

No-Trade in the Laboratory

Marco Angrisani, Antonio Guarino, Steffen Huck and Nathan Larson*

July 2008

Abstract

We test the no-trade theorem in a laboratory financial market where subjects can trade an asset whose value is unknown. Subjects receive clues on the asset value and then set a bid and an ask at which they are willing to buy or to sell from the other participants. In treatments with no gains from trade, theory predicts no trading activity, whereas, in treatments with gains, trade becomes theoretically possible. Our experimental results show that subjects fail to reach the no-trade equilibrium by pure introspection, but they learn to approach it over time, through market feedback and learning. (JEL C92, D8, G12, G14)

1 Introduction

Trade in financial markets is typically attributed to two types of source: liquidity and hedging reasons on the one hand, and informational or speculative reasons on the other. The first concerns traders who have a private reason, such as hedging risk, portfolio rebalancing, or a sudden need for cash, to buy or sell an asset. The second refers to situations in which the value of an asset may be common but uncertain, and trade is generated as individuals attempt to profit from private information about that asset value. Although the informational explanation for trade strikes many market observers as plausible, it is at odds with some celebrated results in the theory of financial economics, known collectively as “no-trade

*Angrisani: Department of Economics, University College London (e-mail: m.angrisani@ucl.ac.uk); Guarino: Department of Economics and ELSE, University College London (e-mail: a.guarino@ucl.ac.uk); Huck: Department of Economics and ELSE, University College London (e-mail: s.huck@ucl.ac.uk); Larson: Department of Economics, University of Virginia (e-mail: nl2a@virginia.edu). We thank Ken Binmore, Douglas Gale, David Levine, Andrew Schotter, Georg Weizsacher and Bill Zame for useful comments and discussions. We also thank Chris Tomlinson, who wrote the experimental program, and Tom Rutter, who helped to run the experiment. Guarino and Huck gratefully acknowledge the financial support of the ESRC (World Economy and Finance program). We are responsible for any errors.

theorems.” These theorems state broad conditions under which rational agents cannot trade with each other on the basis of private information alone.

While there are many variants of the no-trade result—see, for instance the classical contributions of Milgrom and Stokey (1982), Tirole (1982), and Sebebius and Geanakoplos (1983), as well as Aumann’s (1976) famous theorem on the impossibility of agreeing to disagree—they share an underlying logic that is easy to explain. Suppose that two agents consider trading an asset that has the same uncertain value to each of them. If the agents begin with a common prior about that value and then receive additional private information that causes their beliefs to diverge, one might expect to see the agent with more pessimistic news sell to the agent with more optimistic news. However, a rational prospective seller must update her beliefs to account for the information revealed by her trading partner’s willingness to buy, and *vice versa*. If both agents are rational this Bayesian updating has an inexorable conclusion. A trade at some price p would make it common knowledge that the agents disagree about whether the expected asset value is greater or smaller than p , and this in turn would imply that the agents had failed to update as completely as rational agents should. As a result, trade cannot occur in equilibrium.¹

Market practitioners tend to be skeptical about the empirical validity of these theoretical no-trade results. One criticism is that the supporting logic is too subtle for individual investors to grasp by introspection. Another critique is that these theorems require assumptions, such as common priors and common knowledge of rationality, that may not be satisfied in real financial markets (even though they are conventional in the theory literature).² Others argue that the volume of trade in financial market is so high that it would be difficult to explain on the basis of liquidity and portfolio rebalancing motives alone (for some empirical evidence on this, see, e.g., Dow and Gorton, 1997). As Ross (1989) puts it “It is difficult to imagine that the volume of trade in security markets has very much to do with the modest amount of trading required to accomplish the continual and gradual portfolio balancing inherent in our current intertemporal models.” While these critiques of the theory may have

¹We neglect for now the knife-edge case of trade at a price equal to the expected asset value, conditional on both agents’ information.

²While we will not attempt a full discussion of the theoretical literature, we should note that there is a considerable body of more recent work addressing the robustness of the no-trade result. For instance, Morris (1994) provides necessary and sufficient conditions for no-trade results to persist even in the presence of non-common priors. Sonsino (1995) and Neeman (1996) study the case in which common knowledge is replaced by quasi common knowledge or common p -belief. Blume et al., (2006) characterize belief restrictions that ensure a no trade result in competitive equilibrium and show that trade occurs generically when these restrictions are not satisfied. Eliaz and Spiegel (2007) apply a mechanism-design approach in order to examine the extent to which non-common priors create a barrier to speculative bets. Serrano-Padial (2007) provides conditions that rule out all trade, including trades at the expected asset value conditional on pooled information.

merit, they are difficult to investigate with field data since they hinge on details about the preferences and beliefs of investors that we cannot hope to observe.

In this paper, we intend to provide a direct test of these no-trade results using experimental data. By creating an artificial financial market and observing subjects trading behavior in the laboratory, we can control for many of the unobserved influences that plague field data. We can also control the presence or absence of Pareto gains from trade. In markets where mutual gains from trade are possible, the no-trade results do not apply. Our main contribution is to run side-by-side markets with and without gains from trade. By comparing trading activity under the two conditions, we can assess how effectively, and through which channels, the no-trade logic is incorporated into the behavior of our subjects.

In our laboratory setting, subjects first spend some time learning about a noisy relationship between an asset value and two clues. Then they enter a trading phase in which each subject receives a signal (typically one of the two clues) and has the opportunity to set a bid price at which she is willing to buy one unit of the asset and an ask price at which she is willing to sell one unit. A bilateral trade occurs if her bid is higher than her trading partner's ask, or *vice versa*, with the price set equal to the midpoint of the buyer's bid and the seller's ask. In markets with no gains from trade, the buyer's payoff is equal to the realized asset value minus the price, while the seller's payoff is just the negative of this, so the game is clearly zero-sum. Payoffs are similar in the markets with gains from trade, but on top of this, each party to a successful trade earns an additional fixed sum as a "commission." This trading game is repeated for 30 rounds, allowing subjects ample time to learn from experience.

The comparison between markets with and without gains from trade provides us with one theoretical prediction: levels of trade should be lower in the latter than in the former. If subjects reason purely by introspection, this difference should be seen immediately, while if they learn from experience, we would expect levels of trade in these two types of markets to diverge over time. A strict interpretation of the no-trade theorems implies a second prediction: levels of trade in the no-gains markets should be zero. However, this is arguably a problematic benchmark for a laboratory experiment. Although we were careful to keep our instructions neutral, the very act of placing subjects in a market and asking them to set prices probably created some presumption that trade should happen. Also, while subjects did have an incentive to stay for the whole session (they were paid part of their show-up fee on a round-by-round basis), it would not be surprising if some of them found not trading boring and tried to trade for entertainment value alone.³ To take this into

³These concerns are, of course, well known in experimental finance. See, for instance, the considerations of Lei et al. (2001) in a study concerning bubbles in asset markets.

account, we run additional control treatments in which either one or both trading partners' information is made public. In the latter case, subjects can have no illusion about the scope for informational trade, while in the former case, the disadvantaged subject should find his prospects equally dim. In assessing whether informational trade has been wiped out in our main treatments with no-gain, we will focus on the residual trade in these control treatments, not zero, as a benchmark.

Our results offer broad, but not unqualified, support for the theoretical no-trade predictions. Introspection alone does not lead subjects to equilibrium outcomes, as levels of trade are positive and similar with and without gains in the first few rounds of trade. However, over time trade declines substantially in the no-gains markets while holding steady in the markets with gains from trade. This decline is most pronounced when private information explains relatively more of the asset value and white noise explains relatively less. Analysis of these trends reveals an important role for market feedback, particularly losses. Both short run losses and cumulative trading losses induce subjects to price more cautiously (setting larger bid-ask spreads). Declining trade in the no-gains treatments (relative to the treatments with gains) can be explained partly by stronger reactions to these losses, and partly because losses are more frequent when there are no gains from trade.

One corollary of the no-trade results is that in equilibrium it should be impossible to earn positive rents from trading on private information. We test this implication for the no-gains treatments with a counterfactual exercise: we ask whether a hypothetical trader with private information could have made positive expected profits from trade with our subjects. In our no-gains treatments, these profit opportunities are small and decline over time. This decline is driven mainly by the scarcity of willing trading partners, not by market efficiency: conditional on trade, prices impound only modestly more private information over time.

While the no-trade theorem is a crucial result in financial economics, there appears to be relatively little work that tests it. In the experimental literature, to the best of our knowledge, there is only one other direct test, a recent and independent work by Carrillo and Palfrey (2007).⁴ Their treatments study how changes in the trading mechanism and in the deterministic function mapping signals to the asset value affect the frequency of trade. In contrast, we include noise in the asset value function and use variation across treatments to investigate how the ratio of white noise to private information affects trade outcomes and to establish a benchmark level of trade when mutual gains are possible. Thus, the two papers are essentially complementary. Carrillo and Palfrey do report one session, with a bilateral auction and a "sum of signals" form for the asset value, that is similar to our

⁴This paper is roughly contemporaneous with ours. (We started our experiments in November 2003, and finished the study in 2008.)

baseline treatment. They observe a positive initial level of trade that falls by about 50% in the second half of the session. This downward trend is broadly consistent with the decline in trade that we report. However, there is no clear pattern of declining trade in their other treatments, nor can changes over time be linked to learning by subjects. In contrast, we show that learning from market feedback explains declining levels of trade in all of our zero-sum treatments. Finally, the level of trade in their experiment seems to be generally higher than in ours, although a formal comparison is impossible on the basis of a single session.

Our work is also complementary to the substantial literature on experimental asset markets (e.g., Forsythe et al., 1982; Plott and Sunder, 1982, 1988; Copeland and Friedman, 1991; for surveys, see Duxbury, 1995; and Sunder, 1995). For the most part, this literature addresses questions about market efficiency and information aggregation, and the predictive power of rational expectations equilibrium, in settings with gains from trade. While we will comment briefly on the first two issues in discussing our results, our focus is on settings where prices are not observed in equilibrium because no-trade occurs.⁵

Also related to our study is experimental work on the “betting game” by Sonsino et al. (2002) and Søvik (2004). The betting game is a simple zero-sum game in which asymmetric information is generated by giving agents different information partitions. Several steps of reasoning (iterated deletion of dominated strategies) should lead agents not to bet (Sebenius and Geanakoplos, 1983).⁶ Sonsino et al. (2002), however, find that subjects frequently fail in this type of reasoning: betting occurs frequently and slows down over time only slowly. Søvik (2004) has replicated this study, but with changes including higher stakes and use of the strategy method. She finds that dominance violations and betting occur much less frequently than in Sonsino et al. (2002) and that subjects’ behavior is consistent with two to four steps of iterated reasoning. These betting game studies are complementary to our work in that they operate in a setting tailored for very sharp tests of narrow hypotheses about iterated dominance and levels of reasoning. In contrast, our interest is in understanding how market feedback and other factors influence subjects’ pricing decisions (and thus levels of

⁵Plott and Sunder (1988) report one set of treatments in which subjects share the same common value for an asset that is traded in a dynamic double auction. They find that prices generally converge to the correct (pooled private information) asset value in later rounds of the experiment. They also mention in passing that the number of trade declines in later rounds, but the main focus of their paper is elsewhere, and they do not present any formal results on this.

⁶The study by Sebenius and Geanakoplos (1983) is extended by Sonsino (1998) to the case where there exists a small probability that players accept the bet when they should reject it. Generically, the result of Sebenius and Geanakoplos (1983) proves to be robust. Recently, the betting game has been revisited by Jehiel and Koessler (2006) in a framework that allows for bounded rationality. Players are assumed to be boundedly rational in the way they forecast their opponent’s state-contingent strategy: they bundle states into analogy classes and play best-responses to their opponent’s average strategy in those analogy classes. In this setting, betting can be an equilibrium outcome.

future trade) in a setting that is closer to a real-world market.

While there do not appear to be any direct tests of the no-trade theorem using field data in the literature (unsurprisingly, given the difficulties discussed above), large empirical trading volumes have spurred many authors to study why investors trade. Grinblatt and Keloharju (2001) use data from the Finnish stock market to identify determinants of buying and selling activity. They find that investors are reluctant to realize losses, that they engage in tax-loss selling activity, and that past returns and price patterns affect trading. Odean (1999) studies investors with discount brokerage accounts and restricts attention to trades for which liquidity, rebalancing, and tax loss motives can plausibly be ruled out. He finds that these trades generated net losses, even before accounting for transaction costs, and offers various conjectures about why these trades were made. These papers are complementary to our own: we exclude most of the determinants of trade that they study in order to focus more sharply on the role of private information, which their data cannot address. Future work linking field and laboratory data would certainly be valuable.

The paper is organized as follows. Section 2 presents the model and its predictions. Section 3 describes the experiment. Section 4 illustrates the main results. Section 5 analyzes subjects' pricing strategies in the experiment. Section 6 discusses the experimental market efficiency. Section 7 concludes. The Appendix contains some proofs, a description of some auxiliary results of the experiment and the instructions.

2 The Trading Game

Consider the following trading game played by two risk neutral agents. There is an asset worth $V = A + B + X$ to each trader, where A and B are random variables drawn from a joint distribution F and X is drawn independently from F_X , with $E(X) = 0$. We assume that F is atomless and has a density f that is strictly positive on its interior, in the following sense: if $A \leq A' \leq B' \leq B$, and $f(A, B) > 0$, then $f(A', B') > 0$. Each of the signals A and B is observed privately by one of the two traders, while neither observes X . We will assume that A and B each have support \mathcal{S} and are interchangeable in F , so the agents can be considered equally well informed about V .⁷ From now on, we will identify an agent with his signal.

In the trading game, each agent will have the opportunity to buy or sell one unit of the asset from the other. After observing their signals, each agent i submits an order $(b_i, a_i) \in \mathcal{P}^2$ consisting of a bid price b_i and an ask price a_i (where \mathcal{P} is a compact subset of the positive

⁷This symmetry will feature in our experimental design. It is not needed for the no-trade result below, but it is helpful in characterizing some equilibria (of the model with gains from trade) we will present later.

real numbers such that $\text{supp}(V) \subseteq \mathcal{P}$). Whenever one agent's bid price is higher than the other agent's ask price, one unit is sold from the latter to the former at a price equal to the midpoint of the bid and ask. More formally,

$$\begin{aligned} &\text{if } a_A \leq b_B, \quad A \text{ sells to } B \text{ at } p = \frac{1}{2}(a_A + b_B); \\ &\text{if } a_B \leq b_A, \quad B \text{ sells to } A \text{ at } p = \frac{1}{2}(a_B + b_A). \end{aligned}$$

If neither of these conditions holds, there is no-trade between A and B . Alternatively, if both conditions hold, both trades are carried out. In this case, there is no net transfer of the asset, but if the two prices are different, there will be a transfer payment between the agents. Notice that this last case is only possible if $a_i \leq b_i$ for at least one agent, that is, if one of the agents offers to sell at a price below his offer to buy.⁸

Agents maximize expected trading profits, where realized profits are $p - V$ for a sale or $V - p$ for a purchase. Fixing a strategy for agent B , i.e., a map from any signal B into a bid $b_B(B)$ and an ask $a_B(B)$, agent A 's expected payoff from placing order (b_A, a_A) with signal A is

$$\begin{aligned} &\Pr(b_B(B) \geq a_A | A) E(p - V | A, b_B(B) \geq a_A) + \\ &\Pr(a_B(B) \leq b_A | A) E(V - p | A, a_B(B) \leq b_A), \end{aligned}$$

with a similar expression for agent B . If the primitives of the game are common knowledge, then a Bayesian Nash equilibrium of the trading game is a strategy profile for both agents in which each agent's strategy maximizes his expected payoff with respect to correct expectations about his opponent's strategy. We will say that an order strategy is regular if b_i and a_i are continuous, differentiable, strictly increasing functions.⁹ If agent B 's strategy is regular, we can write agent A 's expected payoff as

$$\begin{aligned} \pi_A((b_A, a_A); A) &= \int_{B \geq b_B^{-1}(a_A)} f(B | A) \left(\frac{b_B(B) + a_A}{2} - A - B \right) dB + \\ &\int_{B \leq a_B^{-1}(b_A)} f(B | A) \left(A + B - \frac{a_B(B) + b_A}{2} \right) dB. \end{aligned}$$

A similar expression holds for the expected profit of agent B .

Having described our economy, we can now prove the following proposition:

⁸This case is allowed for the sake of completeness, but it will not turn out to be of any practical importance. As a matter of fact, in the experiment we will not allow subjects to post an ask lower than the bid.

⁹Note that our restriction to regular strategies is imposed mainly to streamline the exposition and could be loosened substantially (at some technical cost).

Proposition 1 *There is no BNE in regular strategies for which trade occurs with positive probability.*

Proof. Suppose, to the contrary, that agent A sells to agent B with positive probability. Then, there exists some A^* , with $a_A(A^*) = p^*$, such that $\Pr(b_B(B) \geq p^* | A^*) > 0$. Define B^* to be the signal for agent B for which $b_B(B^*) = p^*$. (We cannot have $b_B(B) > p^*$ for all $B \in \mathcal{S}$ since agent A would then be able to raise his ask price slightly and continue to win with probability one but at a higher price. Continuity then implies that B^* exists.)

Agent A 's first order condition for the optimality of $a_A(A^*) = p^*$ is then

$$-\frac{1}{b'_B(B^*)}f(B^* | A^*)(p^* - A^* - B^*) + \frac{1}{2}(1 - F(B^* | A^*)) = 0.$$

The second term is equal to $\frac{1}{2}\Pr(b_B(B) \geq p^* | A^*)$, which is strictly positive by hypothesis. Given strictly monotonic strategies, this implies that $p^* - A^* - B^* > 0$ and $f(B^* | A^*) > 0$. However, agent B 's first order condition for the optimality of $b_B(B^*) = p^*$ is

$$\frac{1}{a'_A(A^*)}f(A^* | B^*)(A^* + B^* - p^*) - \frac{1}{2}F(A^* | B^*) = 0.$$

Since $f(B^* | A^*) > 0$ implies $f(A^* | B^*) > 0$, this first order condition cannot be satisfied unless $A^* + B^* - p^* \geq 0$, which is a contradiction.

An identical argument rules out a positive probability of sales from agent B to agent A , completing the proof. ■

The intuition of the proposition follows the standard line of no-trade results. Conditional on a trade at price p^* , neither agent can expect a strictly positive gain, because then the other agent would have to expect a loss. Trade in which both agents expect zero profits cannot survive either. If there were signals A^* and B^* and a price $p^* = E(V | A^*, B^*)$ such that agents with these two signals traded at p^* , then the seller must be in the position of trading with all buyers with signals $B \geq B^*$ and earning zero profits when the buyer's signal is B^* . But then he could strictly improve his payoff by raising his ask price: he would lose nothing by dropping the sales to $B = B^*$ buyers and he would improve his sale price versus the $B > B^*$ buyers. The same logic implies that a buyer making zero profit purchases at the margin would be better off if she were to reduce her bid price. Thus, the agents' exercise of market power rules out even zero expected profit trades.¹⁰

¹⁰In rational expectations competitive equilibrium versions of the no-trade theorem, individual agents cannot influence the price as they do here, but risk aversion often plays a similar role in shutting down zero-profit trade. We expect that adding risk aversion should only strengthen the no-trade result for our trading

The restriction to regular strategies simplifies the proof of the proposition considerably but does not appear to be essential. For example, if we relax strict monotonicity to allow bid and ask functions that pool at particular prices, the no-trade result does not change.¹¹ Note that equilibria without trade always exist. One such equilibrium is the strategy profile in which both agents always ask for the highest price in \mathcal{P} and always bid the lowest price in \mathcal{P} , but there are, of course, many others.

2.1 The Trading Game with Gains from Trade

Our baseline treatments will replicate the trading game just described in the laboratory. It would be unreasonable to expect experimental subjects to align perfectly with theory, so in practice we expect to see some positive frequency of trade. If we observe, for example, that trade occurs in 10% of subject pairings, we will need a method to evaluate whether this should be viewed as a failure or a qualified success for theory. Toward that end, we now amend the trading game to include gains from trade. Thus revised, the trading game *does* admit equilibria with positive levels of trade. By running treatments with small gains from trade, we construct a frame of reference for the baseline results: if the baseline treatments generate levels of trade near the “upper bound” of the gains treatments, this is bad news for the theory. Alternatively, if the baseline treatments diverge from the gains treatments in the direction of zero trade, this supports the theory.


The trading game with gains from trade is identical to the baseline model described above with one exception: now, whenever a trade takes place, the seller and buyer each receive a fixed amount c in addition to their trading profits. One could think of c as a type of fixed-fee commission. Consequently, every trading opportunity is associated with total gains from trade equal to $2c$. One can imagine introducing gains from trade in many alternative and perhaps more realistic ways, but this formulation has two appealing features: it is very easy for subjects to grasp and it will permit us to compute equilibria.

This game typically has many equilibria, including the trivial no-trade equilibrium which survives the presence of gains from trade. Since we are interested in upper bounds on the game. We have not attempted a proof of this since risk neutrality is a relatively reasonable assumption for the payoffs at stake in our experiment.

¹¹The logic is roughly the following. Suppose that there is pooling at some price p^* : that is, there are intervals $[\underline{A}, \bar{A}]$ and $[\underline{B}, \bar{B}]$ such that $a_A(A) = b_B(B) = p^*$ for all $A \in [\underline{A}, \bar{A}]$ and $B \in [\underline{B}, \bar{B}]$. Then we must have $E(V | \bar{A}, B \in [\underline{B}, \bar{B}]) \leq p^*$ and $E(V | \underline{B}, A \in [\underline{A}, \bar{A}]) \geq p^*$, otherwise sellers with signal \bar{A} and buyers with signal \underline{B} would not be willing to trade at p^* . But because V is strictly increasing in the two signals, we have $E(V | \underline{B}, A \in [\underline{A}, \bar{A}]) < E(V | \underline{B}, \bar{A}) < E(V | \bar{A}, B \in [\underline{B}, \bar{B}])$, contradicting the pair of inequalities above. Loosely, the most and least enthusiastic sellers in one pool of signals, and the most and least enthusiastic buyers in another pool of signals cannot all be kept indifferent about the prospect of trade with each other.

level of trade that should be observed in the laboratory, we focus on identifying the equilibria that generate the highest levels of trade. Figure 1 presents symmetric equilibrium strategies for some parameter specifications that we will use in two treatments of the experiment, later called treatments GT1 and GT2. In these treatments, the commission c is equal to 5 (which is equivalent to 10% of the average asset value). The parametric distributions of A , B , and X used for these treatments can be found at the beginning of Section 3.3. For each treatment, we show the equilibrium bid and ask functions for which the probability that a pair of agents trades is highest. In Table 1 we present summary statistics for these equilibria, which requires defining some terms.¹² Given bid and ask strategies (b, a) and a signal realization S , we define an agent's *spread* to be $a(S) - b(S)$. His average spread is just his expected spread over all signal realizations. We define the *value transferred* between two agents to be $|V - p|$ if trade occurs and 0 otherwise. The average value transferred is the unconditional expectation of this transfer payment (including events in which trade does not occur).¹³ Conditional on trade, the *absolute price prediction error* is defined to be $|\varepsilon_p| = |p - E(V | A, B)|$. The percentage of trade is just the equilibrium probability that a pair of agents trades.

Figure 1: Symmetric Equilibrium Strategies with Gains from Trade



We will develop these concepts more as we go forward, but for now note that an agent

¹²In the table, for the sake of comparison, we also present the equilibrium statistics for the cases of $c = 2$ and $c = 8$.

¹³Value transferred is typically on the order of a few pence per match. Later, in the results, we will report the total value transferred over five rounds of a treatment. With five sessions per treatment and 16 matches per round, this represents the sum over 400 matches, so the values there will be on the order of several pounds.

with a larger spread will tend to trade more rarely and at more advantageous prices. The value transferred measures how much one party gains and the other loses on average, net of any commissions. Value transferred is a reasonable measure of average profits and losses, but because the asset value contains noise that neither trader observes, value transferred is not a good measure of price efficiency. For this, we have the price prediction error which captures (in the event of trade) how closely the price tracks the best feasible forecast of V , given the agents’ pooled information.

It is also worth noting that the equilibrium outcomes are quite sensitive to the value of the commission. Even a small change (e.g., from 5 to 8) changes the percentage of trade and the other values considerably. In the remainder of this section we provide a relatively informal discussion of how these equilibria were computed. Readers who wish to skip to the description of the experiment should feel free to do so, while those who are interested in the more technical details will find them in the appendix.

Table 1: Equilibria with Gains from Trade: Summary Statistics

	c	% Trade	Value Transferred	$ \varepsilon_p $	Average Spread
GT1	2	2%	0.03	0.42	60.59
	5	14%	0.24	1.07	33.28
	8	43%	1.09	2.04	25.51
GT2	2	8%	0.86	0.42	24.80
	5	62%	6.33	1.98	5.68
	8	97%	11.49	2.02	2.89

Value transferred, absolute price prediction error and average spread are expressed in pence per match.

Let us say that a bid function b is *separating* on a price interval P if there is a signal interval $\mathcal{S}' \subseteq \mathcal{S}$ such that $P = b(\mathcal{S}')$ and b is strictly increasing and differentiable on \mathcal{S}' . Alternatively, say that b *pools* at a price p if the set of signals S for which $b(S) = p$ has positive measure. The same terminology will apply to ask functions. A strategy (b, a) is fully separating if each function is separating on its entire image. Our equilibrium analysis will be directed toward constructing symmetric equilibria in fully separating strategies, if they exist. Where they do not, we look for equilibria that are “almost fully separating.” There are two practical reasons for this approach. First, our main objective from the analysis is a theoretical upper bound on the level of trade that we can take to our experimental treatments. Many equilibria of this trading game typically exist, ranging from the trivial no-trade equilibrium to equilibria with various patterns of pooling and separation. Characterizing all of these

equilibria is unrealistic, but extensive numerical exploration indicates that the highest levels of trade are achieved in equilibria with separating (or almost separating) strategies. Second, equilibria with more extensive pooling require buyers and sellers to coordinate expectations at particular, and sometimes arbitrary, prices. *A priori*, delicate coordination of this sort among experimental subjects seems unlikely, and we see no evidence of clustering at particular prices in our data.

To fix ideas, consider a prospective symmetric, fully separating equilibrium (b, a) with inverse order functions $\beta = b^{-1}$ and $\alpha = a^{-1}$. Consider the bid \tilde{b} of a buyer¹⁴ who observes signal B and faces an opponent who observes A and uses strategy (b, a) . The expected payoff associated with this bid is:

$$\pi(\tilde{b} | B) = \int^{A \leq \alpha(\tilde{b})} f(A | B) \left(A + B - \frac{1}{2}a(A) - \frac{1}{2}\tilde{b} + c \right) dA. \quad (1)$$

In a symmetric equilibrium, this payoff will be maximized at $\tilde{b} = b(B)$. Notice that here, because of the commission, the profit from the marginal purchase (when $A = \alpha(\tilde{b})$ and therefore $p = \tilde{b}$) can be positive even when the expected value of the asset is below the sale price ($A + B - \tilde{b} < 0$). The fact that this marginal trade can be simultaneously profitable for both the buyer and the seller is, of course, what allows trade to survive. The profitability of inframarginal purchases ($A < \alpha(\tilde{b})$) depends on the $(A - \frac{1}{2}a(A))$ portion of the integrand. If ask prices fall relatively quickly with declining A , then buying from sellers with low signals can be profitable. Alternatively, if ask prices drop slowly as A declines, then the standard lemons logic applies: expected profits can be negative even if the marginal purchase is profitable.

These considerations help to illuminate the optimality conditions for equilibrium. If the ask function is strictly increasing and differentiable near $\alpha(\tilde{b})$, then local optimality dictates a first order condition for (1) that can be written as a differential equation for α' . When α' is too large, a slightly higher bid suffices to attract additional sellers with much higher signals, which tends to be profit improving. Conversely, when α' is too small, the buyer can reduce its bid and purchase from essentially the same pool of sellers, but at a better price. The first order condition ensures that neither of these deviations is attractive. Similar conditions apply to best response asks, leading to a first order condition involving β' . Together, these comprise a pair of coupled differential equations in α and β that must be satisfied to meet local optimality.¹⁵

¹⁴Of course, every agent can act as both a buyer and a seller. The shorthand term “buyer” (or “seller”) should be taken to mean “an agent acting in his capacity as a buyer (or as a seller).”

¹⁵Of course, local optimality requires checking a second order condition as well; this condition is satisfied.

A candidate solution to these differential equations must be checked to ensure that the equilibrium strategy is also a global best response to itself. For example, a bid may be locally optimal and nonetheless generate an expected profit in (1) that is negative; in this case, the buyer would be better off submitting a high bid that has no chance of trading. The potential adverse selection problem here—that only agents with bad news about V are willing to sell—is standard. The twist here is that imposing global optimality will tend to place additional limits on the slopes of b and a : as noted above, larger b' and a' tend to mitigate the adverse selection problem while smaller slopes tend to exacerbate it. For a buyer, the idea is that as the pool of sellers (A) deteriorates, the adverse selection problem is less severe if the price (linked to $a(A)$) also declines quickly and more severe if $a(A)$ declines slowly. In practice, these global conditions must be checked on a case by case basis. For some parameter values, they are satisfied, and a fully separating equilibrium exists. In other cases they fail, usually for a subset of the highest and lowest signals. In these cases, it is generally possible to construct a patched equilibrium with separation over some range $[p_l, p_h]$ of prices. Sellers (buyers) who would be instructed by the first order conditions to ask (bid) for more than p_h will instead place non-competitive asks (pool at p_h). The situation at p_l is just the reverse.

3 The Experiment

We ran the study in the Experimental Laboratory of the ELSE Centre at the Department of Economics at UCL between November 2003 and May 2008. We recruited subjects from the ELSE experimental subjects pool, which includes mainly UCL undergraduate students across all disciplines. They had no previous experience with this experiment. Overall, we recruited 280 subjects to run 7 treatments. For each treatment we ran 5 sessions and each session involved 8 subjects.

The sessions started with written instructions given to all subjects (reported in Appendix). We explained to participants that they were all receiving the same instructions. Subjects could ask clarifying questions, which we answered privately.

The experiment consisted of two phases. In a first phase (“learning phase”) subjects could learn how the value of an asset was determined. In a second phase (“trading phase”) they had the possibility of trading the asset.

3.1 The Learning Phase

In this part of the experiment, subjects went through 30 rounds in which they were provided with certain information about the value of the asset and then asked to predict such a value.

Recall that the asset value is $V = A + B + X$. Out of the 8 subjects in each session, four, randomly chosen, were exposed first only to clue A and then to both clues. The other four were exposed first to clue B and then to both clues. Subjects could learn the relation between the clues and the asset value by repeatedly predicting this value on the basis of the clue(s) and receiving feedback afterwards.¹⁶

A more detailed description of this phase is as follows. First, in the instructions subjects were informed that the value of the asset (in the instructions simply labeled as a “good”) was between 0 and 1 pound sterling. They were also told that the asset value depended on various factors and that two of these were “clue A ” and “clue B .” Second, on the computer screen, they were shown a table with a sample of 10 values for the asset and 10 for the corresponding clue (clue A for the first group and clue B for the second).¹⁷ The purpose of the table was to give subjects the chance to begin to form inferences about the relation between the clue and the value. When subjects were ready to proceed, they moved on to a prediction stage. In each round of this stage, the computer generated a new triple (A, B, V) . Each subject was shown his corresponding clue (either A or B) and asked to predict the value of the asset based on the clue they received. After the prediction, the true asset value was revealed. This was repeated for 15 rounds. Then subjects went through a stage of learning with both clues. First, they were shown another sample table on the screen, this time with 10 values of the asset and of both clues. Then, in the 15 last rounds, they observed the value of both clues, made their predictions and received feedback as above.

In order to induce accurate predictions, we used a standard, quadratic, scoring rule: for a prediction with a mistake of x pence a subject obtained $\text{GBP}3.00 - \frac{3}{1000}x^2$. In particular, subjects were paid according to their prediction in two randomly chosen rounds: one selected from rounds 11-15 and another selected from rounds 25-30. Of course, after each prediction, when subjects were informed of the true value of the good, they could also see on the screen the potential payoff for that prediction.

¹⁶In the theoretical model outlined in the previous section, we computed the equilibria under the assumption of common knowledge of the relation between clues and the asset value. Since most facts about the world are not revealed by public announcement, common knowledge is often motivated as the (idealized) culmination of a long period of learning in a stable environment. Our objective here is to take that motivation seriously: subjects had a reasonable, and realistic, chance to approach common knowledge of the model through this learning phase.

¹⁷The table (identical for subjects receiving the same clue) is reported in the Appendix.

3.2 The Trading Phase

In the second phase of the experiment, subjects had the opportunity to trade in a series of 30 rounds. The trading game in each round resembled the theoretical model described in Section 2. In each round subjects received a clue (as in the learning phase) about the asset value and had to submit two numbers: a bid price and an ask price.¹⁸ They were also given the option of selecting a “no-trade” button, which automatically set the bid and ask prices at 0 and 100 pence, respectively.

The subjects, who in the first 15 rounds of the learning phase had seen clue A (B), received clue A (B) also in the trading phase. Each of the four subjects receiving clue A was matched (i.e., had the opportunity to trade) with each of the four subjects receiving clue B , thus generating 16 pairings. Trade occurred between two subjects whenever the bid price of a subject was higher than the ask price of the other subject. In this case, the transaction price was set equal to the average of the buyer’s bid and the seller’s ask. Otherwise there was no exchange of the asset between those two participants. As a subject could trade with up to four other subjects, it could happen that, for instance, a subject bought in one match, sold in another and did not trade in the others.

For each round, we gave subjects an endowment of 20 pence, regardless of whether trade occurred or not. In addition to this, they could earn (or lose) money by trading. For each purchase, the subject earned the difference between the true value of the asset and the price. Similarly, for each sale, he earned the difference between the price and the true value. For the treatments with gains from trade, a subject was given an additional 5 pence for each trade.

After each trading round, subjects received feedback: they could see the true value of the asset, the bid and ask prices set by the four participants of the other group (whom he was matched with), the price of the transactions (if any) and, the payoff for that round.

Finally, before leaving, subjects filled out a short questionnaire, in which they reported some personal characteristics (e.g., gender, age, knowledge of mathematics) and answered questions on their behavior and on their beliefs on other subjects’ behavior in the experiment. Immediately after completing the questionnaire, subjects were paid in private and could leave the laboratory. The payment was equal to the sum of the per-round payoffs of the trading phase, the selected two round payoffs of the learning phase and a fixed fee of £5. The average payment per subject was £16.50 (approximately \$32) for a session lasting about 2.5 hours on average.

¹⁸We imposed the constraint that the bid had to be weakly lower than the ask.

3.3 Experimental Design

We ran seven treatments in total, and our main focus here will be on four of these: two baseline treatments (B1 and B2) with no gains from trade and two gain from trade treatments (GT1 and GT2) with a positive commission. All the treatments differed only in the Trading Phase, whereas the Learning Phase was identical across them.

In all treatments, the asset value was equal to the sum of the two signals plus a noise term: $V = A + B + X$. The components of V were always coordinated so that V would have support on the interval $[0, 100]$. The clues A and B were random variables composed of a common factor z and two independent factors \tilde{A} and \tilde{B} :

$$A = \tilde{A} + z, \quad B = \tilde{B} + z.$$

The main purpose of the common factor z was to camouflage slightly the simple relationship between the clues and the value, presenting subjects with a learning task that is neither too difficult nor trivially obvious. In every treatment, z was distributed uniformly on an interval of length 15. Because z appears twice (once in A and once in B), it explains 30 pence of the potential 100 pence of variation in V .

The remaining 70 pence of variation in V comes from information (\tilde{A}) contained only in clue A , information (\tilde{B}) contained only in clue B , and from noise (X). We study two settings: one in which the ratio of variation due to private information relative to noise is high (B1 and GT1), and one in which that ratio is low (B2 and GT2). In the high information-noise ratio treatments, \tilde{A} and \tilde{B} are distributed uniformly and independently on $[1.5, 33.5]$, while $X \sim U[-3, 3]$. In the low information-noise ratio treatments, \tilde{A} and \tilde{B} are distributed uniformly and independently on $[5, 20]$, and $X \sim U[-20, 20]$. Finally, in treatments B1 and B2, we set $c = 0$ (no gains from trade), while in GT1 and GT2 we set $c = 5$.

The variation in gains from trade (B1 and B2 vs. GT1 and GT2) provides context for evaluating whether our no-trade predictions are close to being satisfied. Theory predicts no-trade in B1 and B2 but allows for positive trade in GT1 and GT2, so if we find (for example) that trade is much rarer in B1 than in GT1, we can take this as support for the theory.

Our interest in the ratio of private information to noise is motivated more informally. The equilibrium prediction of no-trade in B1 and B2 is completely unaffected by the relative importance of \tilde{A} , \tilde{B} , and X . However, in our experiments, we expect subjects to learn about the asset value and the profitability of trade through trial and error. In other words, with due respect to introspection, it seems likely that they learn the consequences of ignoring an opponent's private information from experience. Because in B1 the losses suffered by

ignoring an opponent’s information can be more severe ($|A - B|$ can be larger) and harder to misattribute to bad luck (X is smaller), we might expect subjects to approach the no-trade prediction faster in B1 than in B2. Alternatively, if subjects persistently fail to account for opponents’ private information, then trade should remain higher in B1 than in B2, because disagreements about the asset value (reflected by $|A - B|$) will tend to be larger in B1.

For the two treatments with gains, the variation in \tilde{A} and \tilde{B} *does* play a role in the formal equilibrium predictions, but the white noise X does not. As discussed in Section 2, a higher upper bound level of trade is predicted in GT2 than in GT1, essentially because larger differences in private information in GT1 force agents to trade more cautiously. These predictions, however, only apply if our subjects have learned to model asset values correctly. An alternative conjecture is that more “model uncertainty” will persist in GT2 (because X is larger) and that this will tend to disrupt subjects from coordinating successfully on trade at a price near V .

In addition to these four main treatments we ran two control treatments.¹⁹ We label these control treatments as CE (control-equal information) and CU (control-unequal information). Both these treatments serve as controls for B2 and GT2 since they have in common the same parameter distributions.²⁰ In treatment CE, every subject received both clues A and B . In other words, this was a treatment with symmetric information. In treatment CU, instead, in odd periods, the four subjects who in the first 15 rounds of the learning phase had seen clue A received clue A while the other four subjects received both clues A and B ; in even periods, the inverse was true, i.e., subjects who in the learning phase had observed clue B received clue B and the others received both clues A and B . In other words, in this treatment some subjects were better informed than their counterparts.

These control treatments are intended to measure residual levels of trade that are plausibly unrelated to differences in private information. In CE, there is literally no private information, while in CU, one of the two trading partners always has no private information. Nonetheless, we can imagine other reasons that subjects might trade in these settings. Subjects may trade because they believe themselves to be superior predictors of the asset value, because they are risk-loving or overconfident, or for a variety of other reasons related to non-standard preferences or behavioral biases. Moreover, subjects may trade simply because they find not trading boring. These motives can be very difficult to observe and disentangle,

¹⁹We ran one additional control treatment which was identical to B2, except with lower incentives. Specifically, in each round, a subject traded only with one (randomly matched) subject with the other signal, rather than with all four opponents. The results were generally similar to B2, and we do not report them here.

²⁰We did not run control treatments for B1 and GT1 for reasons that will be clearer later, when we discuss the results.

even in a laboratory setting, and our only objective here is to try to control for them. If one is willing to believe that these sources generate similar amounts of trade in all of our treatments, then we can identify information-based trade off of the differences between our baseline treatments and these controls.

4 Evidence About Strategies and Trading Activity

Because our principal objective is to understand whether and how the no-trade result is confirmed in the laboratory, we will concentrate our analysis on the trading phase of the experiment. Although the learning phase is mainly in a supporting role, we will discuss some of its highlights in the appendix.

4.1 An Overview: Spreads and Trading Activity

4.1.1 B1 and GT1

We begin with an overview of the two treatments in which the ratio of private information to noise was high, B1 and GT1. We focus on three outcome variables: percentage of matches with trade, value transferred, and spreads. The percentage of trade is the outcome that should be zero for the baseline treatments, according to theory.²¹ Note that the percentage of trade does not tell us anything about the *size* of gains and losses from trade. If trade is frequent but always occurs at a price close to V , then the deviation from theory might be regarded as minor; conversely, if trade is rare but always involves large gains and losses, we might hesitate to say that the theory is confirmed. Value transferred, which is equal to $|V - p|$ if trade occurs or 0 if not, lets us assess this. In this section, we report the total value transferred for each treatment over all of the matches in a block of five rounds.²² Both of these outcomes are produced by the interaction of subjects' order strategies. The average spread measures those strategies directly; subjects who set larger spreads will trade less often and at more favorable prices.

Table 2 summarizes these outcomes, grouping the 30 rounds of play into six 5-round blocks. The trend over time is strikingly different in the two treatments. In treatment B1 the average spread starts at approximately 24 and then increases monotonically to reach almost 40 by the end of the experiment. In GT1, instead, it starts at approximately 20

²¹The unit of observation here is a match and not a subject, in order to avoid duplications. Recall that in each round there are 16 matches, since four subjects in group A can trade with four subjects in group B .

²²With five sessions per treatment and 16 matches per session, this means that the number we report is the total value transferred over 400 matches. For comparison, note that the equilibrium value transferred reported in Table 1 refers to one match only.

and remains almost constant for the entire experiment. As a result of these different pricing strategies, the trading activity also looks significantly different.

Table 2: Descriptive Statistics by Treatment: B1 and GT1

	Rounds					
	1-5	6-10	11-15	16-20	21-25	26-30
Average Spread						
<i>Treatment B1</i>	24.03	23.05	28.93	31.33	33.61	39.40
<i>Treatment GT1</i>	19.43	21.15	20.55	20.18	21.10	21.68
% of Matches with Trade						
<i>Treatment B1</i>	30.0	22.5	20.3	17.8	14.2	7.3
<i>Treatment GT1</i>	35.0	20.7	35.7	38.5	31.2	30.7
Value Transferred						
<i>Treatment B1</i>	9.64	7.78	7.63	6.07	4.72	2.07
<i>Treatment GT1</i>	11.78	6.60	11.38	11.77	9.24	7.19

For each treatment, averages are computed over all 5 sessions. Average spread is the average individual spread, expressed in pence. Value transferred is the sum of $|V - p|$, expressed in pound sterling.

In treatment B1 the percentage of trade decreases steadily from 30% to 7.3%.²³ In GT1, instead, there is no clear trend and the frequency of trade is almost always above 30%. The pattern for value transferred is similar: the two treatments begin at similar levels, but value transferred declines monotonically in the treatment with no gains while remaining roughly constant in the treatment with gains. By the end of the experiment the difference between the two treatments is substantial. Tests for the last 15 rounds of trade indicate that these differences are statistically significant at any reasonable significance level: spreads are higher (p -value = 0.000) and the percentage of trade and value transferred are significantly lower (p -value = 0.000 and p -value = 0.001 respectively) in the baseline treatment.²⁴

Mean spreads alone cannot explain all of the nuances of trading outcomes; other factors

²³As mentioned in the introduction, Palfrey and Carrillo (2007) report one session with a structure similar to B1 with the following key differences: (i) their asset value is the sum of signals with no noise, and subjects are told this; (ii) their subjects have fixed roles - sellers cannot buy and buyers cannot sell; (iii) buyers are price-takers; (iv) there are 20 rounds. They report that the level of trade drops from 31.7% in the first 10 rounds to 21.7 in the second ten rounds. If sales to or from either partner had been possible, as in B1, presumably both numbers would be larger. In other treatments that are less similar to ours, they find smaller declines or increases in trade over time. While we reiterate that a formal comparison with this single session is unjustified, it is interesting to note that in our treatment B1, the frequency of trade is 26.1% in the first 10 rounds, 19.0% in the second 10 rounds, and 10.8% in the final ten rounds. Had our experiment ended after 20 rounds, the decline in trade would have looked rather modest.

²⁴Observations are not independent, since the same subjects set the spreads many times, interact within each session and feature in different trading matches per round. We take into account cross-sectional as well as time-series dependence and compute standard errors by bootstrap (1000 replications), clustering at the session level. This remark applies to all our results involving standard errors. If not otherwise specified, we compare treatments pairwise using a one-tailed test.

including subjects' accuracy in predicting the asset value and heterogeneity in spreads also matter. We will return for a more detailed look at the determinants of trade later, but for now, the basic picture seems clear. Behavior starts at similar levels with and without gains, but in the no gains treatment spreads quickly increase, driving trade down to a very low level by the end of the experiment.

4.1.2 B2 and GT2

Now we discuss the results in treatments B2 and GT2. Recall that these treatments differ from the previous ones in that there is less private information and more noise in the composition of the asset value. Table 3 shows the evolution of spreads, percentage of trade and value transferred over time. For B2, the mean spread is steadily increasing over time, starting at 20 and ending at 31. As a result, the percentage of trade falls almost monotonically from 28.6 to 19.6 and so does value transferred, from an initial value of £11.02 to a final one of £6.13. These outcomes start at initial levels that are similar to B1 and trend in the same direction, but the trend is weaker: final spreads are lower, and trade is higher, than in B1.²⁵

Treatment GT2 is qualitatively similar to GT1: our three outcome variables remain fairly constant over time in both treatments. The levels of the average spread and value transferred are not significantly different between GT1 and GT2, but the percentage of trade is significantly lower in GT2 (at the 5% level).²⁶

Significance tests over the last 15 rounds of play confirm that the divergence between the gains and no gains treatments is not as strong for B2 and GT2 as it was for B1 and GT1. The average spread is roughly 29 for B2 versus 19 for GT2, a difference which is significant at the 10% level (p -value = 0.059). The percentage of trade is 19 in B2 and 22.4 in GT2, while value transferred is £7.91 in the former treatment and £11.51 in the latter; neither of these differences is significant at the 10% level (although value transferred is close, with a p -value of 0.119). A natural conjecture is that learning is slower in B2 because it is harder for subjects to distinguish losses due to bad luck from losses due to an incautious strategy when asset values are noisier. We will have more to say about this in Section 5. The fact that realized trade does not diverge much between B2 and GT2, despite the increasing spreads in B2, appears puzzling at first glance. One must bear in mind that these statistics do not control for the degree to which the two private clues “disagree.” If the size of $|A - B|$ affects the likelihood of trade (and we will see in Section 4.2 that it does), then stochastic variation in this clue disagreement across treatments can have a confounding effect when we compare

²⁵These differences between B1 and B2 are not statistically significant over the final 15 rounds, except for value transferred, which is significantly higher in B2 (at the 5% level). For the comparison of B1 and B2 we use a two-tailed test.

²⁶For the comparison of GT1 and GT2 we use a two-tailed test.

their levels of trade. When we control for $|A - B|$ in Section 4.2, the divergence between B2 and GT2 looks substantially sharper.

Table 3: Descriptive Statistics by Treatment: B2 and GT2

	Rounds					
	1-5	6-10	11-15	16-20	21-25	26-30
Average Spread						
<i>Treatment B2</i>	20.04	21.91	25.98	26.74	28.48	31.52
<i>Treatment GT2</i>	20.67	18.74	16.67	16.27	20.30	22.76
% of Matches with Trade						
<i>Treatment B2</i>	28.5	23.7	16.0	24.0	13.5	19.6
<i>Treatment GT2</i>	25.7	21.5	20.7	24.5	23.2	19.5
Value Transferred						
<i>Treatment B2</i>	11.02	12.06	7.21	11.87	5.74	6.13
<i>Treatment GT2</i>	14.74	7.58	8.16	14.28	10.63	9.62

For each treatment, averages are computed over all 5 sessions. Average spread is the average individual spread, expressed in pence. Value transferred is the sum of $|V - p|$, expressed in pound sterling.

4.1.3 Controls for B2 and GT2

The preceding results tell us that trade in both baseline treatments diverges from the gains from trade treatments (albeit less substantially for B2 with respect to GT2) and trends toward zero. Nevertheless, there appears to be a residual level of trade that is not erased with experience (at least not within 30 rounds). This is particularly true for treatment B2. We would like to know whether this residual trade survives because subjects insufficiently account for asymmetric information, or whether some other cause is to blame.

To address this, it is worth looking at the results for our two control treatments, CE and CU (Table 4). Recall that these two treatments had the same parameterization of the signals and asset value as B2 and GT2, but they differed in which information was provided to subjects. In CE there was no private information at all: every subject received both clue A and clue B in every round. In CU, the asymmetry in private information was stark: in each match, one subject received both clues and the other received only one.

In treatment CE, trade cannot be generated by differences in private information, since all subjects have the same information. Nonetheless, trade hovered just above 10% throughout the treatment and sank only to 10.7% over the last five rounds. Similarly, the value transferred remained fairly constant over time and was equal to £6.64 in the last five rounds. Statistical tests show that we cannot reject the hypotheses that the percentage of trade and the value transferred in the last 15 rounds are equal in treatments B2 and CE (p -value

= 0.106 and p -value = 0.248 respectively).

In treatment CU, only one subject had private information, so the imbalance should have been hard for the disadvantaged subject to ignore. Nonetheless, trade remains at 11.5% over the last five rounds, with value transferred of £4.79. Over the last 15 rounds, value transferred is not significantly different between CU and B2 (p -value = 0.278). The percentage of trade appears to be lower in CU than in B2, but the difference is only significant at the 10% level (p -value = 0.098).

Together the evidence from these two control treatments suggests that much of the residual trade in B2, perhaps most of it, should not be attributed to subjects underestimating their opponents' private information. We can suggest at least three explanations for the persistence of a roughly 10% level of trade, even when subjects do not have private information. One is that subjects disagree about how the asset value is related to the clues. Another is that subjects' strategies vary due to experimentation and errors, and this can generate some trade. Third, and importantly, attempting to trade was the only activity in which subjects could engage. It is possible that some subjects tried to trade out of boredom or a sense that this was what they were supposed to do.

Table 4: Descriptive Statistics by Treatment: CE and CU

	Rounds					
	1-5	6-10	11-15	16-20	21-25	26-30
Average Spread						
<i>Treatment CE</i>	21.02	21.49	21.62	30.92	32.41	34.85
<i>Treatment CU</i>	20.06	25.81	22.82	28.81	29.43	34.85
% of Matches with Trade						
<i>Treatment CE</i>	12.5	19.5	14.5	13.0	12.7	10.7
<i>Treatment CU</i>	23.5	17.0	16.7	14.2	13.0	11.5
Value Transferred						
<i>Treatment CE</i>	5.00	7.04	8.44	6.18	5.68	6.64
<i>Treatment CU</i>	13.51	7.98	7.35	6.81	8.24	4.79

For each treatment, averages are computed over all 5 sessions. Average spread is the average individual spread, expressed in pence. Value transferred is the sum of $|V - p|$, expressed in pound sterling.

4.2 Trading Activity and Differential Information

In order to illuminate the parametric and behavioral roots of differences in trade between B1/GT1 and B2/GT2, it is useful to write down a model of how strategies and signals interact to generate trade. Below, we develop a very simple, ad hoc model in which strategies are

linear in signals. While this model is only illustrative and has no theoretical foundation, one can see from Figure 1 that linear strategies are roughly consistent with equilibrium in GT1 and GT2.

Suppose that subject i uses the following bid and ask strategies with a signal S :

$$b(S) = mS - \bar{b} + \varepsilon_{ib} \quad , \quad a(S) = mS + \bar{a} + \varepsilon_{ia}.$$

That is, all agents use the same linear strategies in S , with constants \bar{a} and \bar{b} , and heterogeneity in the form of idiosyncratic zero-mean shocks ε_{ib} and ε_{ia} . Let us endow subjects i and j with signals A and B respectively, and suppose, without loss of generality, that $A < B$. Then i will sell to j with probability

$$\Pr(\varepsilon_{ia} - \varepsilon_{jb} < m|B - A| - \overline{Spread}),$$

where $\overline{Spread} = \bar{a} + \bar{b}$. The event that the agent with higher signal sells to the agent with a lower signal involves a similar expression, but because $A - B$ replaces $|B - A|$ on the righthand side of the inequality, it is less likely. For our illustrative purposes, we will assume that the expression above approximates the total probability of trade.

This formulation helps to unpack some of the determinants of trade. All else equal, trade is less likely when the average spread is higher, and more likely when the subjects' private signals disagree more, or when they put more weight on that information (higher m). Greater heterogeneity across subjects (higher variance for the shocks) has a mixed effect. If the righthand side tends to be negative (making trade unlikely), then higher variance shocks should tend to raise the chance of trade. If the righthand side tends to be positive, then the opposite is true. Perhaps more intuitively, when spreads are large, it takes an unusually high bid and an unusually low ask to get trade—more heterogeneity makes this event more likely. Conversely, when spreads are small, trade will tend to occur unless disrupted by an unusually low bid and high ask, so the effect of heterogeneity goes the other way.

4.2.1 B1/GT1 versus B2/GT2

This framework can inform a comparison between B1/GT1 and B2/GT2. First, let us fix behavior and consider only the parametric differences across the treatments. In treatments B1 and GT1, differences in private information ($|B - A|$) tend to be larger than they are in treatments B2 and GT2; all else equal, this tends to generate relatively more trade in B1 and GT1. Now, consider what happens as subjects adjust their strategies over time. When there is an increase in the average spread (as in B1 and B2), the probability of trade falls, but

the decline will be more dramatic in cases where the difference in signals is more frequently large. Where $|B - A|$ is more often small, the probability of trade is already fairly small to start with. This suggests that the failure of B2 and GT2 to diverge sharply may be in part related to the fact that differences in subjects' signals tend to be small.²⁷ If we condition on realizations in which A and B differ a lot, the divergence between B2 and GT2 may look more pronounced.

This hypothesis is investigated in Table 5. We revisit our tests for cross-treatment differences in means over the final 15 rounds, conditioning only on those rounds with the following signal realizations: “low” $|A - B|$ ($|A - B| \leq 5$), “high” $|A - B|$ ($|A - B| \in [10, 15]$), and “very high” $|A - B|$ ($|A - B| \geq 20$). This grouping is motivated by the fact that 15 is the largest possible difference in treatments B2 and GT2; the last grouping applies only to treatments B1 and GT1, for which $|A - B|$ could be as large as 32.

The differences between B1 and GT1 continue to be significant for all three $|A - B|$ groupings. More interestingly, both percentage of trade and value transferred are significantly lower in B2 than in GT2 (at the 5% level) if we condition on matches where the signal difference is high (but not when the signal difference is low). This is consistent with our hypothesis and suggests that, when there are no gains from trade, experience and market feedback are fairly effective in pruning down the number of trades instigated by large differences in private information about the asset value (in both treatments B1 and B2).

4.2.2 Actual versus Theoretical Outcomes in GT1 and GT2

This framework can also shed light on the differences between our experimental outcomes with gains from trade and the equilibrium outcomes reported earlier. The percentage of trade in GT2 is below the theoretical upper bound reported earlier, while trade in GT1 substantially exceeds its theoretical upper bound. Note that the observed spreads in the two treatments are similar, but we know that large differences in signals are much more common in GT1 than in GT2. Considering the linear model discussed above, it is a simple matter of mechanics that subjects will manage to coordinate on trade more consistently in GT1 than in GT2. Furthermore, subject heterogeneity should, if anything, reinforce this difference. We would expect the noisier environment of GT2 to produce more dispersed estimates of the asset value, which may lead to more dispersed price strategies. As discussed above, greater dispersion in strategies should tend to reduce the chance of trade when spreads are smaller (as they are in GT1 and GT2).

²⁷Of course, a concomitant cause may be the different way subjects adjusted spreads over time as a function of the feedback they received. We will turn to this issue in the next section.

Table 5: Testing Differences across Treatments Conditioning on Differential Information.

	$ A - B $ low	$ A - B $ high	$ A - B $ very high
Spread			
$H_0 : B1 = GT1$ $H_1 : B1 > GT1$	3.66 (0.000)	2.84 (0.002)	3.03 (0.001)
$H_0 : B2 = GT2$ $H_1 : B2 > GT2$	1.61 (0.054)	2.16 (0.015)	
% of Matches with Trade			
$H_0 : B1 = GT1$ $H_1 : B1 < GT1$	-2.61 (0.005)	-4.56 (0.000)	-3.50 (0.000)
$H_0 : B2 = GT2$ $H_1 : B2 < GT2$	-0.48 (0.316)	-1.73 (0.042)	
Value Transferred			
$H_0 : B1 = GT1$ $H_1 : B1 < GT1$	-1.98 (0.024)	-3.78 (0.000)	-1.80 (0.034)
$H_0 : B2 = GT2$ $H_1 : B2 < GT2$	-1.16 (0.123)	-2.29 (0.011)	

Tests refer to cross-treatment differences in means over the final 15 rounds. We report cluster-robust bootstrap t -statistic, and, for the reader's convenience, the corresponding p -value in parentheses. $|A - B|$ low indicates that $|A - B| \in [0, 5)$; $|A - B|$ high indicates that $|A - B| \in [10, 15)$;

$|A - B|$ very high indicates that $|A - B| \geq 20$. For Treatments B2 and GT2, there are no observations for $|A - B|$ very high.

This argument pushes the question back to the level of strategies. In particular, why do spreads in GT1 remain substantially below their equilibrium levels? Some intuition can be gleaned from the equilibrium strategies in Figure 1. Low signal buyers (below about 12) and high signal sellers (above about 38) find they cannot make positive expected profits and drop out of the market. This tends to unravel the market from the ends: if there are no low signal buyers, then a low signal seller faces more severe adverse selection and must set higher prices, but this in turn means that the next tier of buyers get lower seller signals for the same price, forcing them to bid more cautiously, and so on. This process is obviously sensitive to the behavior of subjects with extreme signals.

In practice, subjects with low signals (less than 12) lose 1.58 pence on average when they buy. Their average bid prices with low signals do appear to react to these losses but only slightly, falling from 20.55 in rounds 1-10 to 18.41 in rounds 21-30. Similarly, subjects with high signals (greater than 38) lose 13.75 pence on average when they sell. Again, their ask prices rise only slightly over time, from 72.38 in rounds 1-10 to 76.62 in rounds 21-30. However, buying with a low signal or selling with a high signal are relatively rare events - each occurs less than once per subject per 30 rounds. The profits of subjects with extreme signals are dominated by trades going the other direction (selling with a low signal or buying with a high signal), and these overall average profits are positive. Given this, it may not be too surprising that subjects do not pay sufficient attention to avoiding buying with a

low signal or selling with a high signal. But as long as these traders are in the market, the unraveling described above stalls, leaving the overall level of trade high.

5 Market Feedback and Trading Activity

Let us now examine the extent to which those changes in strategy can be explained as a response to feedback from the market. The strategic variable that we focus on is the spread. Using a regression framework, we study how round-by-round market outcomes affect the spreads that individuals set over time. For this purpose, we assume the following autoregressive data generating process for the individual spread:

$$Spread_{it} = \alpha + \beta Spread_{it-1} + x'_{it-1}\gamma + e_i + u_{it}, \quad (2)$$

where i and t denote the player and the round, respectively.

The column vector x_{it-1} comprises various information that player i observes at round t before choosing $Spread_{it}$. It includes both short run and long run trading outcomes. Short run outcomes are summarized by two binary variables. One of these takes value 1 whenever player i 's total trading profits in round $t-1$ were positive ($Gain_{it-1}$); the other takes value 1 whenever player i 's total trading profits in round $t-1$ were negative ($Loss_{it-1}$).²⁸ Long run market outcomes are summarized by $TProfits_{it-1}$, the total trading profits that player i has earned in all rounds up to and including $t-1$.²⁹ We also want to consider the possibility that subjects learn from or react to each other, so we include $OSpread_{it-1}$, the average spread set by player i 's opponents in round $t-1$.

The error structure in equation (2) features an individual unobservable component e_i , invariant over trading rounds, and an idiosyncratic disturbance term u_{it} . The individual fixed effect e_i allows for heterogeneity in the conditional mean of $Spread_{it}$ across individuals. We assume that, conditional on the unobserved individual specific effect, u_{it} is uncorrelated with past realizations of spread and past market outcomes.³⁰ Namely,

$$E[u_{it}|Spread_{it-1}, x_{it-1}, e_i] = 0. \quad (3)$$

²⁸Of course, since we consider a dummy for gains and one for losses, a dummy for the case in which i did not trade in round $t-1$ is omitted.

²⁹Trading profits include the commission c , in treatments with gains from trade, but do not include the $20p$ per round fixed payment that subjects received.

³⁰As individual current spreads necessarily affect current market outcomes and subjects' market outcomes, there exists an immediate feedback from $Spread_{it}$ to future values of x_{it-1} . This rules out strict exogeneity of the regressors even in a static model.

The Fixed Effects (FE) estimator is biased for a dynamic model such as equation (2). As shown by Nickell (1981), however, the bias decreases with the time dimension of the panel and vanishes asymptotically as time goes to infinity. In our experimental data set, each player is observed for 30 periods, corresponding to the number of rounds in each session of the experiment. As this large number of usable observations per subject is likely to make the FE bias rather small, we estimate equation (2) by FE despite its dynamic structure.³¹

The design of the experiment poses a further econometric issue. Subjects within a session interact with each other over time, and this can introduce correlation across individual observations within a session, although different sessions remain independent. We assume that spatial correlation within sessions is caused by the presence of interaction factors that are unobservable and uncorrelated with the explanatory variables on the right hand side of equation (2). Although the magnitude of such correlations across players due to unobservables is likely to be of a modest order, we take into account the potential presence of cross-dependence in u_{it} and compute standard errors by bootstrap, clustering at the session level. Such a procedure is also robust to heteroskedasticity and serial correlation in u_{it} and has been shown to perform reasonably well when the number of clusters is relatively small (Cameron et al., 2007).

In order to aid comparisons across treatments, we estimate equation (2) pooling all the data together, but allowing each explanatory variable to have a treatment-specific effect. Formally, we run the regression

$$Spread_{it} = \alpha + \sum_j \beta_j D_j Spread_{it-1} + \sum_j (D_j x_{it-1})' \gamma_j + e_i + u_{it}, \quad (4)$$

where D_j is a binary variable taking value 1 when Treatment = j , with $j = B1, GT1, B2, GT2, CE, CU$. The estimation results for equation (4) are presented in Table 6. In Table 7, we report additional results in which the only change is the inclusion of a treatment-specific linear time trend (variable *Round*).³² The time trend provides an indication of whether there is systematic learning about spreads that cannot be explained directly by market outcomes (such as introspection, perhaps).

³¹We also estimate the parameters of equation (2) using a first-differenced GMM approach *à la* Arellano-Bond (1991). Overall, the estimated coefficients for the market outcomes variables are similar to the ones reported below. Nevertheless, estimation accuracy is lower. Since it is very difficult to predict changes in spread and trading outcomes using lagged spreads and outcomes, the use of weak instruments can plausibly cause the GMM estimator to perform poorly in this framework (Bound et al., 1995). For this reason, we prefer the FE estimation results.

³²To make sure that the market feedback variables were not picking up a more complicated, nonlinear time trend, we also tested a specification with dummy variables for each round. The results were not significantly different from those reported below.

Table 6: FE, Dependent Variable $Spread_{it}$ (time trend excluded)

	Treatments					
	B1	GT1	B2	GT2	CE	CU
$Spread_{it-1}$	0.381*** (0.108)	0.243* (0.127)	0.121*** (0.033)	0.282*** (0.099)	0.364*** (0.103)	0.201* (0.104)
$Gain_{it-1}$	3.424*** (1.441)	0.868 (1.651)	-0.457 (2.027)	2.361** (1.024)	1.178 (1.251)	2.379 (1.705)
$Loss_{it-1}$	9.577** (4.149)	3.372** (1.630)	6.451*** (1.306)	6.297*** (2.230)	7.390*** (2.573)	6.360*** (1.143)
$TProfits_{it-1}$	-0.097*** (0.038)	-0.008 (0.011)	-0.060* (0.034)	-0.011 (0.008)	-0.012 (0.040)	-0.063*** (0.019)
$OSpread_{it-1}$	0.213** (0.093)	0.045 (0.040)	0.163* (0.093)	0.102** (0.048)	0.107*** (0.044)	0.226*** (0.054)

Cluster-robust bootstrap standard errors based on 1000 replications in parentheses. *** indicates significance at 1%, ** indicates significance at 5%, * indicates significance at 10%. Sample size: 6936 observations.

The autoregressive coefficient, β , is positive in all treatments and considerably below 1. This indicates that spreads evolve over time with a moderate degree of persistence.

In a nutshell, when it comes to short run market feedback, losses matter, but gains do not. The differential effect of losses with respect to the no-trade event is consistently positive across treatments, indicating that subjects set higher spreads in a round after a trading loss than they do in a round after not trading. The logical interpretation is that such subjects hope to avoid further losses by demanding more favorable prices. The magnitude of this effect is larger in the treatments with no gains from trade than it is when there are gains (9.58 and 6.46 vs. 3.37 and 6.30). Such a difference is negligible for treatments B2 and GT2, but quite substantial for B1 and GT1. This indicates that the decline in trade in B1 relative to GT1 occurs through at least two channels. First, subjects in B1 make losses more often (because they are not rescued by the commission), and this induces them to set relatively higher spreads. In addition, subjects' reactions to losses are stronger when there are no gains from trade. The second channel suggests that introspection and market feedback interact: perhaps the knowledge that mutual gains from trade are possible encourages subjects to remain aggressive even when they make losses. One could conjecture that in a noisier environment (B2 and GT2), subjects are more focused on model uncertainty than on the presence or absence of a commission.

We would expect a subject with a gain in one round to try to continue to trade but improve his profit with a slightly wider spread. The positive coefficients that we see for $Gain$ are consistent with this story, but they are smaller than the $Loss$ coefficients and generally are not significant. Thus it appears that losses are the principal short run channel

through which market outcomes affect future strategy.

We presume that subjects take long run, as well as short run performance into account when revising their strategies. Here, long run performance is measured by $TProfits_{it-1}$. In the baseline treatments, average accumulated profits (across subjects) are identically zero in every round, but in GT1 and GT2, average accumulated profits tend to grow over time (as trading commissions add up). This means that omitting a significant time trend will tend to bias the coefficient on accumulated profits in these two cases, so here we focus on Table 7. Following the prominent role of losses in the short run feedback, our expectation here is that lower cumulative profits (or larger cumulative losses) should induce subjects to set higher current spreads. Beyond the reasons discussed before, subjects who have become relatively wealthier may show more risk tolerance than subjects who have become poorer. This implies a negative coefficient on $TProfits_{it-1}$, which is what we find for our four primary treatments. Just as for short run feedback, we find that cumulative losses push up current spreads the most in treatment B1 and the least in GT1. In the noisier treatments, cumulative losses again have a weaker effect with gains (GT2) than without (B2), but the difference is much smaller. As earlier, the difference between the gains and no-gains treatments could mean that subjects perceive the two environments differently, but in this case there is another possible explanation. $TProfits$ tends to be centered around zero in B1 and B2 (its median values are 2 and 7) but is typically larger in GT1 and GT2 (medians of 67 and 83). If $TProfits$ has a nonlinear effect (e.g., if raising a subject's profits from -50 to 50 affects behavior more than raising them from 50 to 150 would), then our linear specification would tend to fit this by making its effect larger for B1 and B2. To put the coefficients in perspective, at the end of treatment B1, the subject at the 90th percentile of accumulated profits had earned approximately 300 pence more than the subject at the 10th percentile. Had there been a thirty-first round, the coefficient predicts that the less successful subject would have set a spread about 28 points higher than the more successful one.

Subjects also observe their opponents' strategies at the end of each round, so it is possible that a subject's spread reacts to his opponents' past spreads. There are several reasons to expect such a linkage. One, of course, is that opponents' past spreads may predict current spreads, and these are relevant to a subject's optimal strategy. Generally, a larger opponent spread means that at any given price, a more adverse inference must be drawn about V . A best response will typically involve setting a larger spread oneself in order to avoid losses. A second possible connection involves learning: since the trading game is symmetric, imitating opponents may be one technique that subjects use to try to improve their payoffs. Both of these arguments tend to support a positive relationship between $OSpread_{it-1}$ and $Spread_{it}$, and this is generally what we see in Table 6. The average effect (across all six treatments)

is that a 10 pence increase in the average spread set by subject i 's opponents in round $t - 1$ induces a 1.5 pence increase in subject i 's spread in round t . The exception is treatment GT1, where there is no significant effect. Table 7 suggests that these coefficients on $OSpread$ must be interpreted with caution, as most of them lose significance when a treatment-specific time trend is introduced. Much of the link between $OSpread_{it-1}$ and $Spread_{it}$ could simply reflect some other unobserved and omitted variable that causes this time trend.

The treatment-specific time trends in Table 7 (variable *Round*) are all positive and significant, with the exception of GT2. For the four primary treatments, we once again see the pattern that spreads rise more when there are no gains from trade (B1 and B2), and the difference between gain and no-gains treatments is sharper for B1 and GT1. The coefficients imply that holding our other observables fixed, spreads in B1 rise by about 10 points relative to GT1 over the 30 rounds of the experiment. The corresponding difference between B2 and GT2 is about 5 points. Adding the time trend leaves the market feedback coefficients essentially unchanged, suggesting that there are additional forces, separate from market feedback, driving the divergence between the treatments with and without gains. Without more evidence, we must be agnostic about what those forces are. One possibility is that our simple market feedback specification misses some subtle nonlinear or lagged feedback effects that are picked up in the time trend. Another possibility is that subjects are learning about the hazards of the zero sum environments by observing their opponents' outcomes or by introspection at the same time as they learn from direct personal experience. Whatever its source, the effect of the time trend for B1 is to raise average spreads by about 16 points over 30 rounds. This effect is about 50% larger than the effect of having lost money in the previous round, and roughly equal to the effect produced by reducing a subject's end-of-session trading profits from the 80th to the 20th percentile (about a £1.6 reduction in accumulated profits).

Overall, spreads exhibit a fairly low persistence and a constant underlying tendency to move upwards over time that could reflect learning during the experiment as well as strategic interaction and emulation among players. While per round gains appear to be of minor importance for spread setting, unfavorable market outcomes induce a substantial increase in spreads to shield against further losses. Subjects seem prepared to reduce the distance between bid and ask prices whenever the level of accumulated payoffs is large enough. These general patterns appear to be quite consistent across the different experimental designs, although the magnitude of the coefficients varies from one treatment to another. The differences across treatments also fit a fairly consistent pattern. When subjects react by increasing their spreads, thus making trade rarer and on more favorable terms, they do so most strongly in B1 and most weakly in GT1. The responses in B2 and GT2 are of intermediate strength

and the difference between the two is small. The evidence seems to suggest that in a low noise environment, subjects recognize an important difference between the gains and no-gains settings, and this affects how they respond to feedback. In the noisier environment of B2 versus GT2, this differential reaction to feedback is much less pronounced.

Table 7: FE, Dependent Variable $Spread_{it}$ (time trend included)

	Treatments					
	B1	GT1	B2	GT2	CE	CU
$Spread_{it-1}$	0.341*** (0.107)	0.233* (0.127)	0.097*** (0.027)	0.273*** (0.095)	0.330*** (0.110)	0.173 (0.110)
$Gain_{it-1}$	3.960*** (1.059)	0.961 (1.671)	-0.121 (1.614)	2.858** (1.275)	0.412 (1.710)	2.125 (1.383)
$Loss_{it-1}$	10.729*** (3.819)	3.465** (1.628)	6.844*** (1.256)	6.108*** (2.088)	6.933*** (2.504)	6.131*** (1.191)
$TProfits_{it-1}$	-0.092*** (0.039)	-0.027** (0.012)	-0.059* (0.034)	-0.035** (0.016)	-0.009 (0.043)	-0.063*** (0.017)
$OSpread_{it-1}$	0.013 (0.073)	0.057 (0.050)	0.072 (0.058)	0.105** (0.047)	-0.080 (0.059)	0.107* (0.059)
$Round$	0.537*** (0.201)	0.235** (0.107)	0.408** (0.207)	0.232 (0.160)	0.516*** (0.162)	0.396* (0.203)

Cluster-robust bootstrap standard errors based on 1000 replications in parentheses. *** indicates significance at 1%, ** indicates significance at

5%, * indicates significance at 10%. Sample size: 6936 observations.

6 Profit Opportunities and Market Efficiency

So far we have documented the trading activity in our laboratory financial market and studied the trading strategies that our participants used. Now we turn to the economic implications of these strategies. One corollary of the no-trade results is that in equilibrium it should be impossible to earn positive profits from trading on private information. Below, we test this implication in two related ways. First, we estimate how much profit a rational trader with private information could earn by trading against our subjects. If this profit opportunity is close to zero, this suggests that even though trade in our baseline treatments is a deviation from equilibrium, this deviation has minor economic implications. Large profit opportunities would be more worrying, as these would show that some of our subjects persist in using strategies that can be easily exploited. Our analysis, which is based on computing

best responses to the empirical distribution of strategies in each treatment, is presented in the next section.

Second, we ask whether, conditional on trade, the price reflects all available public and private information. This would imply that our subjects tend to trade only at prices for which both the buyer and seller can expect zero profits, which one might regard as a minor departure from the no-trade equilibrium. Because the impounding of private information in prices has been studied extensively, we should point out that our markets differ from those commonly discussed in the market efficiency literature. Because our baseline treatments have no-trade (and therefore no prices) in equilibrium, for these markets our hypotheses are somewhat informal since we cannot make theoretical predictions about whether prices should be efficient. In treatments with gains from trade, we do have theoretical predictions, but equilibrium prices will depend in part on how the gains from trade are split between buyer and seller, and this will not be resolved in our two player game exactly as it would be with a continuum of traders.

6.1 Informational Profit Opportunities

For our baseline treatments, the no-trade theorem implies that positive profits are a disequilibrium phenomenon; thus we expect to see profit opportunities driven toward zero over time. In this section we test that prediction by considering a hypothetical trader who (as discussed above) understands the data generating process and is endowed with the same type of private information (i.e., one signal) as our subjects. For the baseline treatments B1 and B2, we compute the expected trading profit that this trader could have earned by using the order strategy that is a best response to the empirical distribution of play by our subjects. We repeat this exercise twice, computing best responses to empirical play in the sub-samples of Rounds 1-5 and Rounds 26-30. Our hypothesis is that these best response profits should decline toward zero over time for B1 and B2.

For treatments GT1 and GT2, the exercise is slightly more counterfactual. The best response profit of a subject in these treatments mixes the returns to private information with the commission c for a successful trade. Because our interest is in isolating profits attributable to private information, we compute the best response profits of a trader who does not receive c (but in every other respect is the same as one of our subjects). As a theoretical benchmark, we also compute the expected profit that this trader would earn with a best response to the equilibrium strategies for GT1 and GT2 that were derived earlier. In contrast with the baseline treatments, these theoretical returns to private information are strictly positive here. This is possible because the GT1 and GT2 equilibria permit some

“mispricing” to survive (as long as it does not exceed c). Our hypothesis is that the empirical returns to private information should approach these theoretical levels over time. Because GT1 and GT2 have other equilibria with lower levels of trade, this prediction is not as compelling as the one for the baseline treatments.

We should emphasize that this exercise presents a snapshot of profit opportunities at one window in time; if our hypothetical trader were to actually enter the market, he would not necessarily continue to earn these profits because other traders would begin to adjust their strategies. Furthermore, computing a best response with respect to a finite sample of play will tend to overstate expected profits somewhat because the best response will be fine-tuned to the sample. For all of these reasons, the computed profits should be interpreted as upper bounds on the returns that a very skilled subject could earn by using his private information.

6.1.1 Methodology

Consider the best response ask a of a subject with signal A who expects opponents to bid according to a mixed strategy with density $g(b|B)$.³³ (The best response bid is handled identically.) His expected profit from sales can be written as

$$E\pi_{Ask}(a|A) = \int_{B \in \mathcal{S}} \int_{b \in \mathcal{P}} f(B|A) g(b|B) 1(b \geq a) \left(\frac{a+b}{2} - (A+B) + c \right) dbdB,$$

where we set $c = 0$ for all of the treatments, as discussed above. We can alternatively write this expression as

$$\begin{aligned} E\pi_{Ask}(a|A) &= \int_{B \in \mathcal{S}} \int_{b \in \mathcal{P}} R(A, B, a, b) w(A, B) h(B, b) dbdB \\ &= E_{h(B,b)}(R(A, B, a, b) w(A, B)), \end{aligned}$$

where $R(A, B, a, b) \equiv 1(b \geq a) \left(\frac{a+b}{2} - (A+B) \right)$ is the subject’s profit from an opponent signal-bid realization of (B, b) , $h(B, b) \equiv g(b|B) f(B)$ is the unconditional joint distribution of those opponent signal-bid pairs, and $w(A, B) \equiv \frac{f(B|A)}{f(B)}$ is essentially a kernel that adds more weight to opponent signals that are likelier given A .

We (and by assumption, our subject) can compute $R(A, B, a, b)$ and $w(A, B)$ analytically, so if we also know the joint distribution $h(B, b)$ we can compute $E\pi_{Ask}(a|A)$ and find the optimal a . Alternatively, if we only observe a finite sample of draws from $h(B, b)$ (as it is the

³³Assuming a mixed strategy at this stage will allow us to handle heterogeneity in empirical bids more naturally later.

case for our experiments), we can estimate the expectation and compute a best response with respect to that sample. That is, given a sample $\{(B_1, b_1), (B_2, b_2), \dots, (B_N, b_N)\}$ of opponent bids, we estimate

$$E\pi_{Ask}(a | A) \approx \frac{1}{n} \sum_{n=1}^N w(A, B_n) R(A, B_n, a, b_n).$$

For each treatment, we compute best response asks with respect to three sub-samples of opponents' empirical play: all (B, b) pairs in Rounds 1-5, all (B, b) pairs in Rounds 26-30, and the entire sample of (B, b) pairs in Rounds 1-30, in each case pooling the data across all sessions. We proceed similarly to compute a subject's best response bid. His total expected best response profit $\pi_{BR}(A)$ is the sum of his expected profits from buying and from selling. Finally, we can compute the ex ante best response profit of this subject:

$$\pi_{BR} = \int_{A \in \mathcal{S}} f(A) \pi_{BR}(A) dA.$$

6.1.2 Results

These computed best response profits are presented in Table 8. In three of the four treatments, B1, B2, and GT2, the results give strong partial support to our hypotheses: in each case, opportunities to profit from private information are driven substantially lower over the 30 rounds of trading. In contrast, informational profit opportunities in GT1 show no change over time.³⁴ Potential profits in the two baseline treatments start at about the same level, but decline more slowly when the asset value is noisier (Treatment B2). The stronger prediction that potential profits should be driven out entirely in B1 and B2 is not supported in either treatment, but by the end of the experiment even a flawless trader could earn no more than a penny per match in B2, and less in B1. Meanwhile, the profit opportunity in GT2 approaches its theoretical level. While there is no rigorous basis for predicting what would happen with additional rounds of trade, a linear extrapolation indicates that around 20 additional rounds would be needed to wipe out all profits in B1 (and about 45 additional rounds in B2). If (as seems more likely) convergence slows as profits shrink, more time would be required.

For treatment GT1, it is worth remembering that a subject's expected trading profits can be decomposed into the product of two terms: his probability of trading, and his expected profit conditional on trade. We have already seen that trade occurs much more frequently in

³⁴The index of convergence in the last row of the table represents the fraction of excess profits in the first five rounds that are still available in the last five rounds. Values close to zero indicate that most excess profit opportunities have been driven out by the end of the experiment. Values close to one indicate little change in profit opportunities over time.

Table 8: Best Response Profits

	Treatments			
	B1	B2	GT1	GT2
$\pi_{BR}^{(1-30)}$	1.18	1.23	1.68	1.66
$\pi_{BR}^{(1-5)}$	1.69	1.64	1.99	2.10
$\pi_{BR}^{(26-30)}$	0.72	1.05	1.92	1.31
$\pi_{BR}^{Theor.}$	0	0	0.22	1.21
<i>Convergence</i>				
$\frac{\pi_{BR}^{(26-30)} - \pi_{BR}^{Theor.}}{\pi_{BR}^{(1-5)} - \pi_{BR}^{Theor.}}$	0.43	0.64	0.96	0.11

Amounts are in pence per trading opportunity; expected profit per round are larger by a factor of four. Estimates for each subsample ($\pi_{BR}^{(1-5)}$ and $\pi_{BR}^{(26-30)}$) are based on 200 signal-order realizations, while the full sample ($\pi_{BR}^{(1-30)}$) contains 1200 realizations.

GT1 than can be explained by equilibrium theory, and this empirical fact helps to explain the persistently high profit opportunities in GT1.

6.2 Do Prices Aggregate Private Information?

Expected profit opportunities depend on the probability of finding a willing trading partner, and on the distance between the transaction price and the expected asset value conditional on all private information. In this section we discuss the information content of prices. While the previous analysis was hypothetical (how much profit *would* a very talented trader make, given the empirical distribution of signals and strategies), the present analysis is based on the actual realizations of the signals and on the actual transaction prices.

Recall that the expected value of the asset conditional on all private information is $E(V | A, B) = A + B$. We will say that prices are *efficient* if $p = A + B$ always holds. Note that what we are requiring is that the price aggregates all information available in our economy, that is, that the market is efficient in a strong form.³⁵ More realistically, allowing for the possibility that prices are noisy, we will say that that prices are *efficient on average* if $E(\varepsilon_p | A, B) = 0$.³⁶ That is, there may be noise in p (and consequently expected gains and losses from trading at p), but that noise is orthogonal to private information. Alternatively,

³⁵We do not require the price to reveal the realized asset value, since there is a component of noise X that not even the aggregation of private information is able to eliminate.

³⁶The price prediction error $\varepsilon_p = p - E(V | A, B)$ was defined in section 2.1. Note that, more subtly, prices can be consistently wrong and still fully reveal private information. For example, if $p = \frac{1}{2}(A + B)$, then p is sufficient for $E(V | A, B)$ even though the price is too low. We have checked that prices in our treatments are not fully revealing in this sense, but those results are not reported here.

if prices are not efficient on average, we can measure the total variation in ε_p and try to determine how much of the pricing error can be predicted using private information.

Under the (strong form) Efficient Market Hypothesis, we would expect prices to approach efficiency (in both senses) over time, as traders become more experienced. However, it is not clear whether this is the most reasonable prediction for our baseline treatments with no gains from trade. In these treatments, trade should never occur in theory and becomes rarer over time in practice. Thus prices arise from thinner and thinner markets over time, and are generated by the subset of subjects who continue to trade despite incentives not to. One might be skeptical that these conditions would lead to efficient prices. In our gains from trade treatments, markets remain thicker (both in theory and in practice). However, in these treatments there is a $2c$ surplus to split between the buyer and seller, and their equilibrium shares of this surplus depend on V . (Sellers rationally mark up their asks more with low signals (when they are likeliest to sell), and buyers mark down their bids more with high signals (when they are likeliest to buy).) As a result, *equilibrium* prices are neither efficient nor efficient on average, although they cannot be too far from $E(V)$. Thus for treatments GT1 and GT2 we predict that prices should approach these equilibrium levels of information content.

We begin with efficiency. Table 9 reports average values of $|\varepsilon_p|$ for each treatment over the first and last ten rounds.

The average price error starts at roughly similar levels in all of the treatments (roughly 15% of the average asset value) and tends to decline modestly over time (to around 12% of the average asset value in Rounds 21-30). Thus, the market is not much more efficient at the end than it is at the beginning of the experiment.

Table 9: Average Value of $|\varepsilon_p|$

	Treatments			
	B1	GT1	B2	GT2
<i>Rounds 1-10</i>	8.36	7.68	6.58	6.66
<i>Rounds 21-30</i>	7.83	6.81	5.90	5.30
<i>Equilibrium</i>	–	1.07	–	1.98

Amounts are in pence.

Next we ask what fraction of this price error can be attributed to private information that is systematically not impounded in the price by testing whether $E(\varepsilon_p | A, B) = 0$ holds.

For each treatment, we run a simple OLS specification over all of the matches in which trade occurs:

$$\begin{aligned} \varepsilon_p = & \gamma_0 + \gamma_1 \mathit{Rounds}_{[11-20]} + \gamma_2 \mathit{Rounds}_{[21-30]} + \delta_0 v \\ & + \delta_1 (v \cdot \mathit{Rounds}_{[11-20]}) + \delta_2 (v \cdot \mathit{Rounds}_{[21-30]}), \end{aligned} \quad (5)$$

where $\mathit{Rounds}_{[11-20]}$ and $\mathit{Rounds}_{[21-30]}$ are dummy variables for the second and the last ten rounds respectively and $v = A + B - 50$ is the demeaned expectation of V given A and B .³⁷ If the coefficients δ_0 , δ_1 and δ_2 are all zero, then v does not predict the direction of the price error, either early in the experiment or in later rounds. Alternatively, if $|\delta_0|$ is large but, for example, $|\delta_0 + \delta_2|$ is small, then prices are not efficient on average in early rounds, but efficiency improves (v predicts less of the variation in ε_p) towards the end of the experiment. This will tend to be true whenever δ_0 and δ_2 (or δ_1) have opposite signs. If δ_0 and δ_2 (or δ_1) have the same sign, then price efficiency grows worse, not better, over time.

Table 10: A Simple Model for ε_p : OLS Estimates

	Treatments			
	B1	GT1	B2	GT2
<i>Constant</i>	-2.303** (0.95)	-4.311*** (1.56)	-3.908*** (1.36)	-2.321** (1.00)
<i>Rounds</i> _[11-20]	-2.385** (1.04)	1.499 (1.29)	-0.447 (1.77)	4.850*** (1.01)
<i>Rounds</i> _[21-20]	-0.454 (0.87)	2.476** (1.01)	2.072 (1.69)	1.917** (0.94)
<i>v</i>	-0.478*** (0.06)	-0.366*** (0.07)	-0.212*** (0.04)	-0.336*** (0.06)
<i>v · Rounds</i> _[11-20]	-0.165*** (0.03)	-0.012 (0.10)	-0.211* (0.12)	0.072 (0.11)
<i>v · Rounds</i> _[21-30]	0.033 (0.07)	-0.081 (0.09)	-0.037 (0.09)	0.079 (0.07)
<i>R</i> ²	0.55	0.47	0.16	0.22
<i>N</i>	448	768	492	541

Cluster-robust bootstrap standard errors based on 1000 replications in parentheses. *** indicates significance at 1%, ** indicates significance at 5%, * indicates significance at 10%.

The coefficients on v in Table 10 are all negative and significant (at the 1% level), indicat-

³⁷In principle, more flexible specifications in which A and B enter separately could also be studied. In practice, since we find that $E(\varepsilon_p | A + B) = 0$ is violated in a very strong and consistent way, teasing out more subtle interactions between ε_p , A , and B is of limited interest. Demeaning V has no effect on δ_0 and δ_1 , but it allows us to interpret the intercept γ_0 more meaningfully.

ing that prices are not efficient on average in early rounds. A negative coefficient between -1 and 0 means that as pooled private information about the asset value moves away from the *ex ante* expectation of 50, prices move in the same direction, but not as far. The larger (in magnitude) the coefficient, the more sluggishly prices respond to information. Essentially, information is partially but not fully aggregated in prices since the price tends to be anchored to the unconditional expected value more than it should.³⁸ The time interaction on v shows that the amount of information impounded in prices never improves significantly over time. On the contrary, there is in some cases (B1 and B2) an appreciable deterioration.

As noted above (Table 9), the size of the total price error $|\varepsilon_p|$ is similar across the treatments. These results show that in treatments B1 and GT1 (the treatments with more private information and less noise in V), more of the price error tends to be explained by un-impounded private information (large coefficients on v and large R^2). Conversely, in the treatments with less private information and more white noise in V (B2 and GT2), less of the price error is explained by private information, leaving a larger residual that we could loosely attribute to “model uncertainty.” For the gains from trade treatments, we noted earlier that equilibrium prices are not efficient on average. If we run the same regression specification on simulated equilibrium data, the coefficients on v are -0.099 for GT1 and -0.236 for GT2. GT2 comes rather close to this prediction toward the end of the experiment, with $\delta_0 + \delta_2 = -0.257$, but GT1 does not.

In conclusion, weaving in the profit opportunity results for the baseline treatments, we can say that in B1 and B2 spreads generally widen, trade grows rarer, and informational profit opportunities shrink, as predicted by theory. Since conditional on finding a willing counterpart, profits from private information do persist, the reduction of profit opportunities is mainly explained by the difficulty of finding another participant willing to trade.

7 Conclusion

Theoretical no-trade results offer conditions under which informational trade cannot occur *at all* in equilibrium. It is difficult to know whether these theoretical results really fail in the real world (and if so, why), because generally we can neither observe traders’ private information nor confirm that those theoretical conditions are satisfied. At the same time, the consistency violation at the heart of the theory does not shed any light on the process by which real world traders might arrive at a no-trade outcome. Our paper is a first attempt

³⁸This sluggishness in prices has been observed also in other experiments in which subjects (acting as market makers) are explicitly asked to predict the asset value and set the price (Cipriani and Guarino, 2005).

to bridge some of these gaps. In the interest of robustness and external validity, we create a trading game that is recognizable as a financial market (albeit a stylized one) and allow subjects to learn about the data-generating process (asset value and clues) on their own. By gathering data in the laboratory, we are able to observe subjects' private information and exclude many other confounding motives for trade. Because we observe trading outcomes for our subjects from the inception of the market, we are in a position to study the process by which subjects move toward a no-trade outcome or diagnose the reasons that they fail to do so.

We find that market feedback matters: trading losses encourage more conservative strategies, thereby reducing trade. Subjects lose money more frequently when there are no gains from trade, so feedback tends to drive these markets in the direction of no-trade. There may also be some role for introspection: in some cases subjects react to losses more emphatically when they know that trade is zero-sum than they do when mutual gains are possible. This differential response is quite pronounced for parameters that make the role of private information more salient and less strong when the role of private information is obscured by more noise. Overall, our results should offer some vindication to both detractors and defenders of the no-trade theorems' empirical validity. For the latter, the fact that trade is virtually wiped out in a setting (treatment B1) where traders must learn a noisy information structure on the fly demonstrates that the no-trade logic is not as fragile as sometimes supposed. The former can counter that when the level of noise rises (treatment B2), informational trade is driven out much more gradually, and we cannot conclude that it will ever cease entirely. In our view, these results should probably be seen as points on a continuum; perhaps future work can more fully trace out the slope of informational trade with respect to the degree of model uncertainty confronted by traders.

References

- [1] Aumann, R. (1976) "Agreeing to Disagree." *The Annals of Statistics*, 4, 1236-39.
- [2] Arellano, M. (1987) "Computing Robust Standard Errors for Within-Group Estimators." *Oxford Bulletin of Economics and Statistics*, 49, 431-34.
- [3] Arellano, M. and Bond, S. (1991) "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *The Review of Economic Studies*, 58, 277-97.

- [4] Blume, L., Coury, T., and Easley, D. (2006) "Information, Trade and Incomplete Markets." *Economic Theory*, 29, 379-94.
- [5] Bound, J., Jaeger, D. and Baker, R. (1995) "Problems with Instrumental Variables Estimation when the Correlation between the Instruments and the Endogenous Explanatory Variable is Weak." *Journal of the American Statistical Association*, 90, 443-50.
- [6] Cameron, C., Gelbach, J. and Miller, D. (2006) "Robust Inference with Muly-way Clustering." *NBER Technical Working Paper*, No. 327.
- [7] Cameron, C., Gelbach, J. and Miller, D. (2007) "Bootstrap-based Improvements for Inference with Clustered Errors." *NBER Technical Working Paper*, No. 344.
- [8] Carrillo, J. and Palfrey, T. (2007) "No Trade." *CEPR Discussion Paper*, No. DP6554.
- [9] Cipriani, M. and Guarino, A. (2005) "Herd Behavior in a Laboratory Financial Market." *American Economic Review*, 95, 1427-43.
- [10] Copeland, T. and Friedman, D. (1991) "Partial Revelation of Information in Experimental Asset Markets." *Journal of Finance*, 46, 265-95.
- [11] Dow, J. and Gorton, G. (1997) "Noise Trading, Delegated Portfolio Management, and Economic Welfare." *Journal of Political Economy*, 105, 1024-50.
- [12] Duxbury, D. (1995) "Experimental Asset Markets within Finance." *Journal of Economic Surveys*, 9, 331-71.
- [13] Eliaz, K. and Spiegel, R. (2007) "A Mechanism-Design Approach to Speculative Trade." *Econometrica*, 75, 875-84.
- [14] Forsythe, R., Palfrey, T. and Plott, C. (1982) "Asset Valuation in an Experimental Market." *Econometrica*, 50, 537-67.
- [15] Grinblatt, M. and Keloharju, M. (2001) "What Makes Investors Trade." *Journal of Finance*, 56, 589-616.
- [16] Jehiel, P. and Koessler, F. (2006) "Revisiting Games of Incomplete Information with Analogy-based Expectations." *Games and Economic Behavior*, forthcoming.
- [17] Lei, V., Noussair, C. and Plott, C. (2001) "Nonspeculative Bubbles in Experimental Asset Markets: Lack of Common Knowledge of Rationality vs. Actual Irrationality." *Econometrica*, 69, 831-59.

- [18] Milgrom, P. and Stockey, N. (1982). “Information, Trade and Common Knowledge.” *Journal of Economic Theory*, 26, 17-27.
- [19] Morris, S. (1994). “Trade with Heterogeneous Prior Beliefs and Asymmetric Information.” *Econometrica*, 62, 1327-47 .
- [20] Neeman, Z. (1996) “Common Beliefs and the Existence of Speculative Trade.” *Games and Economic Behavior*, 16, 77-96.
- [21] Nickell, S. (1981) “Biases in Dynamic Models with Fixed Effects.” *Econometrica*, 49, 1417-26.
- [22] Odean, T. (1999) “Do Investors Trade Too Much?” *American Economic Review*, 89, 1279-98.
- [23] Plott, C. and Sunder, S. (1982) “Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models.” *Journal of Political Economy*, 90, 663-98.
- [24] Plott, C. and Sunder, S. (1988) “Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets.” *Econometrica*, 56, 1085-1118.
- [25] Sebenius, J. K. and Geanakoplos, J. (1983) “Don’t Bet on It: Contingent Agreements with Asymmetric Information.” *Journal of the American Statistical Association*, 78, 424-26.
- [26] Serrao-Padial, R. (2007) “On the Possibility of Trade with Pure Common Values under Risk Neutrality.” Mimeo, University of Wisconsin-Madison.
- [27] Sonsino, D. (1995) “Impossibility of Speculation Theorems with Noisy Information.” *Games and Economic Behavior*, 8, 406-23.
- [28] Sonsino, D. (1998) “Geanakoplos and Sebenius Model with Noise.” *International Journal of Game Theory*, 27, 111-30.
- [29] Sonsino, D., Erev, I., and Gilat, S. (2002) “On Rationality, Learning and Zero-Sum Betting—An Experimental Study of the No-Betting Conjecture.” Mimeo, Technion University.
- [30] Søvik, Y. (2004) “Strength of Dominance and Depths of Reasoning - An Experimental Study.” Mimeo, University of Oslo.

- [31] Sunder, S. (1995) “Experimental Asset Markets: a Survey.” in Kagel, J. and Roth, A., eds. *The Handbook of Experimental Economics*, Princeton University Press.
- [32] Tirole, J. (1982) “On the Possibility of Speculation under Rational Expectations.” *Econometrica*, 1, 1163-81.

Appendix

A Equilibrium with Gains from Trade

In this section, we characterize a class of symmetric equilibria of the trading game with gains from trade for which trade occurs with positive probability. A symmetric equilibrium consists of a pair of bid and ask functions $(b(S), a(S))$, where S is the signal received by the subject. Signals are drawn from the range $[S_l, S_h]$ according to the joint pdf $f(A, B)$. Since the signals are identically distributed, we will make the convention that a subject's signal is labeled A when we are considering his optimal ask price (hence selling to an opponent with signal B), whereas his signal is labeled B when we are considering his optimal bid price (thus, buying from an opponent with signal A).

We will say that a bid function is *separating* on an interval $[S_1, S_2]$ if $b(S)$ is strictly increasing on that interval and *pooling* if $b(S)$ is constant on that interval. (The same terminology will also apply to ask functions.) The analysis below is designed to identify equilibria in symmetric, separating strategies. Sometimes no fully separating equilibrium exists due to “unravelling” at very low and very high signals; in these cases we can typically amend the analysis slightly to identify semi-separating equilibria with pooling regions at high and low signals. The general approach that we take to constructing these equilibria is as follows:

1. The first order conditions for a separating equilibrium form a system of differential conditions. We solve this system numerically to identify candidate separating equilibria. There is a degree of freedom in the initial conditions, and therefore there is typically a range of candidate equilibria.
2. Second order conditions are checked to ensure local optimality of the solution.
3. The candidate $b(S)$ and $a(S)$ functions from (1) may not span the entire signal space. If so, we extend those functions, respecting any conditions imposed by equilibrium. This can introduce pooling at high signals (for $b(S)$) or at low signals (for $a(S)$).
4. We check global optimality of the candidate strategy. This is necessary because we cannot guarantee that an agent's profits are quasi-concave in his strategy. The main concern is that for some signals the candidate strategy might lose money (even though it is locally optimal). Agents with these signals would do better by placing non-competitive orders.

5. Step 4 tends to be an issue, if at all, for very high and very low signals. We adjust the candidate solution by expanding the set of signals for which there is pooling (and shrinking the set for which there is separation) to find the minimal level of pooling for which global optimality is satisfied.

An order strategy constructed according to Steps 1-5 will represent a Bayesian Nash equilibrium of the trading game. Our main objective will be to examine equilibria of this form for the parameters in our experiments. We make no claims regarding the general existence of such equilibria or the non-existence of equilibria with more exotic patterns of pooling and separation. Below, we elaborate on each of these steps.

A.1 Step 1

A.1.1 First order conditions for a separating equilibrium

Fix a conjectured symmetric equilibrium strategy $(b(\cdot), a(\cdot))$ with the following piecewise form

$$\begin{aligned} b(S) &= 0 \quad \text{for } S < B_l \\ b([B_l, B_h]) &\text{ is strictly increasing, differentiable, and maps onto } P = [p_l, p_h] \\ b(S) &= p_h \quad \text{for } S > B_h, \end{aligned}$$

and

$$\begin{aligned} a(S) &= p_l \quad \text{for } S < A_l \\ a([A_l, A_h]) &\text{ is strictly increasing, differentiable, and maps onto } P = [p_l, p_h] \\ a(S) &= 100 \quad \text{for } S > A_h. \end{aligned}$$

We will discuss bidding for the highest and lowest signals later. For now we focus on the intermediate signals for which b and a are conjectured to be increasing. We will establish necessary conditions for $(b(\cdot), a(\cdot))$ to be a best response to itself on this range. Under risk neutrality, a subject's strategy regarding purchases does not interact with his strategy regarding sales, so we can examine these decisions one at a time; we start with the best response bid. Consider a subject's optimal bid b when he sees signal B and his opponent is using the conjectured equilibrium strategy. Bidding $b > p_h$ cannot be a best response (it trades exactly as often as bidding $b = p_h$, but at a higher price). Meanwhile, bidding $b < p_l$ earns 0 with certainty. Bidding $b = p_l$ is a subject that we will return to later. For now, consider possible bids $b \in (p_l, p_h]$. Such a bid trades iff.

$$a(A) \leq b \quad \text{or equivalently} \quad A \leq \alpha(b),$$

where $\alpha = a^{-1}$ is increasing and differentiable on P .

In the event that he buys when his opponent's signal is A , his payoff is

$$E(v | A, B) - \left(\frac{1}{2}a(A) + \frac{1}{2}b \right) + c = A + B - \frac{1}{2}a(A) - \frac{1}{2}b + c.$$

His expected payoff from placing bid b is therefore

$$\pi(b | B) = \int^{\alpha(b)} f(A | B) \left(A + B - \frac{1}{2}a(A) - \frac{1}{2}b + c \right) dA.$$

If placing a bid $b \in (p_l, p_h]$ is optimal, it must satisfy the first order condition:

$$\begin{aligned} \frac{d\pi(b | B)}{db} &= \alpha'(b) f(\alpha(b) | B) (\alpha(b) + B - b + c) - \frac{1}{2}F(\alpha(b) | B) = 0, \quad \text{or} \\ \alpha'(b) &= \frac{1}{2} \frac{F(\alpha(b) | B)}{f(\alpha(b) | B)} \frac{1}{\alpha(b) + B - b + c}. \end{aligned}$$

Next, proceed we similarly for an agent's best response ask a , given signal A and equilibrium bidding by the opponent. An agent whose ask satisfies $a \in [p_l, p_h)$ will sell iff.

$$a \leq b(B) \quad \text{or equivalently} \quad B \geq \beta(a),$$

where $\beta = b^{-1}$.

In the event that he sells when his opponent's signal is B , his payoff is

$$\left(\frac{1}{2}a + \frac{1}{2}b(B) \right) - E(v | A, B) + c = \left(\frac{1}{2}a + \frac{1}{2}b(B) \right) - A - B + c,$$

so his expected payoff is

$$\pi(a | A) = \int_{\beta(a)} f(B | A) \left(\left(\frac{1}{2}a + \frac{1}{2}b(B) \right) - A - B + c \right) dB.$$

A necessary condition for an ask $a \in [p_l, p_h)$ to be optimal is therefore the following FOC:

$$\begin{aligned} \frac{d\pi(a | A)}{da} &= -\beta'(a) f(\beta(a) | A) (a - A - \beta(a) + c) + \frac{1}{2}(1 - F(\beta(a) | A)) = 0, \\ \text{or} \\ \beta'(a) &= \frac{1}{2} \frac{1 - F(\beta(a) | A)}{f(\beta(a) | A)} \frac{1}{a - A - \beta(a) + c}. \end{aligned}$$

For these necessary conditions to identify a symmetric equilibrium, they must be satisfied when we plug in the conjectured equilibrium strategies; that is, when we plug in $\beta(b) = B$ in the first FOC and $\alpha(a) = A$ in the second. Doing this, and relabeling the argument in each equation as p , we have the following pair of differential equations:

$$\begin{aligned}\alpha'(p) &= \frac{1}{2} \frac{F(\alpha(p) | \beta(p))}{f(\alpha(p) | \beta(p))} \frac{1}{\alpha(p) + \beta(p) - p + c}, \\ \beta'(p) &= \frac{1}{2} \frac{1 - F(\beta(p) | \alpha(p))}{f(\beta(p) | \alpha(p))} \frac{1}{p - \alpha(p) - \beta(p) + c}.\end{aligned}$$

This system of equations characterizes candidate symmetric equilibrium strategies on (p_l, p_h) . Solving these equations analytically appears to be intractable, but solutions can be computed numerically using standard techniques, as described next. To compute a solution, we must choose two initial value conditions for the two differential equations. This amounts to a choice about the range of prices $[p_l, p_h]$ over which trade may occur. Because our trading game treats buying and selling with complete symmetry, we impose the additional assumption that $\frac{p_l + p_h}{2} = E(V) = 50$; this eliminates one of the initial conditions. However, one degree of freedom (loosely corresponding to $|p_h - p_l|$) remains, and this creates the potential for multiple equilibria with different levels of trade. This symmetry assumption and the computation of candidate strategies are detailed in the next section.

A.1.2 Numerical Computation of Equilibrium

For the sake of brevity, rewrite the system of differential equations as

$$\begin{aligned}\alpha' &= \frac{1}{2} \frac{F(\alpha | \beta)}{f(\alpha | \beta)} \frac{1}{\alpha + \beta - p + c}, \\ \beta' &= \frac{1}{2} \frac{1 - F(\beta | \alpha)}{f(\beta | \alpha)} \frac{1}{p - \alpha - \beta + c}.\end{aligned}$$

With suitable initial or boundary conditions, this system could be solved without further discussion. Unfortunately, an initial condition approach is problematic for several reasons. First of all, we have no *a priori* basis for determining the price interval on which equilibrium bids and asks overlap. Without knowing either p_l or p_h , we don't have a starting point for the solver. In principle, this problem is not insurmountable, as we could adopt a forward-shooting approach of "guessing" an appropriate p_l , solving the differential equations, checking for consistency at p_h , and repeating until a p_l yielding consistent solutions was found.

Unfortunately, this approach is unsuitable here because this system of differential equations is unstable at its endpoints. To see this, observe that $\beta(p) > \alpha(p)$ (because $b(S) < a(S)$), so p_l is defined by $\alpha(p_l) = S_l$, and p_h is defined by $\beta(p_h) = S_h$. (In other words, the endpoints of the overlapping region are defined by the ask price at the lowest possible signal and the bid price at the highest possible signal. For the α' equation, this means that at p_l we have $F(\alpha|\beta)|_{p=p_l} = F(S_l|\beta(p_l)) = 0$, so unless $\alpha + \beta - p + c$ also vanishes at p_l , the ask function $a(S)$ has infinite slope at S_l .³⁹ For similar reasons, $\beta(S)$ potentially approaches S_h with infinite slope. Consequently, starting the integration from either endpoint is likely to amplify numerical errors and any misspecification of initial conditions, making convergence to a consistent solution difficult to obtain.

The solution is to appeal once more to symmetry. Note that in our experiments, signals are distributed symmetrically about a mean of 25, and the mean value for the item is 50. (We are using concrete numbers here for clarity only – the argument does not depend on the particular numbers.) Consider the function $h(\delta) = \alpha(50 - \delta) + \beta(50 + \delta)$. The solution to the differential equations is symmetric about the point $(p, S) = (50, 25)$ if $h(\delta) = 50$ for all δ . Furthermore, if the following two conditions are true – (a) $h(\delta) = 50 \Rightarrow h'(\delta) = 0$, and (b) there exists some δ^* such that $h(\delta^*) = 50$ – then we must have $h(\delta) = 50$ for all δ . (That is, symmetry at one point, plus (a), implies symmetry everywhere.) Demonstrating (a) is simply a matter of differentiating h , plugging in the equations for α' and β' , and properly appealing to the symmetry of f and F – the details are omitted. Assumption (b) follows if, for example, we are willing to assume that the highest and lowest sale prices in equilibrium are symmetric about the mean value of the object (that is, $\frac{p_l + p_h}{2} = 50$). Since the buy and sell sides of the market are treated completely symmetrically in this model, (b) is a reasonable assumption, and henceforth we adopt it.⁴⁰

The benefit of being able to assume a symmetric solution is that we now have a natural alternative to “shooting” trial solutions from the endpoints, where the derivatives are badly behaved – we can shoot trial solutions forward and backward from the point of symmetry at $p = 50$. There are two virtues to this: first, neither of the derivatives is degenerate at $p = 50$, and second, symmetry at $p = 50$ means that we only need to guess one parameter. To see this, notice that if we start integrating the differential equations at any arbitrary p_0 , we need to specify two parameters: $\alpha_0 = \alpha(p_0)$ and $\beta_0 = \beta(p_0)$. However, if we start integrating at $p = 50$, we have $h(0) = \alpha(50) + \beta(50) = 50$, so once we fix one of the parameters, say α_0 , the other one is nailed down by symmetry. Reducing the dimension of the parameter space

³⁹Actually, for our signal structure, $f(\alpha|\beta)$ also vanishes, but at a slower rate than $F(\alpha|\beta)$, so the argument is still valid.

⁴⁰It is entirely possible that (b) is also a necessary assumption (i.e., that equilibria treating the buy and sell sides of the market asymmetrically do not exist). However, we don’t have a proof of this.

that must be searched to find a consistent solution dramatically simplifies the computational burden.

A.2 Step 2: Second Order Conditions

We confirmed analytically that the second order conditions are satisfied. Since the calculation is routine but tedious we omit it here, but the details are available from the authors upon request.

A.3 Step 3

For some parameters and some initial conditions, the solutions $\alpha(p)$ and $\beta(p)$ map onto \mathcal{S} . More typically, either $\alpha(P)$ or $\beta(P)$ or both will be a strict subset of \mathcal{S} . In this case, we need to extend one or both of the strategies. The first criterion for these extensions is that they must be best responses. If there are multiple best responses, we use the freedom to choose extensions that do not “break” the candidate equilibrium constructed at Step 1. A few examples may help to illustrate.

Suppose that A and $B \in [10, 40]$ (as in treatment GT2) and that our differential equation solution has $\alpha(p^*) = A_h < \beta(p^*) = 40$ for some p^* . Extending the differential equation solution above p^* is inappropriate since B is never greater than 40, but we still need to assign ask strategies for signals $A \in (A_h, 40]$. Since (by assumption) type A_h finds it optimal to trade with probability zero, so does any $A > A_h$. These types are indifferent among any ask prices weakly greater than p^* . Were we (for example) to set $a(A) = p^*$ for all of these types, we would tempt buyers who are prescribed bids just below p^* to poach this pool of sellers, disrupting the separating part of the equilibrium. By instead setting $a(A) = 100$ for $A \in (A_h, 40]$, we remove this temptation.

Alternatively, suppose that we have $\beta(p^*) = B_h < \alpha(p^*) = 40$. Thus buyers with signal B_h buy with probability 1. We need to extend $b(\cdot)$ to cover buyers with signals $B \in (B_h, 40]$. Standard arguments indicate that these buyers prefer bidding p^* rather than bidding any lower price (since they are more optimistic about V than B_h buyers). Furthermore, they strictly prefer bidding p^* over any higher bid (since a higher bid simply means paying more to win with probability 1). Thus, in this case we must set $b(B) = p^*$ for $B \in (B_h, 40]$.

B The Learning Phase

The learning phase of our experiments plays a supporting but critical role in our study. Our analysis of the trading phase presumes that subjects have had at least some success

in learning how the two clues and the asset value are related. Here we present results on subjects’ accuracy in predicting the asset value from one or two clues during the learning phase.

Recall that this phase consisted of two sets of rounds, 1-15 and 16-30. For the first 15 rounds, we presented a subject with one clue (A for 15 rounds or B for 15 rounds), then asked her to predict V , and finally showed her the actual value of V . Subjects exposed to A and to B would have seen different sample paths of (clue, value) pairs, but of course all of these pairs were generated from the same symmetric distribution. The second 15 rounds were the same, except that all subjects were shown both clues. A table of sample data was presented before round 1 (sample pairs (A, V) or (B, V)) and before round 16 (sample triples (A, B, V)) to give subjects some “preloaded” information before they had to make predictions. Subjects were rewarded for prediction accuracy according to a quadratic loss function.

We define a subject’s prediction error to be the absolute value of the difference between her prediction and the true value of the asset, $|V_t - pred_{it}|$, where $pred_{it}$ is subject i ’s prediction in round t . Remember that the distribution of the asset value and clues differed between treatments B1 and GT1 versus B2 and GT2, but between B1 and GT1, the learning phase was completely identical. (Similarly, the learning phase for B2 and GT2 was identical.) For this reason, we form two sets of pooled data: B1/GT1 and B2/GT2.⁴¹ Table B1 presents trends in the mean prediction error for these two groups, aggregated into five-round blocks. Prediction errors begin at around 15 in both groups and decline with experience. The decline is dramatic for the treatments with low noise, but only modest for the treatments with high noise. Furthermore, treatments B1 and GT1 show a sharp improvement at round 16, when both clues are introduced, that is completely absent in B2 and GT2. This is sensible: since V varies more with each clue in B1/GT1, the improvement in prediction from a second clue is greater.

Table B1: Mean Prediction Error in the Learning Phase

	Rounds					
	1-5	6-10	11-15	16-20	21-25	26-30
<i>Treatment B1/GT1</i>	14.48	10.72	11.69	3.89	3.04	2.93
<i>Treatment B2/GT2</i>	15.26	12.93	11.07	14.77	12.48	11.33

These cross-treatment comparisons need to be qualified because some of the difference

⁴¹The procedures were identical in all treatments. Therefore, the only difference between the two pools consisted in the distribution of the asset value and the clues. The learning phase in CE and CU was totally identical to that of B2/GT2. We do not discuss it here but only note that the results are not significantly different from those of B2/GT2.

between B1/GT1 and B2/GT2 reflects different levels of white noise in V . A perfect forecaster who sees both clues and has learned that $V = A + B + X$ will make a prediction equal to $E(V | A, B) = A + B$, resulting in a prediction error of $|X|$. Since X is larger in B2/GT2 ($U[-20, 20]$ versus $U[-3, 3]$), even a perfect forecaster will perform worse in these treatments relative to B1/GT1. A better cross-treatment comparison would be to ask how well subjects perform relative to the best feasible prediction with both clues, $E(V | A, B) = A + B$. To examine this question, we form the adjusted prediction error, $|A_t + B_t - pred_{it}| = |(V_t - X_t) - pred_{it}|$. Table B2 presents the average adjusted prediction errors for the two groups in rounds 1-15 and 16-30. We also compute these averages session by session. With five sessions per treatment, we have ten independent observations for each group (B1/GT1 and B2/GT2) of the average adjusted prediction error. The bottom part of the table gives the results from a Wilcoxon rank sum test of the equality of the mean adjusted prediction error across the two specifications. Equality is strongly rejected: predictions under B1/GT1 are slightly (but significantly) worse than under B2/GT2 when subjects see only one clue, but are substantially better than B2/GT2 when subjects see both clues. The average improvement in the adjusted error with a second clue is positive for both groups, but much larger for B1/GT1 (9.68 (s.e. 0.53) versus 2.30 (s.e. 0.43)).

Table B2: Average Adjusted Prediction Error

	Rounds	
	1-15	16-30
<i>Treatment B1/GT1</i>	12.20	2.52
<i>Treatment B2/GT2</i>	10.06	7.75
Wilcoxon Test		
$H_0 : B1/GT1 = B2/GT2$	-2.65	-3.63
$H_1 : B1/GT1 \neq B2/GT2$	0.008	0.000

The rank sum for the Wilcoxon test is reported. The corresponding p -value is in parentheses.

Table B3: Standard Deviation of the Adjusted Prediction Error

	Rounds	
	1-15	16-30
<i>Treatment B1/GT1</i>	9.84	5.03
<i>Treatment B2/GT2</i>	8.72	7.06
Wilcoxon Test		
$H_0 : B1/GT1 = B2/GT2$	-0.83	-2.12
$H_1 : B1/GT1 \neq B2/GT2$	0.410	0.034

The rank sum for the Wilcoxon test is reported. The corresponding p -value is in parentheses.

We are also interested in how much variation there is in prediction accuracy. For each

session s we compute the standard deviation σ_s of the adjusted prediction error in rounds 1-15 and 16-30. As above, this gives us ten observations per time period for B1/GT1 and ten for B2/GT2. Table B3 presents Wilcoxon estimates for the equality of the average standard deviations in B1/GT1 and B2/GT2. Variation in accuracy declines in both specifications, but again the change is larger for the less noisy pair of treatments. With one clue, adjusted errors are slightly (and not significantly) more dispersed in B1/GT1, but with both clues, the variance in errors is significantly smaller in B1/GT1 relative to B2/GT2.

Overall, the results give the impression that subjects are reasonably successful in learning to use the clues to predict the asset value. In particular, they converge fairly tightly on quite good predictions when they have both clues and there is relatively little noise confounding their feedback. When feedback is noisier (B2 and GT2), the results suggest a certain degree of model uncertainty that subjects never fully overcome: their predictions improve less over time and with a second clue, and tend to be more dispersed than is the case with less noise.

C Instructions

The instructions were divided into two parts: Part I, for the Learning Phase, and Part II, for the Trading Phase. Part I was identical for all treatments. Part II differed for the clues given to subjects and for the absence or presence of gains from trade. Here we report the instructions used for treatments B1 and B2 and the parts changed for the other treatments.

C.1 Instructions (Part I)

Welcome to our experiment!

Please, read these instructions carefully. Take all the time you want to go through them. Make sure you understand everything. If you have a question, please, do not hesitate to raise your hand. We will be happy to come to you and answer it privately. Please, do not ask your neighbours and do not try to look at their screens.

You are participating in an economics experiment in which you interact with seven other participants. Depending on your choices, the other participants' choices and some luck you can earn a considerable amount of money. You will receive the money immediately after the experiment. Notice that all participants have the same instructions.

The experiment

This experiment consists of **two parts**. In the first part you will learn something about how to assess the value of a good. In the second part you will be given the opportunity to trade this good with the other participants. Therefore, what you are going to learn in the first part will be useful in the second part of the experiment.

We will start by explaining the first part. When all of you have read these instructions, we will start running the first part. After that, we will give you further instructions for the second part, and then we will continue and run the second part.

The Experiment – Part I

In this part of the experiment we will ask you to predict the value of a good. We will give you some clues that will help you in your task.

You will have to make your predictions 30 times (in 30 rounds). In each of these rounds, the computer will choose a new value for the good and you will have to predict it. The computer chooses the value of the good in each round afresh. The value of the good in one round never depends on the value of the good in one of the previous rounds. However, the value of the good does depend on several factors. We will call two of these factors “clue A” and “clue B”; before making your prediction, you will have the chance to observe one or both of these clues. Because the value of the good depends on factors other than clue A and clue B, you should not expect to be able to make perfect predictions, even when you can observe both of these clues. Here are some more details about the value of the good and the information you will receive.

First of all, the value of the good will never be lower than 0 or higher than 100.

Second, in each of the first 15 rounds, you will receive one clue – either A or B. In contrast, during the last 15 rounds, you will receive both clue A and clue B. Your clue (or clues) will be a number (or two numbers) that appears on your screen at the beginning of each round. Both of these clues are related to the true value of the good, but we will not tell you how they are related – you must discover this through experience.

What you have to do

You have only one task: try to predict the correct value of the good.

Procedures for each of the first 15 rounds

In each of these rounds the computer will draw the value of the good, which will be a number between 0 and 100. You will not be told this number. However, on your screen you will see your clue, a number related to this value. After you see it, you can input your prediction. Note that, to input this value, you can use the mouse and also (to select the number more precisely) the up and down arrows at the bottom-right of the keyboard.

Procedures for each of the last 15 rounds

In the last 15 rounds you will have to do exactly the same. However, instead of seeing one clue, you will see two of them, the one that you always saw in the first 15 rounds and a new one. Again, after you see them, you can make your decision, that is, you can input your prediction.

You do not have to rush. Take all the time you want to make your decision.

Remark 1

In the first 15 rounds all of you will receive only one clue. Note, however, that four of you will receive (for all 15 rounds) clue A and the other four will receive clue B. Who receives one clue and who the other is decided randomly by the computer. In the last 15 rounds everyone receives both clues.

Remark 2

Remember that the value that the computer chooses in one round is completely independent of the value it chose in previous rounds or will choose in the next rounds. In every round the computer chooses a new value.

What do you earn for your predictions?

Your earnings in this first part depend on how well you predict the value of the good. For the first 10 rounds, given that you are learning, your performance will not affect your payment. In the rounds from the 11th through the 15th, your ability to predict the value will be important. In fact, the computer will randomly choose one of these rounds and we will pay you according to what you did in that round. Notice that the computer will select the same round for all of you. If in that round you predict the value exactly, you will earn £3.00. If your prediction differs from the true value by an amount x , you will earn $£3.00 - 0.003x^2$. This means that the further your prediction is from the true value, the less you

will earn. Moreover, if your mistake is small, you will be penalized only a small amount; if your mistake is big, you will be penalized more than proportionally.

Analogously, your performance in the rounds from the 16th through the 25th will not directly affect your payoff, as you are learning how to use two clues. In contrast, your payoff will depend on your predictions in the rounds from the 26th through the 30th. As above, the computer will randomly choose one of these rounds and we will pay you according to what you did in that round. We will pay you using the same rule explained above.

To make this rule clear, let us see some examples.

Example 1: Suppose the true value is 50.

Suppose you predict 70. In this case you made a mistake of 20. We will give you $\text{£}3.00 - 0.003 \cdot 20^2 = \text{£}3.00 - \text{£}1.20 = \text{£}1.80$. Similarly, if your prediction was 30, again you made a mistake of 20. And again we will give you $\text{£}3.00 - \text{£}1.20 = \text{£}1.80$.

Example 2: Suppose the true value is 65.

Suppose you predict 55. In this case you made a mistake of 10. We will give you $\text{£}3.00 - 0.003 \cdot 10^2 = \text{£}3.00 - \text{£}0.30 = \text{£}2.70$. Similarly, if your prediction was 75, again you made a mistake of 10 and, again, we will pay you $\text{£}3.00 - \text{£}0.30 = \text{£}2.70$.

Example 3: Suppose the true value is 24.

Suppose you predict 55. In this case you made a mistake of 31. We will give you $\text{£}3.00 - 0.003 \cdot 31^2 = \text{£}3.00 - \text{£}2.88 = \text{£}0.12$.

How do you learn?

Before the first round, we will show you a table with a sample of 10 values for the good and for the clue. The computer generated these values in the same way in which it will generate new values for you. Therefore, this may be helpful in predicting the values in the first rounds. Similarly, immediately before the 16th round, we will show you a table with 10 values for the good and the two clues. Again, this will help you before you start predicting the value of the good using two clues.

Moreover, after each round, of course, we will inform you of the true value of the good. Therefore, you will be able to compare your prediction with the truth. This may help you to improve your ability to make good predictions in later rounds. Again, you should not

expect yourself to be able to predict the value of the good perfectly, as this value depends on other factors besides the two clues.

Finally, at the end of this part of the experiment, your screen will display the two rounds that the computer chose for everyone to be paid. You will also see your own payoff for these two rounds.

C.2 Instructions (Part II)

Thank you for participating in the first part of the experiment. Now, we will start the second part. Please, read these instructions carefully and ask us if they are not clear. As for the first part, of course, all participants have the same instructions.

The Experiment – Part II

What you have to do

This part of the experiment consists of a series of 30 rounds. You will interact with the other people who participated in part I. In each round you will have to make a simple decision. You will have to choose a price at which you are willing to sell a good and a price at which you are willing to buy it. Your sell price can be greater than or equal to your buy price, but it cannot be lower. In each round you will be matched with other four participants in the experiment. If the price at which you want to buy is higher than the price at which any of these other participants with whom you are matched wants to sell, then trade will occur: you will obtain the good and will have to pay a price equal to the average of your “buy” offer and his “sell” offer. Similarly, if the price at which you want to sell is lower than the price at which any other of these four participants want to buy, trade will occur: you will obtain a price equal to the average of the other participant’s buy offer and your sell offer.

Notice that, given that you are matched with other four people, it may happen that you trade (buy or sell) with some of them and not with others. Moreover, given that these other participants can choose different prices, it is well possible that you sell to some people and buy from others.

It is also possible that neither you nor any of the other participants you are matched with makes a buy offer that is at least as high as the price at which the other one offers to sell. In that case, there will be no exchange of the good at all.

In conclusion, there are five possibilities: that you trade with no one, or with 1, 2, 3 or 4 participants.

What determines the value of the good

In the previous part of the experiment you have already seen how the value of the good is related to the information that you receive. The rules do not change: the value of the good is determined exactly as it was in the previous part.

What you earn in a round

First of all, we will give you 20 pence per round, regardless of whether trade occurs or not. In addition to this, you earn profits or losses on any trade that you make.

Therefore, in rounds without trade you will just earn the 20 pence that we give you. You will not earn more money and you will not lose it either.

When you buy a good from another participant, your profit or loss will be equal to the true value of the good minus the price you paid for it.

Similarly, when you sell a good to someone, your profit or loss will be equal to the price that you receive minus the true value of the good.

Clearly, given that in a round you can trade with 1, 2, 3 or 4 people, for each of these trades you can earn profits or suffer losses.

With whom you are matched

Remember that there are 8 participants in this experiment. In the first part of the experiment, in the first 15 rounds, four of you received information only on clue A and four only on clue B. Now, in each round, the computer will match you with the four people who received the other clue. For instance, if you are a participant who received information only on clue A in those first 15 rounds, you will be matched with participants who saw only clue B. If you received clue B, you will be matched with those who saw only clue A.

What information you receive

In this part of the experiment, we will let you know only one clue on the value of the good. If in the first 15 rounds of the previous part of the experiment you received only clue A, now you will see clue A only. Similarly, if you received information on clue B, again you will see clue B only.

Procedures for each round

Remember that the experiment is organized into different rounds and that within each round you will have to choose two prices, one for buying and one for selling. So, let us summarize what happens within each round.

1) At the beginning of each round the computer randomly chooses the value of the good as well as values for the two clues A and B. You will see one of the two clues on the screen.

2) Now you make your decision: select on your screen the price at which you would be willing to sell the good and the price at which you would be willing to buy it. There is only one restriction: the price at which you accept to buy cannot be greater than the price at which you want to sell. Notice that, obviously, the higher the price at which you want to sell and the lower the price at which you want to buy, the lower is the possibility that trade will actually occur. In particular, if you want to be sure of not trading, you can just click on a ‘no trade’ button on your screen. By clicking on that button, the prices will automatically go to the extremes, 0 and 100. Obviously, at those prices, trade never occurs, regardless of what prices the other participants select and you will simply keep your 20 pence.

4) The computer will compare your prices to those of the four participants with whom you are matched. Below we will call these participants the “Other Participants.” By comparing these prices the computer will determine whether there will be exchange of the good or not between you and any of them.

5) On your screen, you will see your prices, the Other Participants’ prices, the true value of the good and your payoff.

Once the first round is over, we will repeat the same procedure for the second round. At the beginning, the computer will choose again values for the good and the two clues. You will choose your buy and sell prices. You will be informed of your payoff.

Examples of payoff

Example 1

Suppose that in a round you choose a price for buying of 30 pence and a price for selling of 40 pence. Suppose also that:

1) The first of the Other Participants chooses a price for buying of 50 pence and a price for selling of 65 pence.

- 2) The second chooses a price for buying of 10p and a price for selling of 20p.
- 3) The third, a price for buying of 20p and a price for selling of 85p.
- 4) The fourth, a price for buying of 0p and a price for selling of 100p.

In this case, given that the first ‘Other Participant’ is willing to buy at a higher price than you are willing to sell, there will be trade between the two of you: the Other Participant will buy the asset from you. You will receive a price of 45p (which is exactly the midpoint between 40p and 50p).

After you have made your decisions you will be informed of the true value of the good: suppose it is 38p. In this case, your profit from this trade will be

$$\begin{aligned} (\text{price} - \text{true value}) &= \\ (0.45 - 0.38) &= 0.07 \end{aligned}$$

By selling the asset in this case you make a profit, as you sold for 45p something that was worth only 38p.

The second Other Participant is willing to sell at a lower price than you are willing to buy, therefore, there will be trade between the two of you: the ‘Other Participant’ will sell the asset to you. You will pay a price of 25p (which is exactly the midpoint between 20p and 30p).

Your profit for this trade will be

$$\begin{aligned} (\text{true value} - \text{price}) &= \\ (0.38 - 0.25) &= 0.13. \end{aligned}$$

By buying the asset in this case you make a profit, as you bought for 25p something that was worth 38p.

What happens with the third Other Participant? You do not sell to him, as his buy price is less than your sell price ($20 < 30$). Likewise, he does not sell to you, as your buy price is less than his sell price ($30 < 85$). Similarly, you do not buy and do not sell to the fourth Participant.

Finally, your Per Round Payoff will be:

$$0.20 + 0.07 + 0.13 = 0.40$$

Example 2

Suppose that in a round you choose a price for buying of 40 pence and a price for selling of 60 pence. Suppose also that the first of the Other Participants chooses a price for buying of 70 pence and a price for selling of 75 pence.

In this case, given that the first ‘Other Participant’ is willing to buy at a higher price than you are willing to sell, there will be trade between the two of you: the Other Participant will buy the asset from you. You will receive a price of 65p (which is exactly the midpoint between 60p and 70p).

After you have made your decisions you will be informed of the true value of the good: suppose it is 78p. In this case, your profit from this trade will be

$$\begin{aligned}(\text{price} - \text{true value}) &= \\(0.65 - 0.78) &= - 0.13\end{aligned}$$

By selling the asset in this case you make a loss. You lose 13p because you sold the good at a price below its true value.

What happens with the other participants depends on their prices, as in the example above.

Example 3

Suppose that there is a round in which you do not wish to trade with anyone under any circumstances. You choose a buy price of 0 and a sell price of 100p. Since it is impossible for the Other Participants’ sell price to be less than 0 or for their buy price to be greater than 100, you can be assured that you will not trade in this round with anyone.

Your final payoff

Your total payoff at the end of the experiment will be computed as follows.

Just for taking part in the experiment, you earn a show up fee of £5.00.

You have earned some money in the first part of the experiment.

Finally, you have earned the per-round payoffs in these 30 rounds.

We will just add all these amounts:

Total Payment = £5.00 + Money earned in Part I + Sum of Per Round Payoffs in Part II.

C.3 Changes to Instructions for Treatments GT1 and GT2

In treatments GT1 and GT2 the instructions were changed in regard to the explanation of earnings as shown below. Of course, also the examples of payoffs were amended accordingly.

What you earn in a round

First of all, we will give you 20 pence per round, regardless of whether trade occurs or not. In addition to this, you earn profits or losses on any trade that you make.

Therefore, in rounds without trade you will just earn the 20 pence that we give you. You will not earn more money and you will not lose it either.

When you buy a good from another participant, your profit or loss will be equal to the true value of the good minus the price you paid for it, plus an additional 5 pence which you earn simply because a trade occurred.

Similarly, when you sell a good to someone, your profit or loss will be equal to the price that you receive minus the true value of the good, once again plus an additional 5 pence which you earn simply because a trade occurred.

Clearly, given that in a round you can trade with 1, 2, 3 or 4 people, for each of these trades you can earn profits or suffer losses.

Examples of payoff

Example 1

Suppose that in a round you choose a price for buying of 30 pence and a price for selling of 40 pence. Suppose also that:

- 1) The first of the Other Participants chooses a price for buying of 50 pence and a price for selling of 65 pence.
- 2) The second chooses a price for buying of 10p and a price for selling of 20p.
- 3) The third, a price for buying of 20p and a price for selling of 85p.
- 4) The fourth, a price for buying of 0p and a price for selling of 100p.

In this case, given that the first ‘Other Participant’ is willing to buy at a higher price than you are willing to sell, there will be trade between the two of you: the Other Participant

will buy the asset from you. You will receive a price of 45p (which is exactly the midpoint between 40p and 50p).

After you have made your decisions you will be informed of the true value of the good: suppose it is 38p. In this case, your profit from this trade will be

$$\begin{aligned} & (\text{price} - \text{true value}) + 5 \text{ pence} = \\ & (0.45 - 0.38) + 0.05 = 0.07 + 0.05 = 0.12 \end{aligned}$$

By selling the asset in this case you make a profit, as you sold for 45p something that was worth only 38p. In addition, you get the extra 5 pence because trade occurred with this Participant.

The second Other Participant is willing to sell at a lower price than you are willing to buy, therefore, there will be trade between the two of you: the ‘Other Participant’ will sell the asset to you. You will pay a price of 25p (which is exactly the midpoint between 20p and 30p).

$$\begin{aligned} & \text{Your profit for this trade will be} \\ & (\text{true value} - \text{price}) + 5 \text{ pence} = \\ & (0.38 - 0.25) + 0.05 = 0.13 + 0.05 = 0.18. \end{aligned}$$

By buying the asset in this case you make a profit, as you bought for 25p something that was worth 38p, and, in addition, you get the extra 5 pence because trade occurred with this Participant.

What happens with the third Other Participant? You do not sell to him, as his buy price is less than your sell price ($20 < 30$). Likewise, he does not sell to you, as your buy price is less than his sell price ($30 < 85$). Similarly, you do not buy and do not sell to the fourth Participant.

$$\begin{aligned} & \text{Finally, your Per Round Payoff will be:} \\ & 0.20 + 0.12 + 0.18 = 0.50 \end{aligned}$$

Example 2

Suppose that in a round you choose a price for buying of 40 pence and a price for selling of 60 pence. Suppose also that the first of the Other Participants chooses a price for buying of 70 pence and a price for selling of 75 pence.

In this case, given that the first ‘Other Participant’ is willing to buy at a higher price than you are willing to sell, there will be trade between the two of you: the Other Participant will buy the asset from you. You will receive a price of 65p (which is exactly the midpoint between 60p and 70p).

After you have made your decisions you will be informed of the true value of the good: suppose it is 78p. In this case, your profit from this trade will be

$$\begin{aligned} &(\text{price} - \text{true value}) + 5 \text{ pence} = \\ &(0.65 - 0.78) + 0.05 = -0.13 + 0.05 = -0.08 \end{aligned}$$

By selling the asset in this case you make a loss. You get 5p because trade occurred, but also lose 13p because you sold the good at a price below its true value, for a net payoff of -8p.

What happens with the other participants depends on their prices, as in the example above.

Example 3

Suppose that there is a round in which you do not wish to trade with anyone under any circumstances. You choose a buy price of 0 and a sell price of 100p. Since it is impossible for the Other Participants’ sell price to be less than 0 or for their buy price to be greater than 100, you can be assured that you will not trade in this round with anyone.

C.4 Changes to Instructions for Treatments CE and CU

In Treatments CE and CU the instructions were changed in regard to the information that subjects received. In Treatment CE, the instructions were changed as follows:

What information you receive

In this part of the experiment, you will receive both clues A and B. This is true for every participant, of course.

In Treatment CU, the change was the following:

What information you receive

In this part of the experiment, you will receive different information in odd rounds and even rounds. If in the first 15 rounds of Part I of the experiment you received only clue

A, now in *odd rounds* (1,3,5,7...) you will see clue A only. In contrast, if you received information on clue B only, now you will see both clues A and B. The reverse occurs in even rounds. If in the first 15 rounds of Part I of the experiment you received only clue A, now in *even rounds* (2,4,6,8...) you will see both clues A and B. In contrast, if you received information on clue B only, now you will see clue B only. Clearly, since you are matched with participants who saw the other clue in the first 15 rounds of Part I, whenever you see only one clue, your matches see them both. Whenever you see both, your matches see only one.

C.5 Computer Screen Shots

Figure 2: Table with a sample of 10 asset values and the corresponding clue A

The screenshot shows a window titled "ELSE Centre Experiment : ROBOT1". Inside, there is a red header bar with "ELSE Centre Experiment" and a yellow bar with "Experimental Instructions". Below this, two lines of text state: "Your clues for this experiment are shown below." and "This table will be shown to you at regular intervals during this part of the experiment." A table with two columns, "Clue A" and "Real Value", is displayed. The table contains 10 rows of data. A blue arrow icon is located in the bottom right corner of the window.

Clue A	Real Value
14.34	40.07
18.77	44.93
25.18	61.05
27.23	37.18
12.01	22.51
18.07	35.83
17.10	55.89
29.30	64.62
39.08	47.47
14.13	49.18

Figure 3: Table with a sample of 10 asset values and the corresponding clue B

ELSE Centre Experiment
Experimental Instructions

Your clues for this experiment are shown below. This table will be shown to you at regular intervals during this part of the experiment.

Clue B	Real Value
14.34	40.07
18.77	44.93
25.18	61.05
27.23	37.18
12.01	22.51
18.07	35.83
17.10	55.89
29.30	64.62
39.08	47.47
14.13	49.18

Figure 4: Table with a sample of 10 asset values and the corresponding clues A and B

ELSE Centre Experiment
Phase 1 - Round 16 of 30

You have finished your task for the first 15 rounds. Thank you. In the next 15 rounds you will be able to use two clues for predicting the value of the good. This table will be shown to you at regular intervals during this part of the experiment.

Clue A	Clue B	Real Value
14.34	27.77	40.07
18.77	26.77	44.93
25.18	35.13	61.05
27.23	11.04	37.18
12.01	12.89	22.51
18.07	18.19	35.83
17.10	39.29	55.89
29.30	34.05	64.62
39.08	9.27	47.47
14.13	32.60	49.18