

The Limits of Noise Trading: An Experimental Analysis

By

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

**First Draft: January 2005
This Version: October 2005**

***Robert Bloomfield (rjb9@cornell.edu), Maureen O'Hara (mo19@cornell.edu) and Gideon Saar (gs25@cornell.edu) are from the Johnson Graduate School of Management, Cornell University. Financial support for this project was obtained from New York University's Salomon Center for the Study of Financial Institutions. We would like to thank Joel Hasbrouck, Wei Xiong, and seminar participants at New York University for helpful comments.**

Abstract

We report the results of a laboratory market experiment 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 Kyle [1985], on the other hand, defines noise trading as demand from investors who do not possess private information about the liquidation value of the risky asset. His definition bundles together investors who trade for risk sharing and others who trade due to incorrect beliefs or because they “just like to trade.”

Shleifer and Summers [1990] view noise traders as investors whose demand for risky assets is affected by beliefs or sentiments that are not fully justified by fundamental news. 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 follow the tradition that began with Black [1986] by separating those who need to trade (e.g., to raise cash or for risk sharing) and those who trade just because they

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

want to (e.g., they think they can identify trends in prices or perhaps they “just like to trade”). For that reason, our experiment features three types of traders who interact in the market: informed traders (who possess fundamental information), liquidity traders (who have to trade for exogenous reasons), and noise traders (who do not possess fundamental information and have no exogenous reasons to trade). 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 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 DeLong, Shleifer, Summers, and Waldmann (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. We also 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.

Third, we find little 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 Mendleson [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 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 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

We also provide evidence relevant to 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 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 presumably because trend chasing strategies collapse in the presence of large market movements. We also find evidence that in the presence of a securities transaction tax, noise traders who lose money in a period trade less in subsequent periods, consistent with

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.

them eventually being driven out of the market. These results suggest that while noise trading may persist in actual markets, this could be because of the arrival of new noise traders and not because of the survival of older ones.

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 consider anew the issues of noise trader profits, price discovery, informational efficiency, and market volatility. Section 5 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 a very similar market structure. 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 two and a half minutes 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*

Our experiment seeks to understand the behavior, welfare and market influence of noise traders. To do so, we create markets for securities that each pays a liquidating dividend equal to 50 plus the sum of two random numbers, each of which is uniformly distributed from -25 to 25. The market includes four informed, four liquidity and four noise traders. Before trading begins, two informed traders learn the value of the first random number while the other two informed traders learn the value of the second random number (so that no trader has perfect information).

Four 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 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 liquidity targets are random, with two-thirds of the securities resulting in a non-zero net liquidity demand.⁵

The remaining four traders (the noise traders) have no information and no trading targets. These noise traders thus fit the definition offered by Black of traders who may “simply like to trade”, but such noise trading may also be motivated by traders’ beliefs that they can interpret market information better than others, or that they can earn profits by providing liquidity (a reasonable belief, given the non-zero demand by liquidity traders in two-thirds of the securities).⁶ By creating three classes of traders, we can focus

⁴ The use of trading targets is standard in experimental work (see 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 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 -10. 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 +10. 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

⁶ 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 the roles played by specific participants in the experiment.

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. In actual markets, investors may behave on certain occasions as noise traders while on other occasions as liquidity traders. The experiment enables us by construction to focus on the effects of noise trading without confounding it with other reasons for trading in financial markets.⁷

To clarify the market influence of noise traders, we manipulate whether traders face a Securities Trading Tax (SST) by imposing, for half of the securities, a tax of \$2 on each purchase and each sale.⁸ We expect the SST to substantially reduce the propensity of noise traders to trade, because the benefits to doing so are made substantially smaller.⁹ The direct effect of the tax on the behavior of the informed and liquidity traders is less obvious: informed traders have valuable information, and liquidity traders suffer substantial penalties for failing to meet their trading targets. Evaluating these effects allows us to speak to many of the policy issues associated with such taxes.

We also manipulate the value of the information held by informed traders. To do so, 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. The greater the distance between the realized security value and the expected value of 50, the

⁷ While our experiment enables the intensity of noise and informed trading to adjust endogenously (traders can trade as much as they want), the number of noise traders and informed traders is constant. Our goal in setting those specific numbers was to ensure we have competitive informed traders and to give noise traders the opportunity to have a significant presence in the market. Since it is impossible to know the relative numbers of noise, liquidity, and informed traders in actual markets, our objective is to demonstrate general results that would hold under the conditions of competitive informed trading and in the presence of noise traders.

⁸ 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.

⁹ The SST manipulation also clarifies the noise traders' motivations. Noise traders might trade simply because they feel they have nothing better to do with their time, or because they believe (correctly or incorrectly) that they can earn profits by doing so. Noise traders who trade simply out of boredom are unlikely to be affected by the SST.

more valuable is the private information of the informed traders. Therefore, the extremity categories can be viewed as describing the extent (high versus low) of adverse selection in the market.

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, 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 conversely. 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 SST securities after the non-taxed securities (we found no significant differences connected to SST order).

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). Trader type and cohort membership are manipulated across traders, and all other factors are manipulated 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*

Each security is continuously traded in a trading period that lasts two and a half minutes. Our double auction market is organized like a typical electronic limit order book where traders can enter buy or sell limit orders and the execution of limit orders follows strict price/time priority. This means that a limit order at a better price, say a higher price for a buy order, has priority over a limit order at a worse price. Within each price level, older limit orders are executed first. Traders can execute outstanding limit orders in the book by clicking a “buy 1” button, which allows them to buy one share at the lowest current asking price, or by clicking a “sell 1” button, which allows them to sell one share at the highest current bid price. Executing a limit order in the book is equivalent to submitting a market (or marketable limit) order.

The limit orders to buy or sell a security must have integer prices between 0 and 100. As soon as a trader enters an order, the order is shown on all traders’ computer screens, indicating that an unidentified trader is willing to buy or sell one more share at the posted price.¹⁰ Traders can cancel their unexecuted limit orders in the book at any time during the trading period. All trades are reported immediately to all traders, indicating the price and the trade direction (whether the trade involved a market buy taking a limit sell order or a market sell taking a limit buy order).

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

¹⁰ A trader’s screen includes one chart indicating bids (limit buy orders) and one indicating asks (limit sell orders). 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.

received detailed written instructions.¹¹ The instructions were reviewed 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 no cash or shares. Unlimited negative cash and share balances were permitted, so traders could hold any inventory of shares they desired, including 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.¹²

2.5. *Statistical Analysis*

Our analysis uses a variety of statistical methodologies including a repeated-measures ANOVA, regression analysis, and bi-variate vector autoregression (VAR).

¹¹ These written instructions are available from the authors.

¹² 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.

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 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 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.¹³ 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.

We use regression analysis 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). 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

¹³ 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.

difficult to interpret, the coefficients themselves are still consistent. We then perform t-tests and Wilcoxon signed-rank tests on the eight observations (one from each cohort) of the coefficients. This approach allows us to draw inferences from across cohorts, which can be assumed independent.

We use bi-variate vector auto-regression (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.

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 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 4.¹⁴

3.1. Summary Statistics

Figure 1 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 orders. This evidence suggests that experimental subjects understood the market

¹⁴ 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.

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 discuss specific motivations for noise trader activity below.

3.2. *Trading Profits*

As predicted by market microstructure models, informed traders make money and liquidity traders lose money. The average trading profit of an informed trader is 142.17 lab dollars (henceforth, dollars), while a liquidity trader loses on average 61.04 dollars.¹⁵

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, but the literature on day trader profitability is mixed on this issue.¹⁶ For example, Harris and Schultz [1998] 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 trades of individuals in Finland and reports that their performance is not different from that of a control group. Interestingly, these individuals who engage in day trading are net suppliers of liquidity in that they use limit orders and, therefore, often function as if they were market makers. Jordan and Diltz

¹⁵ 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.

¹⁶ Day traders are often viewed as quintessential noise traders because they do not trade for either information or portfolio reasons.

[2003] and Barber, Lee, Liu, and Odean [2004] find that day traders lose money on average.

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. This result arises from the trading strategies pursued by noise traders, as 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.

Given that profitable strategies could exist, what are noise traders actually doing? To address this issue, we first note that 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”. Since experiment participants are randomly assigned to the different roles, it is hard to argue that those assigned to be liquidity traders are more rational than those assigned to be noise traders. It is well known from behavioral research, however, that individuals often perceive patterns when none actually exist. Perhaps the explanation for the smaller losses of the liquidity traders (who must trade) compared with those of the noise traders (who don’t have to trade) is just that the liquidity traders are too busy trading their targets to trade simply on a mistaken reading of the market environment.

Panel A of Figure 2 shows that there is an interesting relation between 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).

If noise traders were making money when extremity was high, it would have suggested that they can “read” the market by buying at the beginning of an upward trend or selling at the beginning of a downward trend. The result that they lose most when extremity is high while the degree of extremity has little effect on the liquidity traders' losses points to the opposite conclusion. In fact, noise traders 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].

Do noise traders, instead, attempt to make money by adopting the role of liquidity providers? Panel B of Figure 2 shows how much depth trader types contribute to the best bid or offer prices (BBO) in the limit order book.¹⁷ The depth provided by limit orders of the different trader types is a measure of their willingness to supply (as oppose to demand) liquidity. Limit orders submitted at the best bid or offer prices constitute a reliable measure of liquidity provision because depth away from current prices could be comprised mostly of stale limit orders. 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

¹⁷ 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.

liquidity irrespective of the level of extremity (0.66 and 0.65), and this level is lower than that of either informed or liquidity traders. Thus, noise traders play a minor role in liquidity provision. The evidence suggests that noise traders instead focus on 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 experimental methodology is particularly suitable for examining the link between noise trading and efficiency because we know the fundamental value of each security and can directly observe the trading of noise traders. In the following analysis, we first develop measures that quantify whether a groups’ 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, we construct two measures. For Measure A, we sort executed orders into those trades 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.

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 $A = -1$.

Intermediate strategies will lie in the range $[-1, +1]$. 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 to or interference with value discovery only for the “active” side of a trade.

Panel A of Figure 3 shows that different types of traders affect value discovery in different ways ($p < 0.0001$). 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 (which describes the severity of the adverse selection problem) for measure A ($p = 0.0029$). As Panel B of Figure 3 shows, when the realized value of the security is close to the expected value, the value discovery measure of the noise traders is also close to zero (-0.02). But 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 define the pricing error, DEVP, as the average of the distance of the transaction price from the true value. We then run the following OLS regression:

$$\text{DEVP}_i = a + b * \text{NOISETRD}_i + c * \text{INFTRD}_i + \text{error}_i$$

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.

As reported in Panel A of Table 1, the coefficient on NOISETRD is positive and highly statistically significant in both regressions.¹⁸ Thus, 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. Conversely, noise traders appear to slow down the adjustment of prices.

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 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 perform t-tests and Wilcoxon signed-rank tests on the eight observations of each coefficient (which 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

¹⁸ 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.

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 that informed trading advances informational efficiency, while noise trading retards its attainment.

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]). The empirical literature on day trading, however, finds mixed results 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 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 shows that 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

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

with higher volatility of prices. A similar positive and highly statistically significant coefficient on noise trading is found in the regression on PRANGE. Interestingly, 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. Now none of the tests show statistically significant results. The divergence between the pooled results and the cohort controlled results suggests that cohorts with more active noise traders also have more volatility, but that little of the variation in volatility within cohorts is attributable to variation in noise trader activity.

4. Noise Trading and a Securities Transaction Tax

In the previous section we presented evidence that noise trading affects market behavior, particularly through its association with lower informational efficiency of prices. Curbing noise trading, therefore, could potentially 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 oft-suggested way to reduce noise trading 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 is articulated by Stiglitz [1989], who conjectured that the tax is unlikely to discourage trading by the uninformed and informed traders, and so would mainly serve to drive out

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

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 4 shows that 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).²³ With noise traders submitting fewer limit orders relative to market orders, they provide less liquidity to the market.

The limit order submission rate measure includes limit orders that are submitted away from current market prices and that will, in fact, never execute. Such limit orders could be used by traders trying to disseminate false information in the hope of moving market prices in certain directions. A measure of liquidity provision not subject to this problem is the provision of depth at the best bid and offer prices (BBO) in the limit order book. Panel B of Figure 4 shows that 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

²¹ Two theoretical papers examining the implications of securities transaction taxes in an environment characterized by information asymmetry are Subrahmanyam [1998] and Dow and Rahi [2000]. Subrahmanyam finds that with a STT informed traders trade less and liquidity generally worsens. Dow and Rahi 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 varies with parameters in their model, but for some sets of parameters a securities transaction tax is pareto improving.

²² We will not consider other potential effects of the tax that are 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.

²³ The submission rate of limit orders defined as the number of limit orders divided by the total number of market and limit orders traders submit.

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 profits and losses. Table 3 presents gross and net profits by tax level, extremity level, and trader type. Note 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. That noise traders can profit in certain circumstances is important because it suggests that we did not bias our investigation by creating an environment in which noise traders are "forced" to lose money. This profitability is probably due to two effects. First, 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, noise traders 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 also deter some of the speculative trading on trends that caused noise trader losses in high extremity securities. This reduces the noise traders' losses from 210.34 dollars without taxes to a net loss of 136.56 dollars after taxes are incorporated.

Taxes also affect the profits of the informed traders. 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.

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 effect 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, which falls from 39.97 shares without taxes to 25.45 shares with taxes (p -value = 0.0403). Informed trading also falls from 47.96 to 40.92 shares, although this decrease is not statistically significant.

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 may directly increase the inventory costs of liquidity providers, but if reduced noise trading also 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 as the total price impact of market orders (the distance between the transaction price and the midpoint between the best bid and offer prices in the limit order book). This measure is similar to the effective spread used in empirical studies. 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). There

is, however, an interesting and statistically significant interaction between liquidity and extremity, or information value (p -value = 0.0310). The total price impact decreases from 2.22 to 2.00 dollars for low extremity securities when we impose taxes, but decreases in a much more drastic fashion for high extremity securities, going from 2.75 to 1.87 dollars

That transaction costs fall despite overall lower market volume is an intriguing result. To understand why this happens, we borrow from the empirical market microstructure literature and decompose the total price impact into permanent and temporary price impact components.²⁴ The permanent price impact of a market order is attributable to the private information content of the order. The temporary price impact (also known in the empirical literature as the realized spread) is a measure of efficiency of the liquidity provision mechanism. It provides some indication of the profitability of liquidity provision because it measures the price reversal from the trade price.

We find no statistically significant differences in the permanent price impact across the two tax regimes. The pattern in the total price impact is driven mainly by changes to the temporary price impact, which is much smaller in the presence of taxes (2.17 dollars versus 1.60 dollars, p -value = 0.0031). The dominance of 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. Hence, our analysis demonstrates that while taxes result in lower volume, their impact on liquidity can be positive.

4.3. *Volatility and Informational Efficiency*

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

²⁴ For a market order at time t denote the transaction price by P_t , the midpoint between the bid and offer prices in the limit order book (just before the trade) by M_t , and the midpoint between the bid and offer prices prevailing five trades after the market order for which we are computing the price impact by M_{t+5} . The total price impact for a buy (sell) order is defined as $P_t - M_t$ ($M_t - P_t$). The permanent price impact for a buy (sell) order is defined as $M_{t+5} - M_t$ ($M_t - M_{t+5}$), and the temporary price impact is defined as $P_t - M_{t+5}$ ($M_{t+5} - P_t$). Note that the total price impact is the sum of the permanent and temporary components.

Summers [1989]). While some price volatility reflects adjustment to fundamental information, transaction taxes are supposed to rid prices of “excess” volatility that is due to noise trading.²⁵ 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.

Earlier we found that noise trading had little effect on volatility. Introducing a securities transaction tax seems also to have little effect on volatility. Although both the standard deviation of transaction prices (STDP) and the transaction price range (PRAGNE) are lower under the tax regime, the differences are not statistically significant. Similarly, we also find that an STT has little effect on the informational efficiency of prices. Again, both 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) are slightly lower under the tax regime, but these differences are not statistically significant. Therefore, we cannot conclude that a securities transaction tax can be used to lower price volatility or improve the informational efficiency of prices.²⁶

4.4. *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

²⁵ 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.

²⁶ It is interesting to note that in the theoretical framework employed by Scheinkman and Xiong (2003), a securities transaction tax can cause a substantial reduction in speculative trading but has only a limited impact on price volatility.

period and run a bivariate vector autoregression (VAR) of profit and trading separately for markets with and without taxes.²⁷

Panel A of Table 4 presents the results for a VAR with two lags in the case of no taxes.²⁸ 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 4 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. Furthermore, the Granger Causality test for profit is significant ($p = 0.0422$) indicating that profit Granger-cause trading.²⁹ 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.³⁰

²⁷ 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.

²⁸ 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.

²⁹ The Granger Causality test on NOISETRD is also significant, and so the relation goes in both directions.

³⁰ Our experiment does not force losing traders to slow down their trading due to a budget constraint. This may be one of the reasons for not finding a significant effect of past losses on future trading in the regime with no taxes. However, one should still find an effect if traders learn from their mistakes and realize that they are trading on noise and therefore lose, rather than make, money. It is possible that the transaction tax helps learning because it serves as an “awareness” device that makes noise traders pay more attention to the relation between trading and losses.

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 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, nor do they specialize in providing 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 because noise trader strategies change the strategies pursued by other traders. 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 implication of our results is that noise traders may 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.³¹ Indeed, this difficulty underscores a general concern with 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 influence 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.

³¹ For a discussion of the issues surrounding participation effects and securities transaction taxes see O'Hara [2004].

References

- Amihud, Y. and H. Mendleson, 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 Preferencing 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. Waldmann, 1990, Noise Trader Risk in Financial Markets, Journal of Political Economy, 98(4), 703-738.
- DeLong, B., Shleifer, A., Summers, L., and R. Waldmann, 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.

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)

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.

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)

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
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
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 2
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 3
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 4
 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

