Noise Trading in a Laboratory Financial Market: A Maximum Likelihood Approach

Marco Cipriani
George Washington University

Antonio Guarino
University College London and University of Venice[*†]

Abstract

We study the extent to which, in a laboratory financial market, noise trading can stem from subjects' irrationality. We estimate a structural model of sequential trading by using experimental data. In the experiment, subjects receive private information on the value of an asset and trade it in sequence with a market maker. We find that, in the laboratory, the noise due to the irrational use of private information

[*]Acknowledgements: We gratefully acknowledge the financial support of the C.V. Starr Center and the Center for Experimental Social Science (CESS) at New York University. Guarino also gratefully ackowledges the financial support of the ELSE Center at University College London and of the Leverhulme Trust. All errors are ours.

[†]E-mail addresses: Cipriani: mciprian@gwu.edu; Guarino: a.guarino@ucl.ac.uk.

accounts for 35 percent of the decisions. When subjects act as noise traders, they abstain from trading 67 percent of the time. When they trade, the probability that they buy is significantly higher than the probability that they sell. (JEL: C92, D8, G14)

1 Introduction

A standard assumption of the theoretical market microstructure literature is that in financial markets there is a proportion of traders who act as "noise traders." Noise traders buy or sell financial assets with fixed probabilities because of exogenous, unmodeled reasons. They are usually contrasted to "informed traders," who use their private information and maximize expected profits. In markets with asymmetric information, the presence of noise traders, although theoretically unpalatable, is needed for the market not to break down.[1] In the literature, the presence of noise traders is typically justified by agents having other reasons to trade, typically liquidity and hedging reasons, beside informational motives. Noise trading is viewed

[1] For instance, in a dealer's market, if every rational trader were informed, the dealer (or market maker) would refuse to trade since he would be facing agents who are better informed than he is. This is a case in which the no-trade theorem (see, e.g., Milgrom and Stokey, 1982) would apply.

as a simple way to capture these non informational reasons to trade. In financial markets, however, an additional source of noise trading can stem from traders' irrational behavior. While standard financial theory assumes that all agents are rational, in actual markets traders make mistakes, may have some degree of bounded rationality, or may just not follow the expected utility maximization paradigm.

This paper studies the extent to which noise trading can stem from traders' irrationality. We analyze a laboratory financial market where subjects traded in sequence with a competitive market maker. All subjects were informed and were given incentives to maximize their expected profits. They did not have any other reason to trade. We construct a likelihood function for the sequence of trades that we observe in the experiment. We estimate the proportion of noise decisions in the experiment, i.e., the proportion of time when subjects disregarded their private information and acted as noise traders. We also estimate the proportion of time when subjects acting as noise traders decided to buy, sell or abstain from trading.

The estimation of a sequential trade model by maximum likelihood was first carried out by Easley, Kiefer and O'Hara (1997) using NYSE transactions data. Our paper is the first to use the same methodology with experimental data. Easley, Kiefer and O'Hara (1997) estimate the amount of noise trading in the market, but cannot distinguish between noise trading due to hedging or liquidity reasons and noise trading due to agents' irrationality. In our experiment, the only reason for agents to trade is informational. Therefore, the extent to which we observe noise trading can only be accounted for by traders' irrationality.

Another novelty of our study is that, in contrast with Easley, Kiefer and O'Hara (1997), in the estimation we do not impose the restriction that noise traders must buy or sell in the same proportion. We allow for the possibility that they may have a bias for buying or selling. We show that such a bias is indeed present in the laboratory financial market: noise traders are more likely to buy than to sell.

Section 2 describes the experimental data. Section 3 presents the estimation methodology. Section 4 discusses the results. Section 5 concludes.

2 The Experiment

We use the data collected by Cipriani and Guarino (2004) in an experimental study with undergraduate students of all disciplines at New York University. In particular, we use the data of the "Flexible price" and "No history" treatments.

For each treatment, the experiment was repeated for four sessions. In each session there were 13 participants, one acting as subject administrator and 12 acting as traders. The procedures of the experiment replicate a simple model of sequential trading in a dealer's market. Here we only summarize them and refer the reader to Cipriani and Guarino (2004) for a detailed illustration.

This was a paper and pencil experiment. At the beginning of each session, the experimenters gave written instructions (available on request) to all subjects. The instructions were read aloud in an attempt to make the structure of the game common knowledge.

Each session consisted of ten rounds. In each round all subjects could trade one after the other. The sequence of subjects for each round was chosen randomly. Each subject was called to trade only once in each round.

Before each round, an experimenter determined the value of the asset by flipping a coin: if the coin landed tail, the value of the asset for that round was 100, otherwise it was 0. Traders were not told the outcome of the coin flip.

During the round, an experimenter acted as market maker, setting the price at which people could trade. The other experimenter was outside the room with two bags, one containing 30 blue and 70 white chips and the other 30 white and 70 blue chips. The two bags were identical. Before trading, each subject drew a chip from one of the two bags (with replacement). If the asset value was 100, the subject drew the chip from the first bag. If it was 0, he drew it from the second.[2] Therefore, the chip color was a signal on the asset value. Subjects were instructed not to tell anyone what the chip color was. Therefore, neither the market maker nor the other subjects knew the realization of the signal drawn by the subject.

After observing the chip color, the subject entered the room and decided whether he wanted to buy, to sell or not to trade. In one treatment, the subject administrator recorded all subjects' decisions on the blackboard, where he also recorded the prices at which subjects could trade the asset. Hence, in this treatment, each subject knew not only his own signal, but also the history of trades and prices. In the other treatment, traders' decisions were not made public. Therefore, a trader knew only his signal and the price at which he could trade.[3]

At the end of each round, i.e., after all 12 participants had traded, the re-

[2] Of course, the subject did not know from which bag he was drawing the chip.

[3] Despite this difference, Cipriani and Guarino (2004) report that subjects' behavior was not significantly different in the two treatments. This agrees with the prediction of economic theory according to which, in the two treatments, subjects should behave identically since the price is a sufficient statistics for the history of trades. In this paper we use data from both treatments in order to have more precise parameter estimates.

alization of the asset value was revealed and subjects were asked to compute their payoffs. All payoffs were in a fictitious currency called lira.[4] Payoffs were computed as follows. If a subject bought, he earned 100+Value-Price lire. If he sold, he earned 100+Price-Value lire. If he decided not to trade, he earned 100 lire.

During the experiment the price was updated after each trade decision in a Bayesian fashion. According to economic theory, rational subjects should always follow their signal, i.e., they should buy after seeing a positive signal and sell after seeing a negative one. No one should decide not to trade, as private information allows the traders to make money by trading with the market maker. Therefore, when a subject decided to buy, the price was updated assuming that he had seen a positive signal. Similarly, when a subject decided to sell, the price was updated assuming that the subject had observed a negative signal. Finally, in the case of a no trade, the price was kept constant.

[4] At the end of the experiments, lire were converted into dollars at the exchange rate of 1/65.

3 Methodology

The main objective of our work is to estimate the parameters of the sequential trading model on which the experiment is based.

In each round of the experiment the asset value V_d (for $d = 1, 2, 3...80$) was either 0 or 100 with equal probability.[5] Given that the asset values were the results of coin flips, they were independently drawn. In the experiment all subjects were informed and were given incentives to maximize expected profits. We allow, however, for the possibility that they deviated from rational behavior. In particular, we assume that only a proportion of time μ did subjects trade the asset to maximize their expected profit based on their private signal. The remaining proportion of time $(1 - \mu)$ they behaved irrationally as noise traders. We assume that they bought, sold or did not trade with fixed probabilities, ε_B, ε_S and $(1 - \varepsilon_B - \varepsilon_S)$. Note that, in the estimation, we assume that subjects' decisions were independent of one another.

As we said in the Introduction, the standard models of sequential trading (see, e.g., Glosten and Milgrom, 1985) assume that some agents are informed and maximize expected profits, whereas the others are noise traders acting

[5] We have data for two treatments, each consisting of four sessions. Given that in each session the experiment was repeated for 10 rounds, we have a total of 80 rounds.

for unmodeled liquidity or hedging reasons. In contrast, in our setup, all subjects are informed, but we allow for the possibility that they sometimes trade irrationally.

The parameters that we want to estimate using the experimental data are the proportion of rational decisions (μ), and the probability that a subject acting as a noise trader decides to buy (ε_B) or sell (ε_S). To this purpose, we use a methodology first used in the empirical market microstructure literature by Easley, Kiefer and O'Hara (1997). We construct the likelihood function for the sequence of trades that we observe during the rounds. To construct such a function, let us first consider the probability of a sequence of trades in a round in which the asset value was 100. Given that traders act independently, the probability of B_d buys, S_d sells and N_d no trades in round d conditioned on $V_d = 100$, is

$$Pr(\{B_d, S_d, N_d\}|V_d = 100, \mu, \varepsilon_B, \varepsilon_S) = \quad (1)$$

$$K[0.7\mu + (1-\mu)\varepsilon_B]^{B_d}[0.3\mu + (1-\mu)\varepsilon_S]^{S_d}[(1-\mu)(1-\varepsilon_B-\varepsilon_S)]^{N_d},$$

where $K = \frac{(B_d+S_d+N_d)!}{B_d!S_d!N_d!}$ is the number of permutations of B_d buys, S_d sells and N_d no trades. Note that in this expression the proportion of the time μ in which subjects act rationally is multiplied by 0.7 or 0.3. Recall that subjects received a signal X with precision 0.7, i.e., $\Pr(X = x|V_d = x) = 0.7$ for $x = 0, 100$. Therefore, even in rounds when the asset value was 100, rational subjects would buy it only 70 percent of the time, and sell it the remaining 30 percent of the time.

The probability of B_d buys, S_d sells and N_d no trades in a round in which the value of the asset was 0 is computed in a similar manner:

$$\Pr(\{B_d, S_d, N_d\}|V_d = 0, \mu, \varepsilon_B, \varepsilon_S) = \tag{2}$$

$$K[0.3\mu + (1-\mu)\varepsilon_B]^{B_d}[0.7\mu + (1-\mu)\varepsilon_S]^{S_d}[(1-\mu)(1-\varepsilon_B-\varepsilon_S)]^{N_d}.$$

Note that in order to compute these probabilities we only need the number of buys, sells and no trades in each round. Indeed, the probability of a sequence of decisions does not depend on the order in which they are taken.

Given that in each round the asset value was 0 or 100 with equal probabilities 1/2, the unconditional probability of B_d buys, S_d sells, and N_d no trades in a given round is the weighted average of the previous two expressions, with equal weights 1/2, i.e.,[6]

[6] One half is the probability of $V_d = 0$ or $V_d = 100$ in a round. In our sample, the frequency was almost the same as the probability (52.5 percent of the time the asset value was 100).

$$\Pr(\{B_d, S_d, N_d\}|\mu, \varepsilon_B, \varepsilon_S) = \tag{3}$$

$$\frac{1}{2}\Pr(\{B_d, S_d, N_d\}|V_d{=}0, \mu, \varepsilon_B, \varepsilon_S) + \frac{1}{2}\Pr(\{B_d, S_d, N_d\}|V_d{=}100, \mu, \varepsilon_B, \varepsilon_S).$$

Finally, in order to compute the likelihood function, we need to compute the joint probabilities of the history of trades in all rounds. The asset value draws at the beginning of each round were independent. Therefore, the probability of a history of trades over multiple rounds can be written as the product of the probability of the histories of trades in each single round:

$$\Pr\left(\{B_d, S_d, N_d\}_{d=1}^{80}|\mu, \varepsilon_B, \varepsilon_S\right) = \Pi_{d=1}^{80}\Pr(\{B_d, S_d, N_d\}|\mu, \varepsilon_B, \varepsilon_S). \tag{4}$$

For simplicity, in our estimation we will maximize the logarithm of the likelihood function, which, after dropping a constant term, is

$$\sum_{d=1}^{80}\log\frac{1}{2}([0.7\mu + (1-\mu)\varepsilon_B]^{B_d}[0.3\mu + (1-\mu)\varepsilon_S]^{S_d}[(1-\mu)(1-\varepsilon_B-\varepsilon_S)]^{N_d})$$
$$+\frac{1}{2}([0.3\mu + (1-\mu)\varepsilon_B]^{B_d}[0.7\mu + (1-\mu)\varepsilon_S]^{S_d}[(1-\mu)(1-\varepsilon_B-\varepsilon_S)]^{N_d}). \tag{5}$$

By maximizing the log likelihood function we will estimate the proportion of time in which subjects acted rationally (μ), and the proportion of time in which subjects acting as noise traders bought (ε_B) or sold (ε_S) the asset.

4 Results

To find the maximum of the log-likelihood function we used the Nelder-Mead method. Table 1 summarizes the results of the estimation.

The estimated value of μ is 0.67. This is equivalent to say that the proportion of noise decisions in the experiment is about one third of the total. As we will see afterwards, some of these noise decisions consisted in buy or sell orders. Therefore, the irrationality of subjects' behavior generated noise trading in a market where there were only informational reasons to trade. Traders' irrationality may be one of the reasons (beside hedging and liquidity reasons) why financial markets may not collapse even if agents differ only in the information that they hold.

TABLE 1

Parameter	Value	Standard Deviation
μ	0.67	0.02
ε_B	0.26	0.04
ε_S	0.03	0.04

It is interesting to compare our findings with those of the empirical literature. Easley, Kiefer and O'Hara (1997) use transaction data of an actual stock (Ashland Oil) in the New York Stock Exchange for their estimation. They estimate a proportion of informed traders, μ, equal to 0.17. The proportion of noise trading that they find in the actual market is much higher than the proportion of irrationality that we find in our experiment. This seems to suggest that irrationality can explain only a fraction of noise trading observed in actual financial markets and that non informational reasons to trade, such as hedging and liquidity reasons, do play an important role.

Note, however, that the difference between our result and that of Easley, Kiefer and O'Hara (1997) could also be partly due to another reason. In their estimation, they assume that traders have perfect information on the asset value, i.e., that the precision of private information is equal to 1. If we made such an assumption in our estimation (i.e., since we know that the precision of the signal in the experiment was 0.7, if we maximized an incorrect likelihood function) we would obtain a much higher proportion of noise decisions (73 percent) than we do when using the correct signal precision. This indicates that incorrect assumptions on the precision of the signal may bias the estimated proportion of informed and noise traders in the market. Therefore, trying to estimate such a precision in actual financial markets could be an important development of the empirical literature.[7]

[7]Easley, Keifer and O'Hara (1997) cannot estimate the precision of the signal, as they

Now let us analyze the noise decisions in the experiment. The majority of irrational decisions are not buy or sell orders, but no trades. In particular, the proportion of no trade decisions, $(1 - \varepsilon_B - \varepsilon_S)$, is 70 percent. Abstention from trading is therefore an important reason why private information cannot be correctly aggregated in financial markets.

When noise traders decided to trade, most of the time they bought. In particular, we estimate a value of ε_B equal to 27 percent, whereas ε_S was only 3 percent. There is a statistically significant bias towards buying the asset: the null hypothesis that ε_B and ε_S are equal is strongly rejected by the data (LRatio test = 3.8). This is at odds with the assumption usually made in the empirical literature, i.e., that noise traders buy and sell with equal probability. Note, however, that, if we impose that ε_B equals ε_S, we estimate only a slightly higher proportion μ of rational decisions (68 percent), i.e., the bias stemming from imposing an incorrect restriction on the parameters is quite small. Moreover, in actual financial markets, there are other reasons for noise trading beside traders' irrationality, like hedging

only use data on the total number of buys and sells in each trading day. To estimate the precision of the private information that traders receive would require an analysis of the informational content of the particular sequence of buys, sells and no trades during the trading day.

and liquidity needs. Therefore, the overall noise trading activity may be more balanced than what we observe in the experiment.

5 Conclusions

We estimated the parameters of a sequential trading model by using experimental data. We found that about one third of subjects' decisions were irrational. Most of these irrational decisions were no trades. Nevertheless, some decisions were sell and, more often, buy orders. Hence, irrationality by market participants seems to explain some noise trading in financial markets.

References

[1] Cipriani, Marco and Antonio Guarino (2004). "Herd Behavior in a Laboratory Financial Market." Mimeo, UCL and GWU.

[2] Easley David, Nicholas Kiefer, and Maureen O'Hara (1997). "One Day in the Life of a Very Common Stock." *Review of Financial Studies*, 10 (3), pp. 805-835.

[3] Glosten, Lawrence and Paul Milgrom (1985). "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders."

Journal of Financial Economics, 14 (1), pp. 71-100.

[4] Milgrom, Paul and Nancy Stokey (1982). "Information, Trade and Common Knowledge." *Journal of Economic Theory*, 26 (1), pp. 17-27.