

Transaction Costs and Informational Cascades in Financial Markets: Theory and Experimental Evidence

Marco Cipriani and Antonio Guarino*

February 2007

Abstract

We study the effect of transaction costs (e.g., a trading fee or a transaction tax, like the Tobin tax) on the aggregation of private information in financial markets. We analyze a financial market à la Glosten and Milgrom, in which informed and uninformed traders trade in sequence with a market maker. Traders have to pay a cost in order to trade. We show that, eventually, all informed traders decide not to trade, independently of their private information, i.e., an informational cascade occurs. We replicated our financial market in the laboratory. We found that, in the experiment, informational cascades occur when the theory suggests they should. Nevertheless, the ability of the price to aggregate private information is not significantly affected. (JEL C92, D8, G14)

Keywords: Informational Cascades, Herd Behavior, Trade Costs, Tobin Tax

*Cipriani: Department of Economics, George Washington University; E-mail: marco.cipriani@gwu.edu. Guarino: Department of Economics and ELSE, University College London; E-mail aguarino@ucl.ac.uk. The work was in part carried out when Cipriani visited the ECB as part of the ECB Research Visitor Programme; Cipriani would like to thank the ECB for its hospitality. We are especially grateful to Douglas Gale and Andrew Schotter for their comments. We gratefully acknowledge the financial support of the C.V. Starr Center and the Center for Experimental Social Science (C.E.S.S.) at New York University. Guarino gratefully acknowledges the support of the Economic and Social Research Council (ESRC) through ELSE. All errors are ours.

1 Introduction

There is a long and widespread debate on the role of transaction costs in the functioning of financial markets. Such costs can be imputed to different reasons, like brokerage commission fees, time involved in acquiring knowledge and record keeping, transaction taxes and so on. It has been argued that transaction costs can have negative effects on the process of price discovery. The presence of even very small costs makes rebalancing expensive. Therefore, valuable information can be held back from being incorporated into prices.

In recent years, such a debate has gained new strength, following the proposals by many policy makers around the world to introduce security transaction taxes (e.g., a Tobin tax),¹ especially in emerging markets. Opponents have stressed how these taxes can reduce the efficiency with which markets aggregate information dispersed among their participants. In contrast, proponents have argued that such taxes only reduce excessive trading and volatility, and can prevent the occurrence of financial crises.²

In this paper, we will contribute to the debate on the role of transaction costs in financial markets by developing a theoretical analysis and testing its results in a laboratory experiment. We will focus, in particular, on the role of transaction costs in the process of information aggregation.

We will first present a theoretical model similar to that of Glosten and Milgrom (1985), in which traders trade an asset with a specialist (market maker). The specialist sets the prices at which traders can buy or sell. The prices are efficiently set according to the order flow, i.e., to the sequence of trades. Traders can buy or sell one unit of the asset or abstain from trading. Some traders trade to exploit their private information and others do it for non-informational (e.g., liquidity) motives. If traders decide to trade, they have to pay a transaction cost.

We will show that the presence of transaction costs has a significant effect on the ability of the price to aggregate private information dispersed among different agents. Transaction costs cause “informational cascades,” i.e., situations in which all informed traders neglect private information and abstain

¹In 1978, Tobin proposed that all major countries should introduce a tax on foreign-exchange transactions.

²For a review of this debate, see Habermeier and Kirilenko (2003) and the comment by Forbes (2003).

from trading.³ Such blockages of information can occur when the price is far away from the fundamental value of the asset. Therefore, transaction costs can cause long lasting misalignments between the price and the fundamental value of a security. Not only informational cascades are possible in our economy, actually they occur with probability one. Eventually, the trade cost overwhelms the importance of the informational advantage that the traders have on the market maker and, therefore, informed agents prefer not to participate in the market, independently of their private information.

Our results contrast with those of Avery and Zemsky (1998), who show that informational cascades cannot occur in financial markets where trade is frictionless. In their work, agents always find it optimal to trade on the difference between their own information (the history of trades and the private signal) and the commonly available information (the history of trades only). For this reason, the price aggregates the information contained in the history of past trades correctly. Eventually, the price converges to the realized asset value. With trade costs, in contrast, the convergence of the price to the fundamental value does not occur, since trading stops after a long enough history of trades. With positive probability no-trade cascades occur when the price is far from the fundamental value of the asset, even for a very small transaction cost.

To test our theory, we have run an experiment. A laboratory experiment is particularly fit to test our theory, since in the laboratory we observe the private information that subjects receive and we can study how they use it while trading. In our laboratory market, subjects receive private information on the value of a security and observe the history of past trades. Given these two pieces of information, they choose, sequentially, if they want to sell, to buy or not to trade one unit of the asset at a price efficiently set by a market maker. By observing the way in which they use their private information, we directly detect the occurrence of cascades. The experimental results are in line with the theoretical model. Indeed, cascades form in the laboratory as the theory predicts, i.e., when the trade costs overwhelms the gain to trading.

It is worth mentioning that the Avery and Zemsky (1998) model has recently been tested and has found support in experimental studies. Cipriani and Guarino (2005) and Drehman et al. (2005) have found that informational cascades rarely occur in a laboratory financial market with no financial fric-

³For a review of the literature on informational cascades, see Gale (1996), Hirshleifer and Teoh (2003) and Chamley (2004).

tions. We relate our experimental findings to these studies. We find that, even if with trade cost cascades do arise, the ability of the price to aggregate private information is not significantly affected. This happens because, with transaction costs, there is a lower incidence of irrational behavior, and, in particular, of trading against one's own private information. Indeed, in our experiment, the level of rationality is very high, higher than in Cipriani and Guarino (2005). Such a higher level of rationality stems from a sharp reduction in the frequency by which traders disregard their private information (for instance to act as contrarians and trade against the market). This higher level of rationality makes the price reflect private information more accurately. This explains why the overall impact of transaction costs on the market efficiency is very small.

Our theoretical analysis lends credibility to the arguments against the introduction of a security transaction tax (like a Tobin tax): it is true that financial frictions impair the ability of prices to aggregate private information by making informational cascades possible. The occurrence of cascades is confirmed by the laboratory experiment. As proponents of the tax have suggested, however, financial frictions also reduce the occurrence of irrational behavior, which improves the ability of the price to reflect subjects' private information. These two effects offset each other in the laboratory, so that the transaction cost does not significantly alter the financial market informational efficiency.

The structure of the paper is as follows. Section 2 describes the theoretical model and its predictions. Section 3 presents the experimental design. Section 4 illustrates the experimental results. Section 5 concludes.

2 The Theoretical Analysis

2.1 The model structure

Our analysis is based on a Glosten and Milgrom (1985) type of model. In our economy there is one asset traded by a sequence of traders who interact with a market maker. Time is represented by a countably infinite set of trading dates indexed by $t = 1, 2, 3, \dots$

The market

The fundamental value of the asset, V , is a random variable distributed on $\{0, 100\}$, with $Pr(V = 100) = p$. At each time t , a trader can exchange

the asset with a specialist (market maker). The trader can buy, sell or decide not to trade. Each trade consists of the exchange of one unit of the asset for cash. The trader's action space is, therefore, $\mathcal{A} = \{buy, sell, no\ trade\}$. We denote the action of the trader at time t by X_t . Moreover, we denote the history of trades and prices until time $t - 1$ by H_t .

The market maker

At any time t , the market maker sets the prices at which a trader can buy or sell the asset. When posting these prices, he must take into account the possibility of trading with agents who have some private information on the asset value. He will set different prices for buying and for selling, i.e., there will be a bid-ask spread. We denote the ask price (i.e., the price at which a trader can buy) at time t by a_t and the bid price (i.e., the price at which a trader can sell) by b_t .

Due to unmodeled potential competition, the market maker makes zero expected profits by setting the ask and bid prices equal to the expected value of the asset conditional on the information available at time t and on the chosen action, i.e.,⁴

$$\begin{aligned} a_t &= E(V|h_t, X_t = buy, a_t, b_t), \\ b_t &= E(V|h_t, X_t = sell, a_t, b_t). \end{aligned}$$

We also define the "price" of the asset at time t as the market maker's expected value of the asset *before* the time- t trader has traded, i.e., $p_t = E(V|h_t)$.

The traders

There are a countably infinite number of traders. Traders act in an exogenously determined sequential order. Each trader, indexed by t , is chosen to take an action only once, at time t . Traders are of two types, informed and uninformed (or noise). The trader's type is not publicly know, i.e., it is his private information. At each time t , an informed trader is chosen with probability μ and a noise trader with probability $1 - \mu$. Noise traders trade for unmodeled (e.g., liquidity) reasons: they buy, sell and do not trade

⁴We use capital letters for random variables and small letters for their realizations. For instance, h_t is a particular history realized at time t , while H_t denotes any possible history until that time. Furthermore, to simplify the notation, we write $E(.|y)$ to mean $E(.|Y = y)$, i.e., the expected value conditional on the realization y of the random variable Y .

with fixed, positive probabilities. Informed traders have private information on the asset value. If an informed trader is chosen to trade, he observes a private signal on the realization of V . The signal is a random variable S_t distributed on $\{0, 100\}$. We denote the conditional probability function of S_t given a realization v of V by $\sigma(s_t|v)$. We assume that, conditional on the asset value v , the random variables S_t are independent and identically distributed across time. In particular, we assume that

$$\sigma(0|0) = \sigma(100|100) = q > 0.5.$$

In addition to his signal, a trader at time t observes the history of trades and prices and the current price. Therefore, his expected value of the asset is $E(V|h_t, s_t)$.

Trading is costly: if a trader decides to buy or sell the asset, he must pay a transaction cost $c > 0$. Every trader is endowed with an amount $k > 0$ of cash. His payoff function $U : \{0, 100\} \times \mathcal{A} \times [0, 100]^2 \rightarrow \mathbf{R}^+$ is defined as

$$U(v, X_t, a_t, b_t) = \begin{cases} v - a_t + k - c & \text{if } X_t = \textit{buy}, \\ k & \text{if } X_t = \textit{no trade}, \\ b_t - v + k - c & \text{if } X_t = \textit{sell}. \end{cases}$$

The trader chooses X_t to maximize $E(U(V, X_t, a_t, b_t)|h_t, s_t)$. Therefore, he finds it optimal to buy whenever $E(V|h_t, s_t) > a_t + c$, and sell whenever $E(V|h_t, s_t) < b_t - c$. He chooses not to trade when $b_t - c < E(V|h_t, s_t) < a_t + c$. Finally, he is indifferent between buying and no trading when $E(V|h_t, s_t) = a_t + c$ and between selling and no trading when $E(V|h_t, s_t) = b_t - c$.

2.2 Predictions

Before moving to the main analysis, it is worth remarking that in this economy there exist a unique ask and a unique bid at which the market maker makes zero profits in expected value. As long as there is adverse selection, the market maker sets a bid-ask spread (Glosten and Milgrom, 1985).

Proposition 1 *At each time t , there exist a unique bid and a unique ask price. Moreover, $b_t \leq p_t \leq a_t$.*

Proof. See the Appendix.

In order to discuss the model's predictions, let us introduce the concept of informational cascade. An informational cascade is a situation in which it

is optimal for a rational agent to make a decision independently of his private information (i.e., to ignore his own private signal).

Definition 1 *An informational cascade arises at time t when*

$$\Pr[X_t = x|h_t, S_t = s] = \Pr[X_t = x|h_t]$$

for all $x \in \mathcal{A}$ and for all $s \in \{0, 100\}$.

In an informational cascade, the market maker is unable to infer the trader's private information from his actions and to update his beliefs on the asset value. In other words, in an informational cascade trades do not convey any information on the asset's fundamental value.

>From a behavioral point of view, an informational cascade can, potentially, correspond to three different trading behaviors, which are of interest in the analysis of financial markets:

1) Traders may disregard their private information and *conform* to the established pattern of trade (e.g., after many buy orders, they all buy irrespective of their signal). In this case, we say that the traders are *herding*.

2) Alternatively, traders may disregard their private information and trade *against* the established pattern of trade (e.g., after many buy orders, they all sell irrespective of their signal). In this case, we say that the traders are engaging in *contrarian behavior*.

3) Finally traders may simply *abstain from trading independently of their private information*. In this case we say that there is a *no-trade cascade*.

We now prove that, in equilibrium, the first two situations never arise: traders will never trade against their signal, i.e., engage in herding or contrarian behavior. In contrast, an informational cascade of no trades arises almost surely as t goes to infinity.

Proposition 2 *In equilibrium, a) a trader either trades according to his private signal or abstains from trading; b) an information cascade in which a trader does not trade independently of his signal occurs almost surely as t goes to infinity.*

Let us first discuss point a. To decide whether he wants to buy, to sell or not to trade the asset, an agent computes his expected value and compares it

to the ask and bid prices. If at time t he receives a signal of 100, his expected value will be

$$E(V|h_t, S_t = 100) = 100 \Pr(V = 100|h_t, S_t = 100) =$$

$$100 \frac{q \Pr(V = 100|h_t)}{q \Pr(V = 100|h_t) + (1 - q)(1 - \Pr(V = 100|h_t))} >$$

$$100 \Pr(V = 100|h_t) = E(V|h_t) = p_t.$$

Similarly, if he receives a signal of 0, his expected value will be

$$E(V|h_t, S_t = 0) < E(V|h_t) = p_t.$$

Given that, according to the previous result, $b_t \leq p_t \leq a_t$, an agent will never find it optimal to trade against his private information, i.e., to buy despite a signal 0 or to sell despite a signal 100.⁵

Let us now consider point b: as t goes to infinity, agents decide not to trade with probability 1. We refer the reader to the Appendix for a formal proof. Here let us note that, over time, as the price aggregates private information, the informational content of the signal becomes relatively less important than that of the history of trades. Therefore, after a sufficiently large number of trades, the valuation of any trader and of the market maker will be so close that the expected profit from trading will be lower than the transaction cost. Every trader, independently of his signal realization, will decide not to trade.

[Insert Figure 1 about here]

⁵In equilibrium, a trader's expectation conditional on receiving a signal equal to 100 (0) is not only greater (smaller) than the price, it is also greater (smaller) than the ask (bid) price, i.e., $E(V|h_t, S_t = 0) < b_t \leq p_t \leq a_t < E(V|h_t, S_t = 100)$. If at time t a trader receives a signal of 100, his expected value will be $E(V|h_t, S_t = 100) = 100 \Pr(V = 100|h_t, S_t = 100)$. On the other hand, $a_t = E(V|h_t, X_t = buy) = 100 \Pr(V = 100|h_t, X_t = buy)$. Clearly, $\Pr(V = 100|h_t, X_t = buy) = \Pr(V = 100|h_t, S_t = 100) \Pr(S_t = 100|h_t) + \Pr(V = 100|h_t) \Pr(Noise|h_t)$, where, with an abuse of notation, $\Pr(S_t = 100|h_t)$ is the probability that the traders in t is informed *and* has received a signal of 100, and $\Pr(Noise|h_t)$ is the conditional probability that the trader in t is noise. Since all these probabilities are strictly positive and $100 \Pr(V = 100|h_t) = p_t < 100 \Pr(V = 100|h_t, S_t = 100)$, it is immediate to see that $E(V|h_t, S_t = 100) > a_t$. By a similar argument, $E(V|h_t, S_t = 0) < b_t$.

Figure 1 illustrates this situation. In an economy with no trade cost, in equilibrium the market maker would set a bid higher than the expected value of a trader with a negative signal and an ask lower than the expected value of a trader with a high signal. As a result, any trader would find it optimal to act according to his private information. The trade cost eliminates such an incentive. The gain from trade of a trader with a positive or negative signal is lower than the trade cost. Therefore, any informed trader decides not to trade. In equilibrium, since trades are uninformative, the market maker sets both the bid and the ask equal to the price, i.e., $b_t = p_t = a_t$.

Note that, when an informational cascade arises, it never ends. This happens because, during a cascade, a trade does not convey any information on the asset value. After a cascade has started, all the following traders have the same public information as their predecessor (i.e., nothing is learned during the cascade). Therefore, the subsequent traders will also find it optimal to neglect their own private information and decide not to trade, thereby continuing the cascade.

This result on traders' behavior has an immediate implication for the price during a cascade. The market maker will be unable to update his belief on the asset value and, as a consequence, the price will not respond to the traders' actions. The price will remain stuck for ever at the level it reached before the cascade started. Note that such a level may well be far away from the fundamental value of the asset. Therefore, in a market with transaction costs, long-run misalignments between the price and the fundamental value arise with positive probability. This suggests that, as opponents of security transaction taxes have claimed, introducing such taxes can significantly impair the process of price discovery.⁶

The impact of transaction costs in financial markets has been discussed theoretically also by Lee (1998) in a setup different from that of Glosten and Milgrom (1985). Lee studies a sequential trade mechanism in which traders are risk averse. They receive signals of different precisions. Traders can trade more than once with the market maker and can buy or sell different quantities (shares) of the asset. The market maker, however, is allowed to set one price only, i.e., in contrast with our setup, there is no bid-ask spread, nor

⁶Securities transaction taxes are usually levied as a percentage tax on the amount transacted and not as a lump-sum tax. It is easy to show, however, that the theoretical results that we present in this section would also hold if the transaction cost were modeled as a percentage. We chose a lump-sum transaction cost only because it is easier to implement in the laboratory.

can the market maker set different prices according to the trade size. In Lee (1998), transaction costs may trigger quasi-cascades that result in temporary information blockages followed by informational avalanches in which previously hidden private information is suddenly revealed. Eventually, however, a complete stop of information occurs. The paper by Lee is very important, as it first discusses the role of transaction costs in the process of information aggregation. In particular, the idea that a gradual buildup of private information is suddenly revealed is obtained as an equilibrium outcome. It should be noted, however, that his results depend on the market maker ignoring the information contained in the current trade. As noted by Chamley (2004), in the absence of this assumption, these results are unlikely to hold. Such an assumption is difficult to reconcile with trading between rational agents. In our model, in contrast, as in any model à la Glosten and Milgrom, the market maker does take into account the information potentially contained in the present trade, by setting a bid-ask spread. Despite this, an informational cascade occurs almost surely.

2.3 A simplified model

In the remaining part of the paper we will discuss whether the predictions of the model are supported or not in a laboratory financial market. To implement the model in the laboratory, we made two important simplifications. We only had informed traders in the market (i.e., there were no noise traders) and we allowed the market maker to set one price only (instead of a bid price and an ask spread).⁷ These changes were aimed at simplifying the experiment, without losing any important characteristic of the theoretical model.

Of course, since every trader was informed and the market maker could not set a bid-ask spread, he would make losses in expected value. This was not a problem in the experiment, since an experimenter played the role of the market maker. He chose the price in order to minimize the expected losses. It is easy to show that the equilibrium price at each time t is just the expected value of the asset conditional on the information available at time t , i.e., p_t .

It is important to note that Proposition 2 holds true in this simplified version of the model too. Indeed, we have proved above that $E(V|h_t, S_t =$

⁷In the absence of noise traders, had we allowed for a bid-ask spread, the equilibrium spread would have been so wide that any incentive to trade for informed subjects would have been eliminated.

$0) < p_t < E(V|h_t, S_t = 100)$, which implies that a trader should never trade against private information when facing a price equal to the expected value of the asset conditional on the history of trades. Furthermore, also in this simplified version of the model, the expected profit from buying or selling the asset is eventually overwhelmed by the trade cost and no-trade cascades occur almost surely.

In the experiment we also set particular values for the relevant parameters of the model. In particular, we assumed the probability of the asset value being 100, p , equal to $\frac{1}{2}$, and the precision of the signal, q , equal to 0.7. Finally, we set the trade cost, c , equal to 9, which corresponds to 18% of the unconditional expected value of the asset. This level of the transaction cost is obviously greater than actual transaction costs, or than reasonable taxes on financial transactions.⁸ We chose such a high level so that the probability of the transaction cost becoming binding was high enough to offer a sufficient number of observations in our analysis. In theory, the transaction costs become binding with probability 1 irrespective of their size, as the number of trading dates goes to infinity. In the experiment, the number of trading dates is, of course, finite and this dictated the choice of c .⁹

3 The Experiment and the Experimental Design

3.1 The experiment

This was a paper and pencil experiment. We recruited subjects from undergraduate courses in all disciplines at New York University. They had no previous experience with this experiment. In total, we recruited 52 stu-

⁸According to Domowitz, Glen, and Madhavan (2000), in the U.S. stock market, trading costs in the period 1990–1998 were equal to 2.2% of the mean returns. Habermeier and Kirilenko (2003) report that, in Sweden, between 1984 and 1991, taxes on equity trading were levied at 0.5% and taxes on option trading were levied at 1%. These figures are lower than our cost. One needs, however, to consider that, following an informational event, the number of trades on a stock is very high. As a result, even low levels of transaction costs or taxes can trigger a cascade.

⁹In our experiment, each sequence of trades consists of 12 decisions, a relatively large number compared to similar studies in the literature.

dents to run 4 sessions.¹⁰ In each session we used 13 participants, one acting as subject administrator and 12 acting as traders. The procedure was the following:

1. At the beginning of the sessions, we gave written instructions (available from the authors on request) to all subjects. We read the instructions aloud in an attempt to make the structure of the game common knowledge to all subjects. Then, we asked for clarifying questions, which we answered privately. Each session consisted of ten rounds of trading. In each round we asked all subjects to trade one after the other.
2. The sequence of traders for each round was chosen randomly. At the beginning of the session each subject picked a card from a deck of 13 numbered cards. The number that a subject picked was assigned to him for the entire session. The card number 0 indicated the subject administrator. In each round, the subject administrator called the subjects in sequence by randomly drawing cards (without replacement) from this same deck.
3. Before each round, an experimenter, outside the room, tossed 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, the same experimenter stayed 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. Each subject, after his number was called, had to go outside the room and draw a chip from one bag. If the coin landed tail the experimenter used the first bag, otherwise he used the second. Therefore, the chip color was a signal for the value of the asset. After looking at the color, the subject put the chip back into the bag. Note that the subject could not reveal the chip color to anyone.
4. In the room, another experimenter acted as market maker, setting the price at which people could trade. After observing the chip color, the subject entered the room. He read the trading price on the blackboard

¹⁰Subjects were recruited by sending an invitation to a large pool of potential participants. For each session of the experiment, we received a large number of requests to participate. We chose the students randomly, so that the subjects in the experiment were unlikely to know each other.

and, then, declared aloud whether he wanted to buy, to sell or not to trade. The subject administrator recorded all subjects' decisions and all trading prices on the blackboard. Hence, each subject knew not only his own signal, but also the history of trades and prices.¹¹

5. At the end of each round, i.e., after all 12 participants had traded once, the realization of the asset value was revealed and subjects were asked to compute their payoffs. All values were in a fictitious currency called lira. In each round students were given an endowment of 100 lire. Each time a student decided to buy or sell the asset, he would have to pay a trade cost of 9 lire. Therefore, students' payoffs were computed as follows. In the event of a buy, the subject obtained $100 + (v - p_t) - 9$ lire; in the event of a sell, he obtained $100 + (p_t - v) - 9$ lire; finally, if he decided not to trade he earned 100 lire. After the tenth round, we summed up the per round payoffs and converted them into dollars at the rate of $\frac{1}{65}$. In addition, we gave \$7 to subjects just for participating in the experiment. Subjects were paid in private immediately after the experiment and, on average, earned \$22.50 for a 1.5 hour experiment.

3.2 The Price Updating Rule

The price was updated after each trade decision in a Bayesian fashion. When a subject decided to buy, the price was updated up, assuming that he had seen a positive signal. Similarly, when a subject decided to sell, the price was updated down, assuming that the subject had observed a negative signal. In the case of no trade, the price was kept constant. The rationale for this rule is very simple. In equilibrium, when the trade cost was smaller than the expected profit from buying or selling the asset, subjects should have always followed their signal, i.e., they should have bought after seeing a positive signal and sold after seeing a negative one. On the other hand, not trading was an equilibrium decision when the expected profit from buying or selling the asset was not higher than the trade cost. Therefore, in equilibrium, a buy would reveal a positive signal, a sell a negative signal, and a no trade would be uninformative.

¹¹Subjects were seated far away from each other, all facing the blackboard. No communication was allowed in the room. The entrance was in the back of the classroom. When making his decision, the subject was facing the blackboard, but not the other participants.

It is important to remark that we chose a trade cost of 9 to avoid the possibility that the cost could be binding only upon receiving a negative signal and not upon receiving a positive one (or vice versa). Let us consider the case of a series of buy orders. With our updating rule, after a series of buys, the price moved from 50 to 70, 84, 93, 97,...At a price of 84, a subject would be (theoretically) indifferent between buying and not trading upon receiving a positive signal and would strictly prefer to sell upon receiving a negative one.¹² Had we chosen another level of the trade cost (for instance 10), theoretically a no trade would have clearly revealed that the subject received a positive signal (since he could never decide not to trade with a negative signal). To be consistent with the theoretical framework, we should have updated the price. This would have made the updating rule quite complicated and difficult to explain to the subjects.¹³

Of course in the experiment we could observe some decisions off the equilibrium path: for instance, a decision not to trade when (theoretically) the

¹²This happens because, when the price is 84 (i.e., after two buys), a trader with a negative signal has an expected value of 70 and therefore would make a $84 - 70 - 9 = 5$ expected profit by selling the asset. In contrast, a trader with a positive signal has an expected value of 93 and therefore would make a $93 - 84 - 9 = 0$ expected profit by buying the asset.

¹³Keeping the price constant after a no trade is optimal for the market maker only if we impose some assumptions on the subjects' behavior in the case of indifference. Not updating the price after a no trade at the price of 84 is equivalent to assuming that an indifferent subject in this case buys the asset, so that a no trade always is an off the equilibrium path decision and does not convey any information on the signal that the subject received.

Similarly, at a price of 93, a subject with a negative signal would value the asset 84 and, therefore, would be indifferent between selling and no trading. A subject with a positive signal would value the asset 97 and, because of the trade cost, would strictly prefer to abstain from trading. Not updating the price after a no trade at a price of 97 is equivalent to assuming that an indifferent subject in such a case does not trade, so that a no trade does not convey any information on the private signal.

After a sequence of sell orders, the price moves from 50 to 30, 16, 7, 3....Similar considerations to those above can be made for no trade decisions at prices of 16 or 7.

The assumptions we made in the cases of indifference, theoretically legitimate, turn out to be fairly consistent with actual behavior in the laboratory. Indeed, over the whole experiment, the frequency of a no trade conditional on receiving a bad signal was very close to the probability of a no trade conditional on receiving a good signal (52% and 51% respectively), which implies that a rational market maker would have not updated the price after a no trade. Given subjects' behavior, for a price of 84, 93, 16 and 7, after a no trade, a market maker should have updated the price to 86, 93, 20 and 7, very close to our updating rule.

trading cost was not binding; or a decision to buy, or sell, when the trade cost (theoretically) overwhelmed the gain to trading. We assumed that if someone bought despite the trade cost, he must have received a positive signal, and, accordingly, updating the price up was a sensible decision. If someone sold despite the trade cost, he must have received a negative signal, and, again, updating the price down was reasonable. Finally, if someone decided not to trade even though the trade cost was not binding, we did not update the price. This updating rule is identical to the one used by Cipriani and Guarino (2005) to analyze a financial market where there are no frictions; therefore, it also facilitates the comparison between our results and those of that experiment.

It is worth mentioning that a great advantage of our setting is that, as already mentioned, the price moved through a grid. There were only few values at which the price was set during the entire experiment.¹⁴ In the first three rounds (which we do not consider in our data analysis), subjects had the opportunity to see exactly how the price moved in response to the order flow.

In the next section we describe the results. The results refer to the last seven rounds of each session only.¹⁵ We do not take the first three rounds into account for two reasons. First, although the experiment was very easy and subjects did not have problems in understanding the instructions, we believe that some rounds were needed to acquaint subjects with the procedures. By considering only the last 7 session, we concentrate on the decisions taken after the learning phase. Second, considering only the last seven rounds makes easier to compare our results to those presented in Cipriani and Guarino (2005).¹⁶

¹⁴See the previous note.

¹⁵In each round, the 12 subjects were asked to trade in sequence. Therefore, the results refer to 336 decisions.

¹⁶We replicated the analysis considering all ten sessions and we obtained very similar results (available on request) to those described in the remainder of the paper.

4 Results

4.1 Trading Behavior

Now we present the results of our experiment. We will find it useful to compare our results to those obtained by Cipriani and Guarino (2005). Cipriani and Guarino (2005) ran several treatments, studying a financial market similar to that described in this paper, but with no transaction costs. The treatment that they label “flexible price” is identical to ours, with the exception that there is no transaction cost. Therefore, we can use that as a clear benchmark to evaluate the impact of transaction costs on subjects’ behavior and on the resulting price paths. We use the label CG-FP to refer to the flexible price treatment in Cipriani and Guarino (2005).

Let us start by discussing the average level of rationality in the experiment. Given the sequential structure of the game, to classify a decision as rational or irrational, we need to make some assumptions on each subject’s belief on the choices of his predecessors. Following Cipriani and Guarino (2005) we assume that each subject believes that all his predecessors are rational, that all his predecessors believe that their predecessors are rational and so on. Under this assumption, a rational subject should always behave as predicted by the theoretical model.¹⁷

The level of rationality is high (see Table 1): 82% of the overall decisions in the laboratory were rational, i.e., not in contrast with the theoretical model, while only 18% of actions were irrational. Such level of rationality is higher than the 65% in CG-FP.¹⁸ This increase in the level of rationality is mainly due to the drop in the proportion of irrational no trades. In CG-FP subjects should have always traded, in order to exploit their informational

¹⁷In the experiment, subjects sometimes made decisions off the equilibrium path, i.e., decisions that could not be the outcome of a rational choice. An important issue is how to update subjects’ beliefs after they observe such decisions. If the decision is a no trade, we assume that the following subjects do not update their beliefs (which is consistent with our price updating rule). If the decision is a buy (sell), we assume (again, consistently with our price updating rule) that the following subjects update their beliefs as though it publicly revealed a positive (negative) signal (i.e., the signal that implies the lower expected loss). This last assumption is, in fact, quite innocuous. For instance, if we assumed that trades that cannot be the outcome of a rational choice do not convey any information, our results would be virtually identical since, in the entire experiment, we observed only 4 of such trades (out of 336 decisions).

¹⁸According to the Mann-Whitney test, the proportion of rational trades is significantly different in the two experiments (p-value=0.03).

advantage on the market maker. Nevertheless, they decided not to trade in 22% of the cases, which added to the level of irrationality. In our experiment, the proportion of no trades was significantly higher (51%).¹⁹ These no trade decisions, however, happened mostly when they should have occurred according to theory: indeed, 79% of no trades occurred when the difference between the trader's expected value and the price was not higher than the trade cost and, therefore, trading could not be profitable.

¹⁹The p-value for the null that the proportion of no trades is the same in the two experiments equals 0.03 (Mann-Witney test).

Table 1: Rational and irrational decisions

Transaction Costs			
<i>Rational Decisions</i>			<i>82%</i>
	Buying or Selling	41%	
	No Trading	41%	
<i>Irrational Decisions</i>			<i>18%</i>
	Buying or Selling	8%	
	No Trading	10%	
<i>Total</i>			<i>100%</i>
CG-FP			
<i>Rational Decisions</i>			<i>65%</i>
<i>Irrational Decisions</i>			<i>35%</i>
	Buying or Selling	13%	
	No Trading	22%	
<i>Total</i>			<i>100%</i>

In order to understand better the increase in no trade decisions that we observe with trade costs, we computed the proportion of no trade decisions for different levels of the absolute value of the trade imbalance. The trade imbalance at time t is defined as the difference between the number of buys and the number of sells taken by subjects from time 1 until time $t - 1$. As the trade imbalance increases in absolute value, the expected profit from trading becomes smaller and smaller, thus reducing trader's incentives to trade upon their information.

Table 2:
**Proportion of no-trade decisions for different levels of the
absolute value of the trade imbalance**

Transaction Costs	
Trade Imbalance	No Trades
0-1	21%
2-3	67%
>3	78%
CG-FP	
Trade Imbalance	No Trades
0-1	16%
2-3	22%
>3	33%

Both in our experiment and in CG-FP, the frequency of no trades increases monotonically with the absolute value of the trade imbalance. In CG-FP, however, this increase is modest, the proportion of no trades going up from 16% with a trade imbalance between -1 and 1 to 33% with a trade imbalance greater than 3 or smaller than -3 . In contrast, in our experiment the proportion of no trades jumps from 21% when the absolute value of the trade imbalance is at most 1 to 67% when the absolute value is 2 or 3 . The proportion of no trades is even higher when the trade imbalance is higher than 3 or lower than -3 . Since for a trade imbalance of at most 1 not trading was irrational, while it was rational after a trade imbalance of 2 , these results confirm that no trade decisions occurred mainly when they were rational. In a nutshell, the trade cost had a significant impact on subjects' decisions exactly when theory suggests that it should become binding.

Rational no trade decisions are not the only reason for the increase in the level of rationality in our experiment as compared to CG-FP. The increase in the level of rationality is also due to a different proportion of trading against private information. In CG-FP, 13% of all decisions are irrational buys and sells against private information. In contrast, in our experiment only 6% of decisions are irrational trades against the private signal.²⁰

²⁰Using the Mann-Whitney test, the p-value for the null that the proportion of trades against private information is the same in the two experiments equals 0.03. Note that the

Cipriani and Guarino (2005) explain part of the irrational buying or selling decisions in their experiment by invoking contrarian behavior and a modest proportion of herding behavior. Both contrarianism and herding can occur for a level of the trade imbalance which is high in absolute value. Herding refers to the situation in which a trader with a signal in contrast with the past history of trades (i.e., a positive signal after many sells or a negative signal after many buys) decides to disregard his private information and follow the market trend.²¹ In contrast, contrarianism refers to the situation in which a trader with a signal which reinforces the past history of trades (i.e., a positive signal after many buys or a negative signal after many sells) decides to disregard both his private information and the market trend.²² In Table 3 we report the proportion of herding and contrarian behavior. In CG-FP, subjects herded 12% of the time in which herding could arise, and they acted as contrarians in 19% of the time in which contrarianism could occur. In contrast, in our experiment, when the absolute value of the trade imbalance was greater than 2, most of the time subjects preferred not to trade:²³ this explains why both contrarianism and herding are even less pronounced in our experiment than in CG-FP, which in turn explains why the proportion of irrational trades is lower.²⁴

6% indicated in the text differs from the percentage (8%) in Table 1 since 2% of trades were irrational not because they did not agree with private information, but because the trade cost was greater than the expected profit from trading.

²¹In particular, following Cipriani and Guarino (2005), we say that in period t there is a situation of potential herd behavior when the trade imbalance (in absolute value) is at least 2 and the subject receives a signal against it. If the subject trades against the signal we say that he herds.

²²In particular, following Cipriani and Guarino (2005), we say that in period t there is a situation of potential contrarian behavior when the trade imbalance (in absolute value) is at least 2 and the subject received a signal agreeing with it. If the subject trades against the signal we say that he engages in contrarian behavior.

²³Note that, according to the Mann-Whitney test, the proportion of contrarianisms is significantly different in the two experiments (p -value = 0.03), but the proportion of herding is not (p -value=0.26). This result confirms the finding by Cipriani and Guarino (2005) that in the presence of the price mechanism herd behavior rarely arises.

²⁴It is also interesting to note that, when the trade imbalance (in absolute value) was at most 1 (and, therefore, the trade cost was not binding), irrational decisions to buy or sell amount to 4%, whereas they amount to 9% in CG-FP. Although this difference is not significant (p -value=0.14), it seems to suggest that, with trade costs, subjects are more careful in making their decisions (and, therefore, they use their information more efficiently). This may have added to the rationality in the experiment.

Table 3: Herding and Contrarian Behavior

Herd Behavior	Transaction Costs	CG-FP
Herding	7%	12%
Following Private	29%	46%
No Trade	64%	42%
Relevant cases	70	66
Contrarian Behavior		
Contrarian behavior	7%	19%
Following private	23%	63%
No trade	69%	18%
Relevant cases	147	132

4.2 No-trade Cascades

Table 4 shows the sequences of traders' decisions during the experiment. We highlighted in gray those periods in which a subject took a decision facing a price greater than or equal to 93 or smaller than or equal to 7 and, therefore, could have rationally chosen not to trade independently of his private information. In bold we indicated those decisions that were indeed no trades.

The table clearly illustrates the pervasiveness of no trade decisions. In many rounds (19 out of the 28) there were indeed long sequences of no trades. In almost all the cases, the no trade decisions started when the price reached the level of 7 or of 93. In four rounds only ($IV - 4$, $IV - 5$, $IV - 7$, $IV - 8$) sequences of no trades started at a price of 16 or 84 and, even more rarely, at prices closer to 50, the unconditional expected value of the asset.²⁵

[Insert Table 4 about here]

We must stress that our experiment offers a particularly *tough* test to the prediction that, when trade is costly, cascades of no trades occur in

²⁵The roman numeral denotes the session. The Arab numeral indicates the round within each session. For instance, $II - 5$ refers to the fifth round of the second session. Note that we are reporting only the last 7 rounds of each session of the experiment. This is why, for instance, the first round in session II is labeled $II - 4$.

the financial market. At a price of 7 or 93, the expected payoff from trading becomes equal to or lower than zero, depending on the signal that the subject receives. At the price of 93, a rational agent receiving a negative signal has an expected value of 84 and is *indifferent* between selling and not trading (since the trade cost is set at 9).²⁶ Analogously, when the price is 7, a rational agent receiving a positive signal has an expected value of 16 and is *indifferent* between buying and not trading.²⁷ Therefore, in our setup, deviations from a cascade of no trades are not necessarily irrational. A no trade cascade, as defined in Section 2, theoretically occurs at prices of 7 and 93 under the assumption that an indifferent agent always decides not to trade. But, on the other hand, in the laboratory, we may expect some subjects to follow a different (and, still rational) strategy, thus breaking the cascade. Furthermore, it must be noticed that at a price of 84 a rational agent receiving a positive signal is indifferent between buying and not trading. Therefore, theoretically, the price of 93 can only be reached if an indifferent agent buys with positive probability. Similarly, a no trade cascade can arise at the price of 7 only if an indifferent agent at the price of 16 sells with positive probability. In summary, in our set up, cascades are more difficult to form and easier to break.

What happened in the experiment? Given the signal realizations, if all subjects had been rational and followed the rules indicated in the previous section in the cases of indifference, a no trade cascade would have occurred in 25 of the 28 rounds.²⁸ In fact, in the experiment sequences of no trades did occur in 19 out of these 25 rounds. Sequences of no trades often arose in the laboratory and were almost never broken. Therefore, the theoretical prediction of the model finds strong support in the laboratory.

4.3 The Price Path

Figure 2 presents the price paths in the different rounds of the experiment. Recall that we updated the price up after a buy, down after a sell, and

²⁶In contrast, a rational agent receiving a positive signal has an expected value of 97 and, therefore, his optimal choice is to abstain from trading.

²⁷In contrast, a rational agent receiving a negative signal has an expected value of 7 and, therefore, his optimal choice is to abstain from trading.

²⁸In particular, this is true under the assumption that an indifferent agent buys (sells) with probability 1 when the price is 84 (16) (which is the assumption used to compute the equilibrium price).

kept it constant after a no trade. Therefore, the no trade cascades which we described above manifest themselves as sequences of periods in which the price remained constant. These sequences, as it is easy to observe, were quite pervasive during the experiment.

Let us now discuss how this trading behavior affected the ability of prices to aggregate private information. In order to do so, we analyzed the distance between the final actual price (i.e., the price after all subjects have traded) and the full information price (i.e., the price that the market maker would have chosen if the signals had been public information). Figures 3 and 4 show this distance and contrast it to that in CG-FP. By comparing our results with those in CG-FP, two interesting differences emerge. First, in our experiment we have one instance in which distance between the final price and the full information price was greater than 50. This was the instance in which a misdirected cascade arose (i.e., the full information price, 84, was above the unconditional expected value, whereas the actual price, 16, was below it). This never happened in CG-FP. Second, in CG-FP 50% of the time the difference between the final price and the full information price was less than 5 lire, whereas this happens only 21% of the time in our experiment. Since with trade costs subjects often stopped trading whenever the price was 7 or 93, it was less common for the difference between the full information price and the final price to become very small.

The average distance between final actual price and the full information price is 14.5 lire. Had subjects behaved in a perfectly rational manner, such a distance would have been 14 lire.²⁹ This reflects the similarity between the behavior observed in the laboratory and that predicted by the theory.

Cipriani and Guarino (2005) report a distance between last actual prices and full information prices of 12 lire. In their case, however, theoretically the distance should have been 0. Therefore, in CG-FP, there was a misalignment of the price with the fundamental value of the asset (which is expressed by the full information price) due to the irrational trades in the laboratory. In contrast, in our experiment, the inability of the price to aggregate private information completely cannot be attributed to a significant discrepancy between theoretical and actual behavior: indeed the level of irrationality (and,

²⁹This theoretical distance is obtained using, for the cases of indifference, the same assumptions explained above when we discussed informational cascades. An alternative way of computing the last price would be to assume that, when indifferent, subjects always followed their signals. Under this assumption, the average distance would have been 10 lire.

more specifically, the proportion of trades against private information) is significantly lower than in CG-FP. The distance between actual and full information prices is due to the presence of transaction costs, that reduced the incentive of subjects to reveal their private information by placing orders on the market.

In conclusion, the overall effect of trade costs on the informational efficiency of prices is ambiguous: on the one hand, trade costs reduce the incentive for subjects in the laboratory to trade irrationally against their private information; on the other hand, they increase (both theoretically and experimentally) the incentive to abstain from trading altogether. In our experiment, these two forces offset each other and, as a result, the ability of prices to aggregate private information is not significantly different from the case of a frictionless market studied in the laboratory by Cipriani and Guarino (2005).

These results help to shed some light on the possible effect of a tax on financial transactions, like a Tobin tax. It has been argued, by opponents of the tax, that such a tax would generate misalignments of asset prices with respect to the fundamentals. Our analysis supports this view: by introducing a wedge between the expectations of the traders and of the market maker, a tax on financial transaction may prevent the aggregation of the private information dispersed among market participants. On the other hand, as proponents of the tax have suggested, a tax on financial transaction reduces the incidence of irrational trading by market participants, and in particular, of herding and contrarianism. These two effects offset each other in the laboratory so that the introduction of a security transaction tax does not significantly alter the ability of the price to aggregate private information.

5 Conclusions

We have analyzed the effects of transaction costs in financial markets through a theoretical model, and have tested the predictions of the model through a laboratory experiment. We have shown that transaction costs cause informational cascades in which all traders abstain from trading, independently of their private information. This theoretical result is supported by experimental evidence: we observed cascades in our laboratory market when the theory predicts that they should indeed arise. Informational cascades impair the process of information aggregation, and may create a misalignment between

the price and the fundamental value of an asset. In this sense, our theoretical and experimental results highlight the negative effect of transaction taxes (e.g., the Tobin tax) on the process of price discovery. Our experimental results, however, suggest that one should be cautious in concluding that introducing a transaction tax would have a strong effect on the informational efficiency of the financial market. In fact, by comparing the results of our experiment to previous experimental results on markets without frictions, we found that the presence of a transaction cost does not affect the convergence of the price to the fundamental value in a significant way. This is due to the fact that transaction costs reduce the frequency by which agents irrationally trade against their private information.

References

- [1] Avery, Christopher and Zemsky, Peter. "Multidimensional Uncertainty and Herd Behavior in Financial Markets." *American Economic Review*, September 1998, 88(4), pp. 724 – 748.
- [2] Chamley, Christophe. *Rational Herds. Economic Models of Social Learning*. Cambridge: Cambridge University Press, 2004.
- [3] Cipriani, Marco and Guarino, Antonio. "Herd Behavior in a Laboratory Financial Market." *American Economic Review*, 2005, 95(5), pp. 1427 – 1443.
- [4] Domowitz, Glen, and Madhavan, Ananth. "Liquidity, Volatility and Equity Trading Costs across Countries and over Time." *International Finance*, 2001, 4(2), pp. 221 – 55.
- [5] Forbes, Kristin. "Comment on "Securities Transaction Taxes and Financial Markets"." *IMF Staff Papers*, 2003, 50(Special Issue), pp. 181 – 186.
- [6] Gale, Douglas "What Have We Learned from Social Learning?" *European Economic Review*, 1996, 40, pp. 617 – 628.
- [7] Glosten, Lawrence and Milgrom, Paul. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics*, March 1985, 14(1), pp. 71 – 100.

- [8] Habermeier, Karl and Kirilenko, Andrei. “Securities Transaction Taxes and Financial Markets.” *IMF Staff Papers*, 2003, 50(Special Issue), pp. 165 – 180.
- [9] Hirshleifer, David and Teoh, Siew H. “Herd Behaviour and Cascading in Capital Markets: a Review and Synthesis.” *European Financial Management*, March 2003, 9(1), pp. 25 – 66.

6 Appendix

6.1 Proof of Proposition 1

First, we prove the existence of the ask price. Because of unmodeled potential Bertrand competition, the ask price at time t , a_t , must satisfy the condition

$$a_t = E[V|h_t, X_t = 1, a_t, b_t].$$

Let us define I_t as a random variable that takes value 0 if the agent at time t is noise and 1 if he is informed. The expected value of the asset at time t , given a buy order at the ask price a_t is

$$\begin{aligned} E[V|h_t, X_t = 1, a_t, b_t] &= \\ E[V|h_t, X_t = 1, a_t, b_t, I_t = 1] \Pr[I_t = 1|h_t, X_t = 1, a_t, b_t] &+ \\ E[V|h_t] \Pr[I_t = 0|h_t, X_t = 1, a_t, b_t]. & \end{aligned}$$

Let us consider the correspondence $\psi : [0, 100] \rightrightarrows [0, 100]$ defined as $\psi(a_t) := E[V|h_t, X_t = 1, a_t, b_t]$, and let us make the following observations:

1) If $a_t > E[V|h_t, S_t = 100] - c$, $\Pr[I_t = 1|h_t, X_t = 1, a_t, b_t] = 0$ and $E[V|h_t, X_t = 1, a_t, b_t] = p_t$.

2) If $a_t = E[V|h_t, S_t = 100] - c$ then an informed trader receiving $S_t = 100$ can randomize between buying and no trading. Therefore,

$$\begin{aligned} E[V|h_t, X_t = 1, a_t, b_t] &= \\ E[V|h_t, X_t = 1, a_t, b_t, I_t = 1] \Pr[I_t = 1|h_t, X_t = 1, a_t, b_t] &+ \\ E[V|h_t] \Pr[I_t = 0|h_t, X_t = 1, a_t, b_t], & \end{aligned}$$

where the following facts should be noted: if an informed trader receiving $S_t = 100$ buys with zero probability, then $\Pr[I_t = 1|h_t, X_t = 1, a_t, b_t] = 0$ and $E[V|h_t, X_t = 1, a_t, b_t] = E[V|h_t]$. If he buys with probability 1, then $E[V|h_t, X_t = 1, a_t, b_t] = E(V|h_t, S_t = 100) \Pr(I_t = 1, S_t = 100|h_t, y, b_t) + E(V|h_t) \Pr(I_t = 0|h_t)$, where y is an ask price strictly lower than $E[V|h_t, S_t = 100]$ (i.e., the probability $\Pr(I_t = 1, S_t = 100|h_t, y, b_t)$ is computed assuming that an informed trader buys with probability 1). Finally, if he buys with any probability belonging to $(0, 1)$, then $E[V|h_t, X_t = 1, a_t, b_t]$ takes any value between these two levels.

3) If $E[V|h_t, S_t = 0] - c < a_t < E[V|h_t, S_t = 100] - c$, then $E[V|h_t, X_t = 1, a_t, b_t] = E(V|h_t, S_t = 100) \Pr(I_t = 1, S_t = 100|h_t, a_t, b_t) + E(V|h_t) \Pr(I_t = 0|h_t)$, where, again, $\Pr(I_t = 1, S_t = 100|h_t, y, b_t)$ is computed assuming that an informed trader buys with probability 1).

4) If $a_t < E[V|h_t, S_t = 0] - c$, then $E[V|h_t, X_t = 1, a_t, b_t] = E[V|h_t]$.

5) Finally, if $a_t = E[V|h_t, S_t = 0] - c$, then $E[V|h_t, X_t = 1, a_t, b_t]$ takes any value between the two levels found in observations 3 and 4 (the argument to show this is identical to the argument used in observation 2).

Observations 1, 2, 3, 4, and 5 imply that the correspondence ψ is piecewise constant. Furthermore, for $a_t = E[V|h_t, S_t = 100] - c$ and $a_t = E[V|h_t, S_t = 0] - c$, $\psi(a_t)$ takes all the values belonging to the intervals indicated above in observations 2 and 5. Therefore, it is immediate to see that the correspondence $\psi(a_t)$ is non empty, convex-valued and has a closed graph. By Kakutani's fixed point theorem, the correspondence has a fixed point. Furthermore, recalling that $E(V|h_t, S_t = 0) < p_t < E(V|h_t, S_t = 100)$, by using the same observations, it is straightforward to check that such a fixed point is unique.

The proof of the existence and uniqueness of the bid price is analogous.

The proof that $b_t \leq p_t \leq a_t$ follows immediately from the proof in Glosten and Milgrom (1985, p. 81).

6.2 Proof of Proposition 2

To prove the Proposition, we first prove the following lemma.³⁰

³⁰Recall that we use capital letters to indicate random variables and small letters to indicate their realizations.

Lemma In each market, A_t , B_t and P_t converge almost surely to the same random variable (i.e., the bid-ask spread converges almost surely to 0).

Proof of Lemma (based on Glosten and Milgrom, 1985)

First, note that

$$\begin{aligned} \text{Var}[P_t] &= \text{Var} \left[\sum_{k=2}^t (P_k - P_{k-1}) \right] = \\ &= E \left[\left(\sum_{k=2}^t (P_k - P_{k-1}) \right)^2 \right] = E \left[\sum_{k=2}^t (P_k - P_{k-1})^2 \right], \end{aligned}$$

where we use these two facts: i) $P_1 = E[V]$ and, therefore, $\text{Var}[P_1] = 0$; ii) the sequence of prices $\{P_t, t = 1, 2, \dots\}$ is a martingale with respect to the history of trades and prices, and the increments of a martingale have zero expected value and are uncorrelated. Let us define the expected squared increment of the price at time k conditional on information available at time k as $\Delta_k := E[(P_{k+1} - P_k)^2 | H_k, A_k, B_k]$. Then, by the tower property of conditional expectation, $\text{Var}[P_t] = E \left[\sum_{k=2}^t (P_k - P_{k-1})^2 \right] = E \left[\sum_{k=1}^{t-1} \Delta_k \right]$.

Furthermore, given that P_t is the expected value of V conditional on H_t , $\text{Var}[V] \geq \text{Var}[P_t]$, which implies that $\text{Var}[V] \geq E \left[\sum_{k=1}^{t-1} \Delta_k \right]$.

Now, let us define the random variable

$$\varphi_k := \Pr[X_k = 1 | H_k, A_k, B_k] \Pr[X_k = -1 | H_k, A_k, B_k].$$

Note that in the event of a buy at time k the price at time $k+1$ will be equal to the ask at time k ; in the event of a sell, it will be equal to the bid at time k ; and in the event of a no trade it will be in between these bid and ask prices. Using this fact, by simple algebra, we can show that

$$\Delta_k \geq \varphi_k [(A_k - P_k)^2 + (B_k - P_k)^2].$$

By assumption, noise traders buy and sell with strictly positive probabilities. Therefore, φ_k is bounded away from zero by a positive number φ . Hence,

$$\begin{aligned} \Delta_k &\geq \varphi [(A_k - P_k)^2 + (B_k - P_k)^2] \\ \text{and } \sum_{k=1}^{t-1} \Delta_k &\geq \varphi \sum_{k=1}^{t-1} [(A_k - P_k)^2 + (B_k - P_k)^2]. \end{aligned}$$

Finally,

$$\text{Var}[V] \geq E \left[\sum_{k=1}^{t-1} \Delta_k \right] \geq \varphi E \left[\sum_{k=1}^{t-1} [(A_k - P_k)^2 + (B_k - P_k)^2] \right].$$

Given that the variance of the fundamental is bounded and that φ is strictly positive, then $\lim_{t \rightarrow \infty} E \left[\sum_{k=1}^{t-1} (A_k - P_k)^2 \right] < \infty$

and $\lim_{t \rightarrow \infty} E \left[\sum_{k=1}^{t-1} (B_k - P_k)^2 \right] < \infty$, which implies, as is easy to show,³¹ that A_t and P_t converge almost surely to the same random variable and, similarly, B_t and P_t converge almost surely to the same random variable.³²

Now, we prove Proposition 2 in two steps.

Step 1

In this step, we show that if at time t , for all $s \in \{0, 100\}$,

$$E[V|h_t] - c < E[V|h_t, S_t = s] \quad (\text{A1})$$

and

$$E[V|h_t, S_t = s] - c < E[V|h_t], \quad (\text{A2})$$

then, in equilibrium, $\Pr[X_t = x|h_t, S_t = s] =$

$\Pr[X_t = x|h_t]$, for all $s \in \{0, 100\}$ and for all $x \in \mathcal{A}$, i.e., an informational cascade occurs.

Proof of step 1

Suppose that at time t (A1) and (A2) hold and the market maker sets $a_t = b_t = E[V|h_t]$. Then, the conditional probability of a buy order given the trader's information is

$$\Pr[X_t = 1|h_t, S_t = s] = (1 - \mu)\varepsilon_B,$$

where ε_B denotes the probability of a buy order by a noise trader. On the other hand, the conditional probability of a buy order conditioned on the

³¹Using Markov's inequality, one can show that, if $\lim_{t \rightarrow \infty} E \left[\sum_{k=1}^{t-1} (A_k - P_k)^2 \right] < \infty$, then, for any $\varepsilon > 0$, $\lim_{t \rightarrow \infty} \sum_{k=1}^{t-1} (\Pr[|A_k - P_k| \geq \varepsilon]) < \infty$, which implies that A_t converges almost surely to P_t .

³²Note that, if private information were completely uninformative (i.e., the signals were independent of the realization of the asset value), the bid-ask spread would equal 0 starting from time 1, since there would be no asymmetric information between the traders and the market maker.

history of trades and prices only is

$$\Pr[X_t = 1|h_t] = (1 - \mu)\varepsilon_B.$$

Finally, note that, since $\Pr[X_t = 1|h_t, S_t = s] = \Pr[X_t = 1|h_t]$, the expected value of the asset conditional on $X_t = 1$ and the public information is indeed $E[V|h_t]$. Therefore, $a_t = E[V|h_t]$ is the equilibrium ask price.

An identical analysis shows that, when (A1) and (A2) hold,

$$\Pr[X_t = -1|h_t, S_t = s] = \Pr[X_t = -1|H_t]$$

and $b_t = E[V|h_t]$ is the equilibrium bid price.

$$\text{Finally, } \Pr[X_t = 0|h_t, S_t = s] = \Pr[X_t = 0|h_t] = (1 - \mu)(1 - \varepsilon) + \mu.$$

Step 2: Now we prove that there exists a time T such that for any $t > T$, (A1) and (A2) will be satisfied almost surely.

Proof of step 2

Let us define the random variable U_t which takes value 1 if at time t an uninformed trader arrives in the market and value 0 if at time t an informed trader arrives in the market. We can now prove step 2 by contradiction.

Assume that there does not exist a T such that (A1) and (A2) hold almost surely for all $t > T$. This means that, for any t , $E[V|h_t] - c \geq E[V|h_t, S_t = 0]$ or $E[V|h_t, S_t = 100] - c \geq E[V|h_t]$ or both.

Recall that, in equilibrium,

$$\Pr[S_t = 0, U_t = 0|h_t, X_t = 1] = 0$$

and

$$\Pr[S_t = 100, U_t = 0|h_t, X_t = -1] = 0.$$

Therefore, the conditions for equilibrium ask and bid prices at time t are:

$$\begin{aligned} (a_t - E[V|h_t]) \Pr[U_t = 1|h_t, X_t = 1] + \\ (a_t - E[V|h_t, S_t = 100]) \Pr[S_t = 100, U_t = 0|h_t, X_t = 1] = 0. \end{aligned}$$

$$\begin{aligned} (E[V|h_t] - b_t) \Pr[U_t = 1|h_t, X_t = -1] + \\ (E[V|h_t, S_t = 0] - b_t) \\ \Pr[S_t = 0, U_t = 0|H_t, X_t = -1, A_t, B_t] = 0. \end{aligned}$$

As we have proved in the Lemma, B_t , $P_t = E[V|H_t]$ and A_t converge almost surely to the same random variable. For almost every history and for any $\tilde{\varepsilon} > 0$, there exists a time \tilde{T} such that for any $t > \tilde{T}$, $(A_t - E[V|H_t]) < \tilde{\varepsilon}$ and $(E[V|H_t] - B_t) < \tilde{\varepsilon}$. Hence, for almost every history and for every $\hat{\varepsilon} > 0$, there exists a time \hat{T} such that for any $t > \hat{T}$,

$$-(A_t - E[V|H_t, S_t = 100]) \Pr[S_t = 100, U_t = 0 | H_t, X_t = 1] < \hat{\varepsilon},$$

and

$$-(E[V|H_t, S_t = 0]_\omega - B_t) \Pr[S_t = 0, U_t = 0 | H_t, X_t = -1] < \hat{\varepsilon}.$$

Moreover, since the ask and the bid converge almost surely to the price, for almost every history and for every $\varepsilon > 0$, there exists a time $T \geq \hat{T}$ such that, for any $t > T$,

$$-(P_t - E[V|H_t, S_t = 100]) \Pr[S_t = 100, U_t = 0 | H_t, X_t = 1] < \varepsilon$$

and

$$-(E[V|H_t, S_t = 0] - P_t) \Pr[S_t = 0, U_t = 0 | H_t, X_t = -1] < \varepsilon.$$

Summing up both sides of the two inequalities, we have:

$$\begin{aligned} & (E[V|H_t, S_t = 100] - P_t) \Pr[S_t = 100, U_t = 0 | H_t, X_t = 1] \\ & + (P_t - E[V|H_t, S_t = 0]) \Pr[S_t = 0, U_t = 0 | H_t, X_t = -1] < 2\varepsilon. \end{aligned}$$

Now, suppose (A2) is violated. Then, $\Pr[S_t = 100, U_t = 0 | H_t, X_t = 1] =: \tilde{\eta} > 0$. From the above inequality, $E[V|H_t, S_t = 100] - P_t < \frac{2\varepsilon}{\tilde{\eta}}$. By choosing $\varepsilon \leq \frac{1}{2}\tilde{\eta}c$, this implies that (A2) holds, a contradiction. A similar contradiction arises assuming that (A1) is violated.

Figure 1: Informational cascades arise in equilibrium

























Figure 2: The price path
















Figure 3: Distance between the theoretical and the actual last prices in CG-FP




Figure 4: Distance between the theoretical and the actual last price




Table 4: No Trade Cascades

Session I												
IV	S	B	B	B	S	S	B	B	S	B	S	S
V	NT	S	S	NT	NT	B	S	B	S	B	B	B
VI	B	B	NT	S	S	NT	B	B	NT	NT	B	NT
VII	B	B	B	NT	NT	NT	NT	NT	NT	NT	NT	NT
VIII	B	NT	B	B	B	S	S	NT	B	NT	NT	S
IX	NT	B	S	B	B	S	NT	S	B	B	NT	NT
X	NT	B	B	B	S	B	NT	B	NT	NT	NT	NT
Session II												
IV	B	NT	NT	B	S	NT	S	NT	S	S	NT	NT
V	B	B	B	NT	NT	NT	NT	NT	NT	NT	NT	NT
VI	B	B	B	NT	NT	NT	NT	NT	NT	NT	NT	NT
VII	B	S	NT	NT	S	S	S	NT	NT	NT	NT	NT
VIII	S	S	NT	S	NT	NT	NT	NT	NT	NT	NT	NT
IX	B	B	B	NT	NT	NT	NT	NT	NT	NT	NT	NT
X	B	B	B	NT	NT	NT	NT	NT	NT	NT	NT	NT
Session III												
IV	S	S	S	NT	NT	NT	NT	NT	NT	B	B	B
V	S	B	S	S	S	NT	B	S	NT	NT	NT	NT
VI	S	S	B	S	B	S	S	NT	NT	NT	NT	NT
VII	S	S	S	NT	NT	B	NT	S	NT	NT	NT	NT
VIII	B	B	B	NT	S	B	NT	NT	NT	NT	NT	NT
IX	S	S	S	B	S	NT	NT	NT	S	NT	NT	NT
X	S	S	NT	S	NT	NT	NT	NT	NT	B	NT	S
Session IV												
IV	NT	B	B	B	NT	S	B	S	NT	NT	NT	NT
V	S	S	S	S	B	B	NT	NT	NT	S	NT	NT
VI	NT	NT	NT	NT	S	S	B	S	S	NT	NT	NT
VII	S	B	NT	B	NT	S	S	S	NT	NT	B	NT
VIII	B	B	NT	NT	NT	NT	NT	NT	NT	B	NT	NT
IX	B	NT	S	NT	NT	B	NT	B	B	S	NT	B
X	NT	B	S	B	B	B	S	NT	NT	B	S	B

Addendum to “Transaction Costs and Informational Cascades in Financial Markets: Theory and Experimental Evidence”

Note: The following are the instructions for the experiment. The instructions were distributed to the subjects and read aloud. The experiment was run at New York University in 2001.

1 Instructions

1.1 Introduction

You are about to engage in an experiment in market decision making. Various research institutions have provided funds for this experiment and, if you make appropriate decisions, you may earn a good monetary payment.

1.2 The experiment

The experiment consists of a series of 10 rounds. In each round you will have the opportunity to make a trade. You will perform this experiment with 12 other people. When it is your turn to decide to trade you will be offered a price at which you can either buy one unit of a good, sell one unit, or decide not to trade. The good you are trading will ultimately be worth either 0 or 100 units of an experimental currency called “lire.” These lire will be converted into dollars at the end of the experiment at the rate of $\$1 = 65$ *lire*. Whether the good will be worth 0 or 100 lire will be determined at the beginning of each round by flipping a coin: if the coin lands head the value will be 0 lire and if the coin lands tail it will be worth 100 lire. Hence, there is a $\frac{1}{2}$ chance of the good having a value of 0 or 100. The outcome of the coin flip will not be revealed to you, so when it is your turn to trade you will not know the value of the good. You will, however, receive some information about which value is more likely to have been chosen.

Before starting the experiment we will choose one of you as the subject administrator. The subject administrator will help us to run the experiment: in particular, he will toss the coin that determines whether the good is worth 0 or 100 lire. The subject administrator will not trade during the experiment and will receive a payoff equal to the average payoff received by all of you.

Moreover, we will assign a number (from 1 to 12) to everyone, to help us to keep track of your trades during the experiment. *Write the number assigned to you on your worksheet, after your name.*

Procedures for each round.

First, the subject administrator will flip a coin. The outcome of this flip will determine the value of the good: if the coin lands head the value will be 0 lire and if the coin lands tail it will be 100 lire. Of course, you will not be told the outcome of the coin flip.

Second, we will start calling you to trade, using your identification number. The order in which you are called is randomly decided by the subject administrator by using a deck of cards.

Third, when you are called to trade, you will draw a chip from one of two boxes. While both boxes contain only two types of chips, white chips and blue chips, the proportion of the two types of chips in each box is different (it will be verified by the subject administrator of the experiment). One box contains 70 blue and 30 white chips and the other box 30 blue and 70 white chips. If the value of the good is 0 (as determined by the coin flip), we will ask you to draw the chip from the box that contains 70 blue and 30 white chips. If the value is 100, we will ask you to draw the chip from the box that contains 70 white and 30 blue chips. To recap:

- If the value is 100, then there are more WHITE chips in the box.
- If the value is 0, then there are more BLUE chips in the box.

Therefore, the chip that you choose will give you some information about the value of the good. After you draw the chip and look at its color you will put it back into the box, so that the number of chips in the box never changes. *The color of this chip is your private information – do not share it with any other subject.*

Forth, immediately after you draw the chip from the box, we will ask you if you want to buy, to sell or not to trade at a given price. Everyone can observe the choices that have been made by the previous players and the price at which they traded (we will write them on the blackboard). Therefore, when you make your decision (sell, buy or no trade), you know what the people who traded before you decided to do. *You will record the price at which you trade and your decision (check one among Buy, NT=No trade, and Sell) in columns 2 and 3 of your worksheet in the row referring to Round*

1. Before you make your decision we will inform you of the price at which the buy or sell transaction will take place. Remember that at that price you can buy, sell or decide not to trade. However, if you decide to buy or sell, you have to pay a cost of 9 lire; you do not have to pay the 9 lire if you decide not to trade. The price offered to you will be computed as if it were the price set in a perfectly competitive market given the buy and sell decisions made in the past. The price offered to the first trader will be 50 lire. This price is set at 50 because it is half way between 0 and 100 and each of these two values has a $\frac{1}{2}$ chance of being correct. After the first player makes his/her decisions, we will update the price. If this person buys, we will increase the price and the person who comes second will face a higher price. If the first trader sells, we will revise the price downward and the second person will face a price lower than 50 lire. The same will happen for all of the following trades: when a person buys we will increase the price and when he sells we will lower it. The reason for this price setting is simple. At the beginning 50 lire is the best guess of the value of the good, given that the values 0 and 100 are equally likely. When a person buys, we calculate that the value of the good is higher than the posted price and therefore we increase the price. When there is a sell, we calculate that the value of the good is lower than the posted price and therefore we decrease the price. The exact amount by which we change the price depends on the working of competitive markets: the price at each turn is the best guess of the good value given the decisions taken by the subjects in the previous turns.

At the end of the round, after all subjects have made their buy, sell, or no-trade decisions, you will be informed of the value of the good in that round and of your payoff. *You will write this value and your payoff in columns 4 and 5 of the row referring to Round 1.*

Your per-round payoff

Your payoff in any round will be determined as follows.

You will first be given 100 lire just for trading in the round. Therefore,

- *If you decide to buy at the price p you will get*

$$100 + Value - p - 9.$$

- *If you decide to sell at the price p you will get*

$$100 + p - Value - 9.$$

- *If you decide not to trade* you will earn 100 lire.

Why these payoffs? If you buy, you get the value of the good and you have to pay its price. If you sell, you have to borrow the good in order to sell it. You get the price at which you sell and have to repay the value of the good. In both cases of buy and sell you have to pay a trade cost of 9 lire.

Therefore, when you buy, you will gain lire if the value of the good is higher than the price plus 9 lire and lose them if it is lower. When you sell, you will gain lire if the price is higher than the value plus 9 lire and lose them if it is lower. Always remember that the value of the good can be either 0 or 100. Notice also that at most you can lose 9 lire in a round, given that we give you 100 lire for trading.

Examples of the per-round payoff.

Suppose when it is your turn to trade the price is 70. There are two possibilities: either the good is worth 0 or it is worth 100. If you decide to buy and the good is worth 0, then you get $100 + 0 - 70 - 9 = 21$. In this case you lose your money, because you bought for a price of 70 something that is worth 0. If you buy and the good is worth 100, then you get $100 + 100 - 70 - 9 = 121$. In this case, you earn more money because you bought for only 70 lire a good that is worth 100 lire. If you decide to sell and the good is worth 0, then you get $100 + 70 - 0 - 9 = 161$. In this case you earn more money, because you sold for a price of 70 something that is worth 0. If you sell and the good is worth 100, then you get $100 + 70 - 100 - 9 = 61$. In this case, you lose your money because you sold for 70 lire a good that is worth 100 lire. Finally, if you decide not to trade, you will simply keep your initial 100 lire.

After the twelfth subject has made his/her decision, the first round is over. We will then proceed to the second and third rounds where the same procedures of the first round will be repeated. At the beginning of the new round, the value of the good will be determined by a new coin flip. When your turn comes, you will draw another chip and trade the good, according to the same rules described for the first round. You will also be given 100 new lire. You will proceed to the Round 2 and 3 rows of your worksheet and record there the color of your chip and your trade decision.

How your final payment is determined. For the simple fact that you show in time for the experiment you earn \$7. The rest of the payment depends on how you perform. First, we will sum up your payoffs in lire for

all the 10 rounds. We will then convert these lire into dollars at the rate of $\# \text{ dollars} = \frac{\# \text{ lire}}{65}$. For instance, if after 10 rounds you made 1000 lire you will receive $\frac{1000}{65}$ dollars, that is, \$15.40. Your final payment will be equal to this amount plus the \$7.