

Trading Volume and Public Information in an Experimental Asset Market with Short-Horizon Traders

August 26th, 2019

(preliminary and incomplete – please do not cite without permission)

Abstract

We examine the joint impact of investors' trading horizons and public information on trading volume. We hypothesize that public information leads to relative homogenization in the traders' beliefs about the fundamental value of an asset and this reduces their disagreement regarding the fundamental value. Since the long-horizon traders' trade is motivated by the fundamental value, such reduced disagreement leads to a reduction in trading volume. We further hypothesize that public information leads to polarization in the traders' beliefs about other traders' beliefs about the fundamental value and this polarization increases disagreement regarding other traders' beliefs about the fundamental value. Since short-horizon traders' trade is motivated by other traders' beliefs about the fundamental value, such increased disagreement leads to an increase in trading volume. We test these hypotheses in an experimental asset market and find strong evidence in their support.

Keywords. Higher-order Beliefs, Abnormal Trading Volume, Public Information, Short-horizon Traders, Single-period Security Market

1. Introduction

The distinction between short-horizon and long-horizon traders could potentially lie at the heart of many puzzling phenomena. Long-horizon traders hold their positions until the fundamentals of the firm become public, while short-horizon traders move frequently in and out of stocks in anticipation of short-term price fluctuations in the market. This difference in trading strategies implies that long-horizon traders form beliefs about the fundamentals of a firm, while short-horizon traders form beliefs of *other traders' beliefs* of the firm's fundamentals. The latter beliefs are known in the literature as higher order beliefs. Higher order beliefs and fundamental beliefs don't always coincide, so it is important to study how differences between them impact observable phenomena and social welfare. For example, it is widely believed that corporate myopia is due to a large presence of short-horizon traders in the market. While this phenomenon has been empirically documented, it remains puzzling because it is unclear why a sequence of short-horizon traders, each concerned only with the price at which they could sell to the next generation of short-horizon traders, would result in short-term prices that do not reflect long-term cash flows.

In this paper, we provide experimental evidence regarding the role of short-horizon and long-horizon traders in another puzzling phenomenon: the enormous trading volume that is observed in capital markets¹. It is highly unlikely that the observed trading volume is entirely due to consumption and savings needs. Neither can it be explained by the mere volatility of fundamentals, because if a shock to fundamentals caused all traders' beliefs to be revised in the

¹ Early evidence on abnormal trading volume around earnings announcement comes from Beaver (1968), Bamber (1987), Kandel and Pearson (1995). More recent results re-affirm this using data from high-frequency trading (Fleming and Remolona (1999), Green (2004), Evan and Lyons (2008), Chae (2005), Krinsky and Lee (1996).

same way then stock prices would change without much accompanying trade. So, the prevailing wisdom is that a large part of the observed trading volume is caused by disagreement among traders and is speculative in nature. But the disagreement that explains trading volume is unlikely to be disagreement about firms' fundamentals because such a story is inconsistent with the robust empirical finding that public releases of news (such as earning announcements) are followed by large increases in trading volume. Public information about fundamentals would homogenize beliefs about fundamentals and reduce disagreement rather than increase it.

So, we hypothesize that a significant part of the observed trading volume is due to the presence of short-horizon traders who are concerned not directly with predicting changes in the fundamentals but with changes in the beliefs of other traders. What causes disagreement among short-horizon traders about the beliefs of other traders? How does the release of public information result in increased disagreement about the beliefs of other traders while at the same time causing fundamental beliefs to converge?

To illustrate what is involved, consider a situation where traders i and j form beliefs about trader k 's (Susan's) beliefs, with Susan concerned only with beliefs about the fundamentals of a firm. Traders i and j are short-horizon traders who will need to ultimately liquidate their holdings by trading with Susan, who is a long-horizon trader concerned about the firm's fundamentals. If these short-horizon traders disagree about Susan's beliefs, they will also disagree about the price at which Susan would be willing to buy or sell to them, and therefore they would trade among themselves before liquidating their holdings to Susan. Let θ represent the firm's fundamentals and suppose that prior beliefs about θ are described by the improper uniform distribution over the entire real line. Consider noisy signals of the form $\tilde{x} = \theta + \tilde{\varepsilon}$, where $\tilde{\varepsilon}$ is distributed Normal with mean 0 and precision β , so that $E(\tilde{\theta}|x) = x$. Now, suppose that each of the three traders receives

idiosyncratic, but *i.i.d.*, draws of \tilde{x} : $\tilde{x}_i = \theta + \tilde{\varepsilon}_i$, $\tilde{x}_j = \theta + \tilde{\varepsilon}_j$ and $\tilde{x}_k = \theta + \tilde{\varepsilon}_k$. Then $E_l(\tilde{\theta}|x) = x_l$, $l = i, j, k$, so that traders have different assessments of the firm's fundamentals. But, more importantly for our discussion, trader i 's and trader j 's beliefs about Susan's beliefs of the fundamentals are described by $E_i[E_k(\tilde{\theta}|x_k)|x_i] = E(\tilde{x}_k|x_i) = E(\tilde{\theta}|x_i) = x_i$ and, similarly, $E_j[E_k(\tilde{\theta}|x_k)|x_j] = x_j$. So, there is disagreement in higher order beliefs measured by $|x_i - x_j|$. But, now let us introduce a public signal $\tilde{y} = \theta + \tilde{\eta}$, where $\tilde{\eta}$ is independent of $\tilde{\varepsilon}$ and is distributed Normal with 0 mean and precision α . How does the public signal affect the disagreement in higher order beliefs? The fundamental beliefs of Susan are now described by $E(\tilde{\theta}|x_k, y) = \frac{\alpha y + \beta x_k}{\alpha + \beta}$, so that the difference in higher order beliefs of traders i and j is:

$$\begin{aligned} \left(\frac{\beta}{\alpha + \beta}\right) |E(\tilde{x}_k|x_i) - E(\tilde{x}_k|x_j)| &= \left(\frac{\beta}{\alpha + \beta}\right) |E(\tilde{\theta}|x_i) - E(\tilde{\theta}|x_j)| = \left(\frac{\beta}{\alpha + \beta}\right)^2 |x_i - x_j| \\ &< |x_i - x_j| \end{aligned}$$

Thus, in a setting where all traders receive conditionally independent signals of the fundamentals, public information decreases the disagreement in both fundamental beliefs and higher order beliefs. Such information environments cannot explain the increase in trading volume following the release of public information, even in the presence of short-horizon traders.

But, consider the following information environment. Let \tilde{s} be the private signal received by Susan (the long-horizon trader), let \tilde{x}_i and \tilde{x}_j be the private signals of the two short-horizon traders and let \tilde{y} be the public signal observed by all traders. Let \tilde{x} be generic notation for the signals received by short-horizon traders. Suppose that, absent the public signal, $Cov(\tilde{x}, \tilde{s}) = 0$. This implies that the information received by short-horizon traders contains no information about Susan's valuation of the asset and therefore would cause no disagreement and no trade among short-horizon traders. But now suppose the public signal is such that $Cov(\tilde{x}, \tilde{s}|y) \neq 0$, i.e., the

public signal induces a relationship between the private signals of short-horizon and long-horizon traders. Then the release of the public signal will cause short-horizon traders to disagree about Susan's beliefs, while in the absence of the public signal there will be no such disagreement. A particularly interesting case arises when $Cov(\tilde{x}, \tilde{s}|y) < 0$. Then short-horizon trader i could be more optimistic than short-horizon trader j about the firm's fundamentals, but more pessimistic about Susan's beliefs of the fundamentals. Thus, given the public signal, trader i would purchase from trader j if these traders are long-horizon traders, but trader i would sell to trader j if they are short-horizon traders. Also, in the absence of public information there would be no trade if traders are short-horizon traders, even though they disagree in their beliefs about the firm's fundamentals. But, if traders are long-horizon traders, there will be trading volume under exactly the same informational conditions.

Since beliefs are not directly observable, but trading volumes are a direct consequence of differences in beliefs, experimental tests built around trading volumes could yield insights into how economic agents could come to disagree in terms of fundamental beliefs and higher order beliefs, and how fundamental beliefs and higher order beliefs could diverge sharply from each other. Such experiments would also provide insights into the role of short-horizon traders in driving trading volumes. Our goal, in this paper, is to provide such insights.

2. Theory and Hypothesis

Our experimental design is based on the theory of trading volume developed by Kondor (2012). We outline the relevant part of Kondor's theory below and identify testable hypothesis. Kondor assumes that there are two groups of traders, A-traders and B-traders. A-traders are long-horizon traders and arrive late in the market after B-traders have finished trading among

themselves. A-traders hold their shares until the firm pays out an uncertain liquidating dividend, and are therefore concerned with assessing the amount of the liquidating dividend. B-traders are short-horizon traders who can trade among themselves in a market in which A-traders are absent, but who must ultimately liquidate their holdings to A-traders. They do not have the option of holding their shares until the firm liquidates, and therefore the value that a B-trader assigns to the firm depends on his/her beliefs of A-traders' average valuation, and therefore upon his/her beliefs of A-traders average beliefs of the firm's liquidating dividend.

The liquidating value of the firm (its fundamental value) is the sum of two independent components: $\tilde{\theta} = \tilde{\theta}_A + \tilde{\theta}_B$. Assume that each component is distributed Normal with means of μ_A and μ_B , variances σ_A^2 and σ_B^2 , respectively, and that $Cov(\tilde{\theta}_A, \tilde{\theta}_B) = 0$. Let Susan be a representative A-trader and let i and j be representative B-traders. Susan receives private information on θ_A , while i and j traders receive private signals on θ_B . Let $\tilde{s} = \theta_A + \tilde{\gamma}$ be Susan's private signal, where $\tilde{\gamma}$ is independent of $\tilde{\theta}_A$ and $\tilde{\theta}_B$ and is distributed Normal with zero mean and variance σ_γ^2 . Thus, conditional on her private signal, Susan's belief of the firm's fundamental value is:

$$E(\theta|s) = E(\theta_A|s) + \mu_B = \tau s + (1 - \tau)\mu_A + \mu_B, \text{ where } \tau = \frac{\sigma_A^2}{\sigma_A^2 + \sigma_\gamma^2}.$$

Let $x_i = \theta_B + \varepsilon_i$, $x_j = \theta_B + \varepsilon_j$ be the private signals of the representative short-horizon traders i and j , where ε_i and ε_j are *i.i.d.* draws from a Normal distribution with zero mean and variance σ_ε^2 and are independent of $\tilde{\gamma}$. Then,

$$E(\theta|x) = E(\theta_B|x) + \mu_A = \alpha x + (1 - \alpha)\mu_B + \mu_A, \text{ where } \alpha = \frac{\sigma_B^2}{\sigma_B^2 + \sigma_\varepsilon^2}$$

Thus, if traders i and j were to behave as long-horizon traders, the presence of private signals would cause disagreement about the firm's fundamentals, with the magnitude of disagreement described by:

$$|E(\theta|x_i) - E(\theta|x_j)| = \frac{\sigma_B^2}{\sigma_B^2 + \sigma_\varepsilon^2} |x_i - x_j| \quad (1)$$

But, suppose traders i and j are short-horizon traders who cannot hold their positions until the firm pays out its liquidating dividend and have to settle up their trades with Susan. Thus, because θ_A is independent of θ_B , $cov(\tilde{s}, \tilde{x}) = 0$, so the private signals observed by traders i and j are uninformative about Susan's information signal and therefore there is no disagreement about Susan's beliefs. More precisely,

$$E[E(\theta|s)|x_i] - E[E(\theta|s)|x_j] = \tau[E(s|x_i) - E(s|x_j)] = \tau[E(\theta_A|x_i) - E(\theta_A|x_j)] = 0 \quad (2)$$

Thus, there is no reason for these short-horizon traders to speculate about Susan's valuation and no reason to trade among themselves before liquidating their holdings to Susan. The above analysis implies the following hypothesis for the information environment described here:

Hypothesis 1: Private information about fundamentals causes disagreement and trade if traders have long-term horizons, but no disagreement and no trade if traders have short-term horizons.

Now, suppose that in addition to the private signals described above, all traders receive a public signal. The public signal provides noisy information about the aggregate quantity θ , i.e. information about the sum of θ_A and θ_B . Conditional on θ , θ_A and θ_B are no longer independent: a higher θ_A implies a lower θ_B and a lower θ_A implies a higher θ_B , i.e., $Cov(\tilde{\theta}_A, \tilde{\theta}_B|\theta) < 0$, while the unconditional covariance $Cov(\tilde{\theta}_A, \tilde{\theta}_B) = 0$. This negative covariance carries over to noisy signals. Let $\tilde{y} = \theta + \tilde{\eta}$ be the public signal, where $\tilde{\eta}$ is distributed Normal with mean 0 and

variance σ_η^2 and is independent of all the other noise terms. Then, while the unconditional covariance $Cov(\tilde{s}, \tilde{x}) = 0$, the conditional variance $Cov(\tilde{s}, \tilde{x}|y) < 0$, as shown below:

$$Cov(s, x|y) = Cov(s, x) - \frac{Cov(s, y)Cov(x, y)}{Var(y)} = -\frac{Cov(\theta_A, \theta)Cov(\theta_B, \theta)}{\sigma_A^2 + \sigma_B^2 + \sigma_\eta^2} = -\frac{\sigma_A^2 \sigma_B^2}{\sigma_A^2 + \sigma_B^2 + \sigma_\eta^2} < 0$$

Public information of this kind connects the private signals of short-horizon traders to the private signals of long-horizon traders. So, in the presence of the public signal, traders i and j can make inferences about Susan's value for the asset, while in the absence of public information no such inferences are possible.

Proposition 1: The presence of public information causes short-horizon traders to disagree about the beliefs of long-horizon traders.

Proof: Susan's value for the asset conditional on her private information and the public signal is:

$$\begin{aligned} E(\theta|s, y) &= E(\theta) + b[s - E(s)] + c[y - E(y)] \\ &= cy + (1 - c)(\mu_A + \mu_B) + b(s - \mu_A), \end{aligned}$$

where

$$b = \frac{Cov(s, \theta)Var(y) - Cov(y, \theta)Cov(y, s)}{Var(y)Var(s) - Cov^2(y, s)} = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_B^2}$$

and,

$$c = \frac{Cov(y, \theta)Var(s) - Cov(s, \theta)Cov(y, s)}{Var(y)Var(s) - Cov^2(y, s)} = \frac{\sigma_B^2}{\sigma_\eta^2 + \sigma_B^2}$$

Using the fact that $b + c = 1$, $E(\theta|s, y)$ can be expressed as:

$$E(\theta|s, y) = b(s + \mu_B) + (1 - b)y \quad (3)$$

Then the disagreement among short-horizon traders about Susan's valuation is:

$$\begin{aligned}
|E[E(\theta|s, y)|x_i, y] - E[E(\theta|s, y)|x_j, y]| &= b|E(s|x_i, y) - E(s|x_j, y)| \\
&= b|E(\theta_A|x_i, y) - E(\theta_A|x_j, y)|
\end{aligned}$$

But, $E(\theta_A|x, y) = E(\theta_A|y) + \frac{Cov(\theta_A, x|y)}{var(x|y)} [x - E(x|y)]$

Cancelling common terms, the disagreement about Susan's valuation is:

$$|E[E(\theta|s, y)|x_i, y] - E[E(\theta|s, y)|x_j, y]| = b \left(\frac{Cov(\theta_A, x|y)}{var(x|y)} \right) |x_i - x_j| \quad (4)$$

It follows that if the private signals of traders i and j do not coincide, the presence of public information will cause disagreement about Susan's valuation, causing them to trade with each other, while in the absence of public information traders i and j will not trade because they agree about Susan's valuation. Additionally, since $Cov(\theta_A, x|y) = Cov(\theta_A, \theta_B|y) < 0$, if $x_i > x_j$ trader i would be more optimistic than trader j about the fundamentals of the firm but more pessimistic about Susan's assessment of those fundamentals. Thus, trader i would purchase from trader j if both traders were long-horizon traders, but trader i would sell to trader j if both traders are short-horizon traders.

Now, suppose that traders i and j trade as if they are long-horizon traders concerned only with the firm's fundamentals. We show below that in this case, public information causes disagreements to decline, rather than to increase.

Proposition 2: $|E(\theta|x_i, y) - E(\theta|x_j, y)| < |E(\theta|x_i) - E(\theta|x_j)|$ (5)

Proof: Proof in Appendix A.

Hypothesis 2 below follows directly from Propositions 1 and 2.

Hypothesis 2: The presence of public information in settings where all traders also receive private signals will cause trading volume to decrease if traders have long-term horizons, but will cause trading volume to increase if traders have short-term horizons.

In the Kondor (2012) paper, traders extract information from the equilibrium price in the capital market in addition to the information provided by the signals described above. Since the equilibrium price aggregates the information of all traders it could possibly become a sufficient statistic for all the idiosyncratic information that traders receive, in which case beliefs would become homogenous with or without public information. However, Kondor shows that with the introduction of independent supply noise, Hypothesis 1 and 2 are essentially preserved even when traders condition on equilibrium prices. In our experimental design, we do not explicitly introduce supply side noise, but we believe that, as a practical matter, there will always be some naturally occurring noise in any experimental implementation due to unpredictable differences in experimentation and learning by the participants in the study. Given this claim, we interpret the data from our experiment as if participants do not condition on equilibrium prices. Since we are not testing predictions about the magnitudes of trading volumes, but only about the ordering of trading volumes, such interpretation is valid in view of the Kondor (2012) result that when supply noise is present, the qualitative nature of the results are the same regardless of whether traders condition or don't condition on equilibrium prices.

3. Experimental Parameters

For experimental purposes, we make three simplifications to the theory described above. First, we provide perfect information, rather than noisy information, about θ_A to Susan, i.e. in the

experiment, $s \equiv \theta_A$. Second, we provide the public signal $\tilde{y} = \theta + \tilde{\eta}$ to all B-traders, but hide this information from Susan. Thus, Susan's valuation of the firm is simply: $E(\theta|\theta_A) = \theta_A + \mu_B$. This is equivalent to forcing $b = 1$ in the more general expression $E(\theta|s, y) = b(s + \mu_B) + (1 - b)y$ that was derived in (3). Substituting $b = 1$ in equation (4), the disagreement, caused by public information, among A-traders about Susan's valuation of the firm becomes $\left(\frac{\text{Cov}(\theta_A, x|y)}{\text{var}(x|y)}\right) |x_i - x_j|$. Thus, by hiding the public signal from Susan we have magnified the disagreement among short-horizon traders and have increased saliency in the experiment. The third simplification we make is to replace Susan by a computer program that simply pays B-traders the value of $\tilde{\theta}_A$ that is realized in that trial of the experiment plus the constant μ_B in exchange for any shares they may wish to sell to Susan, after trading among themselves. All three simplifications either reduce the complexity of the inferences that participants in the experiment need to make or increase saliency in payoffs, without damaging the hypotheses that we wish to test.

We use the following parameter values in our experiment:

$$\mu_A = \mu_B = 50$$

$$\sigma_A^2 = \sigma_B^2 = 100$$

$$\sigma_\varepsilon^2 = \sigma_\eta^2 = 25$$

We preserve Normalcy of all random variables, as specified in the theory.

Given these parameter values, we calculate below the disagreement parameters in each of our settings:

$$\alpha = \frac{\sigma_B^2}{\sigma_B^2 + \sigma_\varepsilon^2} = 0.8$$

$$\text{Cov}(\theta, x|y) = \text{Cov}(\theta, x) - \frac{\text{Cov}(\theta, y)\text{Cov}(x, y)}{\text{var}(y)} = \sigma_B^2 - \frac{(\sigma_A^2 + \sigma_B^2)\sigma_B^2}{\sigma_A^2 + \sigma_B^2 + \sigma_\eta^2} = 11.1112$$

$$\text{Var}(x|y) = \text{Var}(x) - \frac{\text{Cov}^2(x,y)}{\text{Var}(y)} = \sigma_B^2 \left[1 - \frac{\sigma_B^2}{\sigma_A^2 + \sigma_B^2 + \sigma_\eta^2} \right] + \sigma_\varepsilon^2 = 80.5555$$

$$\text{Cov}(\theta_A, x|y) = \text{Cov}(x, \theta_A) - \frac{\text{Cov}(\theta_A, y)\text{Cov}(x, y)}{\text{Var}(y)} = -\frac{\sigma_B^2 \sigma_A^2}{\sigma_A^2 + \sigma_B^2 + \sigma_\eta^2} = -44.4444$$

Therefore:

The disagreement among long-horizon traders i and j with private information only is $\alpha|x_i - x_j| = 0.8|x_i - x_j|$.

The disagreement among short-horizon traders i and j with private information only is 0.

The disagreement among long-horizon traders i and j with private and public information is $\left(\frac{\text{Cov}(\theta, x|y)}{\text{Var}(x|y)} \right) |x_i - x_j| = (0.1379)|x_i - x_j|$.

The disagreement among short-horizon traders i and j with private and public information is $\left(\frac{\text{Cov}(\theta_A, x|y)}{\text{Var}(x|y)} \right) |x_i - x_j| = -(0.5517)|x_i - x_j|$.

The above parameters imply that $x \in 50 \pm 2\sqrt{125} = 50 \pm 22.36$, with 95% confidence. Therefore, the difference in i and j traders' private signals ≤ 44.72 . This, in turn, implies that the first-order disagreement without public information $\leq 0.8 \times 44.72 = 35.78$ while the first-order disagreement with public information $\leq 0.1379 \times 44.72 = 6.17$. Similarly, the second-order disagreement without public information is 0 while the second-order disagreement with public information $\leq 0.5517 \times 44.72 = 24.67$.

4. Experiment Design

In the experiment, a group of ten participants traded shares of one stock in thirteen independent trading periods. Orders were restricted to a single share. The stock paid a liquidating dividend at the end of the trading period. The participants did not know the amount of the dividend

until after the trading period, but they received signals (clues) about the dividend before trading in the market commenced.

The timeline for the experiment is given in Figure 1.




Figure 1 – Timeline of the Experiment

There were *two* treatment variables. The first, public information, was either *Available* or *Not Available*. The second, trading horizon of the human traders, was either *Long* or *Short*. If trading horizon was *Long*, then the shares held by participants at the end of the trading period were liquidated at the realized dividend value. If trading horizon was *Short*, then the shares held by the human participants were liquidated *at the price the computerized trader (Susan) was willing to pay for the shares*. In essence, short-horizon traders were required to sell their shares to the computerized trader at the end of each trading period. Depending upon the treatment, public information was or was not be available to the long and short-horizon human traders. Note that the computerized trader did not receive a public signal in the Public Information treatments. The 2 X 2 design may be summarized as follows.

		Public Information	
		<i>Available</i>	<i>Not Available</i>
Trading Horizon	<i>Long</i>		
	<i>Short</i>		

Figure 2 – Experimental Treatments

We ran three sessions of each of the four experimental treatments, with each session consisting of 13 trading periods. We generated *three* sets of thirteen securities with dividend, two private signals and a public signal. We used the three sets respectively for the three sessions of each of the four treatments. This ensures that all sessions (even across treatments) were informationally identical. Note that for each trading period, we generated two private signals – one-half of the participants received the first private signal and the other half received the second private signal.

The short-horizon traders were asked to forecast the price the computerized trader would be willing to pay while the long-horizon traders were asked to forecast the dividend value. Soliciting the pre-trading forecasts allows us to directly measure the extent of disagreement among the traders.

The price the computerized trader would pay would be equal to the expected value of the dividend given their information set. Also, the dividend components drawn for a trading period were independent of the dividend components drawn for every other trading period.

At the beginning of each trading period, participants were endowed with a certain number of shares and a certain amount of experimental dollars. Each trading period lasted for four minutes. Trading was organized as a continuous double auction. A participant's shares were liquidated at the end of each trading period at the pre-defined *liquidating value*. Recall that the liquidating value corresponded to the realized dividend in the long-horizon treatments and to the amount the computerized trader was willing to pay in the short-horizon treatments. A participant's *trading profit* in a period corresponded to the difference between their ending cash balance (post-liquidation of shares) and their initial portfolio value (with shares valued at 100).

Cash and stock holdings did not carry over from one period to the next. At the beginning of each trading period, participants received fresh endowments of shares and cash. At any point, participants could not sell more shares than they owned and could not bid for more shares than their cash holding allowed.

The first of the thirteen trading periods was treated as a practice trading period, and participants were not compensated for it. The trading profits and the liquidating value of shareholdings in the last twelve trading periods were converted into cash at a pre-announced exchange rate². The participants were also paid for their trading forecasts per the following formula: $\max \{0, 2500 - 0.25 \times [\text{forecast} - \text{liquidating value}]^2\}$.

After the participants signed in, hard-copy instructions were distributed. The participants read the instructions and then answered a quiz. The answers were reviewed and the participants were paid 25 cents for every correct answer. The participants were recruited from the subject pool at the Economic Science Institute Laboratory at Chapman University. All participants received a show-up payment of seven dollars.

5. Results

We ran three sessions of each of the four experimental treatments. As each session was comprised of twelve trading periods, there are thirty-six observations for each of the four treatments. The results reported below are based on these thirty-six observations. We recognize that the twelve observations from a *single* session are not independent as they involve repeated measures arising from the same participants trading with each other in twelve consecutive trading periods. We control for this in our regression analyses through the use of fixed effects for periods (see below). Trading volume can be measured either as the quantity traded or as the quantity traded

² The study was approved by the Institutional Review Boards at Chapman University and University of California – Irvine.

times the price at which it was traded (dollar trading volume). Tables 1 and 2 report the results using both metrics though the following discussion is restricted to quantity traded for ease of exposition. Recall that each trading period lasted for four minutes. All analyses reported in Tables 1 and 2 were conducted twice – once on full set of data generated from all four minutes and once on the subset of data restricted to the last two minutes of each trading period. Again, for ease of exposition the following discussion focuses on the results generated from using all four minutes of trading period data.

5.1 Trading Volume in Private Information Sessions

Our first hypothesis states that, in the absence of public information, private information causes long-horizon traders to trade, while it induces a no trade scenario among short-horizon traders. This no trade prediction is stark, and we do not expect it to hold true in the laboratory. Thus, armed solely with private information, we expect long-horizon traders to trade more than short-horizon traders. Consistent with this, the aggregate trading volume among long-horizon traders with private information is 3304 units while it drops sharply to 1197 units among short-horizon traders with private information (Table 1, Panel A). To test the significance of this difference in trading volumes, we aggregated the trading volume for the first four trading periods and for the last four trading periods of each sessions. Treating these two sums as independent observations yields six observations per treatment. A Wilcoxon signed rank test on these six observations shows that long-horizon traders react differently to private information than short-horizon traders. Indeed, long-horizon traders generated significantly greater trading volume than short-horizon traders (Table 2, Panel C).

5.2 Trading Volume in Public Information Sessions

Our second hypothesis also suggests a difference in trading behavior between long-horizon and short-horizon traders when they receive a public signal in addition to their respective private signals. While long-horizon traders are expected to engage in fewer trades, short-horizon traders are expected to trade more frequently. Consistent with this, the aggregate trading volume among long-horizon traders drops to 1368 units (Public Information Available treatment) from 3304 units (Public Information Not Available treatment), while the aggregate trading volume among short-horizon traders increases to 1578 units (Public Information Available treatment) from 1197 units (Public Information Not Available treatment) (Table 1, Panel A). Regressing³ the long-horizon trading volume on the availability of public information shows that the coefficient on the Public Information dummy variable (set to 1 if public information is available and zero otherwise) is negative and significant (Table 1, Panel C). However, this coefficient is positive and significant when considering the trading volume of short-horizon traders (Table 1, Panel D). To further assess the difference in trading behavior between long- and short-term traders, we also ran non-parametric Wilcoxon signed rank tests. Similar to the approach described in Section 5.1, we utilized six observations per treatment and found that in the presence of public information, long-horizon traders engage in significantly fewer trades than when public information is not available (Table 2, Panel C). Analogously, short-horizon traders engage in significantly more trades when public information is available than when it is unavailable (Table 2, Panel C).

Traders' reactions to the availability of public information seems to be driven by their (differing) beliefs of the asset's liquidating value. While public information leads to a reduction in long-horizon traders' disagreement regarding the liquidating value, it promotes a greater disagreement among short-horizon traders. This leads to less trading volume among long-horizon

³ We ran a generalized least squares model with a random effect for session and a fixed effect for each of the thirty-six trading periods. Standard errors were clustered at the session level.

traders and greater trading volume among short-horizon traders. We measure this disagreement in two separate ways. First, our experiment allows us to measure this disagreement directly by looking at the participants' forecasts. Recall that there were ten participants in each trading period, and two private signals were drawn for each trading period. Five of the participants received the high private signal while the other five received the low private signal. The forecasts made by the five traders with the low signal were averaged together as were the forecasts made by the five traders with the high signal. We measured the disagreement among traders as the absolute difference between these two averages. When public information is not available, the average disagreement among long-horizon traders is 9.66, and it drops to 4.8 when public information is available (Table 3, Panel A). For short-horizon traders, the average disagreement is 9.39 when public information is not available and increases to 9.42 when it is available (Table 3, Panel A). Note that the availability of public information has a stronger effect on the behavior of long-horizon traders than short-horizon traders.

Our second measure of disagreement is the time-weighted bid-ask spread. As there may be a within-period learning effect, the time weighted bid-ask spread may be a more sensitive measure. When public information is not available, the time weighted bid-ask spread with long-horizon traders is 28.48. When public information is available, this measure drops to 9.53 (Table 3, Panel B). The time weighted bid-ask spread moves in the opposite direction with short-horizon traders. When public information is not available, this spread is 15.59 but it increases to 25.01 when public information is available (Table 3, Panel B).

To assess the significance of these changes in disagreement amongst the long (short)-horizon traders, both the average disagreement as well as the time weighted bid-ask spread measures were regressed against a dummy variable, Public Information, which took value 1 if

public information was available and zero otherwise.⁴ The coefficient on public information in both regressions was negative and significant for long-horizon traders (Table 3, Panel C). This coefficient was positive in both regressions for short-horizon traders, though it was only significant for the more sensitive time-weighted bid-ask spread measure (Table 3, Panel D).

Finally, we regressed⁵ the participant-level trades on the availability of public information. It should be hardest to identify the treatment effect at this level because it obviates any kind of aggregation thereby allowing for individual participant-level idiosyncrasies to have an effect. The signs of the coefficients are as expected (negative for long-horizon traders and positive for short-horizon traders), though it is only significant for long-horizon traders (Table 5, Panels A and B). In summary, we find evidence suggesting long-horizon traders react differently to the availability of public information than short-horizon traders. While the availability of public information leads to less trading by long-horizon traders, it tends to increase the trading activity of short-horizon traders.

5.3 Additional Analyses

In this section we report additional analyses regarding the magnitude of price changes and the efficiency of prices. In the presence of public information we find that both long-horizon and short-horizon traders behave similarly in that the volume-weighted average price is higher for both groups (Figure 3, Panels A and B). This suggests that the magnitude of price change is unambiguously higher in public information. The mean absolute deviation from the liquidating value presents a mixed picture (Figure 4, Panels A and B) in that the availability of public

⁴ We ran a generalized least squares model with a random effect for session and a fixed effect for each of the thirty-six trading periods. Standard errors were clustered at the session level.

⁵ We ran a generalized least squares model with a random effect for subject, a fixed effect for each of the thirty-six trading periods and a fixed effect for each of the six sessions. Standard errors were clustered at the subject level.

information does not appear to significantly impact the efficiency of prices when considering markets solely populated by either long-horizon or short-horizon traders. Note, however, that our primary focus is not on the first moment of prices or forecasts, but rather the second moment. Accordingly, the variance of both prices and forecasts decreases when public information is available for the long-horizon traders but increases when public information is available for the short-horizon traders (Table 4).

6. Conclusion

We provide experimental evidence on the differential impact public information has on the trading volume generated by short-horizon and long-horizon traders. When these traders have only private information, we find that the trading volume generated by long-horizon traders is an order of magnitude higher than the trading volume generated by short-horizon traders. This is so because private information alone causes some disagreement among long-horizon traders about the fundamental value of the asset but causes no disagreement among short-horizon traders about the beliefs of long-horizon traders.

When they have both private and public information, we find that the trading volume generated by long-horizon traders decreases while the trading volume generated by short-horizon traders increases. This is so because public information decreases the disagreement among the long-horizon traders about the fundamental value. However, given our information structure, it increases disagreement among short-horizon traders about the long-horizon traders' beliefs about the fundamental value.

We contribute to a vast stream of literature that attempts to explain the puzzling phenomenon of vast trading volume around earnings announcement and more generally, around release of public information. Two broad groups of theoretical explanations attempt at resolving

this puzzle. The first one (Kim and Verrecchia (1994, 1997)) models public information as a combination of public and private signals. A public announcement stimulates superior information processing by sophisticated investors. This induces or exacerbates information asymmetry and thereby, leads to increased trading volume. The second one (Varian (1989), Harris and Raviv (1993), Kandel and Pearson (1995)) assumes that agents have heterogeneous priors so that even when they process the same public signal they end up with different valuations about the fundamental value of the asset. These differing valuations, in turn, lead to an increased trading volume.

More recently, Kondor (2012) exploits higher order expectations in financial markets to explain increased trading volume around public announcements. This paper combines the common prior assumption from the first group of papers with the characterization of public information as a public signal from the second group of papers. It shows that as long as there are at least some short-horizon traders who focus on the future market price instead of fundamental value of the asset, public announcement increases trading volume because it increases disagreement among short-horizon traders by polarizing their beliefs about the market price. We provide experimental evidence in support of this.

Our paper is related to Gallo (2017). She uses archival data to examine the role of higher-order disagreement in shaping investor beliefs while we use controlled laboratory experiment to examine the role of such disagreement in driving trading volume. More specifically, her primary variable of interest is disagreement among traders while our primary variable of interest is the trading volume such disagreement generates. Trading volume may be generated by disagreement among traders but archival data makes it difficult to disentangle trading volume from disagreement among traders (Banerjee (2011), Fischer, Kim and Zhou (2019)). Our controlled laboratory setting

allows us to disentangle these two and additionally, allows us to directly measure both the disagreement regarding the fundamental value and the disagreement regarding other traders' beliefs about the fundamental value.

Our experiment's design features a fundamental value that is additive in two components. There are two groups of traders and each receives a noisy private signal about only one of the two components. A natural real-world counterpart of this set-up is a firm that operates in multiple geographical locations. For example, a multinational firm may operate in California and in Hong Kong. The investors in California would arguably know more about the firm's operations in California while the investors in Hong Kong would know more about its operations in Hong Kong. Just as in our experiment, the firm's publicly disclosed earnings number would be a noisy signal of the aggregate operation. It can be argued that while our information structure enables a parsimonious operationalization of differential trading horizons in a continuous double auction, it is stylized. Future work can examine the relation between public information, trading horizon and trading volume in a setting with a more general information structure.

References

- Bamber, L. S. (1987). Unexpected earnings, firm size, and trading volume around quarterly earnings announcements. *The Accounting Review*, 510-532.
- Banerjee, S. (2011). Learning from prices and the dispersion in beliefs. *The Review of Financial Studies*, 24(9), pp.3025-3068.
- Beaver, W. H. (1968). The information content of annual earnings announcements. *Journal of Accounting Research*, 67-92.
- Chae, J. (2005). Trading volume, information asymmetry, and timing information. *Journal of Finance*, 60(1), 413-442.
- Evans, M. D., & Lyons, R. K. (2008). How is macro news transmitted to exchange rates? *Journal of Financial Economics*, 88(1), 26-50.

Fleming, M. J., & Remolona, E. M. (1999). Price formation and liquidity in the US Treasury market: The response to public information. *Journal of Finance*, 54(5), 1901-1915.

Fischer, P., Kim, C. and Zhou, F. (2019). Disagreement about fundamentals: measurements and consequences. *Working paper*, University of Pennsylvania.

Gallo, L.A. (2017). The more we know about fundamentals, the less we agree on price: evidence from earnings announcements. *Working paper*, University of Michigan.

Green, T. C. (2004). Economic news and the impact of trading on bond prices. *Journal of Finance*, 59(3), 1201-1233.

Harris, M., & Raviv, A. (1993). Differences of opinion make a horse race. *Review of Financial Studies*, 6(3), 473-506.

Kandel, E., & Pearson, N. D. (1995). Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy*, 103(4), 831-872.

Kim, O., & Verrecchia, R. E. (1994). Market liquidity and volume around earnings announcements. *Journal of Accounting and Economics*, 17(1-2), 41-67.

Kim, O., & Verrecchia, R. E. (1997). Pre-announcement and event-period private information. *Journal of Accounting and Economics*, 24(3), 395-419.

Kondor, P. (2012). The more we know about the fundamental, the less we agree on the price. *Review of Economic Studies*, 79(3), 1175-1207.

Krinsky, I., & Lee, J. (1996). Earnings announcements and the components of the bid-ask spread. *Journal of Finance*, 51(4), 1523-1535.

Kroll, Y., Levy, H. and Rapoport, A. (1988). Experimental tests of the separation theorem and the capital asset pricing model. *The American Economic Review*, pp.500-519.

Varian, H. R. (1989). Differences of opinion in financial markets. In *Financial Risk: Theory, Evidence and Implications* (pp. 3-37). Springer, Dordrecht.

Table 1 – Aggregate Trading Volume for Each Treatment

The columns labeled ‘long-horizon traders’ give the aggregate trading volume generated by long-horizon traders while the columns labeled ‘short-horizon traders’ give the aggregate trading volume generated by short-horizon traders. The aggregation is across 36 trading rounds over 3 sessions for each treatment.

Panel A – Aggregate Trading Volume Measured as Quantity Traded

	Long-Horizon Traders (all 4 minutes)	Short-Horizon Traders (all 4 minutes)	Long-Horizon Traders (last 2 minutes)	Short-Horizon Traders (last 2 minutes)
Private Info Only	3304	1197	1500	500
Private & Public Info	1368	1578	490	745

Panel B – Aggregate Trading Volume Measured as Price Times Quantity Traded

	Long-Horizon Traders (all 4 minutes)	Short-Horizon Traders (all 4 minutes)	Long-Horizon Traders (last 2 minutes)	Short-Horizon Traders (last 2 minutes)
Private Info Only	307343	116347	142430	48956
Private & Public Info	147346	189657	51608	88763

Panel C – Long-Horizon Trading Volume as a Linear Function of Availability of Public Information

Regression summary statistics from the regression of long-horizon trading volume on the availability of public information. Note that public information is a binary variable that is set equal to one if public information is available and is set equal to zero if public information is not available. We have 72 observation of long-horizon trading volume – 36 with public information and 36 without.

	Trading Volume Measured as Quantity Traded (all 4 minutes)	Trading Volume Measured as Price Times Quantity Traded (all 4 minutes)	Trading Volume Measured as Quantity Traded (last 2 minutes)	Trading Volume Measured as Price Times Quantity Traded (last 2 minutes)
Public Information	-53.7778*** (20.6509)	-4444.361** (2031.928)	-28.0556** (13.6112)	-2522.833* (1428.511)
Fixed Effect for Each of the 36 Periods	Yes	Yes	Yes	Yes
Random Effect for Session	Yes	Yes	Yes	Yes
Standard Errors Clustered at Session Level	Yes	Yes	Yes	Yes
R-squared	0.7260	0.6662	0.6546	0.6171
N	72	72	72	72

***, **, * indicate significantly different from 0 at $p < 0.01$, 0.05 and 0.10 levels, respectively. Robust standard error in parentheses. The regressions include a random effect for session and a fixed effect for each of the 36 periods. Standard errors are clustered at the session level.

Panel D – Short Horizon Trading Volume as a Linear Function of Availability of Public Information

Regression summary statistics from the regression of short-horizon trading volume on the availability of public information. Note that public information is a binary variable that is set equal to one if public information is available and is set equal to zero if public information is not available. We have 72 observations of short-horizon trading volume – 36 with public information and 36 without.

	Trading Volume Measured as Quantity Traded (all 4 minutes)	Trading Volume Measured as Price Times Quantity Traded (all 4 minutes)	Trading Volume Measured as Quantity Traded (last 2 minutes)	Trading Volume Measured as Price Times Quantity Traded (last 2 minutes)
Public Information	10.5833*** (3.3405)	2036.389** (842.234)	6.8056*** (2.2967)	1105.75*** (412.3064)
Fixed Effect for Each of the 36 Periods	Yes	Yes	Yes	Yes
Random Effect for Session	Yes	Yes	Yes	Yes
Standard Errors Clustered at Session Level	Yes	Yes	Yes	Yes
R-squared	0.8619	0.7823	0.8369	0.8110
N	72	72	72	72

***, **, * indicate significantly different from 0 at $p < 0.01$, 0.05 and 0.10 levels, respectively. Robust standard error in parentheses. The regressions include a random effect for session and a fixed effect for each of the 36 periods. Standard errors are clustered at the session level.

Table 2 – Aggregate Trading Volume for the First Four Trading Periods and the Last Four Trading Periods of Each Session of Each Treatment

The columns labeled ‘long-horizon traders’ give the aggregate trading volume generated by long-horizon traders while the columns labeled ‘short-horizon traders’ give the aggregate trading volume generated by short-horizon traders. The aggregation is across the first 4 trading rounds and the last 4 trading rounds for each session of each treatment.

Panel A – Aggregate Trading Volume Measured as Quantity Traded

	Trading Periods	Session	Long-Horizon Traders (all 4 minutes)	Short-Horizon Traders (all 4 minutes)	Long-Horizon Traders (last 2 minutes)	Short-Horizon Traders (last 2 minutes)
Private Info Only	First Four Periods	I	390	92	170	30
		II	240	112	100	44
		III	289	186	145	97
	Last Four Periods	I	469	120	171	53
		II	107	70	22	17
		III	665	228	381	88
Private & Public Info	First Four Periods	I	226	178	90	86
		II	177	119	64	55
		III	113	278	52	128
	Last Four Periods	I	165	169	57	93
		II	101	78	28	28
		III	129	235	34	120

Panel B – Aggregate Trading Volume Measured as Price Times Quantity Traded

The columns labeled ‘long-horizon traders’ give the aggregate trading volume generated by long-horizon traders while the columns labeled ‘short-horizon traders’ give the aggregate trading volume generated by short-horizon traders. The aggregation is across the first 4 trading rounds and the last 4 trading rounds for each session of each treatment.

	Trading Periods	Session	Long-Horizon Traders (all 4 minutes)	Short-Horizon Traders (all 4 minutes)	Long-Horizon Traders (last 2 minutes)	Short-Horizon Traders (last 2 minutes)
Private Info Only	First Four Periods	I	31949	7749	13521	2518
		II	25463	10241	11125	4163
		III	27571	18546	14200	9942
	Last Four Periods	I	36125	11706	13201	5175
		II	10898	7541	2179	1823
		III	69537	22801	40241	8945
Private & Public Info	First Four Periods	I	24859	22377	9932	11648
		II	20457	10943	7192	5058
		III	11554	40891	5141	18443
	Last Four Periods	I	16876	18536	5605	10211
		II	10667	8076	2879	2845
		III	14043	24956	3746	12098

Panel C – Results from Wilcoxon Signed Rank Test for Trading Volume for Each Treatment

For each session of treatment, we aggregated the trading volume for first four trading rounds and for last four trading rounds. As we ran three sessions of each treatment, we were able to get six trading volume observations for each treatment. (N = 6)

	Trading Volume Measured as Quantity Traded (all 4 minutes)		Trading Volume Measured as Price Times Quantity Traded (all 4 minutes)		Trading Volume Measured as Quantity Traded (last 2 minutes)		Trading Volume Measured as Price Times Quantity Traded (last 2 minutes)	
Hypothesis	z-score	p-value	z-score	p-value	z-score	p-value	z-score	p-value
Private information only – trading volume among long-horizon traders = trading volume among short-horizon traders	2.201	0.0277	2.201	0.0277	2.201	0.0277	2.201	0.0277
Long-horizon traders – trading volume with private and public info = trading volume with private info only	-2.201	0.0277	-2.201	0.0277	-1.992	0.0464	-1.992	0.0464
Short-horizon traders – trading volume with private and public info = trading volume with private info only	2.207	0.0273	2.201	0.0277	2.207	0.0273	2.201	0.0277

Table 3 – Disagreement Among Traders

Panel A – Average Disagreement Among Traders

For each of the 36 trading periods, five traders received a low private signal while five traders received a high private signal. We averaged the forecast made by the five traders with the low signal and averaged the forecast made by the five traders with the high signal. We measured the disagreement among traders as the absolute difference between these two averages. This table reports the average disagreement across 36 observations. The medians are reported in parentheses.

	Long-Horizon Traders	Short-Horizon Traders
Private Info Only	9.66 (7.4)	9.39 (6.6)
Private & Public Info	4.8 (3.0)	9.42 (8)

Panel B – Time Weighted Bid-Ask Spread

This panel reports the average time weighted bid-ask spread across 36 observations. The medians are reported in parentheses.

	Long-Horizon Traders	Short-Horizon Traders
Private Info Only	28.48 (16.56)	15.59 (11.96)
Private & Public Info	9.53 (8.25)	25.01 (19.33)

Panel C – Disagreement Among Long-Horizon Traders and Time Weighted Bid-Ask Spread Among Long-Horizon Traders as a Linear Function of Availability of Public Information

Regression summary statistics from the regression of disagreement among long-horizon traders on the availability of public information and of time-weighted bid-ask spread among long-horizon traders on the availability of public information. Note that Public Information is a binary variable that is set equal to one if public information is available and is set equal to zero if public information is not available.

	Disagreement Among Long-Horizon Traders	Time Weighted Bid-Ask Spread Among Long-Horizon Traders
Public Information	-4.8587*** (1.338)	-18.9551** (8.93)
Fixed Effect for Each of the 36 Periods	Yes	Yes
Random Effect for Session	Yes	Yes
Standard Errors Clustered at Session Level	Yes	Yes
R-squared	0.6895	0.6991
N	72	72

***, **, * indicate significantly different from 0 at $p < 0.01$, 0.05 and 0.10 levels, respectively. Robust standard error in parentheses. The regressions include a random effect for session and a fixed effect for each of the 36 periods. Standard errors are clustered at the session level.

Panel D – Disagreement Among Short-Horizon Traders and Time Weighted Bid-Ask Spread Among Short-Horizon Traders as a Linear Function of Availability of Public Information

Regression summary statistics from the regression of disagreement among short-horizon traders on the availability of public information and of time-weighted bid-ask spread among short-horizon traders on the availability of public information. Note that Public Information is a binary variable that is set equal to one if public information is available and is set equal to zero if public information is not available.

	Disagreement Among Short-Horizon Traders	Time Weighted Bid-Ask Spread Among Short-Horizon Traders
Public Information	0.0278 (1.8575)	9.4237* (5.5336)
Fixed Effect for Each of the 36 Periods	Yes	Yes
Random Effect for Session	Yes	Yes
Standard Errors Clustered at Session Level	Yes	Yes
R-squared	0.7367	0.8214
N	72	72

***, **, * indicate significantly different from 0 at $p < 0.01$, 0.05 and 0.10 levels, respectively. Robust standard error in parentheses. The regressions include a random effect for session and a fixed effect for each of the 36 periods. Standard errors are clustered at the session level.

Table 4 – Standard Deviation of Prices and Forecasts

	Prices		Forecasts	
	Long-Horizon Traders	Short-Horizon Traders	Long-Horizon Traders	Short-Horizon Traders
Private Info Only	29.01 (N = 3304)	14.98 (N = 1197)	17.98	18.75
Private & Public Info	13.45 (N = 1368)	29.8 (N=1578)	15.81	19.74

Table 5 – Participant-Level Trades as a Linear Function of Availability of Public Information

Panel A – Long-Horizon Participant-Level Trades as a Linear Function of Availability of Public Information

Regression summary statistics from the regression of long-horizon participant-level trades on availability of public information. Note that public information is a binary variable that is set equal to one if public information is available and is set equal to zero if public information is not available.

	Participant-Level Trades Measured as Quantity Traded (all 4 minutes)	Participant-Level Trades Measured as Price Times Quantity Traded (all 4 minutes)	Participant-Level Trades Measured as Quantity Traded (last 2 minutes)	Participant-Level Trades Measured as Price Times Quantity Traded (last 2 minutes)
Public Information	-13** (5.9414)	-804.7833* (427.7645)	-5.25** (2.1692)	-333.0333** (154.6407)
Fixed Effect for Each of the 36 Periods	Yes	Yes	Yes	Yes
Fixed Effect for Each of the 6 Sessions	Yes	Yes	Yes	Yes
Random Effect for Participant	Yes	Yes	Yes	Yes
Standard Errors Clustered at Participant Level	Yes	Yes	Yes	Yes
R-squared	0.2243	0.2041	0.2407	0.2330
N	720	720	720	720

***, **, * indicate significantly different from 0 at $p < 0.01$, 0.05 and 0.10 levels, respectively. Robust standard error in parentheses. The regressions include a random effect for participant, a fixed effect for each of the 36 periods and a fixed effect for each of the 6 sessions. Standard errors are clustered at the participant level.

Panel B – Short Horizon Participant-Level Trades as a Linear Function of Availability of Public Information


Regression summary statistics from the regression of short-horizon participant-level trades on availability of public information. Note that public information is a binary variable that is set equal to one if public information is available and is set equal to zero if public information is not available.

	Participant-Level Trades Measured as Quantity Traded (all 4 minutes)	Participant-Level Trades Measured as Price Times Quantity Traded (all 4 minutes)	Participant-Level Trades Measured as Quantity Traded (last 2 minutes)	Participant-Level Trades Measured as Price Times Quantity Traded (last 2 minutes)
Public Information	2.8667 (2.95)	456.9167 (330.971)	2.1667 (1.4744)	310.3833* (180.3452)
Fixed Effect for Each of the 36 Periods	Yes	Yes	Yes	Yes
Fixed Effect for Each of the 6 Sessions	Yes	Yes	Yes	Yes
Random Effect for Participant	Yes	Yes	Yes	Yes
Standard Errors Clustered at Participant Level	Yes	Yes	Yes	Yes
R-squared	0.1966	0.2691	0.2451	0.2461
N	720	720	720	720

***, **, * indicate significantly different from 0 at $p < 0.01$, 0.05 and 0.10 levels, respectively. Robust standard error in parentheses. The regressions include a random effect for participant, a fixed effect for each of the 36 periods and a fixed effect for each of the 6 sessions. Standard errors are clustered at the participant level.

Figure 3 – Volume Weighted Average Price

Panel A – Volume Weighted Average Price Among Long-Horizon Traders



Panel B – Volume Weighted Average Price Among Short-Horizon Traders






Figure 4 – Mean Absolute Deviation from Liquidating Value

Panel A – Mean Absolute Deviation from Liquidating Value Among Long-Horizon Traders



Panel B – Mean Absolute Deviation from Liquidating Value Among Short-Horizon Traders



Appendix A

Proof of Proposition 2: $|E(\theta|x_i) - E(\theta|x_j)| = \frac{Cov(\theta,x)}{Var(x)} |x_i - x_j|$, and

$E(\theta|x, y) = E(\theta|y) + \frac{Cov(\theta,x|y)}{Var(x|y)} [x - E(x|y)]$. So,

$$|E(\theta|x_i, y) - E(\theta|x_j, y)| = \frac{Cov(\theta,x|y)}{Var(x|y)} |x_i - x_j|$$

The proposition then follows if, $\frac{Cov(\theta,x|y)}{Var(x|y)} < \frac{Cov(\theta,x)}{Var(x)}$

But, $Cov(\theta, x|y) = Cov(\theta_A + \theta_B, \theta_B + \varepsilon|y) = Cov(\theta_A, \theta_B|y) + Var(\theta_B|y)$. So, the desired inequality is equivalent to:

$$\frac{Cov(\theta_A, \theta_B|y)}{Var(\theta_B|y) + \sigma_\varepsilon^2} + \frac{Var(\theta_B|y)}{Var(\theta_B|y) + \sigma_\varepsilon^2} < \frac{Var(\theta_B)}{Var(\theta_B) + \sigma_\varepsilon^2}$$

This last inequality is true due to the fact that $Cov(\theta_A, \theta_B|y) < 0$ and

$$\frac{Var(\theta_B|y)}{Var(\theta_B|y) + \sigma_\varepsilon^2} < \frac{Var(\theta_B)}{Var(\theta_B) + \sigma_\varepsilon^2} \text{ because } Var(\theta_B|y) < Var(\theta_B).$$

Appendix B

Instructions for Short-Horizon Public Information Available Treatment

Instructions Part 1 - Introduction

This is an experiment in the economics of market decision-making. You will be paid in cash at the end of the experiment based upon the decisions you make, so it is important that you understand these instructions. If you have a question, please raise your hand and a monitor will approach you. Otherwise, you should not communicate in any way with anyone else.


In this experiment we are going to simulate a market in which you can buy and sell shares with other participants in the experiment. The currency is called Experimental Currency Units (ECUs) and at the end of the experiment your ECUs will be converted into \$US at the rate \$1 = 2500 ECUs.

Your task

The experiment is broken up into multiple 4-minute trading periods. Each period you can trade (buy and/or sell) shares. At the end of a trading period, each share you own will be sold to a robot trader. This robot is not allowed to buy and sell shares during the period – it will only buy all shares from you at the end of the trading period. The robot trader will pay you its best estimate of the actual end-of-period share value for each share it buys from you. The shares will then expire. In a given trading period, the robot trader will pay all participants the same amount for each share.

During the trading period, you will not know the amount each share is worth, but you will receive two clues about this amount before the market trading period begins. After you have received your clues, you will be given 30 seconds to make a forecast of the amount the robot trader will pay for each share at the end of the period.

At the end of each period, you will be rewarded in ECUs for each of your forecasts. The more accurate your forecasts, the higher your reward will be. A timeline of the sequence of events in each period is as follows.



You will start each year with a balance of \$3,600 ECUs and 4 shares. Any time you buy a share, your share balance increases by 1 and the price you pay is deducted from your ECU balance. Any time you sell a share, your share balance decreases by 1 and the sale price is added to your balance. At the end of the period, your share will be sold to the robot trader, and this amount will be added to your ECU balance.


Additionally, your reward from making the pre-trading forecasts will also be added to your balance. Your earnings for the period are based on your final ECU balance. ECUs and shares do not carry over from one period to the next. This means that if you end period 1 with 100 ECUs and 2 shares, then you will still begin period 2 with \$3,600 ECUs and 4 shares.

How end-of-period share values are determined

The end-of-period share value is the sum of two components, which will be denoted by colored tokens: RED tokens and BLUE tokens. That is, the end-of-period share value is the total number of

RED and BLUE tokens.

The number of RED and BLUE tokens are uncertain and independent of each other. Each of these two quantities is determined by independent random draws from the bell-shaped distribution in the figure below. Although the exact number of RED tokens in the end-of-period share value will vary from drawing to drawing, on average (over many, many drawings) the number of RED tokens will be about 50 and the dispersion around this average is 10.



The above data implies that, in the absence of additional information, you can be 67% confident that the number of RED tokens in the end-of-period share value is in the range of 40 to 60, and you can be 95% confident that it is in the range 30 to 70.

For example, the probability that the number of RED tokens in the end-of-period share value is 50 is approximately 4%, while the probability that the number of RED tokens in the end-of-period share value is 70 is approximately 0.5%.

The number of BLUE tokens contained in the end-of-period share value is also uncertain, but has exactly the same features as the number of RED tokens. Thus, on average the number of BLUE tokens is also 50 and the dispersion around the average is 10. The number of BLUE tokens that is drawn from this distribution is independent of the number of RED tokens drawn from the distribution.

For each trading period, the computer randomly drew the number of RED tokens from the RED distribution and independently drew the number of BLUE tokens from the BLUE distribution. The computer then added the number of RED tokens drawn to the number of BLUE tokens drawn.

This sum is the end-of-period share value that will be announced at the end of the trading period. The end-of-period share value in every trading period is determined independently of previous end-of-period share values.

$$\text{RED tokens} + \text{BLUE tokens} = \text{End-of-Period share value}$$

Summary – Part 1

There will be a short quiz followed by a practice period to allow everyone to become familiar with entering offers and making trades, but before we do let's review the main points of the experiment.

1. Each period you will start off with 3,600 ECUs and 4 shares.
2. The shares you hold at the end of a trading period will be sold to a robot trader. The robot trader will pay you what it believes the actual end-of-period share value is. This value is not necessarily equal to the actual end-of-period share value.
3. You will not know the end-of-period share value, which is the sum of two components, namely RED tokens and BLUE tokens. But, you will receive two clues about this value.
4. After you receive the clues but before trading in the market commences, you will make a forecast of the amount the robot trader will pay you for each share you hold at the end of the period (that is, the robot trader's belief of the end-of-period share value).


Instructions Part 2 - Clues about the end-of-period share value

At the beginning of each trading period each of you will receive two clues. One of the clues will be about the number of RED tokens in the end-of-period share value and this clue will be provided to you confidentially. The other clue will be about the sum of the number of RED tokens and number of BLUE tokens contained in the end-of-period share value and this clue will be publicly provided to all participants.

The robot trader will receive a clue about the number of BLUE tokens contained in the end-of-period share value but it will not receive the public clue about the sum of the number of RED tokens and number of BLUE tokens contained in the end-of-period share value. The clue received by the robot trader is perfect so that it knows the exact number of BLUE tokens contained in the end-of-period share value.

The clues you receive are not necessarily perfect, but are useful in making your own subjective judgments about the end-of-period share value in that trading period. For each trading period the computer draws two confidential clues from the clue distribution for RED tokens that is described below. One of these two clues is randomly given to one-half of the participants and the other clue is given to the other half of the participants.

The confidential clue that you receive each period is equal to the number of RED tokens that was actually drawn from the RED distribution for that trading period plus a random error. The error is a drawing from the bell-shaped distribution in the figure below. On average (roughly 8% of the time) the error is 0 (meaning your clue tells you the exact number of RED tokens) and the dispersion around the average is 5.




To illustrate what this means, suppose:

	<i>Example 1</i>	<i>Example 2</i>
<i>Actual # of RED tokens drawn by computer</i>	<i>40</i>	<i>60</i>
<i>67% chance your confidential clue is in range</i>	<i>35 to 45</i>	<i>55 to 65</i>
<i>95% chance your confidential clue is in range</i>	<i>30 to 50</i>	<i>50 to 70</i>
<i>5% chance your confidential clue is</i>	<i>Greater than 50</i>	<i>Greater than 70</i>
<i>5% chance your confidential clue is</i>	<i>Less than 30</i>	<i>Less than 50</i>

Public clues about the end-of-period share value

Recall that the public clue is a clue about the total end-of-period share value, i.e. about the sum of RED and BLUE tokens. As with the confidential clue, the public clue also contains random error. The public clue is determined as follows. The number of RED tokens drawn by the computer is added to the number of BLUE tokens drawn by the computer and then a randomly drawn error term is added to this total. The random error term in the public clue is independent of the random error term in the confidential clues.

The error in the public clue is drawn from the bell-shaped distribution in the figure below. The average of this error is 0 and the dispersion of the error around the average is 5.



To illustrate what this means, suppose:

	<i>Example 1</i>	<i>Example 2</i>

<i>Actual # of RED tokens drawn by computer</i>	<i>40</i>	<i>80</i>
<i>Actual # of BLUE tokens drawn by computer</i>	<i>60</i>	<i>40</i>
<i>End-of-period share value</i>	<i>100</i>	<i>120</i>
<i>67% chance the public clue is in range</i>	<i>95 to 105</i>	<i>115 to 125</i>
<i>95% chance the public clue is in range</i>	<i>90 to 110</i>	<i>110 to 130</i>
<i>5% chance the public clue is</i>	<i>Greater than 110</i>	<i>Greater than 130</i>
<i>5% chance the public clue is</i>	<i>Less than 90</i>	<i>Less than 110</i>

Pre-trading Forecast Earnings

Your payoff for your pre-trading forecast is determined by a scoring rule. The formula is a little complicated, but what it means is that you maximize your expected payment by correctly guessing the amount the robot trader will pay you for each share at the end of the period. The formula is:

$$\text{Your payoff} = \max\{ 0, 2500 - 0.25 \times [(\text{your forecast}) - (\text{amount robot trader will pay you})]^2 \}.$$

Suppose the amount the robot trader will pay you for each share is 100. If you guess a value of 80, then your payoff would be:

$$\text{Your payoff} = \max\{ 0, 2500 - 0.25 \times [(80) - (100)]^2 \} = 2400 \text{ ECU}.$$

If you correctly guess the amount of 100 that the robot trader will pay you for this example, then your payoff would be

$$\text{Your payoff} = \max\{ 0, 2500 - 0.25 \times [(100) - (100)]^2 \} = 2500 \text{ ECU},$$

which is the maximum payoff you can earn for a forecast. Again, although the formula is a bit complicated it is structured so that you maximize your expected payment by correctly forecasting the actual amount the robot trader will pay you. Also note that your payoff cannot be negative.

Screenshots of the pre- and post-trading forecasting pages are below (assuming you forecasted 80 and the robot’s estimate of the end-of-period share value was 100).

Pre-trading Forecast

Your confidential clue is 40 RED tokens.

The public clue is 105 (sum of RED and BLUE tokens).

Recall: End-of-period share value = RED tokens + BLUE tokens.

Please enter your estimate of the robot trader's belief of the end-of-period share value below and click submit.

Estimate

Submit

Results

You forecasted the robot trader's belief of the end-of-period share value would be: 80

The robot's estimate of the end-of-period share value was: 100

The actual end-of-period share value was: 110

Based on your forecast, you earned: 2,400

Summary – Part 2

There will be a short quiz followed by a practice period to allow everyone to become familiar with entering offers and making trades, but before we do let's review the main points of the experiment.

1. You will receive two clues. You will receive a confidential clue about the number of RED tokens contained in the end-of-period share value. All participants will also receive the same public clue about the end-of-period share value (total number of RED and BLUE tokens). The clues are not necessarily perfect but are informative. The robot trader will receive a perfect clue about the number of BLUE tokens contained in the end-of-period share value so that it knows the exact number of BLUE tokens contained in the end-of-period share value. The robot trader will not receive the public clue.
2. After you receive the clues but before trading in the market commences, you will make a forecast of the amount the robot trader will pay you for each share at the end of the period. This amount is equal to the robot trader's belief of the actual end-of-period share value.

Instructions Part 3 - Computer Interface

Now that you have an overview of the experiment, we will talk in more detail about how you buy and sell shares. A table at the right hand side of your screen will show **Your Holdings** of ECUs and shares in the current period.

Your Holdings	
Shares	4
ECUs	3,600
Cash needed to cover current Offers to Buy	0
Available Cash	3,600
Cumulative Profits	0

The bottom right portion of your screen will display whatever information you have about the value of a share including clues.

Information	
Common Message	The public clue is 105 (sum of RED and BLUE tokens).
Clues	<ul style="list-style-type: none">Your confidential clue is 40 RED tokens.

Buying and Selling Shares

Each period, you can buy or sell shares from one another by making offers to buy or to sell. Every time someone makes an offer to buy a share, a **GREEN** dot will appear in the Market Chart on the left side of your screen. Every time someone makes an offer to sell, an **ORANGE** dot will appear on the graph. Offers are also listed in the **Market Book** to the right of the graph. The offers to buy will be listed in ascending order in **GREEN**, while the offers to sell will be listed in descending

order in **ORANGE**. Once a trade is actually made, the trade will be shown as a **BLACK** dot in the graph. Whoever bought the share will pay the agreed upon price to whoever sold the share. The person who sold the share will receive this payment. The buyer's number of shares will increase by one and the seller's number of shares will decrease by one. You cannot buy shares unless you have enough ECUs in your holdings to pay for it. You cannot sell a share if you do not have one in your holdings. To help you know if you have enough ECUs in your holdings to pay for a share, you may look at the **Available Cash** field. It reflects the difference of your current ECUs holdings and your **Cash needed to cover current Offers to Buy**, which represents the sum of all of your current Offers to Buy. Note that if you cancel an Offer to Buy, then your Cash needed to cover current Offers to Buy will decrease and your Available Cash will increase. Also, your end-of-period earnings are based upon your ECUs holdings – not your Available Cash.



The **Market Actions** section shows you the best prices to buy, or sell, that are currently available on the market. To accept an existing offer from another participant, simply click the **Sell Now** or **Buy Now** button. In this example, the lowest offer to sell is 90 and the highest offer to buy is 72. By clicking on the **Buy Now** button, you buy at the listed price (90); by clicking on the **Sell**

Now button, you sell at the listed price (72).

The screenshot shows a 'Market Actions' panel with two main sections: 'Sell Shares' and 'Buy Shares'. Each section has a text input field, a dark button labeled 'Offer to Sell' or 'Offer to Buy', a light grey box showing a price (72 for sell, 90 for buy), and a dark button labeled 'Sell Now' or 'Buy Now'. Below the 'Sell Shares' section, a note states: 'New offers to sell must be < lowest offer to sell (90)'. Below the 'Buy Shares' section, a note states: 'New offers to buy must be > highest offer to buy (72)'. At the bottom, a text line reads: 'You can cancel your existing offers to buy/sell by clicking on them below.' Below this text are four buttons: three orange buttons labeled 'Offer to sell: 122 ✖', 'Offer to sell: 112 ✖', and 'Offer to sell: 104 ✖', and one green button labeled 'Offer to buy: 54 ✖'.

To propose your own price to buy or sell, you simply type the price at which you would like to buy or sell, in the appropriate box and click the corresponding **Offer to Buy** or **Offer to Sell** button.

When you press **Offer to Buy** or **Offer to Sell** you are agreeing to trade a share at that price. Your offer will appear on the graph and in the Market Book (see Market Chart above).

Your offers to buy or sell will also appear on buttons at the bottom of the **Market Actions** section. Clicking this button will remove your offers from the market (see Market Actions above).

At the end of each Period, the Period Information section will display the number of shares you currently hold, the amount the robot trader will pay you for each share, and your ECU balance, which is calculated as:

$$\text{Ending ECU Balance} = \text{Shares} \times \text{Amount robot trader will pay for each share} + \text{ECUs}$$

Cumulative Profits reflects the sum of your trading profits from the current period as well as all previous periods.

Period 1 is Finished

Period Information	
Shares	2
End-of-Period Share Value	100
ECUs	3,768
Ending ECU Balance $\text{ECUs} + (\text{Shares} * \text{Share Value})$	3,968
Cumulative Profits	3,968

Summary – Part 3

There will be a short quiz followed by a practice period to allow everyone to become familiar with entering offers and making trades, but before we do let's review the main points of the experiment.

1. Each period you will start off with 3,600 ECUs and 4 shares.
2. You can buy and sell shares in the market. Offer to Buy (Offer to Sell) is used to announce a price at which you want to buy (sell) if someone else will accept. Buy Now (Sell Now) is used to instantly buy (sell) a share at a price offered by another trader in the market.
3. Your earnings in a period = your initial ECUs + ECUs you receive from selling shares – ECUs you spend buying shares + amount the robot trader pays you for buying shares you hold at the end of the period + reward for making your pre-trading forecast

Quiz

1. The end-of-period share value is a sum of two components: RED tokens and BLUE tokens. **T / F**
2. You will receive a confidential clue about the number of RED tokens in the end-of-period share value. **T / F**
3. The robot trader will receive a confidential clue about the number of RED tokens in the end-of-period share value. **T / F**
 - a. **Explanation: The robot trader will receive a confidential clue about the number of BLUE tokens in the end-of-period share value. Note that while your confidential clue does not inform you of the exact number of RED tokens in the end-of-period share value, the robot's confidential clue informs it of the exact number of BLUE tokens in the end-of-period share value.**
4. The confidential clue you receive perfectly reveals the number of RED tokens contained in the end-of-period share value. **T / F**
 - a. **Explanation: Your confidential clue does not necessarily inform you of the exact number of RED tokens in the end-of-period share value. It is informative, but does not necessarily perfectly reveal the number of RED tokens in the end-of-period share value.**
5. The confidential clue the robot trader receives perfectly reveals the number of BLUE tokens contained in the end-of-period share value. **T / F**
6. One-half of the participants will receive the same confidential clue. **T / F**
7. The public clue is a clue about the end-of-period share value, i.e. about the sum of RED and BLUE. **T / F**
8. The public clue you receive perfectly reveals the end-of-period share value. **T / F**
 - a. **Explanation: All human traders receive the same public clue regarding the end-of-period share value. This clue is informative but does not necessarily inform you of the precise end-of-period share value. Note that the robot trader does not receive this public clue.**
9. Can you use ECUs from one period for trading in the subsequent periods? **Y / N**
10. Can you buy more shares than your current cash (ECU) holding allows? **Y / N**