THE ECONOMIC JOURNAL

The Economic Journal, **127** (*October*), F518–F544. Doi: 10.1111/ecoj.12464 © 2017 The Authors. The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society. Published by John Wiley & Sons, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

EXPERIMENTS ON PERCOLATION OF INFORMATION IN DARK MARKETS*

Elena Asparouhova and Peter Bossaerts

In dark markets, order submissions are bilateral, and transaction prices are known only to the trading counterparties. Here, we study to what extent the information aggregation theory proposed by Duffie and collaborators predicts outcomes in a laboratory version of such markets. We find that prices aggregate the available information but not in the strict sense of the theory, where prices converge exponentially fast to average private signals. Prices instead fluctuate within bands around this average. The band widths reflect, in the best case, the precision of the average signal and, otherwise, the precision of a single private signal.

> The MiFID II reform means that organised trading of financial instruments must shift to multilateral and well-regulated trading platforms. Strict transparency rules will ensure that dark trading of shares and other equity instruments which undermine efficient and fair price formation will no longer be allowed.
>
> Michel Barnier, European Commission – MEMO/14/15 from 14/01/2014

Since the Great Financial Crisis of 2007–8, dark markets have come under intense criticism and scrutiny, to the point that recent regulation has limited their opacity (the Dodd-Frank act), or even, as in the case of MiFID II reform in the European Union, has effectively outlawed them.[1] Dark markets are decentralised markets where negotiation is private between two parties and, even if a trade is concluded, others are informed about it at best long afterwards (Duffie, 2012).

From a theoretical perspective, one can indeed be sceptical about the merits of dark markets. Since the price at which two parties trade is privileged information only to them, different pairs may trade at vastly different prices. This is in sharp contrast with the uniform price at which everyone trades in equilibrium of a centralised competitive economy. Benchmarking against competitive equilibrium prices is important as the latter delivers Pareto optimal allocations under certain conditions.

Empirical studies have confirmed that, indeed, transaction prices in real-world dark markets can be quite dispersed, especially when retail investors are involved, and that

* Corresponding author: Peter Bossaerts, Eccles School of Business, University of Utah, Salt Lake City, UT 84112, USA. Email: peter.bossaerts@utah.edu.

Financial support from the National Science Foundation (Asparouhova: SES-1061844, Bossaerts: SES-1061824), the Banque de France Foundation, a Moore Foundation grant to Caltech (2006–13) in support of Experimentation with Large, Diverse and Interconnected Socio-Economic Systems, and the Development Fund of the David Eccles School of Business at the University of Utah is gratefully acknowledged. The article benefited from discussions with Darrell Duffie, and from insightful comments from the Editor and an anonymous referee. Data and programs used in the analyses reported here are available at http://uleef.business.utah.edu/DataEJ/.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

[1] See http://europa.eu/rapid/press-release_MEMO-14-15_en.htm.

this dispersion decreases with stricter reporting requirements, for example, through *ex post* reporting of trades; see Goldstein and Hotchkiss (2007).[2]

Market experiments have provided support to the argument that, even in the absence of asymmetric information, decentralised trading mechanisms cannot deliver the competitive equilibrium outcome (Chamberlin, 1948). More generally, since the seminal papers of Smith (1965) and Plott and Smith (1978), the perception in experimental economics has been that a centralised market is needed in order to produce the competitive equilibrium outcome and its welfare merits.

The pessimism about dark markets extends to the case of asymmetric information. When insiders have access to identical signals, Wolinsky (1990) provides a theoretical argument that decentralised markets cannot possibly amplify[3] this privileged information through competition and fully revealing prices like centralised markets do. The ability of centralised markets to amplify information and obtain prices that are fully revealing is a well-established prediction of the rational expectations equilibrium (Radner, 1979). Early experiments with centralised financial markets have confirmed this capability (Plott and Sunder, 1988).

Taking dispersed information as the basis informational environment, Duffie, Manso and collaborators have developed an approach to modelling the trading in dark markets that is based on the idea of 'percolation' in stochastic process theory.[4] When the market participants have observable roles of natural buyers and natural sellers and the information is dispersed among them, dark markets can eventually achieve the informational efficiency of competitive equilibrium (Duffie and Manso, 2007; Duffie *et al.*, 2010) and do so with exponentially fast conversion rate (Duffie *et al.*, 2009, 2014).

Closer inspection of the structure of the economy in this 'theory of information percolation' (TIP) of decentralised markets identifies two possible sources that explain why it produces predictions that are vastly at odds with previously received wisdom.

First, in TIP, it is crucial that everyone knows the other participants' idiosyncratic incentives to trade. Notably, this aspect of the theoretical design, as well as the design of our experiments, differs from that of early market experiments, where decentralised markets were observed to fail (Chamberlin, 1948) while public markets performed well (Smith, 1965). In those markets, participants had private components for their valuations of the traded assets that were not common knowledge and, hence, when two agents met privately, they could not easily tell how aggressively to bargain, thus affecting price discovery. This, of course, is related to the well-known result that there does not exist a mechanism that always guarantees *ex post* efficient trade between two parties when they do not know each other's private valuations (Hurwicz, 1977; d'Aspremont and Gérard-Varet, 1979).

[2] That trading costs decrease when opacity is reduced through mandatory trade reporting is also reported in Bessembinder *et al.* (2006); Goldstein *et al.* (2007). Real-world dark markets exhibit additional interesting features that will not be of concern here, such as a puzzling inverted relation between volume and execution cost (Edwards *et al.*, 2007; Green *et al.*, 2007), a hierarchical core-periphery network structure (Li and Schürhoff, 2012) and a drop in industry concentration ratios when transparency is forced up (Bessembinder *et al.*, 2006). Mandatory trade reporting appears to have the unintended effect of asset substitution towards securities that are outside the reach of regulation (Bessembinder and Maxwell, 2008).

[3] 'Amplify' is used here to convey the idea that, while identical information has been distributed only to a few participants, the market behaves as if the information has been distributed to many more.

[4] The approach is also related to gossip formation in networks; see, e.g. Bénézit *et al.* (2008, 2010*a,b*).

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

Second, agents in TIP (and in our experimental design) are given different, conditionally uncorrelated signals about the unknown common value of the asset. As such, when two agents meet privately, they are not in direct competition as far as information is concerned. Better even, they ought to be interested in learning from each other, so that they can use the acquired knowledge in future meetings with other participants. This is in contrast with Wolinsky (1990) where decentralised markets fail to reveal information when insiders all receive the same information.

This brings up a hitherto little appreciated distinction between various forms of asymmetric information in financial markets. In one setting, that of Wolinsky (1990), informed agents receive the same piece of information and the issue in terms of asset pricing is whether prices reveal the information as if more agents received privileged information than factually did. One should refer to this pricing feature as information amplification.[5] In a second setting, informed agents receive different bits of information and the issue there is whether prices reveal the union of the information sets. One ought to refer to this situation as information aggregation.[6]

TIP is about information aggregation only and our experimental design reflects this: every informed participant gets a different piece of information. Our goal is to test the theory of information percolation in a controlled setting. In our experiment, we aim at verifying the correctness of the core prediction of this theory, namely, that if information is freely available, prices aggregate the information and everyone eventually trades at the same, fully revealing, price.

TIP is a very stylised model. Unlike the general equilibrium models underlying earlier markets experiments on information aggregation (Plott and Sunder, 1988), the model spells out in detail what traders know about each other, how traders meet and how/what trade potentially takes place. Despite its stylised nature, the authors meant this model to apply to real-world dark markets. These markets do not even closely follow the same rules and involve far less common knowledge. Still, real-world dark markets share a key feature with TIP: bidding and trading is private. Our experiment, too, starts from this feature, as well as other crucial features of TIP. However, we do relax the rules of engagement, to make markets look more like real-world dark markets, and change the information structure to be more in line with previous experiments on information aggregation.

We could have taken an alternate route and tested the theory directly, without changes to the rules of engagement while ensuring literally the same information structure and matching technology as in the model. However, a literal test of TIP would have amounted to a test of whether subjects understand game theory, which others have been concerned with (Camerer, 2003), and would have most surely obtained a negative answer given the complex setup of the model. Like the original developers of TIP, we view TIP as a stylised model of the actual institutions out there.

[5] Another example of information amplification is in Huber *et al.* (2008). There, centralised markets (organised either as a call market or a continuous double auction) are studied. Insiders within a group receive the same information (about the liquidation value of the asset) and groups are differentiated by quality of information (information is nested, with better-quality signals encompassing lower-quality signals). How laboratory centralised markets amplify information is studied in detail in Plott and Sunder (1988) and in many follow-up papers, including Bossaerts *et al.* (2014).

[6] For an experimental study see Plott *et al.* (2003).

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

Because those institutions never work exactly as in the theory, the issue becomes: does TIP make valid predictions in a type of market loosely structured as in the theory?[7] This view relates to that of Gilboa *et al.* (2014), where the value of stylised ('stark') economic models is considered to be in their ability to isolate one or a few important factors, while successful economic forecasting depends on economists' ability to determine which models – though often, combination of models – are closest to describing the salient features of the real-world economic situation at hand. Here, we emulate salient features of real-world decentralised markets and ask to what extent TIP, a stylised, game-theoretic model that leaves out a large number of institutional details of real-world markets, can make valid predictions.

Because there are differences between the games in the theory and in the experiment, it is possible that the theory's predictions may not come about, not because the theory is false, but because the experiment deviated from the theory in a crucial way. Ideally, discrepancies should be anticipated and this is what we attempt to do here. Specifically, unlike in TIP, we do not insist that convergence to the fully revealing price is exponential and we spell out the reasons why. Instead, we hypothesise that transaction prices will fluctuate within bands of the fully revealing price. These bands are determined, either in a strong sense, by the precision of the average signal, or, in a narrow sense, by the precision of a private signal. Precision is defined as the maximum possible distance between the signal (average or private) and the true payoff of the security at hand and, therefore, loosely related to precision as statisticians define it, namely, the inverse variance of the posterior.

Besides the rules of engagement, we changed the information structure in order to be in line with recent experiments on information aggregation (Bossaerts *et al.*, 2014). TIP assumes aggregate risk in an economy populated by risk neutral agents. Yet, by now there is ample evidence that subjects in laboratory financial markets are risk averse (Bossaerts and Zame, 2008) in the absence of knowledge of subjects' risk aversion, an unknown risk premium would affect prices and, hence, market beliefs cannot readily be inferred from them. Risk neutral pricing can still be obtained, however, by eliminating aggregate risk. Unfortunately, the information structure in TIP is such that there is always aggregate risk, so we had to alter the information structure, in ways we make clear later.

The experiment provides qualified evidence in support of the theory. While transaction prices do not converge, let alone converge exponentially, they fluctuate within a band around the theoretical fully revealing price, which in this case is the average signal. Somewhat surprisingly, prices do so from the very first transaction. With

[7] There is a precedent for experiments that only have a loose link to the underlying game theory. Bossaerts *et al.* (2014) studies information aggregation in a continuous, centralised double auction institution and demonstrates that a highly stylised, game-theoretic price discovery model explains various aspects of price dynamics. This obtains despite the fact that the rules of the game in the theoretical model are much stricter while common knowledge is far more extensive, compared to the experimental setting. The use of game theory to shed light on empirical phenomena in a much looser setting is actually more prevalent in applied economics than in experimental economics. For instance, the prisoners dilemma and battle of the sexes games have been used to explain real-world advertising budgets (Rao *et al.*, 1995) and the public goods provision game is appealed to, so free-riding in alliances such as NATO can be understood (Sandler and Hartley, 2001). In these instances, the situation is far more complex than in the original games. Nevertheless, game theory predicts the essence of the observed phenomena.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

the exception of outliers – as traditionally defined in statistics – transaction prices stay within the weak band around the average signal. The interquartile range (25% percentile to 75% percentile) falls within the stronger bounds in nine out of 14 replications. As such, prices within the interquartile range reflect the greater precision of the average signal, not that of a private signal.

TIP is a theory of informational efficiency of prices. It predicts that prices will aggregate (exponentially fast) the dispersed information in the marketplace. Finance scholars focus on informational efficiency because prices are signals to aid managers in capital budgeting. Economists, on the other hand, are primarily interested in allocative efficiency. This concerns the extent to which gains from trade are exhausted. The issue of allocational efficiency comes about in TIP due to the different private valuations that buyers and sellers have for the traded asset. In our experiment, on the other hand, gains from trade emerge because of subjects' natural aversion to risk. Because there is no aggregate risk in our experiment, everyone in principle can trade to a risk-free position, and hence, allocational efficiency can be measured in terms of the reduction in risk subjects manage to obtain through trading.

We find that changes in efficiency range from as low as 0 and as high as 49% (measured as the reduction in risk and, in particular, as the reduction in the imbalance of a two-stock portfolio). Significantly, there appears to be no relationship between increases in allocational efficiency and the extent of information aggregation in pricing. Evidently, informational efficiency appears to be neither necessary nor sufficient for allocational improvements. We contrast these findings with those obtained in an experiment where everything is the same, except that participants trade in a centralised market. The centralised market generates even less allocational efficiency: it ranges from a low of −28% to a high of only 30%. How could one explain these low levels of allocational efficiency? We shall argue that they are the consequence of the well-known 'Hirshleifer effect' (Hirshleifer, 1971), whereby in complete markets with initially uninformed traders, the receipt of a public signal prior to trading impairs risk-sharing and therefore reduces welfare (in comparison to the no-public signal case).

The remainder of this article is organised as follows. The next Section introduces the experimental design. Section 2 discusses predictions. Section 3 presents the results. Section 4 provides a concluding discussion.

1. Experimental Design

1.1. *Sessions Overview*

We conducted five experimental sessions (approved by the local Institutional Review Boards). The first four constitute the main treatment and are labelled sessions 1 to 4. The last session, called session B, is an additional session that was conducted as a benchmark. Session B is discussed later on. Here we focus on the main treatment.

Each session consisted of three or four replications of the same situation. The participants in sessions 1 and 2 were undergraduate and graduate students from the California Institute of Technology; those in sessions 3 and 4 were from the University of Utah. Between 14 and 22 participated. Each participant had a experimental ID used

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

to log into the trading software. IDs were of the form U followed by a number from 1 to I, I being the number of participants in a session. Participants received a sign-up reward of \$5, which was theirs to keep no matter what happened in the experiment. Each session lasted approximately 2.5 hours in total and the average earnings from participation were \$55 per person, with a minimal and maximal payments of \$31.50 and \$78 for sessions 1 and 2, and \$33.10 and \$73.30 for sessions 3 and 4.

Upon arrival at the experimental laboratory, participants were asked to take seats in front of computer terminals and received a set of written instructions. The session was divided into several sections, the instruction section, the practice trading, and the actual trading replications section. We discuss each in turn.

1.1.1. *Instruction*

Each session began by the experimenter reading the instructions aloud, while also projecting them on a large screen. During the instruction period, if participants had any questions, they were asked to raise their hands and the experimenter would answer the questions. No other oral communication was allowed. Participants communicated with one another via the computer terminals through the trading software.

1.1.2. *Practice*

Following the instruction, there was one practice replication, where participants familiarised themselves with the rules of trading and the trading software (described in more detail below). The instruction and practice trading periods lasted approximately one hour.

1.1.3. *Replications*

After a short break, the participants were asked to log into the trading software for the actual trading rounds. Each experimental session had three or four identical replications. At the start of a replication, participants were endowed with an initial portfolio of securities and a private signal. They engaged in sending offers to other participants through the market software described in detail below. After conclusion of the trading within a replication, the payoff-relevant information was made public and participants' accounts immediately reflected the payout from that replication. As mentioned above, two of the three or four replications were chosen at random and participants received their experimental participation reward of \$5 in addition to the payoff from the two randomly chosen replications.

Below we describe in detail the market institution.

1.2. *Markets*

The experimental markets were designed as follows. In each replication, participants were endowed with five units of one of two types of securities, called stock X and stock Z. Participants could trade one of these securities, namely stock X, during a pre-determined time span (fifteen minutes). Short sales were allowed for up to 10 units of X. After market closure, both securities paid a liquidating dividend, determined by the random drawing of a number x between 0 and 10 (drawn from the discrete uniform

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

distribution over \$0 to \$10, with increments of \$0.10). The payoffs on X and Z were complementary: stock X paid x dollars, while Z paid 10-x dollars (all accounting was done in US dollars). Half the participants started out with only stock X, while the other half started out with only stock Z.

In the first two sessions, a third security was also available and, like stock X, it could be freely traded. This security, called Note, always paid a liquidating dividend of \$5 and, hence, was risk free. All agents started with a zero endowment of Notes. Short sales (of up to 10 units) of Notes were also allowed. Short sales of Notes allowed participants to obtain cash if they needed it to purchase stock X. In these sessions, participants initially holding only Z were also given a cash allocation of \$50. In the second set of (two) sessions, we simplified matters and eliminated the Notes. Instead, all participants, those holding only Z as well as those holding only X initially, were endowed with cash (\$50), so there was no need to short sell Notes in order to acquire cash. As a result, we could shorten the length of trading and run a fourth replication in Sessions 3 and 4.

In all sessions, participants' payment per replication depended on the change of their cash holdings. Notice that participants started with initial cash and securities and ended with cash only (because securities paid a liquidating dividend). That is, participant earnings at the end of each trading replication were determined by the liquidating payoffs on their final holdings of X, Z and Notes (if available), as well as by the change in cash through trading.

Because there was an equal number of X and Z in the (laboratory) economy, there was no aggregate risk and hence, according to standard asset pricing theory, prices in equilibrium should equal expected payoffs even if participants were risk and/or ambiguity averse. This design feature was added with the purpose of readily inferring how much information was revealed in prices; in equilibrium, prices should equal conditional expectations, without adjustment for risk.

Because some participants started out with only X and others with only Z, there were natural incentives to trade because of risk and/or ambiguity aversion. Those initially endowed with the non-traded asset Z would have to purchase stock X in order to reduce their exposure to risk/uncertainty; those initially endowed with the traded asset X would have to sell X in order to reduce their exposure. Participants were told explicitly about these incentives.

Each participant was given a private signal about x, generated as follows. First, a common signal S, within one dollar of x was generated but not disclosed to the participants. Based on S, private signals $S_i, i \in \{1, 2, \ldots, I\}$, were generated independently across the I participants, and uniformly between $S - 1$ (dollars) and $S + 1$ dollars, with increments of \$0.10. Thus, each participant i, based on his/her own private signal S_i only, could infer that x was within \$2 of S_i. Had participants been given the common signal S, they could have inferred that x was within one dollar of that signal. That is, the common signal S was more precise than each of the private signals.

Trade took place online, through an exchange platform called Flex-E-Markets (see http://www.flexemarkets.com). In sessions 1 to 4, the markets for X were all organised as bilateral private markets. In a single session, called session B and serving as a benchmark, the market for X was organised as a centralised exchange. In this

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

centralised public exchange setting, everyone had access to the book of all limit orders. When an order was submitted it was visible to all market participants. All transaction prices were made immediately available to everyone, through a table of transaction prices and also through a graph.

In sessions 1 to 4, we exploited the ability of our electronic market platform to allow for the arrangement of multiple private exchanges organised as continuous double auctions (as opposed to a single organised public exchange). In the private exchanges for security *X*, limit orders were addressed to a specific counter-party, known only by an online trading acronym such as U3. Acronyms revealed whether the trader was a net buyer of *X* (endowed with only *Z*) or a net seller of *X* (endowed with only *X*).[8] Only the counter-party was able to see the offer and react to it, either by submitting a counter-offer (a limit order that did not cross the original offer) or by accepting the offer (through a limit order that crossed the original offer). Limit orders and transaction prices remained privileged information only to the participants who submitted them and those who received them. The market in Notes in sessions 1 and 2 was a standard public continuous double auction. As mentioned before, participants could not trade stock *Z*; the market in stock *Z* remained closed throughout the duration of each replication of all sessions.

Participants were barred from communicating with each other in any other way than through order submission. As such, we emulated real-world dark markets in a highly stylised way (and in line with the TIP). Because order submission is fully electronic, all details of negotiations and trades are available to us, the experimenters.

In online Appendix A, we reproduce the instruction set presented to the participants (sessions 1 and 2; those for sessions 3, 4 and B were analogous).

2. Predictions

As mentioned in the Introduction, our experimental implementation of dark markets has a looser set of rules of engagement than in TIP. Details of the discrepancies can be found in online Appendix B. Here, we first state the main prediction to come out of TIP and then we discuss to what extent we expect differences in our experimental setting, and why.

In TIP, agents infer other agents' private signals (S_i), eventually recovering (and hence trading on) the average, $\bar{S}$, of the private signals, $\bar{S} = \sum S_i/I$. The private signals are within two dollars of the true outcome *x*, while the common signal is far more precise: it is within one dollar of the true *x*. The average private signal, $\bar{S}$, is a good estimate of the common signal *S*, so if we ignore the (small) estimation error, $\bar{S}$ is also within one dollar of the true *x*. In TIP, prices should eventually reflect $\bar{S}$. Thus, if one takes the theoretical model at heart, the following strong prediction can be made:

[8] For instance, odd numbered 'U's (U1, U3, . . .) would all be net buyers, while even-numbered 'U's (U2, U4, . . .) would all be net sellers.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

HYPOTHESIS 1. *The mean/median trade price converges at an exponential rate to the conditional expectation of x given* $\bar{S}$.[9]

Exponential convergence may be unrealistic because, in TIP, it hinges on a specific property of the information structure. Specifically, unknown outcomes (the 'parameter' in Bayesian parlance) is binary, so agents' posterior distributions can be summarised by a scalar (the chance that the parameter takes one of two values) and, hence, all an agent has learned can be reflected in the bid or offer (itself a scalar), provided that the bid/offer functions are strictly monotonic.

In contrast, our informational structure had to follow closely that of recent experiments on information aggregation in centralised markets, in order to avoid aggregate risk. There, the 'parameter' takes on several potential values (101, to be precise) and, hence, the posterior distribution can no longer be summarised by a scalar. Unfortunately, because bids/offers are scalars, this means that traders may at best be able to recover only one aspect of counterparties' posteriors, say, the posterior mean, leaving them to guess other aspects. Importantly, they need to know their counterparties' precision in order to update their own beliefs correctly (Bayesians weigh different signals using their respective precision). To assess their counterparties' precision, traders need to know how often they traded before, because each trade causes priors to be updated and therefore the precision of the posterior to increase. In a dark market, traders have no access to the trading experience of their partners, by definition.

Nevertheless, traders can build informed guesses of their trading partners' beliefs, as follows. Most conservatively, a trader could assume that his trading partner has never traded before and, hence, that the posterior belief reflected in her offer is based on her private signal only. This conservative approach leads to weak bounds on prices at which agents are willing to trade. Specifically, we hypothesised that in this case prices would fluctuate in a two-dollar band around the aggregate signal. Two dollars is the maximum that the true dividend can deviate from any private signal. That is, +/− two dollars is the 100% confidence interval based on the precision of a single signal. Because it is a 100% confidence interval, it is also a no-arbitrage interval; offers beyond this interval constitute an arbitrage opportunity (a sure way to make money). Within the bounds, trades can be rationalised in terms of some level of risk aversion and some pattern of holdings of risky securities.

This argument gives rise to the following hypothesis.

HYPOTHESIS 2 – WEAK. *Trade prices are within a weak band of two dollars around* $\bar{S}$.

More aggressively, an agent may be confident that he knows (or is faced with a partner who knows) the average private signal $\bar{S}$. Ignoring estimation error of $\bar{S}$ as an estimate of S, then this agent should realise that the true x is within one dollar. In that case, prices within one dollar of the average private signal do not allow for arbitrage opportunities and, as such, prices could at most fluctuate within a narrower bound of

[9] Notice that when x is close to the upper and lower boundaries of its [0, 10] range the conditional expectation of x does not necessarily equal $\bar{S}$.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

one dollar. Trades within the narrower bound could be rationalised in terms of some level of risk aversion and some initial portfolio holdings.

HYPOTHESIS 2 – STRONG. *Trade prices are within a strong band of one dollar around* $\bar{S}$.[10]

Figure 1 provides a graphical depiction of our predictions, plotted against the outcomes of the first replication in the first experimental session (referred to as Session 1-1). The following are shown below:

(*i*) the draw for x (and hence, the payoff on Stock X);
(*ii*) the common signal S from which the private signals were drawn;
(*iii*) the twenty-two private signals S_i;
(*iv*) the average of the private signals, $\bar{S}$;

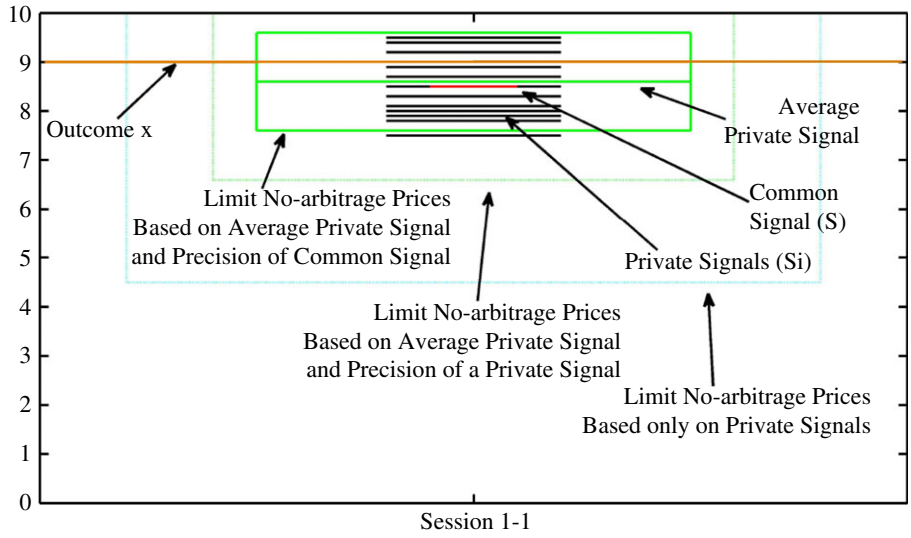

Fig. 1. *Graphical Depiction of Concepts*

Notes. This Figure presents a graphical depiction of concepts in the main hypotheses, using Session 1-1 as reference. The brown line indicates outcome x, around which the common signal S (short red line) was drawn. Individual signals S_i (black line segments) were drawn around S. Solid green box indicates no-arbitrage price bounds assuming precision of the common signal (x is within one dollar of S) and is centred at the average of the individual signals (green line). According to Hypothesis 1, prices should converge to the green line. According to Hypothesis 2 – strong, prices should fluctuate within solid green box. Larger dotted green box depicts no-arbitrage region of prices assuming precision of a private signal (+/− two dollars from average private signal) According to Hypothesis 2 – weak, prices should remain within the dotted green box. Prices beyond the (largest) dotted blue box imply arbitrage opportunities for all traders even if they were only to consider their own private signal (prices are more than two dollars away from any private signal). Colour figure can be viewed at wileyonlinelibrary.com.

[10] If agents take into account the estimation error in $\bar{S}$ as an estimate of the common value S, then they would accept trades at prices beyond the range of average private signal +/− one dollar. As such, our Hypothesis 2 – strong could be relaxed a bit.)

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

(*v*) the conditional expectation of x given the average private signal;
(*vi*) strong (narrow) no-arbitrage price bounds;
(*vii*) weak (wider) no-arbitrage bounds; and
(*viii*) bounds within which the price can be rationalised in terms of some trader's private signal and its accompanying precision (x is within two dollars).

3. Results

Table 1 lists initial allocations. There were two types of traders, distinguished by whether they initially held only the traded security X or the non-traded security Z. There were an equal number of each type and, because payoffs on securities X and Z were complementary, there was no aggregate risk. Subjects within a session rotated types across replications.

Table 2 displays details on the realisations of the dividend (x), the common signal (S) and the average signal ($\bar{S}$), per replication. Numbers in parentheses are:

(*i*) the theoretical standard error of $\bar{S}$ as an estimate of the common signal S (based on the fact that the individual signals S_i are drawn uniformly between $S - 1$ and $S + 1$); and
(*ii*) the sample standard error of $\bar{S}$.

The error of the average private signal as an estimate of S is usually small (at most \$0.17). Theoretically, this means that the average private signal is not exactly within one dollar of the true x but may be a bit further away. Based on standard statistical reasoning, one can state that, in the worst case, there is a 95% chance that the average private signal is within 1.33 dollars ($=1 + 1.96 \times 0.17$) of the true x.

Figures 2–5 display the evolution of offer prices (bids, asks) and trades, across the entire duration of 14 replications over the four sessions, 1 to 4, during the first half of the period that markets were open and during the second half. Prices that conform to Hypothesis 2 – weak lie within the darker shaded area; those that conform to Hypothesis 2 – strong are inside the lighter shaded area. Quick inspection of the plots reveals that, with a few exceptions, trade prices are within the acceptable region for Hypothesis 2 – weak, and in many replications stay within the acceptable region for Hypothesis

Table 1
Summary Statistics

Session	Type	Stock X	Stock Z	Notes	Cash
1 and 2	I	5	0	0	\$0
	II	0	5	0	\$50
3 and 4	I	5	0	NA	\$50
	II	0	5	NA	\$50

Notes. This Table displays the distribution of holdings of Stock, Notes (NA = Not Available) and Cash across the two types of participants. In all replications, there were an equal number of participants of each type. Notes were not available in Sessions 3 and 4, instead cash of \$50 was given to type I of participants.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

Table 2

Signals and Outcomes

Session (number of participants)	Replication	x	S	Average (SE)
1 (22)	1	9.0	8.5	8.6 (0.13/0.17)
	2	8.0	8.0	8.1 (0.13/0.17)
	3	2.0	2.1	1.8 (0.13/0.14)
2 (12)	1	3.2	2.4	2.4 (0.17/0.13)
	2	4.1	4.2	4.4 (0.17/0.13)
	3	6.8	7.7	7.6 (0.17/0.13)
3 (14)	1	9.1	9.2	9.2 (0.15/0.13)
	2	1.3	2.3	2.2 (0.15/0.11)
	3	6.4	7.4	7.5 (0.15/0.10)
	4	1.0	0.4	0.4 (0.15/0.13)
4 (14)	1	8.5	9.0	9.1 (0.15/0.12)
	2	9.4	8.5	8.6 (0.15/0.13)
	3	3.4	3.0	3.1 (0.15/0.11)
	4	2.5	1.6	1.5 (0.15/0.11)

Notes. For all replications across the four experimental sessions, this Table displays the outcomes x (which determined the payoffs on Stocks X and Z), the common signal S, and the average of the private signals, $\bar{S}$, per replication. In parentheses are the standard errors of $\bar{S}$ (left: standard error based on the assumed uniform distribution of private signals; right: sample standard error).

2 – strong except for a few outliers (Session-Replications 1-2, 1-3, 2-2, 2-3, 3-2, 3-4, 4-2, 4-3, 4-4). We turn to more complete evidence later.

Also evident from the Figures is the lack of trend in transaction prices. In other words, there does not seem to be 'convergence'. Instead, the distribution of prices early in a replication is hardly different from that later on. When they stay within the acceptable region for Hypothesis 2 – strong, prices are within this range even for early transactions. Table 3 shows in its first column the number of orders (bid and asks) per trade in each replication of the four sessions and in its second column the number of orders before the very first transaction. On average it takes about 12 orders for each transaction. However, it takes on average 50 orders before the first trade. It is hard to evaluate this number against TIP. While orders do not necessarily generate trades, each order reveals the information that the submitter has. However, in TIP, each trader can send only one order to a counter-party when they meet, while in our experimental setting, she can revise her order multiple times. Unfortunately, there are no clear predictions in TIP about changes over time in the likelihood that orders lead to trades.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

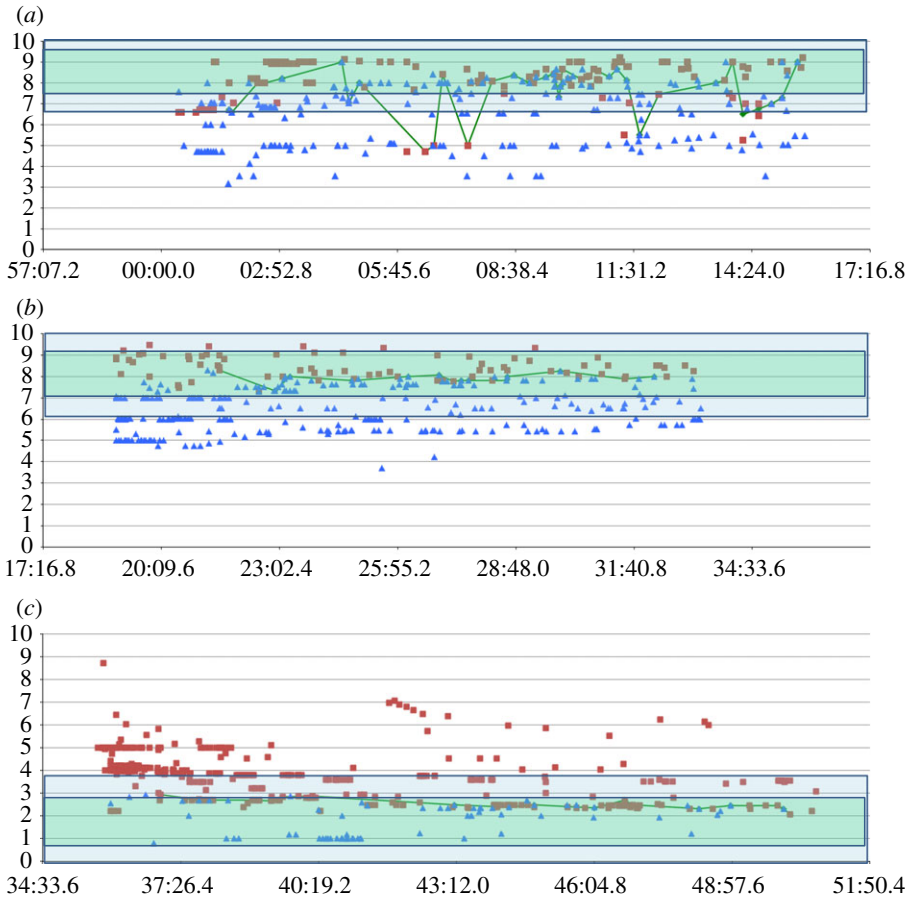

Fig. 2. *Trade Prices in Sessions 1-1 to 1-3*

Notes. This Figure displays the evolution of trade prices (diamonds, connected), bids (arrows) and asks (squares) in three replications (Sessions 1-1 (*a*), 1-2 (*b*) and 1-3 (*c*)). The darker shaded region contains prices that satisfy absence of arbitrage conditional on the average private signal and assuming the precision of the common signal (Hypothesis 2 – strong); the lighter shaded region is the same but assuming the precision of a private signal (Hypothesis 2 – weak). Colour figure can be viewed at wileyonlinelibrary.com.

In addition, not only is there no drift, price observations seem to be drawn independently over time. In the majority of replications, variability fall somewhat as time passes. So, convergence in our data appears at best to be a matter of reduction in price fluctuations. This evidence is contrary to the strict interpretation of TIP as expressed in Hypothesis 1.

The lack of time series dependence in prices also validates the simple test statistics about mean and median prices that we report next. These test statistics require independent observations. Table 4 provides descriptive statistics of trade prices during entire replications and across their first and second halves. For reference, the second column repeats the average private signal $\bar{S}$. Included, in parentheses, are tests of whether the mean trade price (left number) or the median trade price (right number)

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

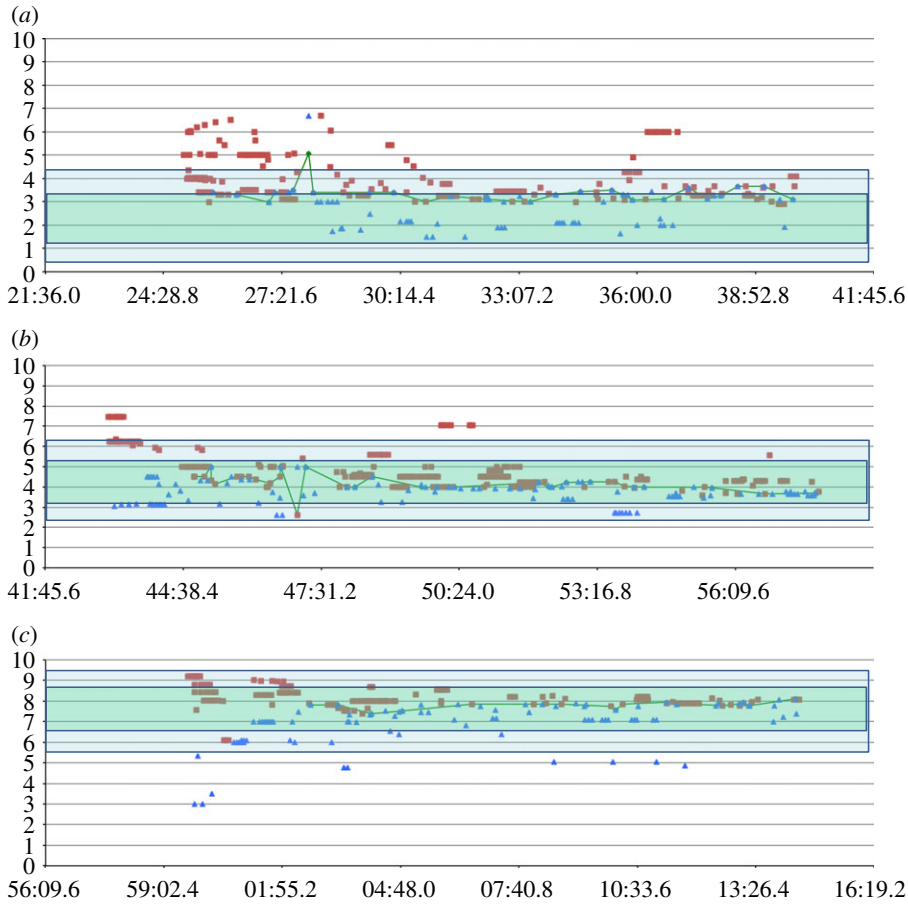

Fig. 3. *Trade Prices in Sessions 2-1 to 2-3*

Notes. This Figure displays the evolution of trade prices (diamonds, connected), bids (arrows) and asks (squares) in three replications (Sessions 2-1 (*a*), 2-2 (*b*) and 2-3 (*c*)). The darker shaded region contain prices that satisfy absence of arbitrage conditional on the average private signal and assuming the precision of the common signal (Hypothesis 2 – strong); the lighter shaded region is the same but assuming the precision of a private signal (Hypothesis 2 – weak). Colour figure can be viewed at wileyonlinelibrary.com.

is significantly different from the average private signal. In the former case, the Table lists the z-statistic of the null that the average trade price equals the average private signal; in the latter case, one asterisk (two asterisks) indicates that the proportion of transactions at prices above the average private signal is significantly different from 0.50 at the 10% (1%) level (two-sided test).

In all replications (and in most halves), mean and median trade prices are significantly different from the average private signal. Even in the second half of replications, mean and median trade prices are significantly different from the average private signal. Exceptions are the 3rd and 4th replications of Session 4. Further confirming evidence against convergence is the finding that mean and median prices are insignificantly different from the average private signal during the first half of 1-2,

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

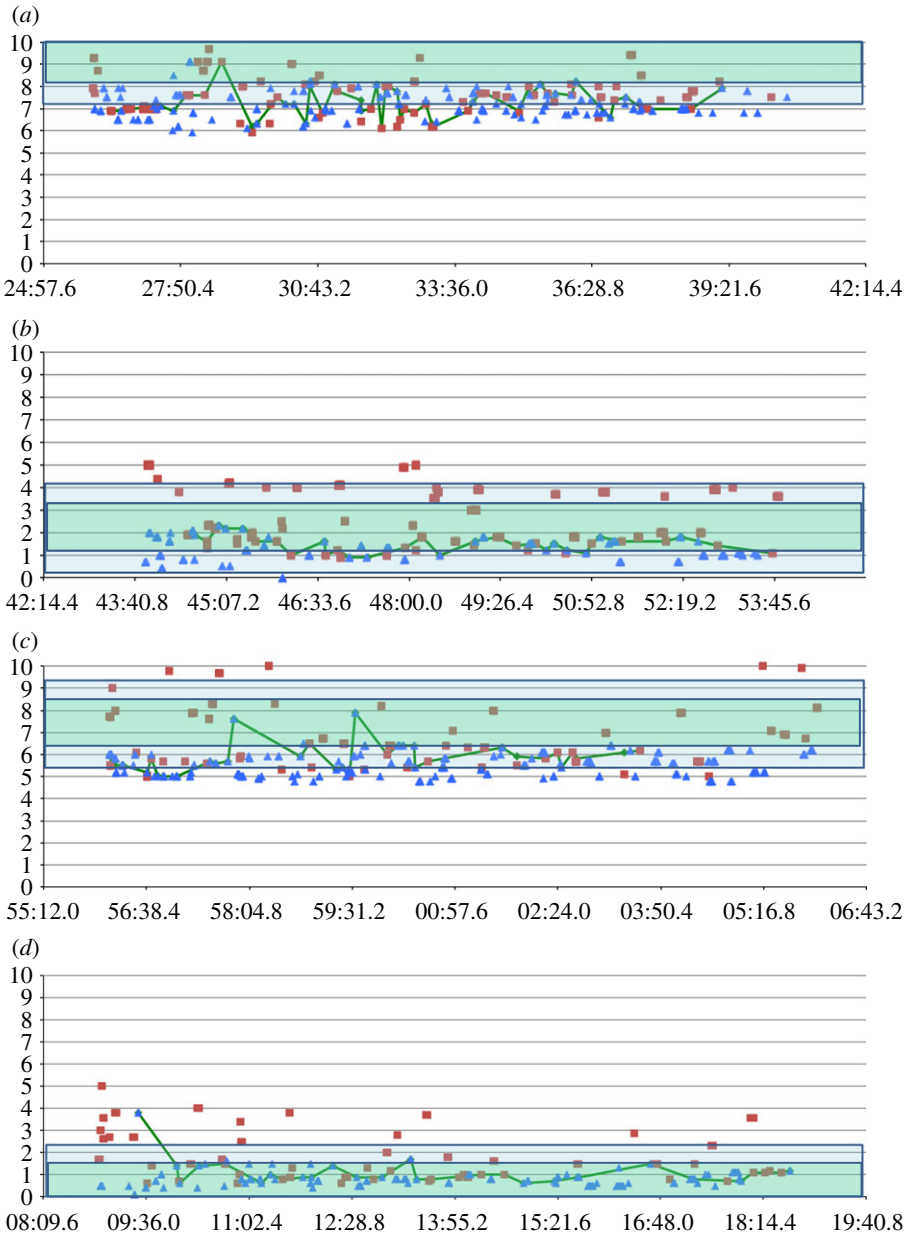

Fig. 4. *Trade Prices in Sessions 3-1 to 3-4*

Notes. The Figure displays the evolution of trade prices (diamonds, connected), bids (arrows) and asks (squares) in four replications (Sessions 3-1 (*a*), 3-2 (*b*), 3-3 (*c*), and 3-4 (*d*)). The darker shaded region contain prices that satisfy absence of arbitrage conditional on the average private signal and assuming the precision of the common signal (Hypothesis 2 – strong); the lighter shaded region is the same but assuming the precision of a private signal (Hypothesis 2 – weak). Colour figure can be viewed at wileyonlinelibrary.com.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

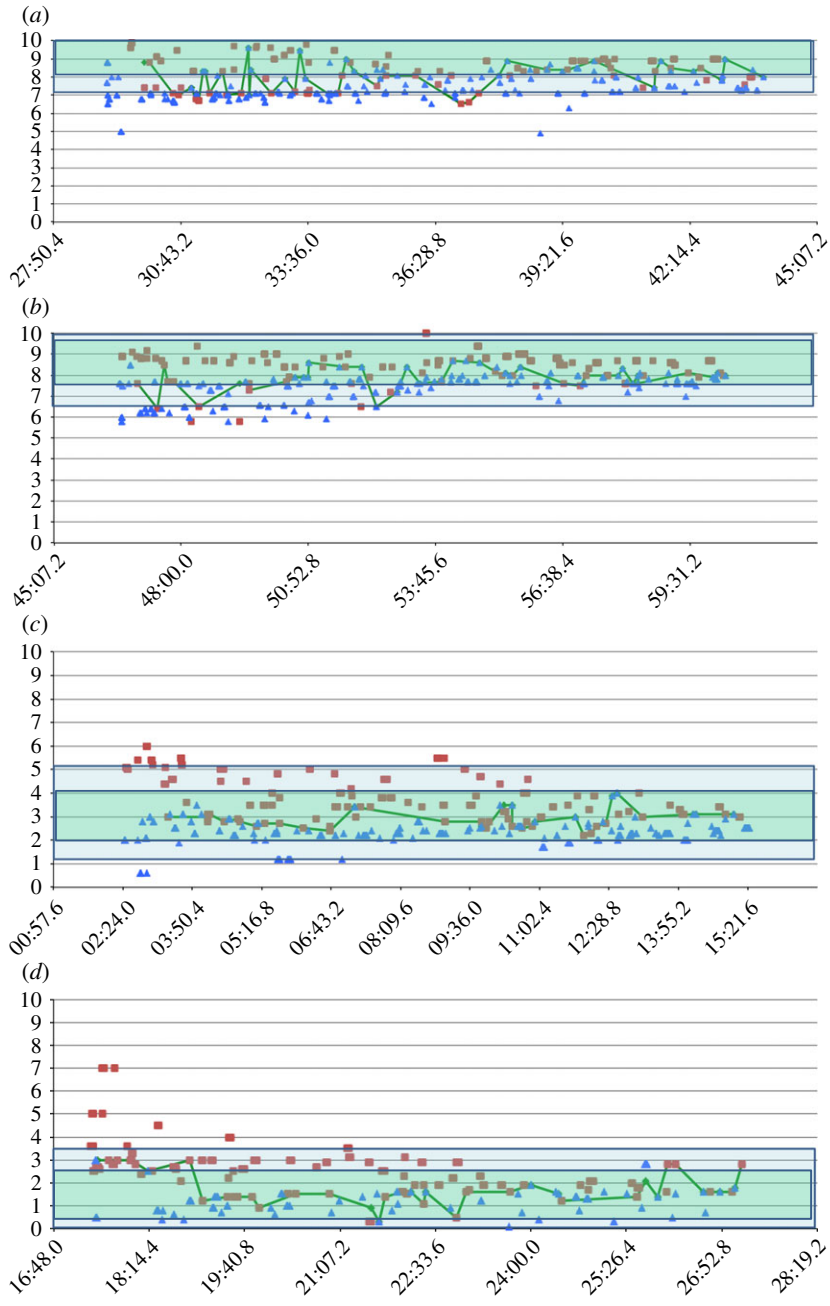

Fig. 5. *Trade Prices in Sessions 4-1 to 4-4*

Notes. This Figure displays the evolution of trade prices (diamonds, connected), bids (arrows) and asks (squares) in four replications (Sessions 4-1 (*a*), 4-2 (*b*), 4-3 (*c*), and 4-4 (*d*)). The darker shaded region contain prices that satisfy absence of arbitrage conditional on the average private signal and assuming the precision of the common signal (Hypothesis 2 – strong); the lighter shaded region is the same but assuming the precision of a private signal (Hypothesis 2 – weak). Colour figure can be viewed at wileyonlinelibrary.com.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

Table 3

Order Volume

Session (number of participants)	Replication	Orders/trade	Orders before first trade
1	1	8.4	38
(22)	2	24.4	112
	3	19.7	88
2	1	9.8	30
(12)	2	8.3	49
	3	13.4	73
3	1	7.9	34
(14)	2	11.3	65
	3	11.1	5
	4	12.4	64
4	1	9.5	55
(14)	2	12.1	29
	3	14.9	64
	4	8.5	15

Notes. This Table displays the number of orders (bids, asks) per transaction and number of orders (bids, asks) before first transaction, per replication.

2-2 and 2-3, while they differ significantly during the second half of those replications. Still, in a majority of replications, the interquartile range is smaller over the second half than over the first half.

Closer inspection of Table 4 does reveal that prices are in line with Hypothesis 2 – weak, and often even Hypothesis 2 – strong. There are only two replications out of 14 (3-1 and 3-3), where the median price is outside the bounds imposed by Hypothesis 2 – strong (median prices are more than one dollar different from the average private signal $\bar{S}$), and in only one of these exceptions (3-1) is the median price outside the bound imposed by Hypothesis 2 – weak (more than two dollars from $\bar{S}$). Not only the median but also most of the distribution of trade prices stays within the bounds of Hypothesis 2 – weak and Hypothesis 2 – strong: in nine replications the interquartile range stays within the bounds of Hypothesis 2 – strong; with the sole exception 3-1, the interquartile range never goes outside the range of acceptable prices for Hypothesis 2 – weak.

Could participants really have inferred the average private signal? Table 4 lists the average number of trades per capita. Across all replications, participants traded 4.6 times on average. Hence, each participant received 4.6 signals on average. Ignoring that there were a finite number of participants, this means that after the last trade, each participant had information that was based on $2^{4.6} \approx 24$ signals on average. So indeed, participants could in principle have inferred the average signal, which was actually not even based on 24 private signals, but only 22 (replications in Session 1) or less (all other replications).

Figures 6 and 7 provide graphical depictions of the evidence in Table 4. Shown are boxplots of all traded prices for each replication, the average private signals (solid line) and price ranges that conform to Hypothesis 2 – strong (smaller boxes). The Figures also display boxes (dashed, larger) that indicate the range of acceptable prices given

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

Table 4

Trade Prices

Session-replication (average trades per subject)	Average private signal $\bar{S}$	Trade price mean/median		
		All	Early	Late
1-1	8.6	7.6/8.0	7.2/8.0	7.9/8.1
(3.9)		(−5.4)/(**)	(−4.2)/(**)	(−3.7)/(**)
		[7,8.2]	[6.2,8.1]	[7.3,8.6]
1-2	8.1	7.9/8.0	7.9/8.0	7.9/7.9
(1.4)		(−2.9)/(**)	(−1.3)/( )	(−3.2)/(*)
		[7.8,8.0]	[7.7,8.1]	[7.8,8.0]
1-3	1.8	2.5/2.5	2.8/2.7	2.4/2.4
(1.9)		(14.7)/(**)	(21.8)/(**)	(14.2)/(**)
		[2.4,2.7]	[2.4,2.9]	[2.3,2.5]
2-1	2.4	3.4/3.3	3.6/3.4	3.3/3.3
(4.3)		(12.1)/(**)	(4.6)/(**)	(18.9)/(**)
		[3.1,3.4]	[3.3,3.4]	[3.1,3.4]
2-2	4.4	4.2/4.2	4.3/4.0	4.0/4.0
(5.8)		(−2.9)/(**)	(−0.5)/( )	(−8.4)/(**)
		[4.0,4.5]	[4.1,4.5]	[4.0,4.2]
2-3	7.6	7.8/7.8	7.7/7.8	7.8/7.9
(2.7)		(4.3)/(**)	(0.8)/( )	(6.0)/(**)
		[7.8,7.9]	[7.5,7.8]	[7.8,7.9]
3-1	9.2	7.3/7.2	7.3/7.1	7.3/7.3
(7.6)		(−22.7)/(**)	(−14.6)/(**)	(−17.3)/(**)
		[6.9,7.7]	[6.9,7.6]	[6.9,7.7]
3-2	2.2	1.5/1.55	1.5/1.6	1.4/1.5
(4.6)		(−10.2)/(**)	(−5.5)/(**)	(−11.6)/(**)
		[1.2,1.8]	[1.0,1.9]	[1.2,1.6]
3-3	7.5	5.9/5.8	5.7/5.7	6.1/6.1
(5.3)		(−17.7)/(**)	(−8.1)/(**)	(−10.7)/(**)
		[7.8,7.9]	[7.5,7.8]	[7.8,7.9]
3-4	0.4	1.1/0.9	1.2/0.9	1.0/0.95
(4.3)		(6.5)/(**)	(4.3)/(**)	(8.5)/(**)
		[0.8,1.175]	[0.8,1.4]	[0.8,1.1]
4-1	9.1	7.9/8.1	7.7/7.4	8.1/8.1
(7.3)		(−10.9)/(**)	(−7.6)/(**)	(−7.8)/(**)
		[7.1,8.4]	[7.075,8.3]	[7.65,8.5]
4-2	8.6	7.8/7.9	7.6/7.7	8.0/8.0
(6.0)		(−9.4)/(**)	(−6.2)/(**)	(−8.5)/(**)
		[7.6,8.1]	[7.6,7.9]	[7.75,8.1]
4-3	3.1	2.9/2.8	2.9/2.8	2.9/2.8
(5.3)		(−3.1)/(*)	(−3.3)/(*)	(−1.8)/( )
		[7.8,7.9]	[7.5,7.8]	[7.8,7.9]
4-4	1.5	1.8/1.6	1.9/1.5	1.7/1.6
(5.7)		(2.5)/(*)	(2.0)/( )	(1.6)/(*)
		[1.4,2.5]	[1.4,3.0]	[1.5,1.8]

Notes. For each replication, this Table displays the mean and median trade prices, over the entire replication, over the first half ('Early') and over the second half ('Late'). In parentheses: z-statistic of the null that the mean trade price equals the average private signal/binomial test of whether the median equals the average private signal (** Significant at p = 0.01; * Significant at p = 0.1; otherwise p > 0.1; all p values are one-sided). In square brackets: inter-quartile range (i.e. 25% quartile and 75% quartile).

private signals. This range stretches from the highest value of x that the subject with the highest private signal would infer, down to the lowest value of x that the subject with the lowest private signal would infer. If a trade happened at a price outside this box,

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

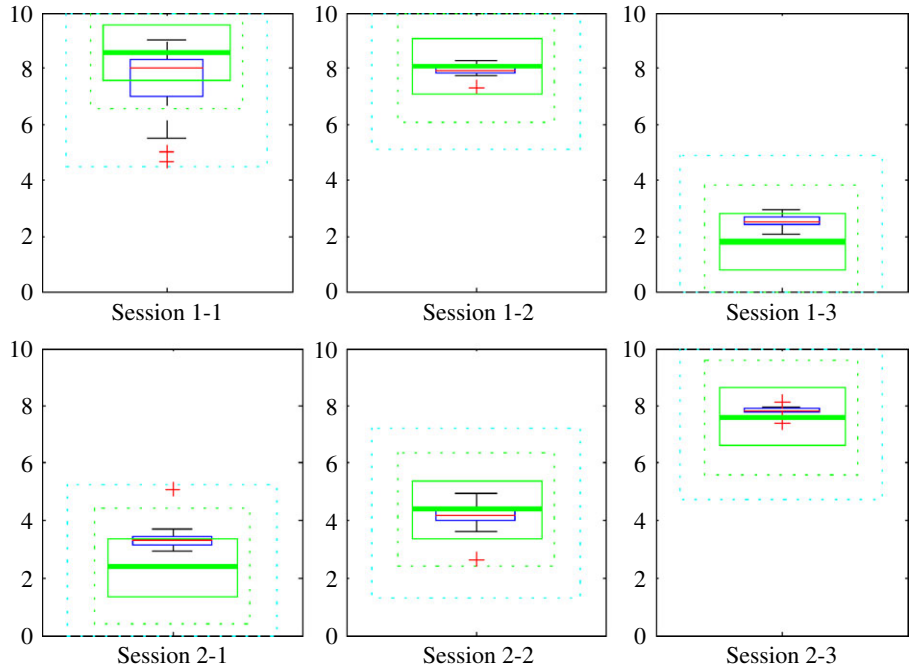

Fig. 6. *Box Plots of Trade Prices in Sessions 1-1 to 2-3*

Notes. This Figure presents the box plots of transaction prices in the six replications from Sessions 1 and 2. Each blue box represents the interquartile range (25th to 75th percentile) with the red line indicating the median transaction price. Black whiskers are drawn such that the range between them covers approximately 99.3% of the data (assuming a gaussian distribution). The data that fall outside the range of the whiskers are considered outliers and indicated with a red plus sign. The bold green line indicates the average private signal. The solid green box represents the range of prices consistent with no arbitrage and assuming that traders know the average private signal and attribute to it the precision of the common signal (the true value of x is within one dollar of the common signal). The larger dotted-green box depicts the same but assuming a precision equal to that of a private signal (the true value of x is within two dollars of a private signal). Finally the largest dotted box depicts the range of prices that are within two dollars from the minimal and maximal of all signal. Colour figure can be viewed at wileyonlinelibrary.com.

one of the two parties to the trade for sure was making a mistake. Notice that there is only one trade of this kind in 3-4, and a small number in 3-1.

Finally, Table 5 displays the percentage increase in allocational efficiency for each replication. Assuming all subjects are risk averse and because there is no aggregate risk, full efficiency is obtained when everyone ends up holding a risk-free portfolio. This requires holders of five (non-traded) Stock Z to buy five Stock X, while participants who were initially endowed with (five) Stock X should sell everything. Efficiency increases can, therefore, be measured as one minus the ratio of average absolute differences in holdings of X and Z at the end of trading over average absolute differences in initial holdings of X and Z. (The latter equals 5). Table 5 shows that efficiency increases lie somewhere between 0% and 49%. Comparison of Tables 5 (efficiency gains) and 4

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

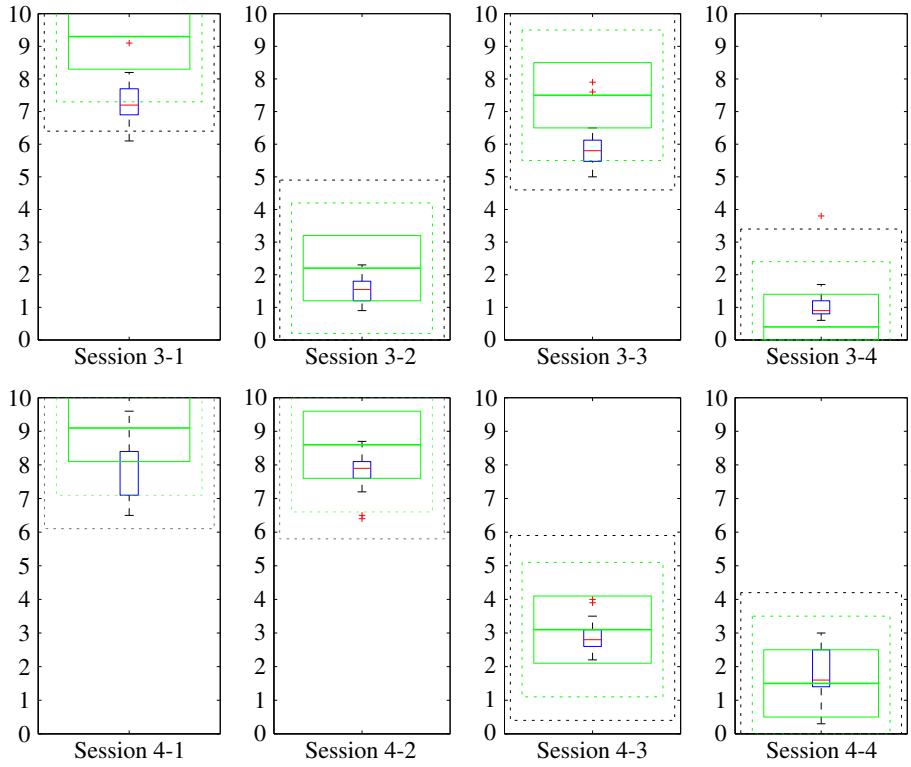

Fig. 7. *Box Plots of Trade Prices in Sessions 3-1 to 4-4*

Notes. This Figure presents the box plots of transaction prices in the eight replications from Sessions 3 and 4. Each blue box represents the interquartile range (25th to 75th percentile) with the red line indicating the median transaction price. Black whiskers are drawn such that the range between them covers approximately 99.3% of the data (assuming a gaussian distribution). The data that falls outside the range of the whiskers are considered outliers and indicated with a red plus sign. The bold green line indicates the average private signal. The solid green box represents the range of prices consistent with no arbitrage and assuming that traders know the average private signal and attribute to it the precision of the common signal (the true value of x is within one dollar of the common signal). The larger dotted-green box depicts the same but assuming a precision equal to that of a private signal (the true value of x is within two dollars of a private signal). Finally the largest dotted box depicts the range of prices that are within two dollars from the minimal and maximal of all signals. Colour figure can be viewed at wileyonlinelibrary.com.

(price quality) suggests that there is no correlation between efficiency gains (i.e. extent to which subjects traded to risk-free positions) and quality of pricing (i.e. the extent to which prices correctly aggregated the average private signal). For instance, the lowest efficiency gains (0%, in Session 3-2) obtain when the interquartile range of trade prices is within the bounds of Hypothesis 2 - strong throughout trading. The highest efficiency gains (49%, in Session 3-1) are recorded when most of the interquartile range of trade prices was outside the bounds of Hypothesis 2 – strong; some prices were even outside the much weaker bounds of Hypothesis 2 – weak.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

Table 5
Allocative Efficiency

	Replication			
Session	1	2	3	4
1	22	30	12	
2	27	40	47	
3	49	0	40	23
4	34	17	31	20

Notes. This Table displays the percentage increase in allocative efficiency from initial to final holdings. Full efficiency is achieved when average absolute differences in holdings of Stock *X* and Stock *Z* reach zero. Efficiency increases are measured as one minus the ratio of average absolute differences in unit holdings of Stock *X* and *Z* at the end of trading over that before trading.

4. Discussion

We reported results from fourteen replications where participants could trade a risky security through 'dark markets'. Negotiations and trade remained privately known only to the parties involved. Everyone had a private signal about the liquidating payoff of the traded security, and everyone knew that, absent these private signals, buyers were in the market because they wanted to offset the risk of a non-traded security they were holding, and sellers were in the market because they wanted to reduce the risk of their allocation of traded securities.

Despite the absence of aggregate risk, mean/median transaction prices did not converge to expected payoff conditional on the average private signal (Hypothesis 1). This may be an indication that prices did not correctly aggregate the dispersed private information. Still, prices fluctuated within narrow no-arbitrage bands around the average of the private signals. That is, prices were benchmarked off the average private signal, which was substantially more informative than a single private signal. For example, in nine of the 14 replications, the interquartile range of transaction prices stayed within the region of prices that are consistent with absence of arbitrage assuming that everyone knows the average private signal and assumes its precision equals that of the common signal it is reflecting (Hypothesis 2 – strong). Prices, therefore, perfectly revealed the average private signal because they stayed within bounds that were far tighter than those implied by traders' private signals alone. With one exception, the interquartile range of trade prices lay within no-arbitrage bounds of Hypothesis 2 – weak. In the four replications where prices stayed outside the bounds of Hypothesis 2 – strong but within the bounds of Hypothesis 2 – weak, trades appeared to be based on beliefs that equalled the average of the private signals, yet with a precision equal to that of a single private signal.

The latter is significant. As mentioned before, TIP makes strong assumptions on the nature of the unknown outcome ('parameter'), which imply that two parties to a trade can infer all aspects about each other's posterior from offers. Not only will trading partners know each other's best estimate of the common value of the traded asset (the posterior mean), but also the precision of their beliefs. In more realistic settings, this would require traders to know how often their counterparties had traded before. Because markets are opaque, it is not obvious how traders could ever acquire this knowledge.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

Evidently, our subjects did manage to read counterparties' information from offers but their trades were mostly consistent with having taken a conservative position: they ascribed low precision to their own posterior, to the point that, after many trades, everyone effectively knew what the average private signal was, yet nobody was aware that it was the average private signal. Instead, prices and offers were consistent with the traders taking the information revealed in the counterparty's offer as reflecting only her own signal, as if she had never traded before. The interquartile range of trade prices revealed a more aggressive stance. Because prices within this range generally fell within narrow no-arbitrage bounds, they reflected the belief that one had learned the aggregate signal.

Altogether, our experiment provides qualified support for TIP. While prices did not converge to conditional expectations given the aggregate information in the marketplace, they fluctuated within narrow no-arbitrage bands around this conditional expectation. One cannot emphasise enough the import of the key findings. At no time could any participant verify that they were actually trading within narrow no-arbitrage price levels given the aggregate information. While everyone traded 4.6 times on average, individual participants could not have been aware that others were trading as intensively and, hence, could not have known that the information revealed by offers of their counter-parties eventually reflected far more than 4.6 private signals.

Participants may at times have attributed incorrect precision to the information they ended up collecting through trading. This may explain why there was no correlation between allocative and informational efficiency. We noted how replications with the highest and lowest increase in allocational efficiency are also replications with low and high level of information aggregation, respectively. Evidently, there is a disconnect between informational and allocational efficiency; we suspect that inability to assess the precision of one's information correctly is the cause.

Overall, allocational efficiency gains are not high (at most 49%). This seems to contrast with the amount of trading: it takes five trades (of one unit) to trade to full efficiency and on average participants traded 4.6 times. Closer inspection of the data, however, revealed that individuals often traded in the opposite direction to that required by allocational efficiency. Presumably this reflects that traders perceived prices not to equal expectations (conditional on their own information and the information revealed in the price), thus incentivising them to move away from establishing risk-free positions.

One should be cautious, however, about our measure of improvements in allocational efficiency. Full efficiency is obtained only if all participants are risk averse in the economic sense of the term (decreasing marginal utility). While risk aversion is typical in laboratory financial markets like ours, there is quite a bit of cross-sectional variation. Subjects almost invariably do avoid risk but their choices often do not reflect the demand for diversification that their risk avoidance would imply (Bossaerts *et al.*, 2007). Consequently, without a better understanding of the nature of risk aversion that subjects exhibit in the laboratory, our measures of improvement in allocational efficiency remains very crude.

Therefore, our experiment, and any other that renders revealing prices, cannot really be used to study allocative efficiency. If it had been common knowledge that prices were fully revealing, then the amount of risk that remained was miniscule. Indeed, the payoff of the risky assets was at most about one dollar above or below trade

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

prices. This should be compared to the risk agents faced before trading started: based on the private signal alone, the true payoff was as much as two dollars above or below the private signal. Of course, since participants evidently were not always aware that trade prices reflected the average of the private signals, risk must have been perceived to be somewhere in between these extremes.

To put things into perspective, we replicated our experiment but instead of forcing trade through private markets, in session B, we opened a standard, centralised market (continuous double-sided open book system, as implemented in Flex-E-Markets). There, everyone could see all traded prices. Across ten replications, prices revealed the average of the private signals equally well as in the private-markets treatment; price variability was much lower though than when markets were private. Figures 8 and 9 plots the evolution of prices and corresponding boxplots in four typical replications. Because everyone could see all traded prices and price volatility was low, participants should have been better aware of how precise prices reflected information. Therefore, the remaining risk was commensurately reduced and, as a result, participants had far less incentive to trade to risk-free positions. Indeed, allocational efficiency gains, as computed above, ranged from a low of −28% to a maximum of only 30%.

The findings about allocational efficiency remind one of the 'Hirshleifer Effect' (Hirshleifer, 1971). When there are strong incentives to trade based on private information but prices would reveal the true value of individual endowments, competitive markets may make many agents worse off. In Hirshleifer's model, agents start with endowments which they do not know the true value of and, being risk averse, they wish to obtain insurance before the true value is revealed. However, if equilibrium prices end up revealing the true value of the endowments, there is no chance anymore for them to insure and, while some may be better off (if they happened to have valuable endowments), others will be far worse off (if their endowments were revealed to be of low value). *Ex ante*, risk-averse agents would like insurance; *ex post*, nobody can insure.

In our experiment, we observe a similar behaviour. Based on private information alone, risk-averse participants would like to insure. Because trading reveals information and, hence, the value of individual endowments, there is less scope for insurance. It is worthwhile to point out that subjects often did complain about the lack of 'fairness' in our setting when they were allocated endowments that were revealed to be of low value before they had the chance to trade away to more 'equitable' positions.

In our experiment, everyone had free access to private information. An interesting extension we plan to work on in future experiments is to investigate what happens when private information is costly to acquire. In centralised markets, Sunder (1992) has shown that this generates the Grossman–Stiglitz paradox (Grossman and Stiglitz, 1980): when information is auctioned off, prices (of information) drop to zero and prices are fully revealing, while when information must be acquired at a fixed cost, prices become more noisy. In contrast, in dark markets, TIP predicts that costly information acquisition does not necessarily stand in the way of full revelation. This is, among others, because information acquisition exhibits strategic complementarities (Duffie *et al.*, 2014): if one agent knows that her counterparty has acquired information, it is in her interest to acquire

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

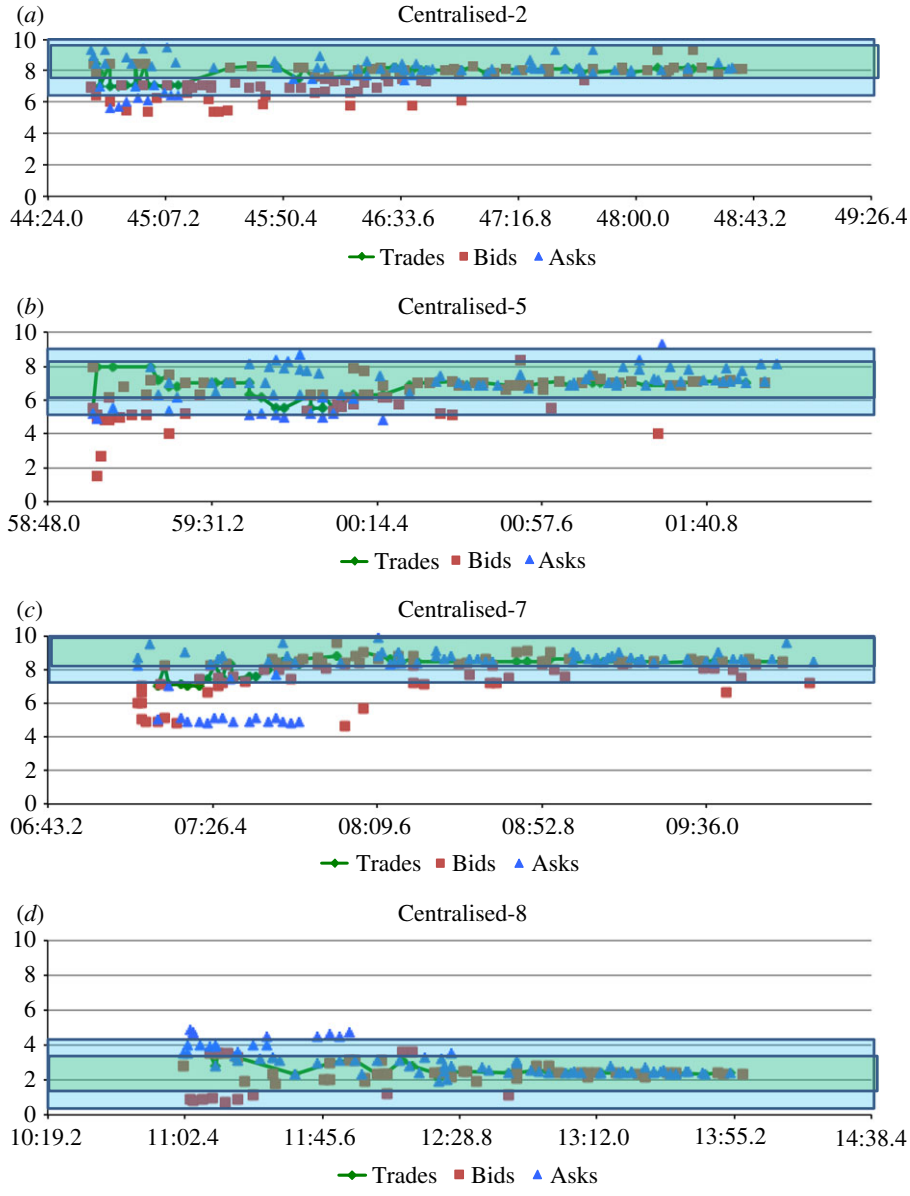

Fig. 8. *Trade Prices in Replications with Centralised Markets*

Notes. This Figure displays volution of trade prices (diamonds, connected), bids (arrows) and asks (squares) in four replications with centralised markets. The darker shaded region contain prices that satisfy absence of arbitrage conditional on the average private signal and assuming the precision of the common signal; the lighter shaded region is the same but assuming the precision of a private signal. Colour figure can be viewed at wileyonlinelibrary.com.

information as well. The fact that prices fluctuate in our experimental dark markets (albeit in narrow bands) should further provide incentives to enhance the informational quality of prices.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

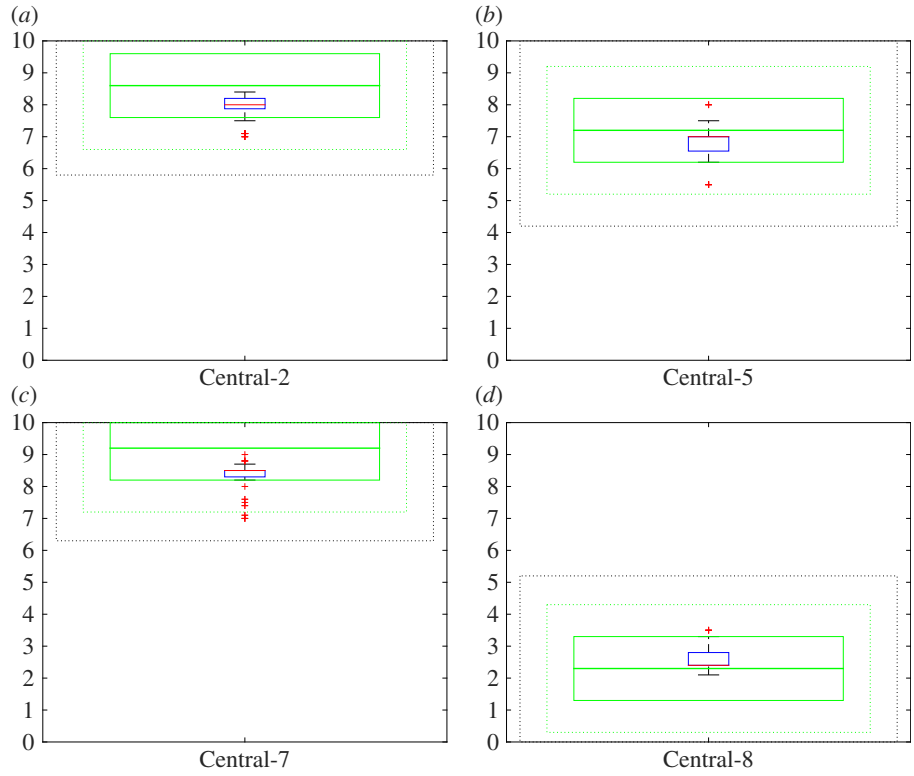

Fig. 9. *Box Plots of Trade Prices in Replications with Centralised Markets*

Notes. This Figure presents the box plots of transaction prices in four replications of the centralised markets treatment. Each blue box represents the interquartile range (25th to 75th percentile) with the red line indicating the median transaction price. Black whiskers are drawn such that the range between them covers approximately 99.3% of the data (assuming a gaussian distribution). The data that fall outside the range of the whiskers are considered outliers and indicated with a red plus sign. The bold green line indicates the average private signal. The solid green box represents the range of prices consistent with no arbitrage and assuming that traders know the average private signal and attribute to it the precision of the common signal (the true value of x is within one dollar of the common signal). The larger dotted-green box depicts the same but assuming a precision equal to that of a private signal (the true value of x is within two dollars of a private signal). Finallly the largest dotted box depicts the range of prices that are within two dollars from the minimal and maximal of all signal. Colour figure can be viewed at wileyonlinelibrary.com.

From a policy perspective, the recent tendency to force all trade onto centralised exchanges could be justified by our finding that prices in centralised markets remain closer to the fully revealing level. Still, dark markets did not fare that badly. Even if they were not aware of it, participants traded more than 75% of the time at prices within narrow bands of the fully revealing price. A more complete evaluation of policy should await our experiments with costly information acquisition. There, centralised markets are known to not provide sufficient incentives, as mentioned before, and it is an open empirical question whether decentralised markets do.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

We emphasise that we are not claiming that dark markets would always aggregate information as effectively as in our experiments. As we discussed at length, the setting of the percolation theory of Duffie, Manso and collaborators is highly specific and our experiment was designed with this specificity in mind. Importantly, traders should know why others are in the market absent private information: they should know each other's private valuations for the asset and thus their 'roles' as natural buyers or sellers of an asset. Likewise, information has to be dispersed: there has to be a level playing field – everyone's private piece of information is as valuable as anyone else's. Then again, our information structure is far looser than in TIP, and our trading protocol is less structured. Yet we observe that our decentralised markets manage to aggregate the dispersed information effectively, as in TIP.

University of Utah
University of Melbourne

Additional Supporting Information may be found in the online version of this article:

Appendix A. Instructions.
Appendix B. Brief Description of TIP Model.

References

Bénézit, F., Blondel, V., Thiran, P., Tsitsiklis, J. and Vetterli, M. (2010*a*). 'Weighted gossip: distributed averaging using non-doubly stochastic matrices', *Proceedings 2010 IEEE International Symposium on Information Theory (ISIT)*, pp. 1753–7.

Bénézit, F., Denantes, P., Dimakis, A.G., Thiran, P. and Vetterli, M. (2008). 'Reaching consensus about gossip: convergence times and costs', University of California San Diego, mimeo, available at http://ita.ucsd.edu/workshop/08/files/paper/paper_253.pdf (last accessed: 8 August 2017).

Bénézit, F., Dimakis, A.G., Thiran, P. and Vetterli, M. (2010*b*). 'Order-optimal consensus through randomized path averaging', *IEEE Transactions on Information Theory*, vol. 56(10), pp. 5150–67.

Bessembinder, H. and Maxwell, W. (2008). 'Markets: transparency and the corporate bond market', *Journal of Economic Perspectives*, vol. 22(2), pp. 217–34.

Bessembinder, H., Maxwell, W. and Venkataraman, K. (2006). 'Market transparency, liquidity externalities, and institutional trading costs in corporate bonds', *Journal of Financial Economics*, vol. 82(2), pp. 251–88.

Bossaerts, P., Frydman, C. and Ledyard, J. (2014). 'The speed of information revelation and eventual price quality in markets with insiders: comparing two theories', *Review of Finance*, vol. 18(1), pp. 1–22.

Bossaerts, P., Plott, C. and Zame, W. (2007). 'Prices and portfolio choices in financial markets: theory, econometrics, experiments', *Econometrica*, vol. 75(4), pp. 993–1038.

Bossaerts, P. and Zame, W.R. (2008). 'Risk aversion in laboratory asset markets', *Research in Experimental Economics*, vol. 12, pp. 341–58.

Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.

Chamberlin, E.H. (1948). 'An experimental imperfect market', *Journal of Political Economy*, vol. 56(2), pp. 95–108.

d'Aspremont, C. and Gérard-Varet, L.A. (1979). 'Incentives and incomplete information', *Journal of Public economics*, vol. 11(1), pp. 25–45.

Duffie, D. (2012). *Dark Markets: Asset Pricing and Information Transmission in Over-The-Counter Markets*, Princeton, NJ: Princeton University Press.

Duffie, D., Giroux, G. and Manso, G. (2010). 'Information percolation', *American Economic Journal: Microeconomics*, vol. 2(1), pp. 100–11.

Duffie, D., Malamud, S. and Manso, G. (2009). 'Information percolation with equilibrium search dynamics', *Econometrica*, vol. 77(5), pp. 1513–74.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.

Duffie, D., Malamud, S. and Manso, G. (2014). 'Information percolation in segmented markets', *Journal of Economic Theory*, vol. 153, pp. 1–32.
Duffie, D. and Manso, G. (2007). 'Information percolation in large markets', *American Economic Review*, vol. 97 (2), pp. 203–9.
Edwards, A.K., Harris, L.E. and Piwowar, M.S. (2007). 'Corporate bond market transaction costs and transparency', *Journal of Finance*, vol. 62(3), pp. 1421–51.
Gilboa, I., Postlewaite, A., Samuelson, L. and Schmeidler, D. (2014). 'Economic models as analogies', Economic Journal, vol. 124(578), pp. F513–33.
Goldstein, M. and Hotchkiss, E. (2007). 'Dealer behavior and the trading of newly issued corporate bonds', AFA 2009 San Francisco Meetings Papers.
Goldstein, M.A., Hotchkiss, E.S. and Sirri, E.R. (2007). 'Transparency and liquidity: a controlled experiment on corporate bonds', *Review of Financial Studies*, vol. 20(2), pp. 235–73.
Green, R.C., Hollifield, B. and Schürhoff, N. (2007). 'Financial intermediation and the costs of trading in an opaque market', *Review of Financial Studies*, vol. 20(2), pp. 275–314.
Grossman, S. and Stiglitz, J. (1980). 'On the impossibility of informationally efficient markets', *American Economic Review*, vol. 70(3), pp. 393–408.
Hirshleifer, J. (1971). 'The private and social value of information and the reward to inventive activity', *American Economic Review*, vol. 61(4), pp. 561–74.
Huber, J., Kirchler, M. and Sutter, M. (2008). 'Is more information always better?: experimental financial markets with cumulative information', *Journal of Economic Behavior and Organization*, vol. 65(1), pp. 86–104.
Hurwicz, L. (1977). 'On the dimensional requirements of informationally decentralized pareto-satisfactory processes', in (K. Arrow and L. Hurwicz, eds.), *Studies in Resource Allocation Processes*, Cambridge: Cambridge University Press.
Li, D. and Schürhoff, N. (2012). 'Dealer networks', CEPR Discussion Paper DP10237.
Plott, C. and Sunder, S. (1988). 'Rational expectations and the aggregation of diverse information in laboratory security markets', *Econometrica*, vol. 56(5), pp. 1085–118.
Plott, C.R. and Smith, V.L. (1978). 'An experimental examination of two exchange institutions', *Review of Economic Studies*, vol. 45(1), pp. 133–53.
Plott, C.R., Wit, J. and Yang, W.C. (2003). 'Parimutuel betting markets as information aggregation devices: experimental results', *Economic Theory*, vol. 22(2), pp. 311–51.
Radner, R. (1979). 'Rational expectations equilibrium: generic existence and the information revealed by prices', *Econometrica*, vol. 47(3), pp. 655–78.
Rao, R.C., Arjunji, R.V. and Murthi, B. (1995). 'Game theory and empirical generalizations concerning competitive promotions', *Marketing Science*, vol. 14(3), pp. G89–G100.
Sandler, T. and Hartley, K. (2001). 'Economics of alliances: the lessons for collective action', *Journal of Economic Literature*, vol. 39(3), pp. 869–96.
Smith, V.L. (1965). 'Experimental auction markets and the walrasian hypothesis', *Journal of Political Economy*, vol. 73(4), pp. 387–93.
Sunder, S. (1992). 'Market for information: experimental evidence', *Econometrica*, vol. 60(3), pp. 667–95.
Wolinsky, A. (1990). 'Information revelation in a market with pairwise meetings', *Econometrica*, vol. 58(1), pp. 1–23.

© 2017 The Authors.
The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society.