

The Limits of Noise Trading: An Experimental Analysis

By

Robert Bloomfield, Maureen O'Hara, and Gideon Saar*

First Draft: January 2005

***Robert Bloomfield (rjb9@cornell.edu) and Maureen O'Hara (mo19@cornell.edu) are from the Johnson Graduate School of Management, Cornell University. Gideon Saar (gsaar@stern.nyu.edu) is from the Stern School of Business, New York University.**

Abstract

In this research we investigate the behavior of noise traders and their impact on the market. We do this in an experimental market setting that allows us to determine not only how noise traders fare in a competitive asset market with other traders, but also how the equilibrium changes if a securities transactions tax (“Tobin tax”) is imposed. We find that noise traders lose money on average: they do not engage in extensive liquidity provision, and their attempt to make money by trend chasing is unsuccessful as they lose most in securities whose prices experience large moves. Noise traders adversely affect the informational efficiency of the market: they drive prices away from fundamental values, and the further away the market gets from the true value, the stronger this effect becomes. With a securities transaction tax, noise traders submit fewer orders and lose less money in those securities that exhibit large price movements. The tax is associated with a decrease in market trading volume, but informational efficiency remains essentially unchanged and liquidity (as measured by the price impact of trades) actually improves. We find no significant effect, however, on market volatility, suggesting that at least this rationale for a securities transaction tax is not supported by our data.

The Limits of Noise Trading: An Experimental Analysis

Noise traders play a ubiquitous role in the finance literature. Fischer Black [1986] dedicated his AFA presidential address to the beneficial effects of “noise” on markets, concluding that “noise trading is essential to the existence of liquid markets.” Shleifer and Summers [1990] and Shleifer and Vishny [1997] identified noise traders as the basis for the limits of arbitrage, arguing that noise trading introduces risks that inhibit arbitrageurs and prevent prices from converging to fundamental asset values. Noise traders in the guise of SOES bandits have been credited with enhancing price discovery (see Harris and Schultz [1998]), while their day trading counterparts have been disparaged for creating speculative bubbles in asset prices both home and abroad (Scheinkman and Xiong [2003]). The excessive volatility allegedly arising from noise trading is also blamed for a variety of economic ills ranging from market crashes to the failure of globalization.¹

Despite this central importance, there remains considerable debate regarding the role of noise trading in financial markets. For example, there is little agreement as to whether noise trading enhances or detracts from informational efficiency; whether noise trading increases or decreases price volatility; or even whether noise traders can survive in financial markets, in either the short or the long run. This disagreement spills over into debates over whether society is well advised to limit noise trading by taxation or other means, or to ignore it altogether due to its inconsequential nature in affecting market outcomes.

Undoubtedly, the confusion surrounding the role of noise traders stems in part from basic disagreements over even what constitutes noise trading. Black, in his original

¹ See, for example, “Day Trading Makes a Comeback” Reuters, December 19, 2004, where day traders “were blamed for adding irrationality to an exuberant market”. Similarly, critics of globalization point to the excessive short-term speculation in foreign currency trading as undermining the economic viability of developing countries (discussion of this issue can be found at www.globalsolidarity.org).

discourse on the subject, was careful to distinguish between noise traders and traders transacting for informational or portfolio reasons. He argued that “People who trade on noise are willing to trade even though from an objective point of view they would be better off not trading. Perhaps they think the noise they are trading on is information. Or perhaps they just like to trade.”² Shleifer and Summers [1990] view noise traders as responding to “pseudo-signals that investors believe convey information about future returns but that would not convey such information in a fully rational model.” These authors note, however, that “when [arbitrageurs] bet against noise traders, arbitrageurs begin to look like noise traders themselves. They pick stocks instead of diversifying, because that is what betting against noise traders requires....It becomes hard to tell the noise traders from the arbitrageurs.” And so, too, does it become difficult to discern the specific effects of noise traders from the more general behaviors characteristic of functioning markets.

In this research we seek to clarify these issues by investigating the behavior of noise traders and their impact on the market. We do this in an experimental market setting that allows us to determine not only how noise traders fare in a competitive asset market with other traders, but also how the equilibrium changes if we impose restrictions, such as a securities transactions tax (STT) (“Tobin tax”) on the market. Our experimental framework allows for traders to pursue a wide variety of strategies, and for market equilibria to exhibit a wide range of outcomes. We analyze how noise traders behave in these markets, and how they prosper and fail as measured by trading profits and losses. We also evaluate how noise trading affects market properties such as informational efficiency, volatility, and the losses and gains of other market participants. And we examine how individual traders and the market fare in the presence of a

² For an excellent analysis of excessive trading volume see Odean [1999].

transactions tax. Taken together, our analysis allows us to delineate the limits of noise trading in asset markets.

Experimental analysis seems particularly well suited for addressing this issue. As noted earlier, there is an extensive theoretical literature on noise trading, but the complexity of evaluating numerous (potentially non-rational) trading strategies imposes limitations on even the most insightful analyses (see, for example, the papers of DSSW [1990; 1991] or the more recent work of Scheinkman and Xiong [2003]). Similarly, there is a large empirical literature investigating the effects of particular trader types on markets (see, for example, Barber et. al. [2004]; Garvey and Murphy [2001]; Linnainmaa [2003]), but again the complexity of analyzing actual market data (when such data is even available) makes inference problematic. Experimental markets hold out the dual benefits of flexibility and control; traders are free to choose whatever strategies they like, while confounding factors are held constant across different market settings. This allows us to differentiate behaviors that are otherwise difficult to disentangle.

Our analysis provides a number of important results, several of which we highlight here. First, we find that noise traders drive prices away from fundamental values, and not towards them as is sometimes alleged. Two aspects of our findings here are particularly intriguing. We find that trading periods in which noise traders are most active are associated with higher pricing errors. And we find that the destabilizing effects of noise traders are greatest when prices are farther away from true values. Thus, noise traders hinder market efficiency exactly when the market needs it most. In the aftermath of the recent Nasdaq market crash, some have questioned the role played by small investors or day traders in contributing to the internet bubble. Our results here provide some support for the notion that these noise traders contributed to the market's inefficiency. We find that the presence of noise traders is not beneficial to market efficiency, and the farther away the market gets from true value, the stronger this effect becomes.

Second, we find that noise traders lose money, and they incur substantially greater losses than do liquidity traders. Because noise traders in our markets have neither an informational advantage nor a portfolio motive for trade, their trading behavior is best characterized as either liquidity provision or trend chasing. We find little evidence of the former, with noise traders instead generally taking rather than making liquidity. Unfortunately, the trend chasing strategy is not successful. We find these traders lose relatively little when prices are close to true values, but they suffer massive losses when new information takes on extreme values. This pattern allows noise traders to occasionally make small profits but ultimately to suffer even greater losses. Numerous authors have investigated whether day traders make profits in actual markets. Although the results of specific studies differ, our results are close to those of Barber, Lee, Liu and Odean [2004] who find substantial losses associated with day trading in Taiwan, where average daily losses are more than twice the size of average daily gains.

Our results here are also relevant for the debate over the survival of noise traders in financial markets. Friedman [1953] first argued that wealth dynamics would ultimately drive noise traders from markets, and a large literature has expanded on this idea (see DeLong, Shleifer, Summers, and Waldman (DSSW) [1991]; Blume and Easley [1992; 2002]; Shefrin and Statman [1994]; Sciubba [1999]; Sandroni [2000]). DSSW offer the counter view that the greater risk taken on by noise traders could permit survival given a positive risk-return trade-off. Our results cast doubt on this conjecture: we find that noise traders do not prosper in markets because trend chasing strategies collapse in the presence of large market movements. We also find some evidence that noise traders who lose money in a period then trade less in subsequent periods, consistent with their eventually being driven out of the market. These results suggest that while noise trading may persist in actual markets it does so because of the arrival of new noise traders, and not because of the survival of older ones.

Third, we find some evidence that a higher level of noise trading is associated with higher price volatility. The notion that short term speculation, presumably the hallmark of noise traders, creates excess volatility in financial markets has long been debated. Many countries impose a Securities Transactions Tax (STT) in order to curb such speculation, and several times over the past two decades such a tax has been proposed in the U.S. as well. While the costs and benefits of the STT have been extensively discussed in the literature (see Stiglitz [1989]; Summers and Summers [1989]; Amihud and Mendelson [1993]; Schwert and Seguin [1993]; Subrahmanyam [1990]) one problem with this discussion is the scarcity of empirical evidence on the effects of imposing such a tax (see Umlauf [1993] on the effects of the STT in Sweden).

In our experimental market we find that imposing the STT indeed affects trader behavior. With such a tax, noise traders submit fewer orders, and lose less money in those securities that exhibit large price movements. Market trading volume falls, but informational efficiency remains essentially unchanged and liquidity (as measured by the price impact of trading) actually improves. We find no significant effects, however, on market volatility, suggesting that at least this rationale for the STT is not supported by our data. Overall, our findings support the notion that restricting noise trading does little harm to the market and may be useful if the effects on individual traders are considered.³

This paper is organized as follows. The next section sets out the experimental design we use in this analysis. In Section 3 we then present our results on the effects of noise traders in markets. Our analysis examines whether noise traders make money by trading; whether noise traders help or hurt price discovery; and the relationship between noise trading and market properties such as informational efficiency and volatility. In Section 4 we introduce a Securities Transactions Tax into our experimental markets, and

³ The SEC moved to curb day trading in 2000 by requiring at least \$25,000 to open a day trading account. This action stemmed from the 1999 incident in which a day trader killed his family and nine people at two brokerages in Atlanta after having sustained day trading losses of \$150,000.

consider anew the issues of noise trader profits, price discovery, informational efficiency, and market volatility. Section 5 then summarizes our results and offers conclusions on the role of noise traders in financial markets.

2. Experimental Design

We now describe the nature of our experiment and the specific features of our markets. This section is based on Bloomfield, O'Hara, and Saar [2005], which uses very similar market structures. As a useful preliminary, we note the following definitions. A *cohort* is a group of twelve traders who always trade together. Our experiment includes eight cohorts, for a total of 96 traders. A *security* is a claim on a terminal dividend, and is identified by the distribution of information and liquidity targets (described below). A *trading period* is a 150-second interval during which traders can take trading actions for a specific security. Only one security is traded in each trading period. Each cohort of participants traded 24 securities sequentially (one in each trading period) over a 90-minute session. Unless otherwise indicated, all prices, values and winnings are denominated in laboratory dollars (\$), an artificial currency that is converted into US currency at the end of the experiment.

2.1. Experimental Goals and Design

The first goal of our experiment is to examine the behavior and welfare of noise traders. To achieve this goal, we construct markets with three types of traders. Four traders (designated the informed traders) are given information about security value that is not given to the remaining traders. Four other traders (the liquidity traders) do not know the true value, but are assigned trading targets they must achieve before the end of trading if they are to avoid a penalty described below. Liquidity traders failing to meet their targets are assessed a penalty equal to \$100 for each unfulfilled share. This penalty is large enough to make trading attractive to the liquidity traders, and once they hit their

targets, liquidity traders can buy or sell as many shares as they please without penalty.⁴ The remaining four traders (the noise traders) have no information and no trading targets, and therefore have no reason to trade other than that they believe that they can interpret market information better than others, or that they can provide liquidity.⁵

By creating three classes of traders, we can focus on how the behavior of noise traders *differs* from that of informed and liquidity traders. This focus on differences allows us to draw clearer inferences about noise trader behavior, because we can use the other types of traders as control groups in our analyses.

One possible concern in any study of noise traders is that traders with no information or liquidity targets may trade simply because they feel they have nothing better to do with their time. Because this could be viewed as the definition of a noise trader, we do not view this as a problem. However, simply observing a baseline level of noise trade would shed little light on the forces that drive such trade and the cost-benefit analyses that influence noise traders' decisions. Consequently, to examine the conditions that influence noise trader behavior and how noise traders affect the market, we also manipulate security value realizations and aggregate liquidity demand.

Specifically, security value realizations are structured as follows. Each security pays a liquidating dividend equal to 50 plus the sum of two random numbers, each of which is uniformly distributed from -25 to 25. Two informed traders learn the value of the first random number before trading begins while the other two informed traders learn the value of the second random number (so that no trader has perfect information). We categorize securities as high-extremity if they have values below 35 or above 65; we categorize securities as low-extremity if they have values between 36 and 64. We allow

⁴ The use of trading targets is standard in experimental work (see, for example, Lamoreaux and Schnitzlein [1997], Cason [2000], Bloomfield and O'Hara [1998; 2000] and Bloomfield, O'Hara and Saar [2005]), and it captures the notion that liquidity traders are transacting for exogenous reasons relating to the need to invest or the need to liquidate positions.

⁵ Each trader knows his or her type, and each knows the populations of informed, liquidity, and noise traders in the market. Traders do not know, however, the roles played by specific participants in the experiment.

noise traders the opportunity to earn profits by providing liquidity. For that purpose, we set the targets of the liquidity traders such that the aggregate number of shares they demand is positive, zero, or negative in different securities.⁶

Our second goal is to understand the market influence of noise traders. This goal presents a challenge in experimental design. Ideally, one could exploit the advantages of the experimental method by comparing markets that are identical in every way, except for the presence of noise traders. However, for markets of finite size, such controls are not possible: a market with noise traders must either have more total traders or fewer traders of other types than a market without noise traders. To avoid this problem, we use two different techniques. First, we use a variety of measured independent variables, such as number of trades involving noise traders. Unlike manipulated independent variables, which are under the complete control of the experimenter, measured independent variables are determined by participant behavior. As a result, we cannot rule out the possibility that associations between market behavior and noise trader behavior are driven by variables that we do not observe, but that are associated with both the dependent and independent variables (similar to the situation in empirical work).

To minimize the risk of this ‘correlated omitted variables’ problem that arises from the use of measured independent variables, we also use a controlled manipulation to assess the influence of noise traders: the Securities Transaction Tax (SST). In a Tax setting, we impose a \$2 fee for each trade. In a no-tax setting, there are no transaction costs. We expect the SST to substantially reduce the propensity of noise traders to trade, because the benefits to doing so are made substantially smaller. The direct effect of the tax on the behavior of the informed and liquidity traders is less obvious: informed traders

⁶ For 8 securities, half of the liquidity traders must buy 20 shares and the other half must sell 30 shares, for an aggregate liquidity demand of -20 . For another 8 securities, half of the liquidity traders must sell 20 shares, while the other half must buy 30 shares, for an aggregate liquidity demand of $+20$. The remaining eight securities have zero aggregate liquidity trader demand: For 4 of those securities, half of the traders must buy 20 shares and the other half must sell 20 shares; for the other 4 securities, half of the traders must buy 30 shares, and the other half must sell 30 shares.

have valuable information to trade on, and liquidity traders suffer substantial penalties for failing to meet their trading targets. In addition to clarifying the influence of noise traders, the SST manipulation also allows us to speak to the policy issues associated with such taxes.

2.2. *Experimental Controls*

To ensure that we can draw clear inferences from treatment differences, we created two groups of 12 securities, each containing four securities from each level of aggregate liquidity demand. Within each group, exactly half of the securities with each level of aggregate liquidity demand (-20 , $+20$ and 0) have extreme value realizations, while the other half do not. We balance the direction of the aggregate liquidity demand with the direction of extreme value realizations, so that just as many securities with extremely positive (negative) values have positive aggregate liquidity demand as negative aggregate liquidity demand, and just as many securities with positive (negative) aggregate liquidity demand have extremely positive values as negative values. All of the cohorts traded both groups of twelve securities in the same order; however, half of the cohorts traded the first group with an SST and the second group without one, while the other half of the cohorts traded the first group without an SST and the second group with one.

In summary, our experiment is a fully factorial repeated-measures design in which each variable is perfectly balanced for order of presentation and for other treatments. Specifically, the design includes trader type (informed, liquidity, noise), SST (present, absent), extremity (high, low), cohort (eight cohorts of twelve traders each), trading order (SST first, SST last), and aggregate liquidity demand (-20 , $+20$ and 0).⁷ Trader type and cohort membership are manipulated across traders, and all other factors are manipulated

⁷ Analysis of results by trading order (SST first, SST last) did not produce any significant patterns, and therefore we do not discuss it further in the paper.

within traders. We count on the random assignment of participants to trader types to minimize the possibility that differences across trader types are driven by individual differences.

2.3. *Trading*

Our double auction market is organized like a typical electronic limit order book with strict price/time priority (older limit orders in the book at each price level are executed first). During the 150 seconds of trading in each security, traders can enter bids (limit orders to buy one share at a chosen price) and asks (limit orders to sell one share at a chosen price). Traders can delete their orders at any time during the period. Traders can enter as many bids and asks as they wish, subject to two rules. First, they cannot enter bids with prices that are equal to or more than the price of the best outstanding ask, or asks with prices that are equal to or less than the price of the best outstanding bid (so that the market is not “crossed” or “locked”). Second, all bids and asks must have integer prices between 0 and 100, inclusive. As soon as a trader enters an order, the order is shown on every trader’s computer screen, indicating that an unidentified trader is willing to buy or sell one more share at the posted price.⁸

Traders may also take other traders’ bids and asks. Traders take an ask by clicking a “buy 1” button, which allows them to buy one share at the lowest current asking price. Traders take a bid by clicking a “sell 1” button, which allows them to sell one share at the highest current bid price. Taking an ask is equivalent to entering a market (or marketable limit) buy order, while taking a bid is equivalent to entering a market (or marketable limit) sell order. All trades are reported immediately to all traders, indicating the price

⁸ A trader’s screen includes one chart indicating bids and one indicating asks. The left side of each chart shows every price at which an order has been posted (shown in green for the highest bid and lowest ask price, and yellow for other prices), and the number of shares posted at that price (shown by a number to the left of the graph). The right side of each chart shows every price at which the trader has personally posted an order, and the number of shares that the trader has posted at that price. The center of each chart also includes a solid red line indicating the highest bid or lowest ask entered by any trader, and a solid green line indicating the highest bid or lowest ask entered by that particular trader.

and the trade direction (whether the trade involved a market buy taking an ask or a market sell taking a bid).

2.4. Subjects, Instructions and Incentives

The experiments were conducted in the Business Simulation Laboratory (BSL) at the Johnson Graduate School of Management at Cornell University. The participants in the experiments were Johnson MBA students. Each experimental 90-minute session involved twelve participants (one cohort). Upon arriving at the BSL, each subject received detailed written instructions, a copy of which is given in the Appendix. The instructions were reviewed in detail by the experiment administrator, who also answered any questions. The administrator then guided participants through the use and interpretation of the computer interface by trading a practice security, which was exactly like the securities to be traded during the experiment, except that trading outcomes did not affect participants' cash winnings.

Traders started trading in each security with an endowment of \$0 in cash and zero shares. Unlimited negative cash and share balances were permitted, so traders could hold any inventory of shares they desired, including large short positions. Traders were told that at the end of trading, shares paid a liquidating dividend equal to their true value, so that their net trading gain or loss for a security would simply be equal to their ending share balance times the value of each share, plus their ending cash balance. Any penalties assessed to a liquidity trader for failing to hit a target are deducted from this trading gain or added to her trading loss.

We determined cash winnings by subtracting a "floor" from each trader's winnings in laboratory dollars, and then multiplying by an exchange rate that converts laboratory dollars into US dollars. The floor and exchange rate were derived from pilot experiments separately for each type of trader, and were designed so that each type would receive average winnings of approximately \$25 per 90-minute session (with an expected

minimum payment of \$10.00). Traders were not told the floor or exchange rate, however, to minimize any gaming behavior.⁹

2.5. *Statistical Analysis*

We use two types of statistical methodologies in Sections 3 and 4: (i) a repeated-measures ANOVA, and (ii) a regression (or VAR) analysis.

From a statistical standpoint, the repeated-measures ANOVA is a conservative and robust procedure for analyzing experimental data. To judge statistical significance, we compute the average of the dependent variable within each cell (defined by the appropriate factors) for each of the eight cohorts. A repeated-measures analysis effectively treats each cohort as providing a single independent observation of the dependent variable. This design therefore reduces the problem, common in experimental economics, of overstating statistical significance by assuming that repetitions of the same actions by the same subject or group of subjects are independent events.

When appropriate, we will use the ANOVA terminology of "main effect," "interaction," and "simple effect" to describe the statistical tests. A main effect examines the influence of one factor averaging over all the levels of the other factors. An interaction is when the effect of one factor is different at different levels of the other factors. A simple effect looks at the influence of one factor holding another factor at a specific level.¹⁰ All the results we report in Figures 3 through 7 and Table 3 are based on

⁹ As trade progresses, traders may become aware that they are incurring trading losses. Traders' payoffs are based upon their trading performance, so poorer performance results in lower payments. One concern is that losing traders may take on excessive risk, a problem sometimes referred to as the "house money" effect. Excessive risk-taking in experiments can be curbed by subtracting trading losses from a "floor" to determine actual payoffs. Because the floor level is unknown to traders, the actual level of their trading losses is also unknown, thereby reducing their tendency for overzealous trading.

¹⁰ For example, a significant Tax main effect without a significant Type*Tax interaction means that the STT exerts a similar influence on the behavior of all trader types. A significant Type*Tax interaction implies that the different types of traders behave differently in the two tax regimes (with and without tax) with respect to the dependent variable under investigation. An example of simple effects is looking separately at the three types of traders to see whether Tax exerts a significant influence on the behavior of each type.

the ANOVA analysis. The ANOVA results are easy to recognize in the text of Sections 3 and 4 because we provide the p -values associated with the different findings.

The second type of statistical procedure we use is a regression analysis. We use regressions to look at the relationships between market-wide measures of informational efficiency or volatility and the trading of informed and noise traders. The unit of observation for the regression analysis is a security (i.e., one trading period). Where appropriate, we investigate the effects of extremity levels or the tax regimes using dummy variables, and report the results of Chow tests for structural change as a means of assessing the significance of these manipulations.

A possible concern with such an analysis is that regression errors for different trading periods of the same cohort of subjects may be correlated. To examine the robustness of our results, we use a methodology that is similar in spirit to the one in Fama and MacBeth [1973]. We run a separate regression for the 24 trading periods of each cohort. While the potentially correlated errors make the standard errors of the coefficients difficult to interpret, the coefficients themselves are still consistent. We then perform t -tests and Wilcoxon signed-rank tests on the eight observations of each coefficient. This approach allows us to draw inferences from across cohorts, which can be assumed independent. The regression results are presented in Tables 1, 2, and 4.

In Section 4.5 (Table 5) we use a bivariate vector autoregression (VAR) to analyze the dynamic relation between noise traders' profits and their trading intensity. The VAR is particularly suitable for analyzing the possibly related behavior of a trader over time in the different securities he or she trades. We use Granger-causality tests to make an inference on whether traders' performance affects their future behavior.¹¹

¹¹ For a standard exposition of vector autoregressions see Hamilton [1994].

3. Noise Traders and their Effect on the Market

The experimental framework detailed in the previous section is designed to investigate how noise traders affect the market. As a useful preliminary, we begin with summary statistics to provide a sense of the behavior of the market and the orders initiated by the traders. We then look at whether noise traders make money by trading. We proceed to examine whether noise traders help or hurt value discovery, and the relationship between noise trading and market-wide measures of informational efficiency and volatility. Discussing the effects of a securities transaction tax (STT) on the market is deferred to Section 4.¹²

3.1. Summary Statistics

Figure 1 presents the evolution over time of volume and bid-ask spreads in our markets. Each panel divides the trading period in a security into ten 15-second intervals. Since trading starts when the book is empty, no executions can take place before there are limit orders in the book. Therefore, traders tend to submit more limit orders in the first 15-second interval and as a result executed trade volume is lower. Panel A shows that volume increases markedly in the second interval, and from that point volume exhibits the usual “U” shape observed in equity markets. Panel B shows that bid-ask spreads also follow the “U” shape over time that is prevalent in equity markets and experimental studies (see, for example, Bloomfield, O’Hara, and Saar [2005]).

Figure 2 provides information on the orders submitted by traders. Panel A shows that all three trader types (noise, informed, and liquidity) use both market and limit

¹² The manner in which we report our results—first in Section 3 on noise traders’ behavior and impact on the market and then in Section 4 on the effects of the securities transaction tax—is for expositional simplicity. The results we present in Section 3 are based on data gathered from trading periods with and without taxes. What enables us to postpone the discussion of taxes is that there are no statistically significant higher-level interactions of trader type, extremity, and taxes in the ANOVA analysis except for profits. As for profits, the directional effect of taxes on noise traders and informed traders is the same (just the magnitude of the effects differs across the tax regimes), and therefore we can still draw clear inferences even before discussing taxes. Table 3 contains a complete breakdown of profits by trader type, extremity level, and tax regime.

orders. This evidence suggests that experimental subjects understood the market mechanism and felt comfortable pursuing various order strategies. Panel B shows the average number of trades by the three types of traders and their breakdown into market and executed limit orders. Liquidity traders trade the most (44.8 shares), followed by informed traders (43.2 shares), and noise traders (32.2 shares).

The heavy trading by the liquidity traders is understandable given their exogenously imposed portfolio targets. Similarly, because informed traders have private information, their trading activity is a natural consequence of their efforts to profit from their informational advantage. Noise traders are given no targets or special information, but they, too, seem to trade a significant number of shares. We provide no specific motivation for this behavior, but note that it could be consistent with a wide variety of arguments advanced in the literature such as overconfidence, trend-chasing, liquidity provision, or simply liking to trade. To evaluate this issue further, we now turn to analyzing the trading behavior of noise traders.

3.2. *Trading Profits*

The standard result in the theoretical market microstructure literature is that informed traders make money and liquidity traders lose money. This is also the case here. The average trading profit of an informed trader is 142.17 experimental dollars (henceforth, dollars), while a liquidity trader loses on average 61.04 dollars.¹³

Less obvious is the question of whether noise traders can make money. Black [1986] posits that most of the time noise traders as a group will lose money. As noted earlier, the literature on day trading is mixed on whether day traders, who can be viewed as one type of noise traders, make or lose money on average. Harris and Schultz [1998]

¹³ These numbers represent trading profits before any payment of taxes. The pattern of profits net of taxes across trader types is similar to that presented here. Section 4 looks in detail at the effect of security transaction taxes on traders' profits.

show that SOES bandits (individual investors who used Nasdaq's Small Order Execution System for day trading) make money on average. They typically attempt to capitalize on short-term momentum in prices and hold a position for only a few minutes. Linnainmaa [2003] looks at day traders in Finland and finds that their performance is not different from that of a control group comprised of non-day-traders. Interestingly, day traders in her sample are net suppliers of liquidity in that they use limit orders and therefore often function as if they were market makers. Jordan and Diltz [2003] and Barber, Lee, Liu, and Odean [2004] find that day traders lose money on average.

The experimental design does not "force" noise traders to lose money. Because there is private information in the market, if noise traders are good at identifying price trends, they could make money by trading on these trends like Harris and Schultz [1998]'s SOES bandits. The experimental design also features imbalances in the buy and sell targets of the liquidity traders, so noise traders can presumably make money by supplying liquidity as in Linnainmaa [2003], smoothing out demand fluctuations and getting compensated for making the market.

The evidence in the experiment is that noise traders realize trading losses, 81.13 dollars per trader on average, that are even greater than those of the liquidity traders who have to trade. In fact, the more noise traders trade, the more money they are likely to lose. The correlation between the trading profit of a noise trader in a trading period and the number of shares he or she trades across all trading periods and all cohorts in the experiment is -0.345 (with $p < 0.0001$ against the hypothesis of zero correlation). This result is reminiscent of Odean's [1999] findings that individual traders trade "too much".

Panel A of Figure 3 shows that there is an interesting pattern with respect to trading profit and the extremity of the realized value, which can be viewed as a (noisy) measure of the value of the private information that the informed traders possess ($p = 0.011$ for the extremity*type interaction). Noise traders lose relatively little (12.08 dollars) when the realized value of the security is not too far from the expected value.

However, they lose a lot of money (150.18 dollar) when the realized value is far from the expected value (the difference in the noise trader's losses is statistically significant, $p = 0.0162$). Informed traders, not surprisingly, make much more money when their information is valuable, 218.55 dollars, than when their information is less valuable, 65.79 dollars. Liquidity traders lose about the same in both cases (53.71 and 68.37).

That extremity has little effect on liquidity traders' losses while greatly affecting noise traders' losses suggests that noise traders are unable to trade profitably on price trends. In fact, they seem to behave like the "mistaken" investors hypothesized by Friedman [1953] who buy high and sell low, or the day traders in Jordan and Diltz [2003] and Barber, Lee, Liu, and Odean [2004].

We also examine whether noise traders attempt to make money by adopting the role of liquidity providers as in Linnainmaa [2003]. Panel B of Figure 3 shows how much depth trader types contribute to the best bid or offer prices (BBO).¹⁴ Informed traders are much less likely to provide depth when the value of their information is high, consistent with the findings of Bloomfield, O'Hara, and Saar [2005]. Noise traders supply about the same level of liquidity irrespective of the level of extremity (0.66 and 0.65), and this level is much lower than that of either informed or liquidity traders. Thus, noise traders play a very minor role in liquidity provision, and focus instead of trend chasing, a strategy that is not particularly successful.

3.3. *Value Discovery and Informational Efficiency*

Black [1986] states that noise trading "actually puts noise into the prices", and concerns that excessive speculation can drive prices away from fundamental values are mentioned in Summers and Summers [1989] and Schwert and Seguin [1993]. An

¹⁴ We take snapshots of the contribution to depth at the BBO of all trader types each time a trader submits or cancels an order. The average contribution to BBO depth of a trader of a certain type is computed as the simple average of these snapshots divided by the number of traders of each type.

experimental methodology is particularly suitable for examining this issue because we know the fundamental value of each security and we are also able to directly observe the trading of noise traders. In the analysis which follows, we first develop measures that quantify the contribution of a trader type to value discovery, or whether their trades move prices closer to or away from the true value. We then look at market-wide measures of informational efficiency (pricing errors), and use regression analysis to relate those to the magnitude of noise trading and informed trading in the market.

To examine the contribution of each trader type to value discovery (or the adjustment of prices toward the true value of the security) we construct two measures. For Measure A, we sort executed orders into those moving prices toward the true value or away from it. In particular, if the true value is higher than the price, we assign +1 to a buy order of a trader that resulted in a trade and -1 to a sell order that resulted in a trade. If the true value is lower than the price, we assign -1 (+1) to a buy (sell) order of a trader that resulted in a trade. The measure is then aggregated for all market and executed limit orders of a trader and divided by the number of his trades.

Measure A therefore has the following properties. First, it is always in the range $[-1, +1]$. If a trader always buys when the price is below the value and sells when the price is above the value, he is helping price discovery and will have measure $A = +1$. If he is always hindering price discovery by buying when the price is above the true value of the security and selling when it is below the value, he will have measure $A = -1$. The more positive (negative) measure A of a trader, the more his trades contribute to (interfere with) value discovery.¹⁵ Measure B of value discovery is constructed in a similar way to measure A, except that we consider only trades resulting from market

¹⁵ We computed two versions of this measure: One using all orders that resulted in a trade, and one that excludes trades that occur when the true value is within the bid-ask spread. The results were similar for the two measures, but since the logic behind the measure is less robust when the true value is within the spread, we present the results of the second version.

orders submitted by the trader. Therefore, it considers contribution to or interference with value discovery only for the “active” side of a trade.

Panel A of Figure 4 shows that the different types of traders affect value discovery in a different way ($p < 0.0001$). As one might expect, the trades of informed traders help prices adjust to the true value: their measure A is equal to 0.23 and measure B is equal to 0.12. If noise traders help price discovery by making price adjust faster as some have claimed in the controversy over the trading of SOES bandits (e.g., Battalio, Hatch, and Jennings [1997]), then the value discovery measures of the noise traders should also be positive. We find the opposite: both measures A and B are negative for noise traders (-0.08 and -0.15 , respectively), echoing the notion in Black [1986] or the concerns in Summers and Summers [1989] that noise traders are driving prices away from fundamental values.

There is also a statistically significant interaction between trader type and extremity for measure A ($p = 0.0029$). As Panel B of Figure 4 shows, when the realized value of the security is close to expected value, the value discovery measure of the noise traders is negative, but close to zero (-0.02) which means that their trading is more neutral with respect to value discovery. On the other hand, when the realized value is further away from the expected value of the security, or when the private information of the informed traders is more valuable, measure A is equal to -0.13 . This means that exactly when prices are in greater need for informational efficiency, when they are further away from the true value and there is valuable private information in the market, noise traders hinder rather than help the adjustment of prices to their full-information level.

Our second analysis relates the average level of informational efficiency in the market to the levels of noise trading and informed trading. For each trading period, we compute the measure DEVP as the average of the distance of the transaction price from

the true value (pricing error).¹⁶ We also record all trades of the noise traders (the sum of their market orders and executed limit orders), and similarly for the informed traders. We then run the following OLS regression:

$$DEVP_i = a + b * NOISETRD_i + c * INFTRD_i + Error_i$$

where NOISETRD are the trades of the noise traders, INFTRD are the trades of the informed traders, and the subscript i denotes all trading periods pooled across the eight cohorts. We also run a similar regression for another pricing error measure that is not influenced by the magnitude of the bid-ask spread, DEVMID. This measure is computed by averaging the distance of the true value from the midpoint between the prevailing bid and the ask prices when a trade takes place.

Panel A of Table 1 presents the results of the two regressions. The coefficient on NOISETRD is positive and highly statistically significant in both regressions. This means that trading periods in which noise trading is higher (holding constant the amount of informed trading) are associated with higher pricing errors or lower informational efficiency. In contrast, the coefficient on INFTRD is negative and highly statistically significant in both regressions. This means that a higher amount of informed trading is associated with lower pricing errors or higher informational efficiency. This negative coefficient is consistent with informed traders putting new information into prices, thereby helping prices adjust to their full-information level. The results suggest, however, that noise traders in fact slow down the adjustment of prices or introduce noise into prices.¹⁷

A potential econometric problem with the above specification is that the regression errors for different trading periods of the same cohort of subjects may be

¹⁶ A trading period constitutes two and a half minutes of trading in which one security is traded before its dividend is publicly revealed and gains and losses are realized.

¹⁷ We also ran slightly different models to see if the results are sensitive to the manner in which we control for the level of trading. We ran univariate and multivariate regressions where we normalized NOISETRD and INFTRD by total volume and the results were similar in the sense that noise trading had a positive and significant coefficient and informed trading had a negative and significant coefficient.

correlated. We examine the robustness of our results using the following procedure: (i) we run a separate regression for the 24 trading periods of each cohort, and (ii) we then perform t-tests and Wilcoxon signed-rank tests on the eight observations of each coefficient. The regression coefficients from each cohort are consistent even in the presence of this econometric problem, and using their (cross-cohort) variability allows us to construct tests using eight observations that can be assumed independent.

Panel B of Table 1 provides the means and medians of the coefficients on NOISETRD and INFTRD together with the associated p -values for the tests. In both models with DEVP and DEVMID, the mean and median of the NOISETRD coefficients are positive and statistically significant, while the mean and median coefficients on INFTRD are negative and statistically significant. These results support our conclusion from the pooled regressions on the effects of noise and informed trading on informational efficiency.

Like the value discovery measures, the analysis of pricing errors also exhibits an interesting interaction with the extremity manipulation. To investigate this interaction we introduce a dummy variable that is equal to zero for low extremity securities and one for high extremity securities. Panel C of Table 1 shows the regressions on DEVP and DEVMID in a full dummy variable format (without an intercept) to look separately at the effects of low and high value extremity on the relation between the informational efficiency measures and both NOISETRD and INFTRD. We also use a structural stability test to investigate the question of whether the relation between informational efficiency and trading is the same across the two subsamples of securities (low and high extremity).

We observe that informed traders are most helpful in lowering pricing errors when the value of their information is higher. In fact, the negative coefficient on INFTRD is not statistically different from zero in the low extremity subsample, but it is highly statistically significant in the high extremity subsample. Noise trading seems to be associated with higher pricing errors mainly in the low extremity case. Chow tests for

structural stability performed on both regressions reject the null hypothesis of no structural change (in both cases with $p < 0.0001$).¹⁸

3.4. *Volatility*

The limits-to-arbitrage literature claims that noise trading could increase the volatility or risk associated with financial securities (e.g., De Long, Shleifer, Summers, and Waldmann [1990] and Shleifer and Summers [1990]). Black [1986] also states that “anything that changes the amount or character of noise trading will change the volatility of prices” (p. 533). Reducing the “excess” volatility of prices is also one of the arguments mentioned in favor of a securities transaction tax (e.g., Summers and Summers [1989], Schwert and Seguin [1993]). The empirical literature on day trading finds mixed results, however, on whether noise trading causes subsequent volatility. Lynch-Koski, Rice, and Tarhouni [2004] use a vector autoregression and find that noise trading (which they proxy by activity on message boards followed by day traders) increases volatility in the following day. Battalio, Hatch, and Jennings [1997] use a different proxy (trading on Nasdaq’s Small Order Execution System) and find higher volatility in the one-minute interval immediately following 1000-share transactions on SOES, but lower volatility over longer periods.

We compute two measures of price volatility in our experimental markets. The first measure is the standard deviation of transaction prices in a trading period (STDP), and the second measure is the price range (high minus low) in a trading period (PRANGE).¹⁹ Our analysis relating volatility to noise trading is similar to the one we

¹⁸ The Chow test looks at the stability of the estimated relation over two subsamples: low and high extremity. The log likelihood ratio statistic is based on the comparison of the restricted and unrestricted maximum of the log likelihood function. Both the restricted and unrestricted log likelihood numbers are obtained by estimating the regression using the whole sample. The restricted regression uses the original set of regressors (intercept, NOISETRD, and INFTRD), while the unrestricted regression adds a dummy variable for each high extremity security.

¹⁹ The daily price range is a measure of volatility that seems particularly relevant to day trading (e.g., Bernstein [1998]).

perform for the market-wide measures of informational efficiency in Section 3.3. For each trading period we compute STDP and PRANGE and record the levels of noise trading (NOISETRD) and informed trading (INFTRD). We then run OLS regressions of each volatility measure on the trading variables.

Panel A of Table 2 presents the results of the two regressions. The coefficient on NOISETRD in the regression on STDP is positive (0.0258) and statistically significant. Thus, trading periods in which noise trading is higher (holding constant the amount of informed trading) are associated with higher volatility of prices. A similar positive and highly statistically significant coefficient on noise trading is found in the regression on PRANGE. The coefficient on informed trading, INFTRD, is negative but not statistically significant in either regression.

Correlated errors may again be a concern, so we run separate regressions for each cohort and conduct tests on the coefficients from these regressions. Panel B of Table 2 reports the means and medians of the coefficients on NOISETRD and INFTRD together with the associated p -values for the tests. None of the tests show statistically significant results. While it is true that these tests utilize only eight observations (one for each cohort) and therefore constitute a rather high hurdle for finding statistical significance, the results in this panel caution us against interpreting too strongly the significant results in the pooled regressions.

In Panel C of Table 2 we look at how trading affects volatility separately in the two extremity levels. We use a dummy variable that is equal to zero for low extremity securities and one for high extremity securities, and present the two regression models in a full dummy variable format (without an intercept). We observe that neither noise nor informed trading seem to affect volatility in low extremity securities. On the other hand, noise trading has a positive coefficient and is highly statistically significant in both the STDP and PRANGE regressions for high extremity securities. This means that when prices are further from expected values, higher levels of noise trading are associated with

higher volatility. This finding seems to suggest that noise traders would exacerbate volatility in times of market stress. The coefficient on informed trading is negative in both regressions, but only in the regression on PRANGE is it statistically significant. We use Chow tests to look at the stability of the estimated relation over the two subsamples (low and high extremity). The tests performed on both regressions reject the null hypothesis of no structural change (with $p < 0.0001$ for STDP and $p = 0.0019$ for PRANGE).

Due to the lack of significance of the robustness tests in Panel B of Table 2, however, one should interpret the results of the pooled regressions with the dummy variables for extremity with caution. While greater noise trading seems to be associated with higher volatility at times when prices move a lot (and controlling for the trading of the informed traders who move the prices), these results cannot be viewed as a strong support for the notion that noise traders cause “excess” volatility.

4. Noise Trading and a Securities Transaction Tax

In the previous section we presented evidence that more noise trading is associated with lower informational efficiency of prices, and perhaps (though the evidence is weak) also with greater volatility. Curbing noise trading could therefore help make prices better reflect the fundamental values of securities, thereby improving risk sharing and the allocation of investment. Since there is no easy way outside of an experimental setting to identify noise traders, one way to reduce noise trading that was suggested in the past is a securities transaction tax (see Tobin [1978], Stiglitz [1989], and Summers and Summers [1989]).²⁰

Proponents of the tax claim that it will reduce stock market volatility and improve the efficiency of prices without harming liquidity. The intuition for why this would occur

²⁰ For a recent list of countries that impose a securities transaction tax see Pollin, Baker, and Schaberg [2002].

is articulated by Stiglitz [1989] who describes a world populated by informed investors, uninformed investors, and noise traders, very much like the one we constructed in the experiment. He conjectures that the tax is unlikely to discourage trading by the uninformed and informed traders, and therefore its effect will mainly be to drive out the noise traders. A survey of arguments on the potential costs and benefits of a securities transaction tax can be found in Schwert and Seguin [1993] and Pollin, Baker, and Schaberg [2002].²¹ Our ability to see who is trading in our experimental markets enables us to test Stiglitz's conjecture, as well as to investigate the effect of the tax on volume, liquidity, informational efficiency, and volatility.²²

4.1. Noise Traders' Strategies and Profit

The imposition of taxes affects noise traders' strategies in terms of liquidity provision. Panel A of Figure 5 looks at the submission rate of limit orders defined as the number of limit orders divided by the total orders traders submit. The submission rate of limit orders by noise traders declines under the tax regime, from 0.71 without taxes to 0.65 with taxes (p -value = 0.0792 for the tax*type interaction). With noise traders submitting fewer limit orders relative to market orders, they provide less liquidity to the market.

²¹ Two theoretical papers that examine the implications of securities transaction taxes are Subrahmanyam [1998] and Dow and Rahi [2000]. Subrahmanyam investigates a securities transaction tax in a model with asymmetrically informed traders (either strategic or competitive). In his framework, a transaction tax makes informed traders trade less. With strategic traders, the tax enhances liquidity if the informed trader is a monopolist, but worsens liquidity with multiple informed traders. With competitive informed traders, the tax tends to reduce liquidity under most parameter values. However, if markets are large and competitive but the mass of informed traders is small, liquidity can increase under the tax regime. Dow and Rahi present a model with informed speculators and uninformed investors with hedging motives, and also find that a transaction tax would make the informed traders trade less. The effects of taxes on the welfare of investors and the informativeness of prices depend on the values of different parameters in their model, but for some sets of parameters a securities transaction tax is pareto improving.

²² We will not be discussing other potential effects of the tax that were mentioned in the literature such as a reduction in "excess" production of private information, a move away from the short-term focus on the part of management, the equilibrium effects on expected returns and the cost of capital, distortions in capital structure, and the flight of trading to markets abroad, as well as issues related to implantation, compliance, and tax revenues.

Panel B of Figure 5 looks directly at the provision of depth at the best bid and offer prices (BBO) by different trader types. Taxes again affect trader types differently (p -value = 0.0214). Noise traders reduce their provision of liquidity at the BBO from 0.82 shares without taxes to 0.50 under the tax regime (this reduction is highly significant with simple effect p -value = 0.0066). On the other hand, informed traders and liquidity traders provide greater liquidity when taxes are imposed (1.02 and 1.23, respectively) compared to the no tax case (0.81 and 1.06, respectively).

As traders' strategies change with the imposition of the tax, so too will their returns. Table 3 presents gross and net profits by tax level, extremity level, and trader type. We observe that noise traders actually make money (45.49 dollars) when there are no transaction taxes and the value of the information of the informed traders is low. In fact, in this cell (no tax, low extremity) they make more money than do informed traders (who make 30.94 dollars). This probably reflects two effects. The first is profitable liquidity provision as we find that noise traders have a higher rate of limit order submission and greater provision of BBO depth in the no tax, low extremity case. Second, noise traders lose less when trading with the informed traders in low extremity securities because adverse selection is low. On the other hand, when informed traders have valuable information (in the high extremity securities), noise traders lose a lot of money, 210.34 dollars.

The situation changes markedly under the tax regime. Noise traders now lose in both low and high extremity securities. In the low extremity securities, we observed before that they provide much less liquidity and so do not realize market making profits. As such, they have gross trading losses of 69.64 and net trading losses (when taxes are considered) of 124.91 dollars. On the other hand, taxes seem to deter some of the speculative trading on trends that caused their losses in high extremity securities when there were no taxes in the economy. Under the tax regime, their gross trading loss is

90.02 dollars (compared with 210.34 dollars without taxes) leaving them with a net loss of 136.56 dollars after taxes are incorporated

Taxes also affect the profits of the informed traders. While they make gross trading profits even when their information is less valuable, their average net profit is just 0.8 dollars. When their information is more valuable in high extremity securities, informed traders seem to trade less aggressively on their information and their gross trading profit goes down from 275.37 (without taxes) to 161.73 (with taxes). This leaves them with only 97.90 dollars of net profit under the tax regime in the high extremity case.

These results suggest that taxes are not a pareto improving intervention in this economy (unlike the result for some parameter values in Dow and Rahi [2000]) as the informed traders make significantly less money. However, taxes do seem to be good for the noise traders under certain circumstances. In particular, when the realized value is further away from the expected value and the market is characterized by more volatile prices and higher adverse selection, taxes deter speculative trading by noise traders and so reduce their losses.

4.2. *Volume and Liquidity*

There are reasons to believe that the imposition of taxes will lower overall volume in the market. Constantinides [1986] argues theoretically how an increase in transaction costs reduces volume, and Umlauf [1993] finds empirically that the imposition of a securities transaction tax in Sweden reduced trading volume. We find this to be the case in our experiment as well. The average number of shares traded per capita in our experiment goes down from 46 shares without taxes to 35.44 shares under the tax regime (p -value = 0.0013). The most marked decrease is in the trading of the noise traders, going down from 39.97 shares without taxes to 25.45 shares with taxes (p -value = 0.0403).

Informed traders also decrease trading, from 47.96 to 40.92 shares, although this decrease is not statistically significant.²³

While both proponents and opponents of securities transaction taxes agree that volume is likely to decline, there is no consensus on the effects of taxes on liquidity. Lower volume (or frequency of transactions) may increase the inventory costs of liquidity providers in the spirit of Amihud and Mendelson [1980]. On the other hand, if reduced noise trading makes prices less volatile, the risk of holding shares would go down (Stoll [1978]). Therefore, inventory holding costs may go either up or down. As for adverse selection costs, these could go up because the probability of informed trading is higher when noise traders reduce their trading by a much larger amount than do informed traders (e.g., Glosten and Milgrom [1985], Easley and O'Hara [1992]). On the other hand, if strategic informed traders trade less aggressively on their information due to decreased noise trading (in the spirit of Kyle [1985]), adverse selection costs would not change.

We measure liquidity here as the total price impact of market orders, which is equivalent to the effective spread used in empirical studies. We also borrow from the empirical market microstructure literature and decompose the total price impact into a permanent price impact and a temporary price impact. The total, permanent, and temporary price impacts for a market order at time t that is characterized by a transaction price P_t and quote midpoint M_t are computed as follow:

$$\begin{aligned} \text{Total Price Impact} &= \begin{cases} P_t - M_t & \text{for market buy orders} \\ M_t - P_t & \text{for market sell orders} \end{cases} \\ \text{Permanent Price Impact} &= \begin{cases} M_{t+5} - M_t & \text{for market buy orders} \\ M_t - M_{t+5} & \text{for market sell orders} \end{cases} \\ \text{Temporary Price Impact} &= \begin{cases} P_t - M_{t+5} & \text{for market buy orders} \\ M_{t+5} - P_t & \text{for market sell orders} \end{cases} \end{aligned}$$

²³ The slight decrease in the level of informed trading that we observe could be consistent with the theoretical models of Subrahmanyam [1998] and Dow and Rahi [2000], but there is too much variability in the data to determine that for sure.

where $t+5$ denotes the midquote prevailing five trades after the trade for which we are computing the price impact measures

Note that these definitions mean that the total price impact is the sum of the permanent and temporary price impacts. The total price impact (effective spread) measures the transaction costs incurred by trader who submits a market order. The permanent price impact of a market order is viewed as a manifestation of the private information content of the order. The temporary price impact (also known in the empirical literature as the realized spread) is viewed as a measure of efficiency of the liquidity provision mechanism. It provides some information about the profitability of liquidity provision because it measures the price reversal from the trade price.

The total price impact of market orders is significantly lower under the tax regime: 1.93 dollars compared with 2.48 dollars without taxes (p -value = 0.0085). While the permanent effect seems to be a bit lower with taxes, the effect is not statistically significant. The large decrease in the total price impact comes from a large decrease in the temporary price impact from 2.17 dollars without taxes to 1.60 dollars with taxes (p -value = 0.0031). Liquidity, as manifested by the price impact of a market order, seems to be significantly better under the tax regime and the lower price impact is not due to a lower information content of market orders but rather to more efficient (and less expensive) liquidity provision, or in other orders, less price reversal.

There is, however, an interesting and statistically significant interaction with extremity of the realized value of the security (p -value = 0.0310 for total impact and p -value = 0.0437 for the temporary impact). Figure 6 shows the permanent, temporary, and total (the sum of the two) impacts separately for high and low extremity securities both with and without taxes. The total price impact goes down a bit from 2.22 to 2.00 dollars for low extremity securities when we impose taxes. This decrease is to a large extent due to a decrease in the temporary impact from 1.95 to 1.65 dollars. Taxes, however, seem to affect liquidity for high extremity securities in a much more drastic fashion. The

temporary price impact is 2.40 dollars without taxes and only 1.54 dollars with taxes. This is reflected in a total transaction cost to demanders of liquidity that goes down from 2.75 to 1.87 dollars when taxes are imposed.

The decrease in transaction costs of market orders despite the fact that overall volume in the market goes down is not a trivial result. The fact that it comes mainly from the temporary price impact rather than the permanent price impact points to an explanation in the spirit of inventory control or the provision of liquidity by risk averse investors (see Grossman and Miller [1988]), rather than adverse selection. We observe that noise traders supply less liquidity but informed and liquidity traders seem to increase their presence at the BBO when taxes are imposed. Our profit analysis shows that noise traders are not good at trading on trends when markets are more volatile as in the case of high extremity securities. It is likely that their trading creates overshooting and reversal of prices in the absence of taxes, and since taxes cause them to curb their trading in high extremity securities, prices tend to exhibit smaller reversals and therefore a lower temporary price impact of trades. Hence, our analysis demonstrates that while taxes result in lower volume, their impact on liquidity can be positive.

4.3. *Value Discovery and Informational Efficiency*

We use the measures introduced in Section 3.3 to investigate the effects of taxes on value discovery and informational efficiency. Figure 7 shows that taxes significantly affect how market orders of different trader types contribute to (or interfere with) the adjustment of prices toward their full-information level (p -value = 0.0369 for the tax*type interaction). It seems as if a transaction tax affects the value discovery function of market orders used by informed traders and liquidity traders much more than it does for the noise traders. The most striking result is that informed traders' market orders contribute much less to the adjustment of prices toward the true value under the tax regime, 0.026, than in the case without taxes, 0.214. On the other hand, liquidity traders

are interfering much less with the adjustment of prices under the tax regime (-0.086) as opposed to their impact on prices when taxes are not imposed (-0.279).

To examine the overall effect of these changes on informational efficiency we examine the market-wide pricing error measures, DEVP (deviations of transaction prices from the true value) and DEVMID (deviations of midquotes prevailing at the time of transactions from the true value). Both measures seem to be slightly lower under the tax regime: 13.76 for DEVP and 13.55 for DEVMID compared with 14.65 and 14.15 (respectively) when there are no taxes in the economy. However, these differences are not statistically significant, indicating that there may be little change in informational efficiency when taxes are imposed.

To examine whether taxes influence the manner in which informed and noise trading affect informational efficiency we run regressions similar in spirit to those in Section 3.3. Specifically, we regress the informational efficiency measures on the amount of noise trading and informed trading with dummy variables for tax and no tax securities. We then use structural stability tests to examine the hypothesis that taxes have no influence on the manner in which noise and informed trading affect informational efficiency.

Table 4 provides the results of these regressions. We observe that in the regression on DEVP, greater amount of noise trading is associated with higher pricing errors when there are no taxes in the economy. On the other hand, the coefficient on noise trading is not statistically significant under the tax regime. The Chow test for structural stability rejects the null hypothesis of no structural change ($p = 0.0008$). A similar picture is also observed in the regression on DEVMID, and the Chow test is significant for this regression as well.

These regressions suggest that while taxes may not necessarily lower pricing errors, they may weaken the link between noise trading and informational efficiency. When taxes are imposed, noise traders trade less and therefore do not affect the market to

the same extent. As such, their trading may have a less reliable effect on the informational efficiency of prices. While the magnitude of the coefficient on noise trading is similar with and without taxes, the loss of statistical significance of the coefficient under the tax regime, together with the significant Chow tests, suggest that noise traders' impact on pricing errors diminishes.

4.4. *Volatility*

Probably the number one goal of a securities transaction tax in the eyes of its proponents is a reduction in the volatility of prices (Stiglitz [1989], Summers and Summers [1989]). While part of the volatility of prices reflects adjustment to fundamental information, transaction taxes are supposed to rid prices of "excess" volatility that is due to noise trading.²⁴ However, there is no consensus on whether a transaction tax would in fact reduce volatility. Kupiec [1996], for example, shows that the volatilities of prices and returns respond in opposite ways to the imposition of a securities transaction tax.

We find that the standard deviation of transaction prices (STDP) without taxes is 4.07 and with taxes it is 3.53. Similarly, the transaction price range (PRAGNE) is 20.21 dollars when no taxes are imposed and 16.84 dollars under the tax regime. However, there is enough variability in the data so that these differences are not statistically significant. Also, in the regressions of STDP and PRANGE on the amount of noise and informed trading, the structural stability tests cannot reject the hypothesis of no structural change with and without taxes. This means that taxes do not seem to affect the manner in which noise trading and informed trading are associated with volatility. Therefore, we cannot conclude that a securities transaction tax can be used to lower price volatility even

²⁴ De Long, Shleifer, Summers, and Waldmann [1990] show how noise traders can introduce volatility into prices beyond the volatility that is due to fundamental risk.

though we document that more noise trading is associated with higher volatility and that taxes make noise traders trade less.

4.5 *Dynamics of Profits and Trading*

Finally, we turn to the interesting question of the “survival” of noise traders in financial markets, or whether noise traders will consistently lose money and therefore be driven out of the market. If Friedman’s wealth dynamic arguments are correct, we would expect to find that noise traders who lose more money subsequently trade less. We therefore record for all noise traders their net profit and amount of trading in each trading period and run a bivariate vector autoregression (VAR) of profit and trading separately for markets with and without taxes.²⁵

Panel A of Table 5 presents the results for a VAR with two lags in the case of no taxes.²⁶ Both coefficients on past profit in the profit equation are positive and statistically significant, indicating that there is some consistency (or skill) involved in trading: traders who lose more in the past are also likely to lose more in the future. On the other hand, the two profit coefficients in the NOISETRD equation are not statistically significant, and the Granger Causality test for profit is also not significant. This means that we cannot find evidence that past profit affects future trading. The coefficients on past trading, however, are highly statistically significant in the NOISETRD equation, providing some evidence that traders adopt a “style” of trading—some simply trade more while others trade less.

Panel B of Table 5 presents a similar VAR applied to the securities that are subject to a transaction tax. The most notable difference is that the coefficient on the first lag of profit in the NOISETRD equation is positive and statistically significant.

²⁵ We pool together the 12 consecutive securities that trade without taxes from each of the eight cohorts into one VAR, and the 12 consecutive securities that trade with taxes from each of the eight cohorts into another VAR.

²⁶ Similar results are obtained if an exogenous regressor is added to the model to pick up learning due to experience with the experiment (i.e., a time variable). Also, we looked at different lag structures (up to five lags) and our conclusions are unchanged.

Furthermore, the Granger Causality test for profit is significant ($p = 0.0422$) indicating that profit Granger-cause trading.²⁷ This result means that noise traders who lose more money in a trading period are likely to trade less in the following period. This effect is consistent with the idea that noise traders will eventually be driven out of the market (or at least trade less), but we observe a significant effect only under the tax regime.

It is possible that the transaction tax serves as an “awareness” device in that it makes noise traders pay more attention to the relation between trading and losses. We documented that noise traders trade less under the tax regime, presumably for the same reason, and here we also see that they take note of losses by cutting down further on their future trading. Our experiment, therefore, suggests another mechanism by which a securities transaction tax may affect noise trading that has not been mentioned previously in the literature. It is not just that the higher costs deter trading because some short-term speculative strategies are no longer profitable when taxes are considered. It could also be the case that the tax forces noise traders to keep track of their trading and profits more vigilantly. It therefore serves to educate them and draw their attention to the relation between their trading and the resulting losses, and by that makes noise traders realize that they trade on noise, and therefore cause them to reduce trading or exit the market.

5. The Limits of Noise Trading

Noise trading is a contentious but important issue in the study of asset markets. We have analyzed the role of noise trading in an experimental market with noise traders, informed traders, and liquidity traders with exogenous trading needs. Our purpose was to delineate both how noise traders fare in markets and how their trading affects market performance, and to these ends we have had some success. Our analysis has provided a

²⁷ The Granger Causality test on NOISETRD is also significant, and so the relation goes in both directions.

number of results on the role and behavior of noise traders. In this final section, we now consider these results in the context of their implications for market behavior. We also discuss the limitations of our analysis, with particular attention given to the aspects of noise trader and market behavior we are not able to examine here.

Perhaps the most important implication of our research is that noise traders, on balance, are not particularly beneficial to market performance. Noise traders do not make prices more efficient, they tend to increase market volatility, and there is little evidence that noise traders provide liquidity to other traders. This view of noise trading is very much at odds with positive declarations that noise traders are the “grease that makes markets work.” Our analysis, instead, suggests that markets work in spite of, and not because of, noise traders.

Proponents of noise trading might argue that our analysis misses an important aspect of noise trading: the more noise traders willing to lose money, the greater benefit to informed trading, and so more information production takes place. While aspects of this logic are correct, our results show that it is too simplistic. Noise trading is not a simple scaling factor. Noise traders change the price process in complex ways: Volatility is increased, prices can overshoot true values, and liquidity production is affected. These changes benefit some traders, but also penalize them as well. While our setting does not explicitly allow for changes in the numbers of informed traders, it does allow informed trading to adjust to market conditions. Our findings on the impact of the securities transaction tax illustrate that the linkage between informed trading and noise trading is complex.

A second significant implication of our results is that noise traders do not prosper in the short-run and they would not thrive (or survive) in the long run, especially in the presence of securities transaction taxes. In our experiments, noise traders eschew the potentially profitable role of liquidity provider and instead concentrate on trend-chasing. DSSW [1990; 1991] argue that such behavior could allow for survival because traders

taking on such risks could potentially earn a return from doing so. But the noise trader analysis of DSSW is limited in some very important ways. For example, DSSW [1991] do not allow noise traders to affect prices; instead, rational traders are assumed to be price-setters and noise traders then free-ride, as it were, on the price process. In actual markets, as in our experimental markets here, all traders affect prices, and when they do, they bear both the costs and benefits of their actions.

Another limitation of DSSW [1991] is that they do not allow for differential information. Unlike the informed traders allowed here, DSSW consider a world in which rational traders have correct beliefs about the distribution of returns on assets, while noise traders do not. This framework is an interesting one for many questions, but it does not incorporate the complexities introduced when some traders (the informed) have better information than other traders do. This latter framework is the basis for microstructure analyses of price formation, and as we show here allowing asymmetric information greatly reduces the viability and survivability of noise traders. The noise trader risk in our markets appears to be greatest for the noise traders themselves.

Finally, we turn to our results on the impact of a securities transactions tax. We find that the STT has many useful effects on markets, although reducing volatility is not one of them. If the goal of such a tax is to limit noise trading, then our results provide some backing for this approach. But we caution here that our analysis is not general enough to provide overall conclusions on securities transaction taxes. We do not, for example, allow such taxes to influence portfolio-based liquidity traders' participation decisions, although such effects are likely to be significant. Should liquidity traders opt not to participate in markets with securities transaction taxes, risk sharing and allocational efficiency in the economy would suffer, and market performance would almost surely be degraded.²⁸ Indeed, this difficulty underscores a general concern with

²⁸ For a discussion of the issues surrounding participation effects and securities transaction taxes see O'Hara [2004].

the unfocused nature of such taxes; while their goal may be to affect particular traders, the incidence of the tax falls on every market participant. More targeted approaches, such as the SEC's minimum wealth requirements for day traders, may be a more effective strategy to limit particular types of trading.

What, then, to conclude about the role of noise trading? We believe our results establish that noise trading is not a panacea for market ills. Increasing noise trading in markets would not improve the efficiency of markets, nor would it necessarily improve the provision of liquidity. Conversely, decreasing noise trading could potentially reduce the complexity of the price formation process and perhaps even lower volatility. This is not to say, however, that the optimum level of noise trading is zero. Noise traders trade for a variety of reasons, including perhaps that they simply like to trade. Noise traders will enter markets drawn by these motivations, and they will exit markets due to the losses they sustain. Market forces influence both events, and thus ultimately the market itself determines the limits of noise trading.

References

- Amihud, Y. and H. Mendelson, 1980, Dealership Market: Market-Making with Inventory, Journal of Financial Economics, 8, 31-53.
- Amihud, Y. and H. Mendelson, 1992, Transaction Taxes and Stock Values, in Modernizing U.S. Securities Regulation, ed. K. Lehn and R. Kamphius (Irwin: NY, NY)
- Atkins, A. and E. Dyl, 1997, Stock Price Volatility, Transactions Costs and Securities Transactions Taxes, Managerial and Decision Economics, 18, 709-718.
- Barber, B., Lee, Y-T, Liu, Y-J, and T. Odean, 2004, Do Individual Day Traders Make Money? Evidence from Taiwan, Working Paper, U.C. Davis.
- Battalio, R. H., Hatch, B. and R. Jennings, 1997, SOES Trading and Market Volatility, Journal of Financial and Quantitative Analysis, 32(2), 225-238.
- Bernstein, J., 1998, The Compleat Day Trader II, McGraw-Hill.
- Black, F., 1986, Noise, Journal of Finance, 41(3), 529-543.
- Bloomfield, R. and M. O'Hara, 1998, Does Order Preferecing Matter? Journal of Financial Economics, 50, 3-37.
- Bloomfield, R. and M. O'Hara, 2000, Can Transparent Markets Survive? Journal of Financial Economics, 55, 425-459.
- Bloomfield, R., O'Hara, M. and G. Saar, 2004, The Make-or-Take decision in an Electronic Market: Evidence on the Evolution of Liquidity," Journal of Financial Economics, 75, 165-199.
- Blume, L. and D. Easley, 1992, Evolution and Market Behavior, Journal of Economic Theory, 58(1), 9-40.
- Blume, L. and D. Easley, 2002, If You're so Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets, Econometrica, forthcoming.
- Cason, T. N., 2000, The Opportunity for Conspiracy in Asset Markets Organized with Dealer Intermediaries, Review of Financial Studies, 13, 385-416.
- Constantinides, G. M., 1986, Capital Market Equilibrium with Transaction Costs, Journal of Political Economy, 94(4), 842-862.
- DeLong, B., Shleifer, A., Summers, L., and R. Waldman, 1990, Noise Trader Risk in Financial Markets, Journal of Political Economy, 98(4), 703-738.

- DeLong, B., Shleifer, A., Summers, L., and R. Waldman, 1991, The Survival of Noise Traders in Financial Markets, Journal of Business, 64(1), 1-19.
- Dow, J., and R. Rohit, 2000, Should Speculators be Taxed? Journal of Business, 73(1), 89-107.
- Easley, D. and M. O'Hara, 1992, Time and the Process of Security Price Adjustment, Journal of Finance, 47.
- Fama, E. and J. MacBeth, 1973, Risk, Return and Equilibrium: Empirical Tests, Journal of Political Economy, 71, 607-636.
- Friedman, M., 1953, Essays in Positive Economics, University of Chicago Press, Chicago.
- Garvey, R. and A. Murphy, 2001, How Profitable Day Traders Trade: An Examination of Trading Profits, working paper, University College Dublin.
- Glosten, L. R. and P. Milgrom, 1985, Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders, Journal of Financial Economics, 14, 71-100.
- Grossman, S. J. and M. H. Miller, 1988, Liquidity and Market Structure, Journal of Finance, 43, 617-633.
- Hamilton, J. D., 1994, Time Series Analysis, Princeton University Press, Princeton, New Jersey.
- Harris, J. H. and P. H. Schultz, 1998, The Trading Profits of SOES Bandits, Journal of Financial Economics, 50, 39-62.
- Jordan, D. J. and J. D. Diltz, 2003, The Profitability of Day Traders, Financial Analysts Journal, 59(6), 85-94.
- Kyle, A. S., 1985, Continuous Auctions and Insider Trading, Econometrica, 53(6), 1315-1335.
- Kupiec, P. H., 1996, Noise Traders, Excess Volatility, and a Securities Transactions Tax, Journal of Financial Services Research, 10(2), 115-129.
- Lamoreaux, C. and C. Schnitzlein, 1997, When It's Not the Only Game in Town: The Effect of Bilateral Search on the Quality of a Dealer Market, Journal of Finance, 52, 683-712.
- Linnainmaa, J., 2003, The Anatomy of Day Traders, working paper, UCLA.

- Lynch-Koski, J., Rice, E. M. and A. Tarhouni, 2004, Noise Trading and Volatility: Evidence from Day Trading and Message Boards, Working paper, University of Washington.
- Odean, T., 1999, Do Investors Trade Too Much? American Economic Review, 89(5) 1279-1298.
- O'Hara, M., 2004, Liquidity and Price Stability, National Bank of Belgium Working Paper, June 2004.
- Pollin, R., Baker, D. and M. Schaberg, 2002, Securities Transaction Taxes for U.S. Financial Markets, Working paper, University of Massachusetts.
- Sandroni, A., 2000, Do Markets Favor Agents Able to Make Accurate Predictions? Econometrica, 68 (6), 1303-42.
- Scheinkman, J. and W. Xiong, 2003, Overconfidence and Speculative Bubbles, Journal of Political Economy, 111(6), 1183-1219.
- Schwert, G.W. and P.J. Seguin, 1993, Securities Transactions Taxes: An Overview of Costs, Benefits, and Unresolved Questions," Financial Analysts Journal, September-October 1983, 27-35.
- Sciubba, E., 1999, Asymmetric Information and Survival in Financial Markets, University of Cambridge DAE Working Paper, n. 9908.
- Shefrin, H. and M. Statman, 1994, Behavioral Capital Asset Pricing Theory, Journal of Financial and Quantitative Analysis, 29(3), 323-349.
- Shleifer, A., and R. Vishny, 1997, The Limits of Arbitrage, Journal of Finance 52, 35-55.
- Shleifer, A., and L. Summers, 1990, The Noise Trader Approach to Finance, Journal of Economic Perspectives, 4(2), 19-33.
- Stiglitz, J. E., 1989, Using Tax Policy to Curb Speculative Short-Term Trading, Journal of Financial Services Research, 3, 101-115.
- Stoll, H. R., 1978, The Supply of Dealer Services in Securities Markets, Journal of Finance, 33(4), 1133-1151.
- Summers, L. and V. P. Summers, 1989, When Financial Markets Work Too Well: A Cautious Case for a Securities Transaction Tax, Journal of Financial Services Research, (3), 261-286.
- Subrahmanyam, A., 1998, Transactions Taxes and Financial Market Equilibrium, Journal of Business, 71(1), 81-118.
- Tobin, J., 1978, A Proposal for Monetary Reform, Eastern Economic Journal, 4, 153-159.

Umlauf, S., 1993, Transactions Taxes and Stock Market Behavior: The Swedish Experience, Journal of Financial Economics, 33, 227-240.

Table 1
Informational Efficiency Regressions

This table presents a regression analysis of pricing errors. For each security (i.e., for each trading period), we compute two definitions of pricing errors. The first definition (DEVP) is the absolute value of the difference between true value of the security and the transaction price, averaged across all trades in a security. The second definition (DEVMID) is the absolute value of the difference between the true value of the security and the midpoint between the limit order book's best bid and offer prevailing at the time of the transaction, averaged across all trades in a security. We also record the number of trades (sum of market orders and executed limit orders) of the noise traders (NOISETRD) and the informed traders (INFTRD). We then run the following OLS regression:

$$DEVP_i = a + b * NOISETRD_i + c * INFTRD_i + Error_i$$

and a similar specification where DEVMID is the dependent variable. Panel A presents the results of these regressions using all securities and all cohorts (24 securities times 8 cohorts). In Panel B we provide robustness analysis on the results of the above regression. We first run a separate regression for each of the eight cohorts (over the 24 securities each cohort traded). We then report the mean and median of the coefficients from these eight regressions alongside a t-test and a Wilcoxon test against the hypothesis of a zero mean/median. This approach allows us to draw inferences from across cohorts, which can be assumed independent. In Panel C we introduce a dummy variable (D) that is equal to zero for low extremity securities and one for high extremity securities, where low (high) extremity means that the realized value of the security is no more than (at least) \$15 from its unconditional value. We interact the dummy variable with the intercept, NOISETRD, and INFTRD, and run a regression using all securities from the eight cohorts. We present the full dummy variable format (without an intercept) to look at the relationship between pricing errors and trading by noise and informed traders separately in each extremity level.

Panel A: Regressions of Pricing Errors on Noise and Informed Trading

Dependent Variable	a Intercept (<i>p</i> -value)	b NOISETRD (<i>p</i> -value)	c INFTRD (<i>p</i> -value)	R² (in %)	No. of Observations
DEVP	13.9023 (0.0000)	0.0953 (0.0016)	-0.0762 (0.0003)	7.37	192
DEVMID	13.8454 (0.0000)	0.0809 (0.0085)	-0.0724 (0.0008)	6.09	192

Panel B: Robustness Analysis of Pricing Errors Regressions

Dependent Variable	NOISETRD		INFTRD	
	Mean (<i>p</i> -value from t-test)	Median (<i>p</i> -value from Wilcoxon test)	Mean (<i>p</i> -value from t-test)	Median (<i>p</i> -value from Wilcoxon test)
DEVP	0.1151 (0.0198)	0.0965 (0.0300)	-0.0762 (0.0255)	-0.0687 (0.0423)
DEVMID	0.1083 (0.0334)	0.0825 (0.0423)	-0.0742 (0.0355)	-0.0582 (0.0423)

Panel C: Pricing Error Regressions with an Extremity Interaction

Dependent Variable	Low Extremity			High Extremity			R ² (in %)	Obs.
	Intercept *(1-D) (<i>p</i> -value)	NOISETRD *(1-D) (<i>p</i> -value)	INFTRD *(1-D) (<i>p</i> -value)	Intercept *D (<i>p</i> -value)	NOISETRD *D (<i>p</i> -value)	INFTRD *D (<i>p</i> -value)		
DEVP	8.1810 (0.0000)	0.0965 (0.0230)	-0.0300 (0.2223)	20.9426 (0.0000)	0.0607 (0.0716)	-0.1541 (0.0003)	30.09	192
DEV MID	7.5505 (0.0000)	0.0914 (0.0312)	-0.0258 (0.2935)	21.7042 (0.0000)	0.0380 (0.2588)	-0.1597 (0.0002)	32.09	192

Table 2
Volatility Regressions

This table presents a regression analysis of price volatility measures. For each security (i.e., for each trading period), we compute two definitions of volatility. The first definition (STDP) is the standard deviation of transaction prices in a security. The second definition (PRANGE) is the price range (high minus low) in a security. We also record the number of trades (sum of market orders and executed limit orders) of the noise traders (NOISETRD) and the informed traders (INFTRD). We then run the following OLS regression:

$$STDP_i = a + b * NOISETRD_i + c * INFTRD_i + Error_i$$

and a similar specification where PRANGE is the dependent variable. Panel A presents the results of these regressions using all securities and all cohorts (24 securities times 8 cohorts). In Panel B we provide robustness analysis on the results of the above regression. We first run a separate regression for each of the eight cohorts (over the 24 securities each cohort traded). We then report the mean and median of the coefficients from these eight regressions alongside a t-test and a Wilcoxon test against the hypothesis of a zero mean/median. This approach allows us to draw inferences from across cohorts, which can be assumed independent. In Panel C we introduce a dummy variable (D) that is equal to zero for low extremity securities and one for high extremity securities, where low (high) extremity means that the realized value of the security is no more than (at least) \$15 from its unconditional value. We interact the dummy variable with the intercept, NOISETRD, and INFTRD, and run a regression using all securities from the eight cohorts. We present the full dummy variable format (without an intercept) to look at the relationship between volatility and trading by noise and informed traders separately in each extremity level.

Panel A: Regressions of Volatility Measures on Noise and Informed Trading

Dependent Variable	a Intercept (<i>p</i> -value)	b NOISETRD (<i>p</i> -value)	c INFTRD (<i>p</i> -value)	R² (in %)	No. of Observations
STDP	3.4920 (0.0000)	0.0258 (0.0165)	-0.0119 (0.1094)	3.02	192
PRANGE	16.4860 (0.0000)	0.1155 (0.0042)	-0.0417 (0.1349)	4.30	192

Panel B: Robustness Analysis of Volatility Regressions

Dependent Variable	NOISETRD		INFTRD	
	Mean (<i>p</i> -value from t-test)	Median (<i>p</i> -value from Wilcoxon test)	Mean (<i>p</i> -value from t-test)	Median (<i>p</i> -value from Wilcoxon test)
STDP	-0.0058 (0.7360)	-0.0063 (0.6241)	-0.0087 (0.4667)	-0.0095 (0.4412)
PRANGE	0.0132 (0.8498)	-0.0095 (0.9442)	-0.0254 (0.5076)	-0.0798 (0.6241)

Panel C: Volatility Regressions with an Extremity Interaction

Dependent Variable	Low Extremity			High Extremity			R ² (in %)	Obs.
	Intercept *(1-D) (<i>p</i> -value)	NOISETRD *(1-D) (<i>p</i> -value)	INFTRD *(1-D) (<i>p</i> -value)	Intercept *D (<i>p</i> -value)	NOISETRD *D (<i>p</i> -value)	INFTRD *D (<i>p</i> -value)		
STDP	3.5003 (0.0000)	0.0045 (0.7945)	-0.0008 (0.9402)	3.5572 (0.0000)	0.0380 (0.0063)	-0.0208 (0.2245)	4.50	192
PRANGE	16.5534 (0.0000)	0.0394 (0.5395)	0.0071 (0.8500)	18.8406 (0.0000)	0.1641 (0.0015)	-0.1534 (0.0164)	6.83	192

Table 3
Trader Profit with and without Taxes

This table reports gross and net profits of the traders by tax regime and extremity level. In the setting with a securities transaction tax (STT), we impose a \$2 fee for each trade. In a no-tax setting, no fees are imposed. Low (high) extremity means that the realized value of the security is no more than (at least) \$15 from its unconditional value. Gross profit is the trading profit a trader makes before taxes. Net profit is gross profit minus the STT paid by the trader. We first compute each profit definition for an individual trader per security, and then average it for a trader type within each of the eight cohorts. The numbers in the table represent the averages across the cohorts.

Tax Regime	Extremity Level	Gross Profit			Net Profit		
		Noise	Informed	Liquidity	Noise	Informed	Liquidity
No Tax	Low	45.49	30.94	-76.43	45.49	30.94	-76.43
	High	-210.34	275.37	-65.04	-210.34	275.37	-65.04
Tax	Low	-69.64	100.64	-31.00	-124.91	0.80	-110.28
	High	-90.02	161.73	-71.70	-136.56	97.90	-152.21

Table 4
Informational Efficiency Regressions with Dummy Variable for Tax

This table presents a regression analysis of pricing errors, which we use as measures of informational efficiency, by tax regime. For each security (i.e., for each trading period), we compute two definitions of pricing errors. The first definition (DEVP) is the absolute value of the difference between true value of the security and the transaction price, averaged across all trades in a security. The second definition (DEVMIID) is the absolute value of the difference between the true value of the security and the midpoint between the limit order book's best bid and offer prevailing at the time of the transaction, averaged across all trades in a security. We also record the number of trades (sum of market orders and executed limit orders) of the noise traders (NOISETRD) and the informed traders (INFTRD). In the setting with a securities transaction tax (STT), we impose a \$2 fee for each trade. In a no-tax setting, no fees are imposed. We define a dummy variable (D) that is equal to zero for trading periods without taxes and one for trading periods when the STT is imposed. We then run an OLS regression using all securities from the eight cohorts of the pricing error measure on an intercept, NOISETRD, and INFTRD interacted with the tax dummy variable. We present the full dummy variable format (without an intercept) to look at the relationship between pricing errors and trading by noise and informed traders separately in each tax regime.

Dependent Variable	No Tax			Tax			R ² (in %)	Obs.
	Intercept *(1-D) (p-value)	NOISETRD *(1-D) (p-value)	INFTRD *(1-D) (p-value)	Intercept *D (p-value)	NOISETRD *D (p-value)	INFTRD *D (p-value)		
DEVP	13.0791 (0.0000)	0.0912 (0.0128)	-0.0560 (0.0336)	15.4599 (0.0000)	0.0889 (0.1494)	-0.1097 (0.0018)	8.56	192
DEVMIID	12.9997 (0.0000)	0.0749 (0.0448)	-0.0503 (0.0611)	15.4227 (0.0000)	0.0797 (0.2059)	-0.1094 (0.0024)	7.36	192

Table 5
Dynamics between Profit and Noise Trading with and without Taxes

This table presents a vector autoregression (VAR) analysis of trading and net profits of the noise traders. In the setting with a securities transaction tax (STT), we impose a \$2 fee for each trade. In a no-tax setting, no fees are imposed. NetProfit is defined as the trading profit of an individual noise trader in a security minus the STT paid by the trader. We also record the number of trades (sum of market orders and executed limit orders) of the individual noise trader (NOISETRD) in a security. In Panel A, we use the 12 consecutive securities in each cohort where traders did not pay taxes, and run a bivariate VAR with two lags over 320 observations: 4 (noise traders) * 12 (securities) * 8 (cohorts), minus 64 observations due to the lag structure. In Panel B we run a similar bivariate VAR for the 12 consecutive securities in each cohort where taxes were imposed. The VAR is estimated using OLS.

Panel A: VAR of Noise Traders' Net Profit and Trading without Tax

Dependent Variables	NetProfit_{t-1} (<i>p</i> -value)	NetProfit_{t-2} (<i>p</i> -value)	NOISETRD_{t-1} (<i>p</i> -value)	NOISETRD_{t-2} (<i>p</i> -value)	Intercept (<i>p</i> -value)	R² (in %)
NetProfit_t	0.1898 (0.0023)	0.2421 (0.0001)	-3.6007 (0.0069)	-1.1220 (0.4166)	144.5022 (0.0368)	23.71
NOISETRD_t	-0.0008 (0.7797)	0.0007 (0.8186)	0.3109 (0.0000)	0.4561 (0.0000)	11.8577 (0.0002)	40.42

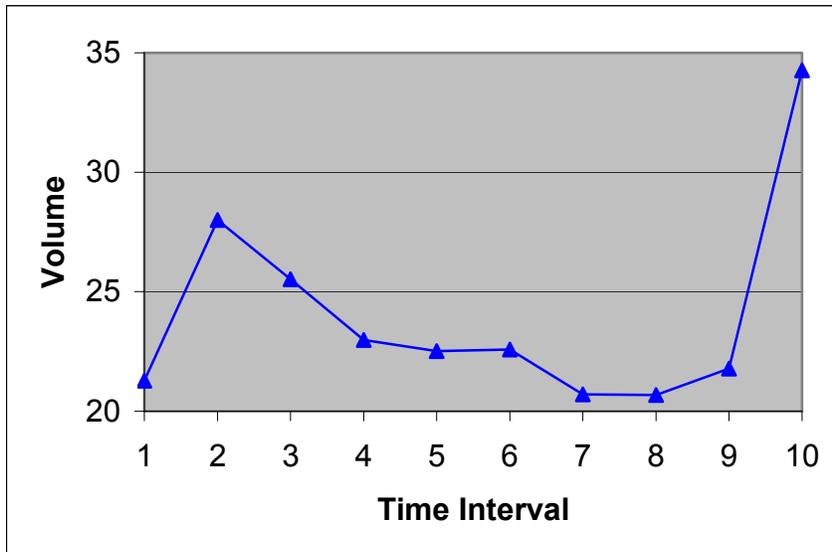
Panel B: VAR of Noise Traders' Net Profit and Trading with Tax

Dependent Variables	NetProfit_{t-1} (<i>p</i> -value)	NetProfit_{t-2} (<i>p</i> -value)	NOISETRD_{t-1} (<i>p</i> -value)	NOISETRD_{t-2} (<i>p</i> -value)	Intercept (<i>p</i> -value)	R² (in %)
NetProfit_t	0.0171 (0.7855)	0.1642 (0.0154)	-4.1581 (0.0003)	-1.5642 (0.2296)	32.5548 (0.3946)	14.20
NOISETRD_t	0.0072 (0.0230)	-0.0040 (0.2313)	0.3851 (0.0000)	0.5105 (0.0000)	5.1826 (0.0071)	47.98

Figure 1
Market-wide Summary Statistics

This figure presents summary statistics for volume and bid-ask spreads over time. Volume is defined as the number of shares traded. Bid-ask spread is the difference between the best price in the book for buying a share and the best price in the book for selling a share. The variables are computed separately for each 15-second interval in a trading period, and then averaged across all trading periods (i.e., securities) for the eight cohorts.

Panel A: Volume



Panel B: Bid-Ask Spreads

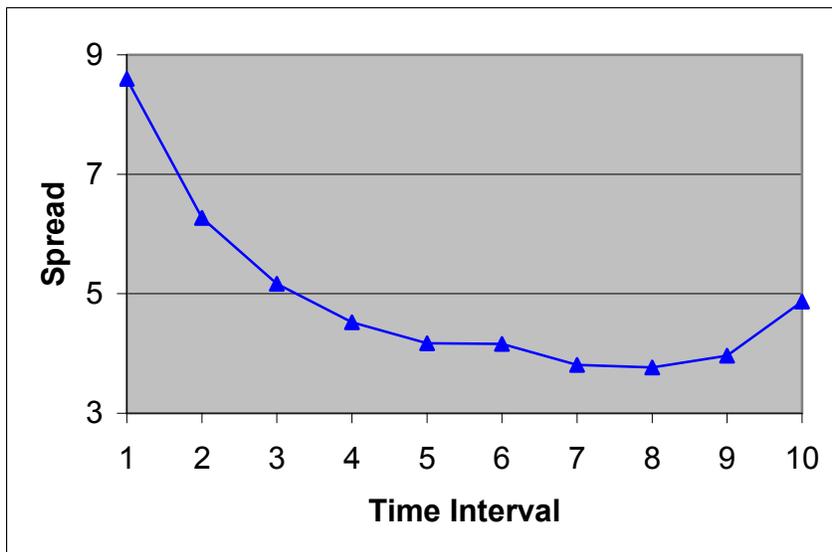
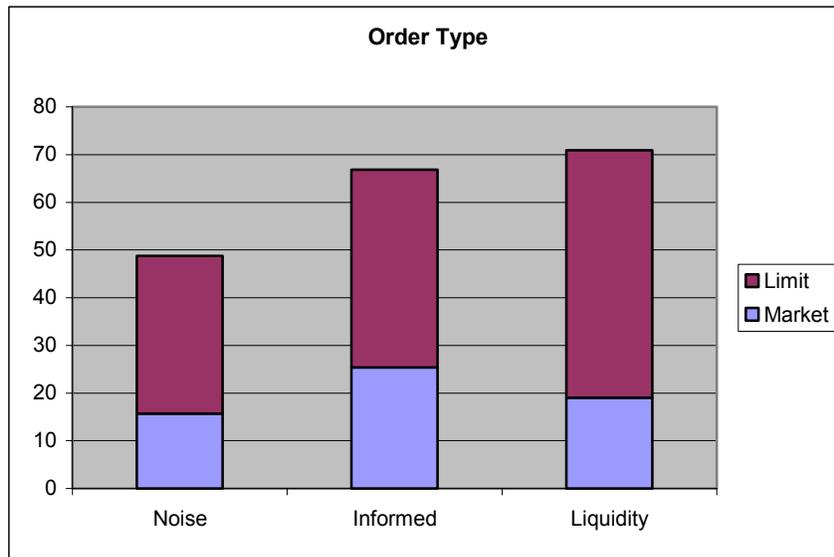


Figure 2
Orders and Trades by Trader Type

This figure presents summary statistics on orders and trades for the different types of traders: noise, informed, and liquidity. Panel A separately reports the average numbers of limit and market orders submitted by the traders. Each limit order submitted by a trader is for one share. Market (or marketable limit) orders are defined as the taking of limit orders at the best prices in the book. Panel B plots the average number of one-share trades executed by a trader who belongs to one of the three types. The number of trades is broken down into trades executed by submitting a market orders and trades that occur as a result of a trader's previously submitted limit order being executed.

Panel A: Market and Limit Orders by Trader Type



Panel B: Trades by Trader Type

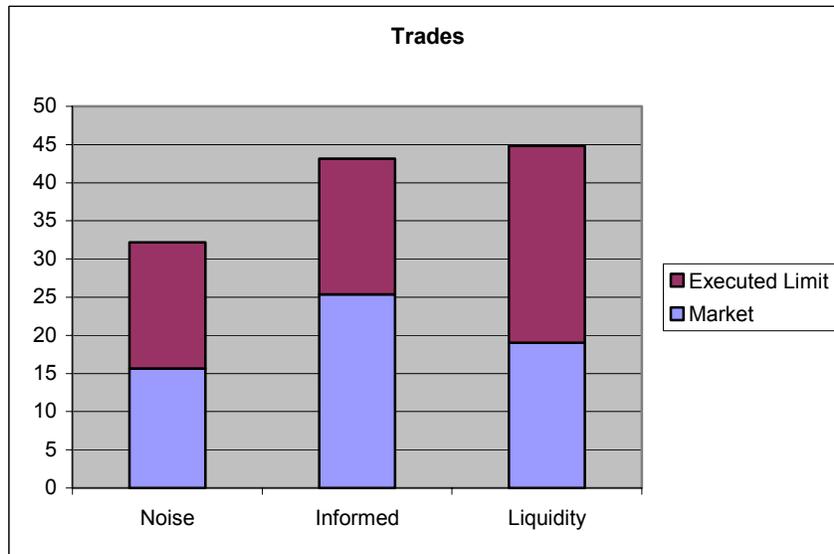
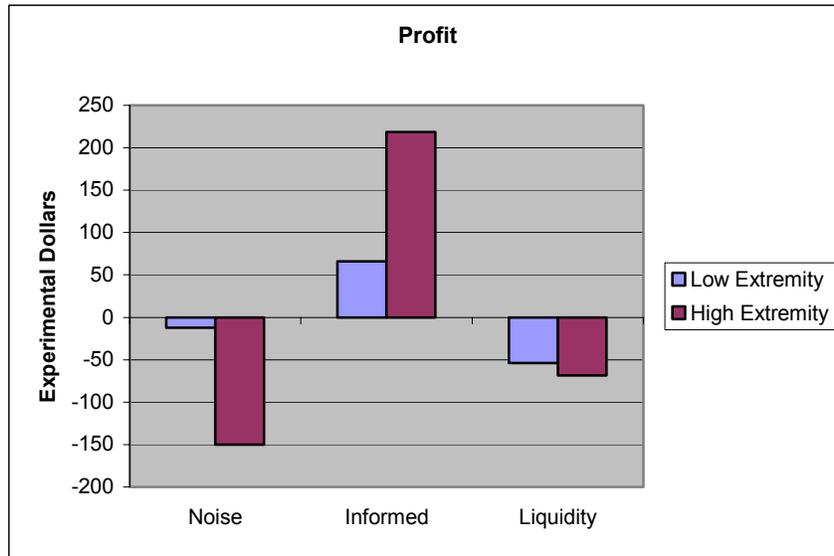


Figure 3
Traders' Profit and BBO Depth Provision

This table reports gross profits of the traders and their contribution to depth at the best bid or offer (BBO) by trader type and extremity level. Low (high) extremity means that the realized value of the security is no more than (at least) \$15 from its unconditional value. In Panel A, gross profit is the trading profit a trader makes before taxes. We first compute the profit for an individual trader per security, and then average it for a trader type within each of the eight cohorts. The numbers in the table represent the averages across the cohorts. Panel B shows how much depth (in number of shares) each trader type contributes to the best bid or offer prices. We take snapshots of the contribution to depth at the BBO of all trader types each time any trader submits or cancels an order. The average contribution to BBO depth of a trader belonging to a certain type is computed as the simple average of these snapshots divided by the number of traders of each type. We then average these numbers across cohorts.

Panel A: Gross Profit in Experimental Dollars by Trader Type and Extremity Level



Panel B: BBO Depth Provision in Number of Shares by Trader Type and Extremity Level

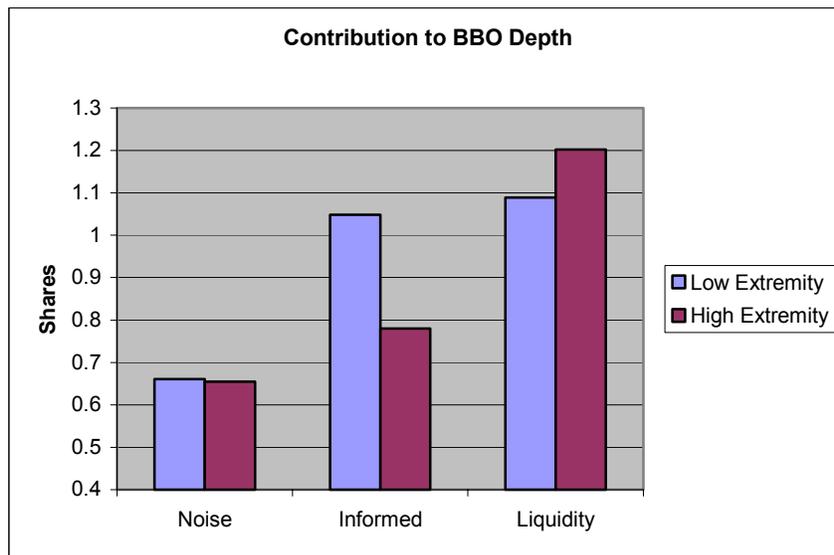
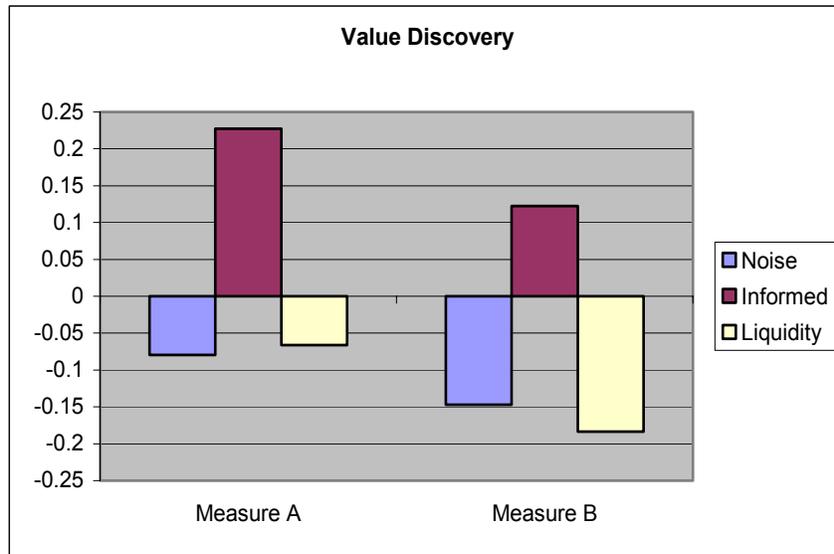


Figure 4
Value Discovery

This figure presents evidence on the contribution of a trader type to value discovery, or whether their trades move prices closer to or away from the true value. In Panel A, we present two measures of value discovery (measure A and measure B) by trader type. For Measure A, we sort executed orders into those moving prices toward the true value or away from it. If the true value is higher than the price, we assign +1 to a buy order of a trader that resulted in a trade and -1 to a sell order that resulted in a trade. If the true value is lower than the price, we assign -1 (+1) to a buy (sell) order of a trader that resulted in a trade. The measure is then aggregated for all market and executed limit orders of a trader and divided by the number of his trades (the measure is therefore always in the range [-1, +1]). The more positive (negative) measure A of a trader, the more his trades contribute to (interfere with) value discovery. Measure B of value discovery is constructed in a similar way to measure A, except that we consider only trades resulting from market orders submitted by the trader. Therefore, it considers contribution or interference with value discovery only for the “active” side of a trade. Panel B presents value discovery measure A by trader type and extremity level. Low (high) extremity means that the realized value of the security is no more than (at least) \$15 from its unconditional value. For both panels, we first compute each measure for an individual trader and then average it for a trader type within each of the eight cohorts. The numbers in the figure represent the averages across the cohorts.

Panel A: Value Discovery Measures A and B by Trader Type



Panel B: Value Discovery Measure A by Trader Type and Extremity Level

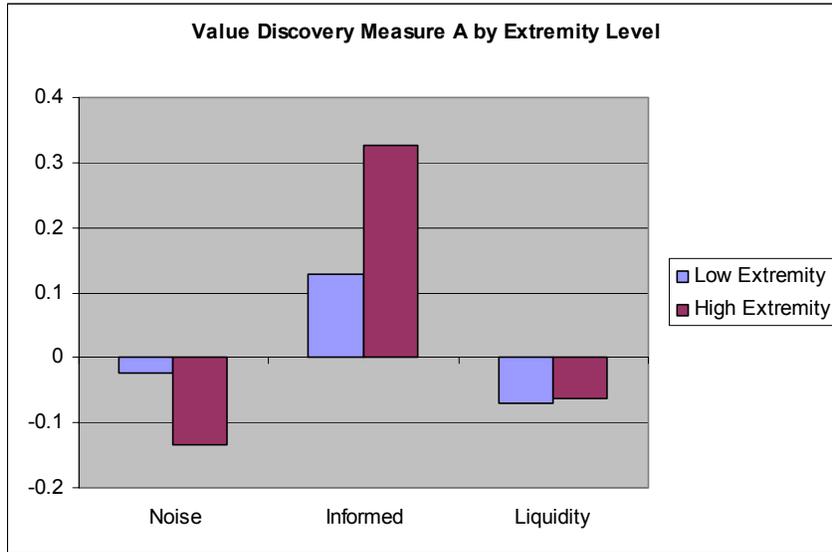
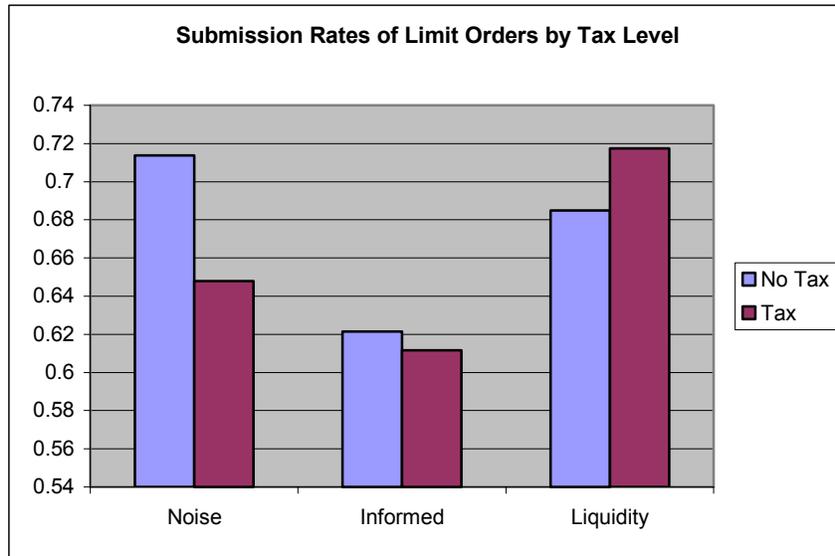


Figure 5
 Trader Strategies with and without Taxes

This figure presents two aspects of the traders' strategies under the different tax regimes. In the setting with a securities transaction tax (STT), we impose a \$2 fee for each trade. In a no-tax setting, no fees are imposed. In Panel A, we report the submission rate of limit orders defined as the number of limit orders divided by the total number of orders a trader submits. We first compute the submission rate for an individual trader and then average it for a trader type within each of the eight cohorts. The numbers in the figure represent the averages across the cohorts. Panel B shows how much depth (in number of shares) each trader type contributes to the best bid or offer prices. We take snapshots of the contribution to depth at the BBO of all trader types each time any trader submits or cancels an order. The average contribution to BBO depth of a trader belonging to a certain type is computed as the simple average of these snapshots divided by the number of traders of each type. We then average these numbers across cohorts.

Panel A: Submission Rates of Limit Orders by Trader Type and Tax Level



Panel B: Depth Provision at the Best Bid and Offer Prices by Trader Type and Tax Level

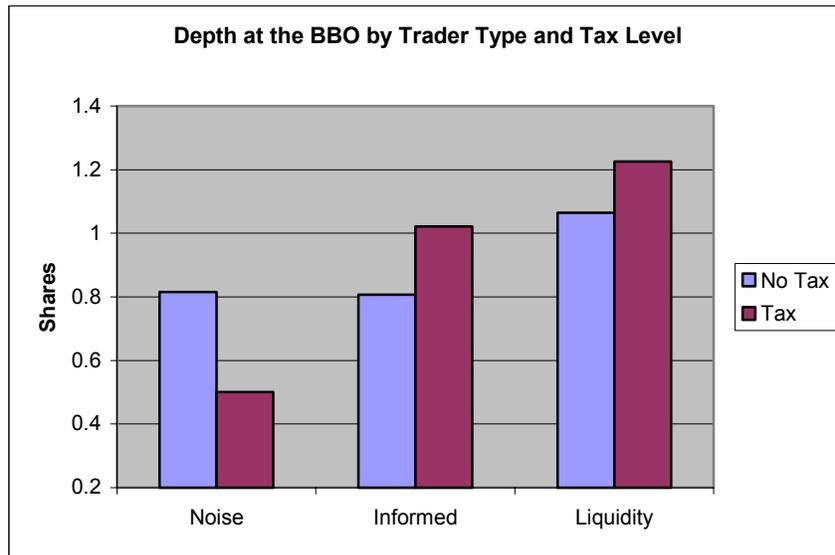


Figure 6
Price Impact of Market Orders by Extremity and Tax Levels

This figure presents evidence on one aspect of market liquidity—the price impact of market orders—by tax regime and extremity level. In the setting with a securities transaction tax (STT), we impose a \$2 fee for each trade. In a no-tax setting, no fees are imposed. Low (high) extremity means that the realized value of the security is no more than (at least) \$15 from its unconditional value. The total price impact (the total height of each column in the figure) is further decomposed into temporary and permanent components. The total, permanent, and temporary price impacts for a market order at time t that is characterized by a transaction price P_t and a prevailing quote midpoint M_t are computed as follows:

$$\text{Total Price Impact} = \begin{cases} P_t - M_t & \text{for market buy orders} \\ M_t - P_t & \text{for market sell orders} \end{cases}$$

$$\text{Permanent Price Impact} = \begin{cases} M_{t+5} - M_t & \text{for market buy orders} \\ M_t - M_{t+5} & \text{for market sell orders} \end{cases}$$

$$\text{Temporary Price Impact} = \begin{cases} P_t - M_{t+5} & \text{for market buy orders} \\ M_{t+5} - P_t & \text{for market sell orders} \end{cases}$$

where $t+5$ denotes the midquote prevailing five traders after the trade for which we are computing the price impact measure. The numbers in the figure represent the averages across the securities and cohorts.

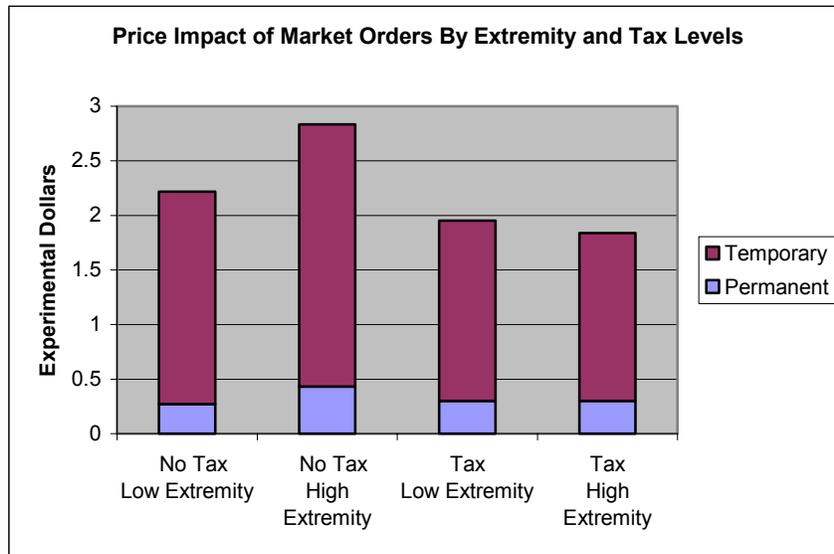
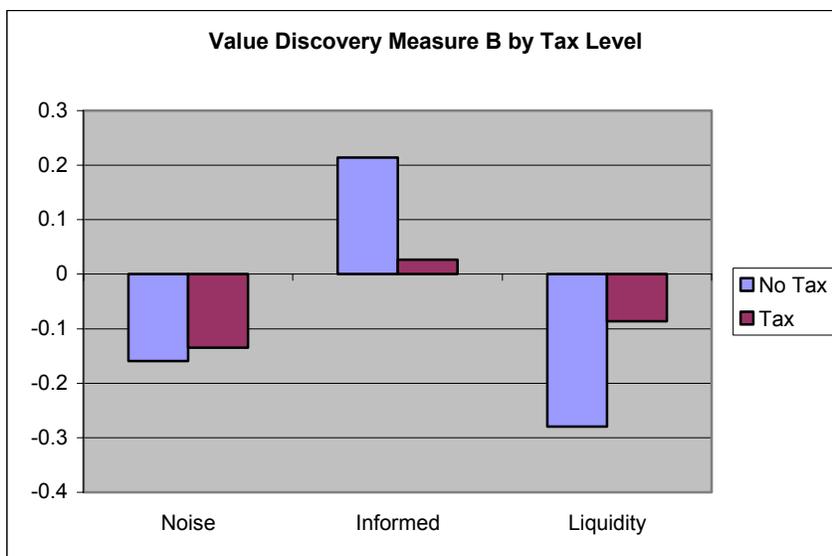


Figure 7
Value Discovery by Tax Levels

This figure presents evidence on the contribution of a trader type to value discovery, or whether their trades move prices closer to or away from the true value, by tax regime. In the setting with a securities transaction tax (STT), we impose a \$2 fee for each trade. In a no-tax setting, no fees are imposed. To construct the value discovery measure (Measure B), we sort a trader's market orders into those moving prices toward the true value or away from it. If the true value is higher than the transaction price, we assign +1 to a market buy order and -1 to a market sell order. If the true value is lower than the transaction price, we assign -1 (+1) to a market buy (sell) order. The measure is then aggregated for all market orders of a trader and divided by the number of market orders (the measure is therefore always in the range [-1, +1]). The more positive (negative) measure B of a trader, the more his market orders contribute to (interfere with) value discovery. We first compute the measure for an individual trader and then average it for a trader type within each of the eight cohorts. The numbers in the figure represent the averages across the cohorts.



Appendix

Instructions for an Open Electronic Limit Order Book (ELOB) Market

Overview

During this session, you will trade many securities that each pay a single liquidating dividend denominated in “laboratory dollars.” We will convert your gains or losses into U.S. dollars to determine your payment. At the end of trading, we will also ask you a series of questions about your experience. The entire session should take about 1.5 hours. Please do not talk with other subjects, look at others’ computer screens, or leave the room without explicit permission from the experiment administrator.

Liquidating Dividends

The dividend paid by every security is equal to 50 plus two random numbers, called “Component 1” and “Component 2.” Each component will each range from -25 and +25, with each realization equally likely, allowing dividends as low as 0 and as high as 100. Note that extreme dividends are less likely than dividends close to 50, because there are many ways that the two random numbers can add up to 0, but only one way that they can add up to the most extreme high or low dividend. (Similar reasoning explains why two rolled dice are more likely to add up to 7 than 2 or 12.)

Types of Traders

The market includes three types of traders:

- **Informed Traders** know one of the two random numbers. There are two who know one number and two who know the other.
- **Liquidity Traders** are forced to end trading with a share balance exactly equal some “target” number of shares, or else they are penalized. Throughout the session, the target is 20 or 30. Two of the traders will have positive targets (they will need to buy shares) and two will have negative targets (they will need to sell shares). The (Sell 20, Buy 30) and (Buy 20, Sell 30) combinations are twice as likely as the (Sell 20, Buy 20) and (Sell 30, Buy 30) combinations.
- **Free traders** are not told either random number, and do not have any liquidity target.

Sequence of Events

For each security, you will follow the same sequence of steps:

- (1) Learn component (if informed trader) or target (if liquidity trader)
- (2) Estimate security dividend
- (3) Trade
- (4) Estimate security dividend
- (5) Learn true dividend of security

Remember that the dividend and market price are not necessarily the same thing. A security's market price is determined by the amount traders are pay or accept, and may change as trading progresses. A security's dividend is determined by the two random numbers selected before trading begins, and never changes.

How to Trade in an Open Electronic Limit Order Book Market

Trading sessions are 150 seconds long (except for the practice security). All traders trade shares by entering orders that others can "take" or by "taking" orders that others have entered. All orders are for one share, but you can enter and take multiple orders at each price.

- **Entering a Bid** A bid is an order to **buy** a share at a stated price. You will buy at that price if someone else chooses to take your bid, and sells a share to you at the price you indicated.
- **Entering an Ask.** An ask is an order to **sell** a share at a stated price. You will sell at that price if someone else chooses to take your ask, and buys a share from you at the price you indicated.
- **Taking a Bid or Ask.** If you click on the "SELL 1" button on the BID column, you will sell a share at the highest bid. If you click on the "BUY 1" button on the ASK column, you will buy a share at lowest ask.
- **Removing a bid or ask.** You can remove (cancel) a bid or ask that you entered, simply by *right-clicking* on it.

Note that the price graph used on the computer screen shows an initial range from \$25 to \$75. However, dividends could be as low as 0 or as high as 100. You can enter and see orders for prices outside the initial range by clicking on the up- and down-arrows between the graphs.

Some Trading Restrictions

The following rules keep you from entering or taking any orders you please.

- ◆ You can't trade with yourself. Requests to take your own order will be rejected.
- ◆ You can never enter a bid at a price greater than your own ask, or an ask at a price less than your own bid. Doing so would be like trying to trade with yourself.
- ◆ You can't enter a bid higher than an existing ask or an ask lower than an existing bid. If you are willing to buy at the lowest ask, simply click the "BUY 1" button. If you are willing to sell at the highest bid, simply click the "SELL 1" button.

Trading Gains and Losses. You start each security with no cash and no shares. However, negative cash and share balances are permitted. Thus, you can buy shares even

if you don't have money to pay for them, and you can sell shares you don't own ("short selling").

After trading a security, the shares you own pay the liquidating dividend. If you have a positive balance of shares, the dividend is added to your cash balance for each share you own. If you have a negative balance of shares, the dividend is subtracted from your cash balance for each share you own. The resulting number is your trading gain (if positive) or trading loss (if negative).

You make money every time you buy a share for less than true dividend or sell a share for more than true dividend. For example, buying a share worth \$30 at a price of \$23 creates a gain of \$7. Selling that share at that price creates a loss of \$7.

Trading taxes. For half of the securities in the experiment, there is a trading tax equal to \$2 for every share you buy or sell. Thus, if you buy a share worth \$30 at a price of \$23 (as in the example above), your net gain after tax is only \$5. Selling that share at that price creates an after-tax loss of \$9. You will always know whether the trading tax is in force.

Penalties for Liquidity Traders. Liquidity traders also may incur penalties for failing to achieve their targets. The penalty is \$100 laboratory dollars for each share you fall short of your target. This penalty is large enough that liquidity traders are always better off trading enough to hit their target, even if they must buy at very high prices or sell at very low prices to do so.

Converting Laboratory Dollars into US Dollars.

Laboratory winnings, as described above, will be converted into US\$ according to the formula

$$\text{US\$ Payment} = (\text{Net Gain/Loss in Laboratory \$} + \text{Adjustment}) \times \text{Exchange Rate.}$$

You are guaranteed a minimum payment of US\$10.

You will not learn the exact adjustment or exchange rates. However, we will tell you that the exchange rate is positive, meaning that the more laboratory dollars you win, or the fewer you lose, the more \$US you take home. The parameters are set so that the average winnings will be approximately \$25, so that no one has an advantage or disadvantage.

How to use the Trading Interface (Order Entry Graphs)

Two ways to BUY:

- **Enter a Bid:** Left-Click a price on the right side of the BIDS graph places an order in the order book. You will buy a share at that price if someone else clicks the “Sell 1” button when you have the highest bid in the book.
- **Take an ask.** Click the “Buy 1” button to buy 1 share at the highest bid price in the book. Your transaction will be executed immediately.

Two ways to SELL:

- **Enter an Ask.** Left-Click a price on the right side of the ASKS graph places an order in the order book. You will sell a share at that price if someone else clicks the “Buy 1” button when you have the lowest ask in the book.
- **Take a Bid.** Click the “Sell 1” button to sell 1 share at the highest bid price in the book. Your transaction will be executed immediately.

Delete a Bid or Ask: Right Click on the right side of the BIDS or ASKS graph to remove an order at that price.

Scrolling the graphs.

Use the buttons between the order entry graphs to

- Scroll higher in both graphs (up arrow)
- Scroll lower in both graphs (down arrow)
- Set the boundaries of both graphs to be equal (“RESET”)

Reading the Information on the Right Side of the Screen

Information about YOUR trades:

- # of Buys indicates the number of times you bought a share, whether by taking someone’s ask or having them take one of your bids
- The average price at which you bought is also reported
- # of Sells indicates the number of times you sold a share, whether by taking someone’s bid or having them take one of your asks
- The average price at which you sold

Information about MARKET trades:

- # of Buys indicates the number of times someone clicked the “Buy 1” button, taking the lowest ask; the average price of those trades is also reported
- # of Sells indicates the number of times someone clicked the “Sell 1” button, taking the highest bid; the average price of those trades is also reported