other words, passing limits through expectations is not always valid).

References

Asparouhova, E., and P. Bossaerts. 2016. Experiments on percolation of information in dark markets. *The Economic Journal* Forthcoming.

Asriyan, V., W. Fuchs, and B. S. Green. 2017. Information aggregation in dynamic markets with adverse selection. *Working Paper* .

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.

Blouin, M. R., and R. Serrano. 2001. A decentralized market with common values uncertainty: Non-steady states. *The Review of Economic Studies* 68:323–46.

Bossaerts, P. 2004. Filtering returns for unspecified biases in priors when testing asset pricing theory. *The Review of Economic Studies* 71:63–86.

Bossaerts, P., W. Yang, and E. Asparouhova. 2017. Why dissenting views gradually become more extreme. *Working Paper* .

Bossaerts, P. L., C. Frydman, and J. Ledyard. 2014. The speed of information revelation and eventual price quality in markets with insiders. *Review of Finance* 18:1–22.

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

Corgnet, B., M. DeSantis, D. Porter, et al. 2015. Revisiting information aggregation in asset markets: Reflective learning & market efficiency. *Working Paper* .

Doob, J. L. 1949. Application of the theory of martingales. *Le calcul des probabilites et ses applications* 23–7.

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.

Golosov, M., G. Lorenzoni, and A. Tsyvinski. 2014. Decentralized trading with private information. *Econometrica* 82:1055–91.

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.

Ostrovsky, M. 2012. Information aggregation in dynamic markets with strategic traders. *Econometrica* 80:2595–647.

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.

Radner, R. 1979. Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica* 47:655–78. ISSN 00129682.

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* 667–95.

Wolinsky, A. 1990. Information revelation in a market with pairwise meetings. *Econometrica: Journal of the Econometric Society* 1–23.

Tables

Table 1. Payoff Structure of the Stock

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

Table 2. 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)

Table 3. Auction Bidding Summary Statistics

This table report the summary of statistics of bidding prices in the information auction. Panel A report the summary statics across all the sessions we run through 2016 to 2017. Panel B report the average bidding price separately for sessions in 2016 and 2017. High treatment is when we have majority of the subjects endowed with information at beginning, while in Low treatment majority of the subjects are not endowed with information at the beginning. In 2016 Design, the number of informed agents varies in multiple levels, and in 2017 Design the number of informed agents vary only in two levels. The details are reported in Table 2.

Treatment	mean	std	median	p10	p25	p75	p90
High	1.08	0.70	1.17	0.10	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

Table 4. Paired T-Test

For each subject who experience both treatments, we compute the average bidding price of the subject per treatment. Then we conduct paired t-test.

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

Table 5. Treatment Effect on Auction Bidding

This table reports the regressions to identify the treatment effect on bidding prices during information auction. Column (1) use indvidual fixed effect and column (2) is estimated using linear mixed models with random intercepts and random coefficients for High treatment at individual level. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

	Dependent variable: bidding price	
	Fixed Effects	*Random Coefficients*
	(1)	(2)
High Treatment	-0.2039^{***}	-0.2034^{***}
	(0.0559)	(0.0513)
Trading Round	0.0144	0.0130^{*}
	(0.0122)	(0.0074)
Intercept		1.1875^{***}
		(0.0855)
Num. obs.	301	301
Adj. R^2	0.750	

Table 6. Average Total Payoff of Uninformed Subjects

This table report the average total payoff of the uninformed subjects per trading round across treatments and information acquisition status.

Panel B. Stock Trading Profit

Treatment	Endowed (1)	Acquired (2)	Failed to Acquire (3)	p-val of (1) minus (2)	p-val of (2) minus (3)
High	2.79	3.47	2.17	*0.63*	*0.39*
Low	3.57	3.37	1.33	*0.89*	*0.07*
p-val of diff.	*0.43*	*0.95*	*0.31*		

Panel A. Total Payoff

Treatment	Endowed (1)	Acquired (2)	Failed to Acquire (3)	p-val of (1) minus (2)	p-val of (2) minus (3)
High	2.79	3.81	4.17	*0.47*	*0.81*
Low	3.57	3.51	3.33	*0.97*	*0.87*
p-val of diff.	*0.43*	*0.86*	*0.31*		

Table 7. Payoff Regression

We regress the payoff the subjects on a set of variables to examine the effect of information acquisition on welfare. The dependent variable in model (1) is the profit extracted from stock trading, i.e. it doesn't include any cash remaining from information auction. The dependent variable in model (2) is the total payoff, including the cash remaining from information auction. "Eventually Informed" is an indicator for the subjects who either receive information endowment or acquired information during information auction. "Acquired" is an indicator for the (initially uninformed) subjects who successfully acquired information. "Failed" is an indicator for the subjects who are uninformed and failed to acquire information. "Endowed" indicates subjects who are endowed with information (for ree). $T_H(T_L)$ is a dummy for High(Low) treatment. "O State" indicates state is Orange. "Seller" is a dummy for sellers. Longshot is a dummy for subjects who received the signals that are opposite of the true state. "Round" is the trading rounds. *p<0.1; **p<0.05; ***p<0.01

Dependent Variable:	Stock Trading Profit			Total Payoff (incl. cash from auction)		
	(1)	(2)	(3)	(4)	(5)	(6)
Comparison Baseline:	Failed	Endowed	Failed	Failed	Endowed	Failed
Eventually Informed	1.6500***			−0.3038		
	(0.5721)			(0.5725)		
Acquired		0.4874			0.7062	
		(0.8472)			(0.8476)	
Failed		−1.5127**			0.5027	
		(0.6201)			(0.6204)	
Acquired $\times T_H$			1.6038			−0.0584
			(1.4429)			(1.4435)
Acquired $\times T_L$			2.1928**			0.3296
			(1.0182)			(1.0187)
Endowed $\times T_H$			0.8322			−1.1674
			(0.8152)			(0.8156)
Endowed $\times T_L$			2.5897***			0.5909
			(0.9992)			(0.9997)
T_H	−0.1420	−0.0186	0.7861	−0.1684	0.0104	0.7849
	(0.5518)	(0.5922)	(0.8908)	(0.5522)	(0.5925)	(0.8912)
O State	0.9835**	0.9814**	0.9824**	0.9898**	0.9868**	0.9883**
	(0.4888)	(0.4891)	(0.4892)	(0.4892)	(0.4893)	(0.4894)
Longshot	−3.1912***	−3.2165***	−3.2312***	−3.1794***	−3.2160***	−3.2351***
	(1.1963)	(1.1977)	(1.1983)	(1.1972)	(1.1983)	(1.1988)
Seller	0.5966	0.5747	0.6086	0.6118	0.5800	0.6145
	(0.4879)	(0.4897)	(0.4903)	(0.4883)	(0.4899)	(0.4905)
Round	−0.0179	−0.0177	−0.0179	−0.0205	−0.0203	−0.0205
	(0.0902)	(0.0903)	(0.0903)	(0.0903)	(0.0903)	(0.0903)
Intercept	1.0220	2.5048***	0.7099	3.0366***	2.4906***	2.7208***
	(0.7949)	(0.8821)	(0.8295)	(0.7955)	(0.8825)	(0.8298)
Observations	695	695	695	695	695	695
Adjusted R^2	0.0197	0.0187	0.0186	0.0116	0.0112	0.0111

Table 8. Bidding Strategy

In this table we regress stock trading bids/asks submitted by traders on a set of variables to understand the determinants of the traders bidding functions. "EV" stands for expected value computed according to traders private signals (see equation 4). "O state" indicates Orange state. Dummy(Period=2 indicates the second period of stock trading. T_L indicates the Low treatment. All regressions include session fixed effect.

	Dependent variable: bidding price					
	(1)	(2)	(3)	(4)	(5)	(6)
EV	0.8372^{***}	0.5095^{***}	0.9466^{***}	0.5048^{***}	0.6467^{***}	0.6325^{***}
	(0.0305)	(0.0396)	(0.0375)	(0.0344)	(0.0497)	(0.0276)
O State	−0.0706	-1.6329^{***}				
	(0.2369)	(0.5532)				
Seller	4.2644^{***}		6.0160^{***}			
	(0.2364)		(0.5117)			
Dummy(period=2)	-0.3757^{*}			-2.0419^{***}		
	(0.1929)			(0.4786)		
Round	0.0164				0.1049	
	(0.0336)				(0.0790)	
T_L	-0.4240^{**}					0.8146
	(0.2099)					(0.5681)
EV × O State		0.0203				
		(0.0547)				
EV × Seller			-0.1863^{***}			
			(0.0499)			
EV × Dummy(period=2)				0.1834^{***}		
				(0.0476)		
EV × Round					−0.0089	
					(0.0080)	
EV × T_L						-0.1344^{**}
						(0.0570)
Observations	1,579	1,579	1,579	1,579	1,579	1,579
Adjusted R^2	0.4351	0.3140	0.4372	0.3068	0.2996	0.3027

Table 9. Informational Efficiency

In this table we regress the deviations from fundamental value of the stock on a set of variables. The dependent variable in the first column is the absolute difference between bid/ask and fundamental value. The fundamental value at the trader level is the payoff of the stock shown in Table 1. The dependent variable in the second column is the absolute difference between transacted price and fundamental value, which is the average payoff value between buyer and seller. The Longshot in the second column will take value of 1 if either buyer or seller is a Longshot.

	Dependent variable:	
	|FV − Bid/Ask|	|FV − Price|
	(1)	(2)
Seller	0.7809***	
	(0.2214)	
Second Period	−0.6471***	−0.7432**
	(0.2213)	(0.3369)
Round	0.0321	0.0861
	(0.0383)	(0.0578)
Longshot	5.7044***	4.5819***
	(0.5021)	(0.5856)
O State	1.0391***	1.0478***
	(0.2215)	(0.3357)
LowInfo	1.3175***	1.5746***
	(0.2391)	(0.3748)
Intercept	4.1796***	2.8424***
	(0.4429)	(0.6633)
Observations	1,579	409
Adjusted R^2	0.1081	0.1837

Table 10. Allocation Efficiency

In this table we examine whether allocation efficiency is improved. The dependent variable in column (1) is number of transaction happened in each trading period. The dependent variable in column (2) is the simulated successful rate of transactions. To derived this measure, we exhaust all possible random matching between buyers and sellers in each trading period. Then we compute the likelihood of a random matching that will result into a successful transaction. It's computed as the number of successful transactions divided by number of possible matching.

	Dependent variable:	
	N. of Transactions	Success Rate of Transactions
	(1)	(2)
LowInfo	-0.7693^{***}	-0.0754^{**}
	(0.2855)	(0.0317)
Second Period	0.5686^{**}	0.1029^{***}
	(0.2690)	(0.0299)
O State	-0.1570	-0.0216
	(0.2692)	(0.0299)
Intercept	4.0589^{***}	0.5154^{***}
	(0.2517)	(0.0280)
Observations	102	102
Adjusted R^2	0.0817	0.1287

Figures

Figure 1: Histogram of Bids in Information Auction

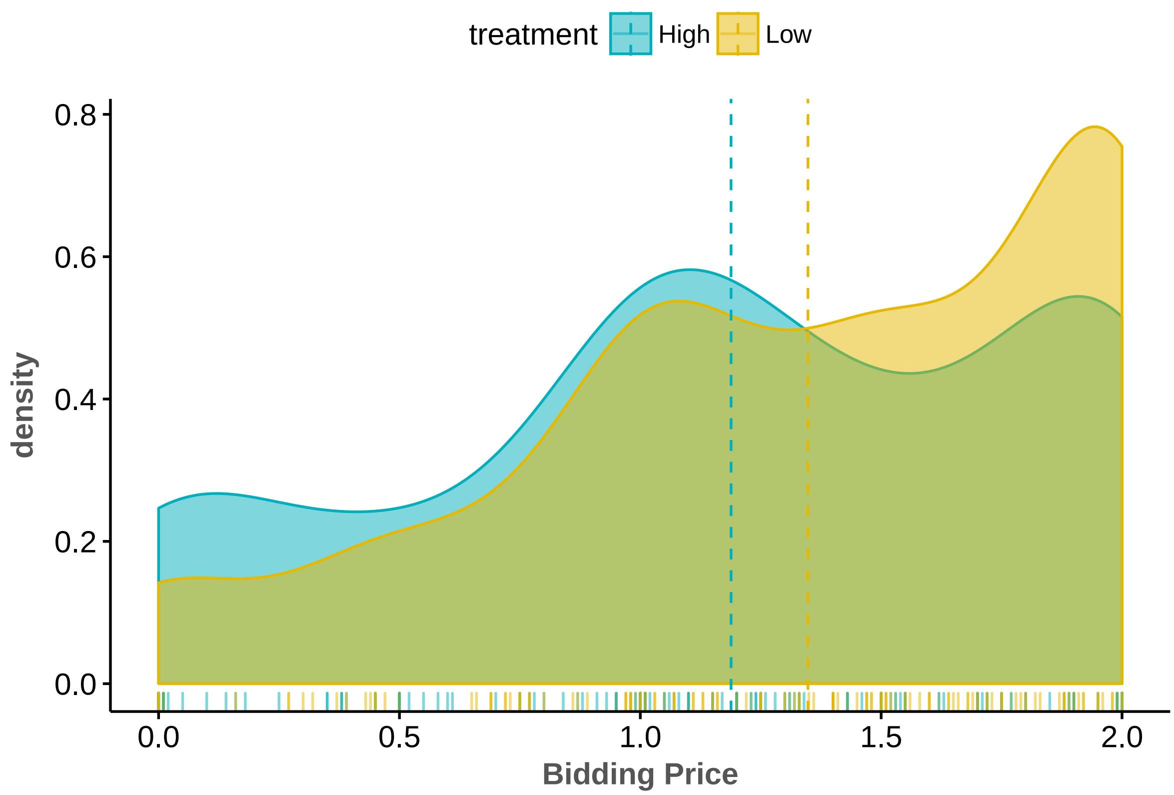

Figure 2: Pairwise Comparison of Auction Bids

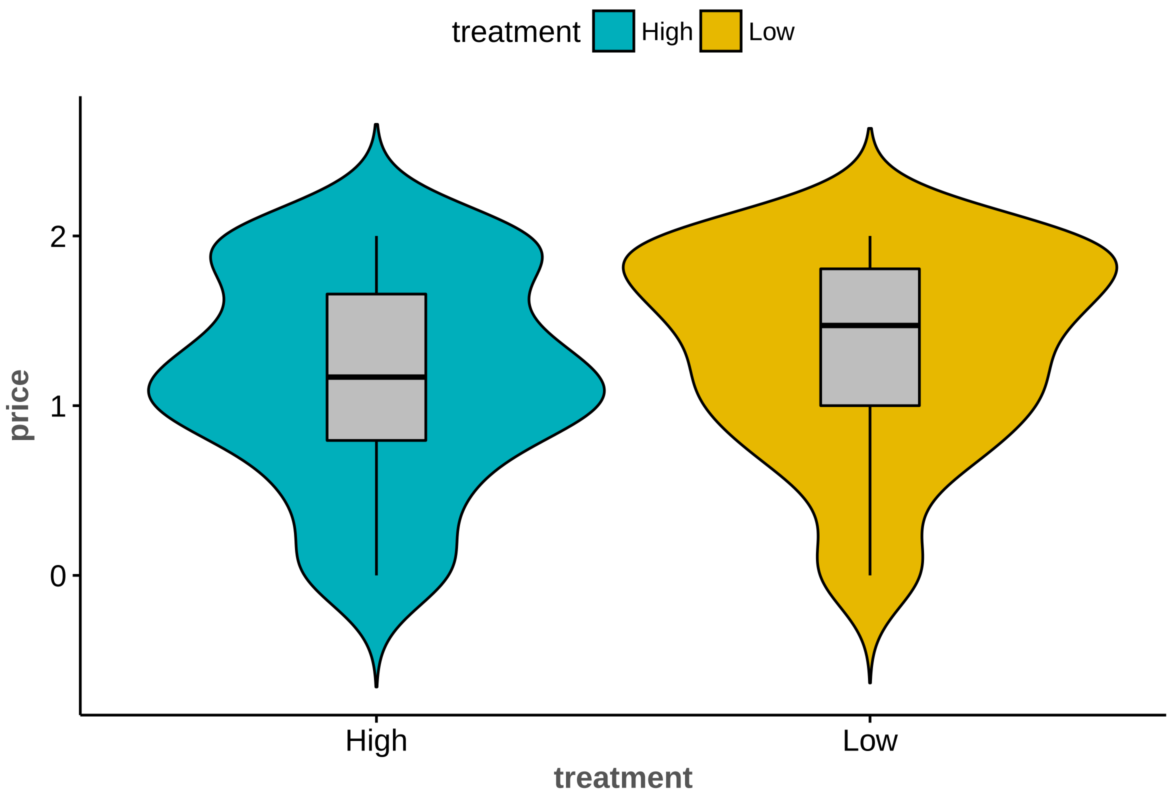

Figure 3: Histogram of Random Coefficients

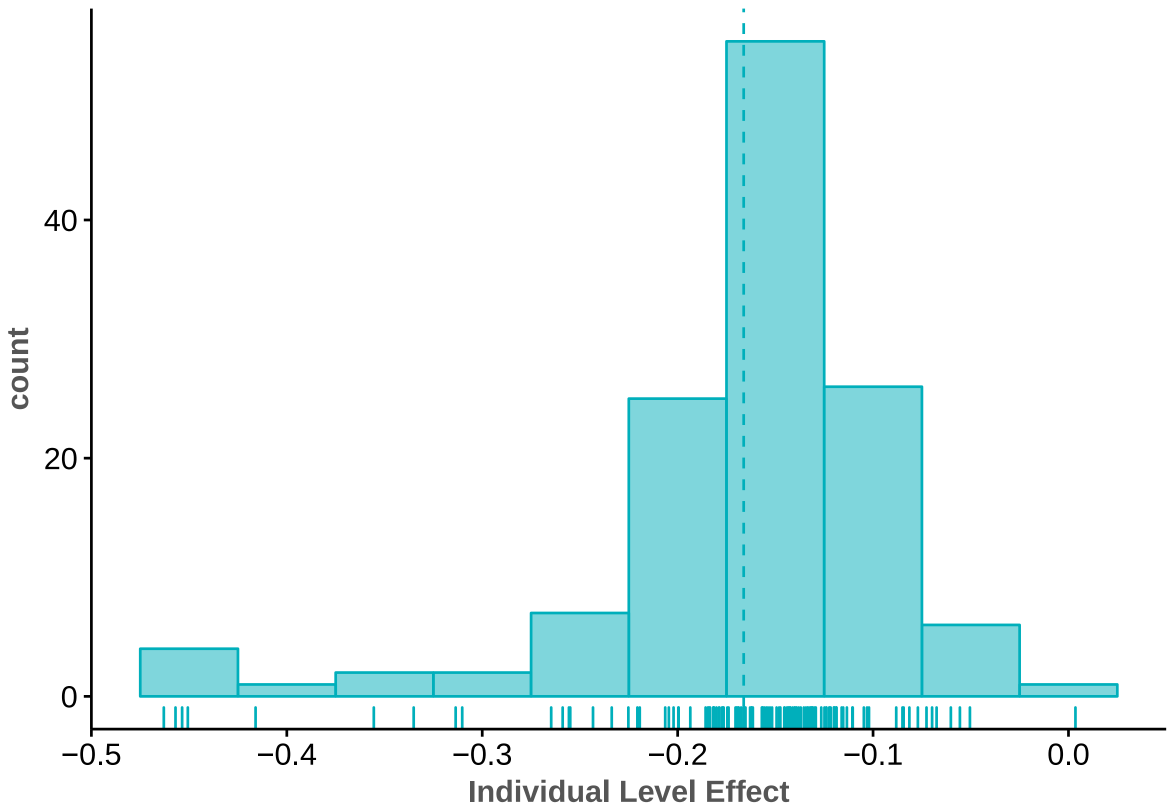

Figure 4: Actual Bids/Asks against Expected Value of the Stock

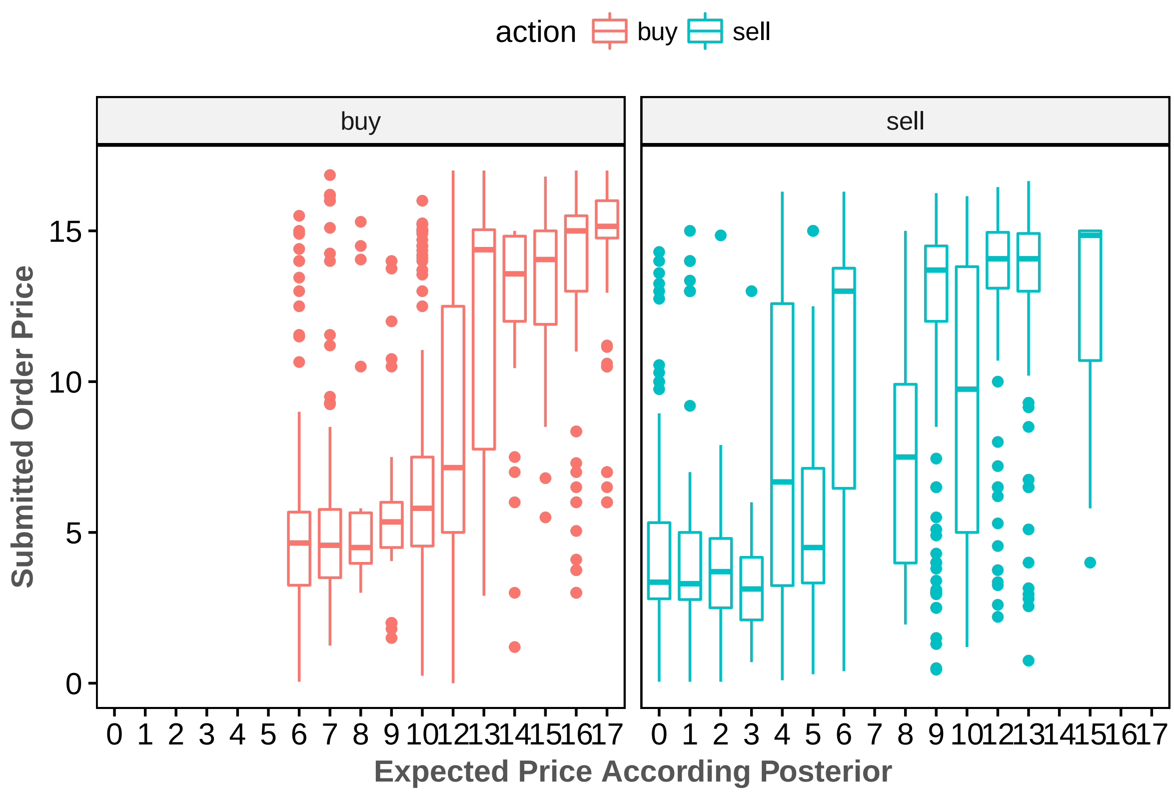

Appendix B: Instructions

In the next few pages, we attached the original experiment instruction that is distributed to each subject before the experiment. In addition, we also construct presentation slides to help the subject comprehend the experimental prcedures. The presentation slides that the experimenter projects on a large screen while reading the instructions can be found here.

INSTRUCTIONS: Information Aggregation Experiment

Summary
In a number of replications (up to 10), you will be given stocks and cash, and you will be invited to trade with others through an online marketplace. Final earnings will be the sum of your earnings in the replications. Stock payoffs are determined by a random draw (think of a coin flip). You will either be given information or will have the ability to purchase information about the outcome of this random draw. Your exact earnings will depend both on the random draw and the prices at which you trade. In addition, you will receive a sign-up reward of $5.

1. Setting
There are two large bins with orange and blue colored balls.

1. In the first bin 2/3 of the balls are orange and the rest (1/3) are blue.
2. In the second bin 2/3 of the balls are blue and the rest (1/3) are orange.

In the former case, we say that the bin is Orange; in the latter case, it is Blue.

1.1. Replications
You will be trading in a series of 10 replications. Every replication will have two trading periods.

1.2. Two Trading Periods Per Replication
In each trading period, you are given some cash, and some units (possibly zero) of a security, called Stock. The Stock's payoff depends on whether the bin is Orange or Blue.

Half of you have been assigned to be **Sellers**.
You will own one unit of Stock at the start of a period.
You will only be able to **sell** the Stock.
You will start each period with 1 unit of Stock.

The other half have been assigned to be **Buyers.**
You will have only cash in the start of a period.
You will only be able to **buy** the Stock.
You will start each replication with $20 and no Stock.

If the bin is Blue, the Stock pays off $17 to the Buyers (if they bought it) and costs $13 to the Sellers (if they sold it).
If the bin is Orange, the Stock pays off $6 to the Buyers (if they bought it) and costs $0 to the Sellers (if they sold it).

Bin	Stock Value	
	Buyer	Seller
Orange	$6	$0
Blue	$17	$13

A replication starts with a coin toss that will not be revealed to you until the end of a replication. The bin that applies for the given replication is the Orange bin if the coin toss is *Heads* and it is the Blue bin if the coin toss is *Tails*.

At the start of each replication, before the first trading period, some of you (possibly all) will be given a sample of nine balls drawn from the selected bin. The vast majority of you will get a sample that consists of mostly orange balls when the bin is Orange and mostly blue balls when the bin is Blue. However, some of you will get a sample that has mostly orange balls when the bin is Blue, or mostly blue balls when the bin is Orange.

Initial Bin Samples
In replications 1 and 6 all of you will receive a sample from the bin. In the rest of the replications, some of you, randomly selected, will have no sample, but will have the option to purchase one via an auction. We will announce at the start of each replication how many have been given a sample (without having to purchase it).

Stock Market
Buyers can submit at most one offer to *buy* a unit of the Stock. The bid should be submitted to the trader called ***mm***, the "market maker."
Sellers can submit at most one offer to *sell* a unit of the Stock. The ask offer should be submitted to the trader called ***mm***.
Then ***mm*** will execute an algorithm, where each buyer is randomly matched with a seller. If the buyer's offer price is higher or equal to the seller's offer price, then there will be a transaction. The transaction price will be the mid-point between the buyer's and the seller's offers. If transaction occurs, only the participating buyer and seller will know the price (or that a transaction took place).

Once a transaction occurs, both the buyer and the seller will get a new updated sample of orange and blue balls that is simply equal the combined sample of the two parties to the transaction.

<u>Example:</u>
If buyer M10 is matched with seller M20, and M10 has submitted offer to buy at $14, while M20 has submitted an offer to sell at $12, then there will be a trade, and the price will be $13 (=1/2*(12+14)). Only M10 and M29 will know that.
If M10 had a signal of 5 blue and 4 orange balls (thinking that more likely than not the bin is Blue but not very sure) and M20 had a signal of 7 blue and 2 orange balls (thinking that the bin is Blue with very high likelihood), then both will move to the second period of trading with a sample of 12 blue and 6 orange balls, inferring that the bin (much) more likely Orange than Blue.

Information Auction
Information will be transmitted to you via the trading interface.
To this end we have created two non-markets called Blue Signal and Orange Signal. Your holdings in those non-markets are your signals.
If you have been (randomly) chosen to have information, then your holdings in the two non-markets will contain the information (see slides for examples). In addition, your cash holdings will be artificially set to -1 during the information auction to prevent you from participating (as you already have a signal). This -1 will not affect your final earnings.
If you have been (randomly) selected to be a trader with no information, you will have the opportunity to purchase a sample in the information auction. You will be given $2 and a spot in the auction.

In every auction there will be one signal for sale. It will go to the highest bidder.

Example: If 5 participants who are uninformed participate in an auction and the bidding prices are 0.89, 0.83, 0.63, 0.62, 0.55, then the signal will be given to the bidder who bid 0.89.

If say, 15 of you are uninformed, we will conduct 3 auctions. Each of you will randomly be chosen into 3 groups of 5 and the above process will repeat for each of the groups. If only five of you are uninformed, we run only one auction. As such, you only ever compete with 4 other people in the information auction stage.

Notice the information auction will happen once in each replication, before the start of the first trading period. And you will learn whether you are a buyer or a seller of the Stock only after the information auction.

Experimental Earnings
If you are uninformed and you did not spend the entire $2 given to you, the rest is yours to keep. Notice that bidding $2 might not be optimal, since if you "save" from your auction money by bidding less than $2, across the 10 replications, the total sum saved may be a significant portion of your earnings.

Buyers are provided with $20 cash at the beginning a trading period. Your earnings per replication are equal to your
Total change of cash (at the end of period 1 and period 2)
PLUS
(value of Stock) X (units of stock you own)
PLUS
(leftover from info auction).
Notice that if you do not trade, your earnings are zero, except for the money not spent in the information auctions.

Sellers are provided with one unit of the Stock in trading period 1 and one unit in period 2.
Your earnings per replication equals your:
Total change of cash (at the end of period 1 and period 2)
MINUS
(cost of Stock) X (units of stock you sold)
PLUS
(leftover from info auction).

Example 1
You are a **seller** and sell 1 unit of stock in Period 1 at $9 and do not sell in period 2.
Your payoff is:
Your change in cash MINUS (value of Stock) X (units of stock you sold) =
If the bin is Orange, then the cost of Stock is $0 = $9 - $0 = $9
If the bin is Blue, then the cost of Stock to you is $13 = $9 - $13 =-$4.
If you did not trade and kept the two units of Stock, your payoff would be $0.

Example 2
You are a **buyer** you start a replication with $40 ($20 in each trading period), and you buy 1 unit of Stock for $9 in the first period.
Your payoff is:
Your change cash PLUS (value of Stock) X (units of stock you own) =
If the bin is Orange, then value of Stock is $6 = -$9 + $6 =- $3
If the bin is Blue, then value of Stock to you is $17 = -$9 + $17= $8
If you do not trade and keep your cash position, your payoff will be $0.

Discussion

- **Information Auction**
  Information is essential in this market. When you are uninformed, you are in a disadvantageous position, because the best guess you have is half-half on the true color of the bin. Thus, as, say, a buyer you might bid too high (in which case you will incur some loss) or too low (in which case you miss opportunities to profit) relative to the true value of the stock. The information auction offers uninformed participants the opportunity to become better informed and formulate better bids.

  Another reason that the uninformed should participate in the information auction is to insure for the case you trade with participants who hold the "wrong" sample, namely, those with majority blue (orange) samples when the bin is orange (blue).
  Here is an example: you are an uninformed seller and the bin is Orange; you are matched with a buyer who holds a sample with majority Blue balls (say 5 blue and 4 orange); you transact one unit of stock – which is likely because the buyer will bid a high price given buyer's sample. You now will inherit the buyer's signal, which is the only signal you have since you yourself started uninformed. Thus you too will be wrongly convinced that the bin is Blue. In the 2nd period, you will then probably ask too high a price because you estimate the stock cost to be \$13 when the bin is Blue. Chances are that most of the other participants will either have the correct signal or no signal at all, and thus your offer will be too high and you will miss on a trading opportunity. If you had purchased your own signal, you'd be with very high likelihood able to avoid such mistakes.

- **Stock Market**
  Notice the value of the Stock is worth \$6 for buyers but costs only \$0 for the sellers when the bin is Orange. The Stock is worth \$17 for buyers and costs \$13 to sellers in the case of a Blue bin. Thus the Stock is always more valuable for buyers than it costs to sellers. When submitting the bids, buyers should keep in mind the following trade off: if you bid a low price, your potential profit will increase, however the probability that a seller will accept your offer will decrease. Conversely, sellers should realize that if they ask a high price, they are less likely to trade.

  Successful transactions are important for all participants because they not only realize the potential gains, and hence, generate your earnings in this experiment, but they also transmit information between the two parties to the trade. After the information is transmitted, the size of your sample of Orange and Blue balls will increase. With a larger sample, you can make more accurate inference about the true color of the bin, and hence formulate better bids and asks for the 2nd period.

Trading

Trading will take place through an electronic trading platform called *Flex-E-Markets*. Log onto http://uleef.business.utah.edu/flexemarkets/ with the trading account name and password given to you. You submit limit orders (orders to buy or sell at a price you determine, or a better price). Transactions take place from the moment a buy order with a higher price crosses a sell order with a lower price. Orders remain valid until you cancel them or the trading period ends.

Why Dissenting Views Gradually Become More Extreme

Peter Bossaerts[a] Wenhao Yang[b] Elena Asparouhova[c]

[a]The University of Melbourne and The Florey Institute of Neuroscience and Mental Health

[b]The University of South Carolina

[c]The University of Utah

November 9, 2017

Here, we invoke a result from mathematical statistics to identify a curious phenomenon, namely: as it grows, dissenting evidence (i.e., evidence against the truth) necessarily becomes more extreme.

Suppose that many agents collect independent samples to test a statistical hypothesis. As they collect more samples, mathematics predicts that the measure of those who believe that the true hypothesis is wrong rapidly goes to zero. However, the few with samples that contain evidence contrary to the truth become far more extreme than the confirming evidence. As a result, the evidence for and against the truth polarizes.

Formally, $(X_n : n \geq 1)$ is a sequence of independent and identically distributed (iid) random variables from a distribution with a density h. Suppose that f is another density that is absolutely continuous with respect to h (i.e., if $h(x) = 0$, then $f(x) = 0$). We define evidence as the odds ratio. This is the ratio of the likelihoods of the observed sample, under the *alternative* hypothesis (*Hypothesis* f, in the numerator) and under the *truth* (*Hypothesis* h, in the denominator):

$$L_n = \prod_{i=1}^{n} \frac{f(X_i)}{h(X_i)}, \quad L_0 = 1.$$

With a slight abuse of notation, we call "a piece of evidence i" the ratio $\frac{f(X_i)}{h(X_i)}$. Our phenomenon is the result of two properties of odds ratios, namely: (i) as n grows there is a sharp increase in the number of sample paths for which the odds ratio declines (towards zero); (ii) the expected effect on the odds ratio from increasing the sample is zero. So, starting from 1, the odds ratio is not expected to change with adding (or subtracting) a piece of evidence (property ii). Nevertheless, it reduces to zero for almost all samples (property i). The only way the two properties can hold simultaneously is if the odds ratios of samples with evidence against the truth increase without bound.

Mathematically, the two properties can be expressed as (i) $L_n \to 0$ almost surely as $n \to \infty$ (ii) $(L_n : n \geq 0)$ is a martingale with respect to $(X_n : n \geq 1)$. The former is often stated in courses in statistics; Doob was the first to point out the latter,[1] though the proof is simple (see Supplementary Material).

The martingale property implies, among other things, that $E(L_n) = 1$. That is, across samples, the unconditional expectation of the odds ratio equals 1.

The Figure below illustrates our point. We generated 10,000 samples of binomial random variables of length N, where N increases from 100 to 1125. The binomial distribution has p (probability of success) = 0.3. For each sample, we compute the odds ratio of $p = 0.4$ against the null (truth) of $p =$ 0.3. The Top Left plot shows the average odds ratio as a function of N. In theory, the average should be 1 (as a result ofthe martingale property), and indeed it is (ignoring sampling bias). This also means that our simulation size (10,000) is sufficiently large to generate the predicted outcome. The Right panel displays the average odds ratio for the samples that generated evidence in favor of the truth (i.e., odds ratio 1). As the sample length increases, the average decreases towards zero. That is, the odds ratio decreases towards zero as evidence in favor of the truth accumulates. Bottom Left is a graph of the average odds ratio for the samples that generated evidence against the truth (i.e., odds ratio >1). Consistent with the theory, the average increases with the sample size, demonstrating that the evidence against the truth gradually becomes more extreme as T increases. The Bottom Right panel depicts the evolution of the maximal odds ratio as a function of the sample length. Again, we observe that the evidence (the odds ratio) quickly becomes radical.

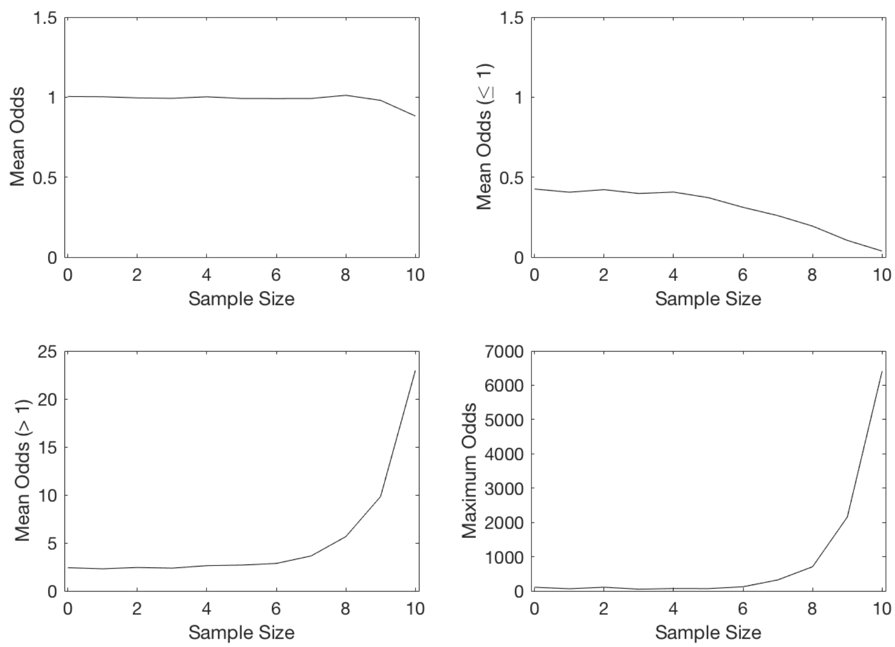

Figure 1. Means and maxima estimated from 10,000 samples of length N. Sample Size $= log_2(N - 100)$.

The comparison of samples with evidence for (Top Right) and against (Bottom panels) demonstrates polarization, consistent with the theory. Importantly, the evidence against the truth radicalizes faster than the evidence in favor of the truth. This of course is the simple consequence of the fact that on average the evidence has to cancel out (the martingale property), but the evidence against the truth becomes rarer, so necessarily will be more extreme than the evidence in its favor.

As such, mathematics generates the curious result that dissenting evidence becomes more extreme as it grows.

This has major implications, among others, for society. Indeed, if members of a society collect large, independent samples in order to determine the veracity of a proposition, then an uninformed person may accidentally run into someone whose sample happens to contradict the truth, and this sample will be extremely convincing. The uninformed person may unfortunately become convinced too. Now consider encounters with those who collected a sample in favor of the truth. Though such encounters are more likely, they will lead in general to a more cautious view of the truth!

Duffie and collaborators[2] recently modeled rational choice in a game where players randomly meet and exchange information in order to determine the truth. If their actions are consistent with the truth, they win; otherwise they lose. The authors cast the game in a financial context (specifically, an over-the-counter market), but that is not essential to their finding, which is that uninformed agents (those without a sample) have an incentive to pay to obtain their own information when they know that the persons they are likely to meet also did.

Our result explains this counterintuitive finding: the uninformed are concerned with encountering someone who accidentally learned to believe the wrong thing; they rationally choose to collect their own sample, to offset the extremism encapsulated in dissenting evidence.

Our result should be interpreted to mean that any single collection of radical evidence must be interpreted with caution. The more convincing the evidence is, the more important it is to spend the effort of collecting another sample. If the extreme evidence happens to be wrong, the new sample will show this with high likelihood, even if it tends to point to the truth in a less convincing way.

References

1. Doob, J. L. Application of the theory of martingales. Le calcul des probabilites et ses applications 2327 (1949).

2. Duffie, D., S. Malamud. and G. Manso, Information percolation in segmented markets. Journal of Economic Theory 153, 132 (2014).

Supplemental Materials for

Why Dissenting Views Gradually Become More Extreme

Proof of the Martingale Property

The proof below is from the notes of Prof. Peter W. Glynn for the course MS&E 321 entitled Stochastic Systems, taught at Stanford for Spring 12-13. Notes are dated June 1, 2013.

Proposition 1. *Let $(X_n : n \geq 1)$ be a sequence of iid random variables with common density h. Suppose that f is another density with the property that whenever $h(x) = 0$, then $f(x) = 0$. Set $L_0 = 1$ and*

$$L_n = \prod \frac{f(X_i)}{h(X_i)}$$

Then, $(L_n : n \geq 0)$ is a martingale with respect to $(X_n : n \geq 1)$.

Proof: $E(L_{n+1}|X_1, ..., X_n) = E(L_n \frac{f(X_{n+1})}{h(X_{n+1})}|X_1, ..., X_n) = L_n E(\frac{f(X_{n+1})}{h(X_{n+1})}|X_1, ..., X_n) =$

$L_n E(\frac{f(X_{n+1})}{h(X_{n+1})}) = L_n \int(\frac{f(x)}{h(x)})g(x)dx = L_n$, since f is a density that integrates to 1.

Proposition 2. $L_n \to 0$

Proof: Note that

$$\log(L_n) = \sum log(\frac{f(X_i)}{h(X_i)}).$$

The strong law of large numbers guarantees that

$$\frac{1}{n}\log L_n \to \int \log(\frac{f(x)}{h(x)})h(x)dx.$$

The right hand side of the above is known as the *relative entropy*. Since log is strictly concave, Jensen's inequality asserts that

$$E\log(\frac{f(x)}{h(x)}) < \log(E(\frac{f(x)}{h(x)})) = 0.$$

As a consequence, L_n converges to zero as $n \to \infty$ a.s., and exponentially fast. This is an example of a sequence of random variables L_n for which $L_n \to 0$ a.s. and yet $EL_n \nrightarrow 0$ as $n \to \infty$ (in other words, passing limits through expectations is not always valid).