

Herding and Price Convergence in a Laboratory Financial Market

Marco Cipriani and Antonio Guarino*

This version: December 15, 2002

Abstract

We study whether herding can arise in a laboratory financial market in which agents trade sequentially. Agents trade an asset whose value is unknown and whose price is efficiently set by a market maker. We show that the presence of a price mechanism destroys the possibility of herding. Most agents follow their private information and prices converge to the fundamental value. This result contrasts with the case of a fixed price, where herding and cascades arise. When the price moves, however, agents may behave as contrarian, i.e., they may trade against the market, something not accounted for by the theory. Finally, we study whether informational cascades arise when trade is costly (e.g. because of a Tobin tax). With trade costs, most subjects rationally decided not to trade and the price was unable to aggregate private information efficiently. (JEL C92, D8, G14)

*Cipriani: Department of Economics, George Washington University; E-mail: mciprian@gwu.edu. Guarino: Department of Economics and ELSE, University College London; E-mail: aguarino@ucl.ac.uk. We are indebted to Andrew Schotter for his guidance. We thank Franklin Allen, Alberto Bisin, Andrew Caplin, Douglas Gale, Steffen Huck and Marco Ottaviani for useful comments. We also thank seminar participants at CREED, New York University, University College London and ESA meeting in Boston. 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. All errors are ours.

1 Introduction

In recent years there has been an increasing interest in herd behavior in financial markets. Especially after the financial crises of the 1990s, many scholars have suggested that herd behavior might be a reason for excess price volatility and financial systems fragility.

The theoretical research on herd behavior starts with the seminal papers by Banerjee (1992), Bickchandani et al. (1992), and Welch (1992).¹ These papers do not discuss herd behavior in financial markets, but in an abstract environment, in which agents with private information make their decisions in sequence. They show that, after a finite number of agents have chosen their actions, all following agents will disregard their own private information and herd. For instance, if some people decide to adopt a technological innovation, then everyone will do the same, even if they have negative private information on the actual value of that innovation. Intuitively, agents' private information is revealed through their actions and becomes public information. Eventually, public information overwhelms private information, and, therefore, agents act independently of their signal. This is an important result, because it gives a rationale for the imitating behavior that we observe in consumers' and investors' decisions. In these first models, however, the cost of taking an action (e.g., the cost of adopting a new technology) is held constant. In other words, these models do not analyze situations in which, when agents make their decisions to buy or sell a good, the price of that good changes. Therefore, they are unsuitable to discuss herd behavior in financial markets, where prices are certainly flexible and react to the order flow.

More recently, Avery and Zemski (1998) have studied herd behavior in a financial market where the price is set by a market maker according to the order flow. They show that the presence of a price mechanism makes an informational cascade (i.e., a situation in which an agent does not use his own information) impossible.² Agents will always find it convenient to trade

¹In this paper we study only informational herding. Therefore, we do not discuss herd behavior due to reputational concerns, as in Sharfstein and Stein (1990), or to payoff externalities. For a survey of herding in financial markets, see Brunnermeier (2001).

²In the literature it has been pointed out (see Smith and Sørensen, 2000) that herd behavior (a situation where all agents do the same thing) and informational cascades (a situation in which agents neglect their private information) are not identical concepts. For an experimental analysis of the difference between herd behavior and informational cascades, see Çelen and Kariv (2002). The distinction between herds and cascades, however, is relevant only when beliefs are unbounded. Given that in our set up beliefs will be bounded,

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). Therefore, it will never be the case that agents neglect their information and imitate previous traders' decisions. For this reason, the price aggregates the information contained in the history of past trades correctly.

It is difficult to test these theoretical models of herding empirically. The existing literature (see, e.g., Lakonishok et al., 1992, and the other papers cited in the survey of Bickchandani and Sharma, 2000) does not test these models directly, but only analyzes the presence of herding in financial markets through statistical measures of clustering. This literature finds that in some cases fund managers tend to cluster their investment decisions. This clustering, however, may or may not be due to herding: for instance, it may be the result of the common reaction to public announcements. The problem for the empiricist is that there are no data on the private information available to the traders and, therefore, it is difficult to understand whether traders decide to disregard their own information and imitate.

This problem can be overcome in an experimental study. Indeed, the advantage of an experimental study is that we can observe variables not available in real market data, in particular the private information that agents have when they make their decisions. In our laboratory market, agents 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, buy or not to trade one unit of the asset. By observing the way in which traders use their private information and react to the decisions of the previous traders we can directly estimate the presence of herding in the market. By testing directly the prediction of the theoretical work, we create a bridge between the existing empirical and theoretical literatures.

The results of our study are quite in line with the predictions of the theoretical models. We compare two cases, one in which the price is fixed and one in which it is efficiently set by a market maker according to the order flow. We find that, when the price is fixed, agents tend to neglect their own information and follow the decisions made by the previous agents. In contrast, when the price is flexible, agents tend to use more their private information and prices converge to the true asset value many times.

By comparing the results of the two treatments, we can understand the role of the price in aggregating private information and preventing herd be-

we use the expressions “herd behavior” and “informational cascade” interchangeably.

havior. The upward movement of the price after a series of buys prevents the next agent to buy if he has negative information. The downward movement of the price after some sells prevents the next agent from selling if he has positive information. The price movement makes the history of trades irrelevant, as the price factors in the decisions of the previous traders.

The fact that agents trade according to their own information and are not affected by the history of trades is confirmed by our third experimental design. In this treatment, the price is flexible but agents cannot observe the decisions made by previous traders. They only know the price at which they can trade. Their behavior does not differ significantly from the behavior that arises in the baseline experiment with flexible prices and observable history. In both treatments, most agents buy when they have a positive signal and sell when they have a negative one.

Avery and Zemski's result, i.e., the fact that prices destroy rational herds, hinges on the assumption that traders trade only for speculative reasons. In particular, they assume that there are no gains from trade and no trade costs in the economy. When these conditions are not satisfied, agents may disregard their own private information and herd (Lee, 1998, Cipriani and Guarino, 2003). When trade is costly, for instance, informational cascades should arise. After some trades, agents should refrain from trading as the trading activity is less beneficial than the trade cost. Because agents refrain from trading, private information is not aggregated. To study the effect of trade costs, we run an experiment in which traders were asked to pay a cost in order to buy or sell the asset. In this treatment, after some trades, agents decide correctly not to follow their signal and not to trade. Given that agents do not trade, the price in many cases fails to aggregate private information.

While theory predicts the effect of price on herding and cascades correctly, it is unable to account for a different phenomenon that we observe in the laboratory. When the price is flexible, agents in some cases neglect their private information in order to trade against the market. They sell, despite a positive signal, when the market price is high and buy, despite a negative signal, when the price is low. This behavior, which we call "contrarian behavior," reduces the ability of the price to aggregate private information efficiently. In our econometric analysis of data we try to give an explanation of this and of other behaviors that we observed in the experiment.

The structure of the paper is as follows. Section 2 describes the theoretical model. Section 3 presents the experimental design. Section 4 illustrates the results of the first three treatments. Section 5 discusses the treatment with

trade costs. Section 6 presents the results of an econometric analysis of the data. Section 7 concludes.

2 The Theoretical Model

2.1 The model structure

Our analysis is based on a simplified version of the Glosten and Milgrom (1985) model. In our economy there is one asset with true value V distributed on $\{0, 100\}$. The asset takes value 100 with probability p and 0 with probability $1 - p$. The asset is traded sequentially. The price of the asset is set by a market maker who interacts with a sequence of traders. At any time t , a trader is randomly chosen to act in the market and can buy, sell or decide not to trade the asset. Each trade consists of the exchange of one unit of the asset. Any trader trades at most once. At any time t , the market maker and the traders know the history of trades, H_t , until time $t - 1$.

The true realization of the asset value V is unknown. All traders, however, receive a private signal on the value of the asset and maximize their expected profit based on that signal. The signal X is drawn from the following conditional probability function: $f(X = x|V = x) = q$ and $f(X = x|V \neq x) = 1 - q$, with $1 > q > \frac{1}{2}$, $x = 0, 100$. After receiving the private information, the traders make their decisions in sequence. Traders are risk neutral, i.e., they maximize their expected profit.³ The expected value of the asset for an informed trader at time t is $E(V|H_t, x)$. On the other hand, the expected value for the market maker before observing the new trader's decision will be conditioned only on the public information available at time t , i.e., it will be $E(V|H_t)$. We consider two different setups: one in which the market maker keeps the price fixed at the unconditional expected value of the asset; and another in which he updates the price, setting it equal to the expected value of the asset conditional on the past history of trades.⁴

³In the original Glosten and Milgrom (1985) model, a proportion of traders are uninformed and trade for exogenous reasons, without respect for their profit. The presence of these noise traders is necessary for the market not to breakdown. Indeed, a market without gains from trade and where agents are risk neutral would collapse in the absence of noise traders, as proven by the no-trade theorem (Milgrom and Stokey, 1982). In our set up, for simplicity all traders are informed and we assume that the market maker is willing to trade even if he make negative profits in expected value.

⁴In the original Glosten and Milgrom model the market maker posts a bid price and

2.2 Predictions when the price is fixed

In our first setup, as in the seminal papers on herding and cascades (e.g., Banerjee, 1992, and Bickchandani et al., 1992) there is no price mechanism. We assume that agents can buy or sell the asset at a given price of 50, the unconditional expected value of the asset. In other words, we do not allow the price to change after agents make their trades. How should agents behave in this case? Let us consider a simple example. Suppose that the first two agents decide to buy the asset. What should the third trader do? If he receives a positive signal, clearly he wants to buy too. But suppose he receives private information that the asset value is 0. From the first two buy orders, he can infer that the first two agents received information that the asset is worth 100. Therefore, he knows that there have been two good signals and one bad signal. Given that all signals have the same precision, he will decide to buy, despite his negative private information. Whenever there is an imbalance of at least two orders in the order flow (i.e., two more buys than sells or vice versa), agents rationally decide to herd and an informational cascade arises.

2.3 Predictions when the price is flexible

Let us discuss now the case in which the price is flexible. The market maker updates the price on the base of the order flow by setting it equal to $E(V|H_t)$. Avery and Zemski (1998) have shown that, in this case, informational cascades cannot arise. The informed trader has an informative advantage on the market maker, as he receives a private signal in addition to knowing the history of trades. He will always find it optimal to use this extra piece of information and buy the asset if he receives a positive signal and sell it if he receives a negative one. Therefore, his action will always reveal his signal, information will always arrive in the market and no one will find it convenient to imitate the previous players.

To illustrate more in detail this point, let us repeat the same example of the previous section and assume, however, that the price is not fixed at 50,

an ask price. We avoid the presence of two prices (the bid and the ask) and assume that the market maker sets the price equal to the expected value of the asset (disregarding that, when a trader comes to buy or sell, this reveals the information he has on the asset value). This is necessary, as in our model there are no noise traders and, therefore, if the market maker set a bid and an ask spread, the market would collapse. The presence of only one price makes our experiment also easier to run. The bid-ask spread has already been studied in the experimental literature, see, e.g., Bloomfield, 1996.

but is set equal to the expected value of the asset conditional on the past history of trades. Let us also assume that the probability that the signal is correct is 70%. After the first two traders buy, the market maker will increase the price from 50 to 84⁵ to take into account that the first two buys came from agents with a high signal. The third agent infers that the previous two traders had a high signal. If his signal is low, his expected value of the asset (given two buy orders and a negative signal) will be 70. Given that he faces a price of 84 he will sell, i.e., he will follow his own private information. By the same argument all traders will always follow their private information. Since the signal that they receive is correct 70% of the time, over time the price will converge to the fundamental value of the asset, thus aggregating the private information dispersed among traders. Therefore, when prices are set efficiently, agents will follow their own private information and the price will aggregate the information spread among traders. Consequently, we should not observe misalignments of the price with respect to the fundamentals.

It is worth noting that the past history of trades is completely irrelevant for the trader's decision. This happens because the price set by the market maker conveys the same information on the expected value of the asset as the past history of trades. A rational subject should act according to his private signal, irrespective of whether he is able to observe the whole history of trades or not. For this reason, we ran a control treatment where we did not allow traders to observe the past history of trades.

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. For each session of the experiment we recruited 13 people, one to act as a subject administrator and the others to act as traders. In total, we recruited 208 students to run 16 sessions (four sessions for each treatment). The experiment was run in the following way:

1. At the beginning of the sessions, we gave written instructions (reported

⁵The price is calculated using Bayes's rule.

in the Appendix) to all subjects. Then we read the instructions aloud and asked clarifying questions. Since the instructions were read aloud and they indicated how the asset value was chosen and how the signal was extracted, the structure of the game was common knowledge to all traders.

2. In each session we performed ten rounds. In each round we asked all students to trade one after the other.
3. The sequence of traders for each round was chosen randomly. At the beginning of the session each student picked a card from a deck of 13 numbered cards. The number that a student picked was assigned to him for the entire session. The card numbered 0 indicated the subject administrator. In each round the subject administrator called the agents in sequence by randomly drawing cards (without replacement) from the same deck.
4. Before each round, the subject administrator went outside the room and tossed a coin: if the coin landed tail, the value of the asset for that round was 100, otherwise it was 0. Of course, no trader was told the outcome of the coin flip.
5. One of us acted in the experiment as the market maker, setting the price at which people could trade. The other 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 looked perfectly identical. Each agent, before trading, had to go outside the room and draw a chip from one bag. If the coin landed tail we used the first bag, otherwise we used the second. Therefore, the chip color was a signal for the value of the asset.⁶ Note that the trader could not share the chip color with anyone. Therefore, neither the market maker nor the other traders knew the realizations of the signal.
6. After seeing the chip color, the agent came inside the room and declared aloud whether he wanted to buy or sell one unit of the asset or decide not to trade. The subject administrator recorded all agents' decisions on the blackboard. On the blackboard he also recorded the history of prices and the new price at which the agent could trade the asset.

⁶Using the notation of the previous section, we chose $q = 0.7$.

Hence, each agent knew not only his own signal but also the history of trades and prices.

7. At the end of each round, i.e., after all 12 students had traded once, the subject administrator revealed the realization of the asset value and we asked students to compute their payoffs. All values were in a fictitious currency called lira. Their payoffs were computed as follows. In the event of a buy, the student obtained $100 + Value - Price$ lire; in the event of a sell, he obtained $100 + Price - Value$ lire; finally, if he decided not to trade he earned 100 lire. This is like to say that for each round we gave 100 lire to agents and they had the option of using them to trade. Given that the price was always between 0 and 100 lire, and that they were given 100 lire at the beginning of each round, students could never lose money.
8. After the twelfth 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 students just for participating in the experiment. On average, students earned \$23 for a one hour and a half experiment.

3.2 The Experimental Design

We ran four different treatments. Here we describe the first three of them, and postpone the discussion of the fourth to Section 5. In a first treatment (“fixed price” treatment), the price was not updated on the basis of the order flow and was fixed at 50. As explained in the previous section, after an imbalance of two decisions in the order flow, an informational cascade should arise, i.e., agents should not use their private information.

In a second treatment (“flexible price”), the price was updated after each trade decision in a Bayesian fashion. Bayesian updating assures that prices were the best guess of the asset value given the history of trades.

The third was a control treatment: in order to understand the effect of history on the behavior of agents we ran an experiment where agents could not observe the decisions of those who traded before them. Indeed, the only difference of this “no history” design with respect to the flexible price treatment was that each agent could not observe what previous traders had done. We asked students not to declare their trade decisions publicly. Each agent, when called, went to the experimenter’s table and wrote his decision on

a piece of paper. Given that decisions were not made public, when a student had to make a decision, he could only read the price at which he could trade on a piece of paper, but did not have the history of trades and of prices written on the blackboard. The only information that he had was his own signal and the price at which he could buy or sell. Although we did not want students to know the past prices, we wanted them to know the mechanism of price formation. In order to make sure that students understood this mechanism, not only we described it in the instructions, but in the first three rounds, we also ran the experiment as in the flexible price treatment. In this way, everyone could observe how the market maker updated the price in reaction to the traders' decisions. Only from the 4th round on, did we not allow agents to see the past history.

In the next section we describe the results of these three treatments. The results refer to the last seven rounds of each session only. We do not take into account the first three rounds for two reasons. First, although the experiment was very easy and agents did not have problems in understanding the instructions, we believe that some rounds were important to make agents acquainted with the procedures. We prefer to distinguish the decisions that agents made in the learning stage from the decisions taken afterward. Second, as explained above, in the “no history” design, we started the experiment with unobservable history only from the 4th round. Therefore, considering the last seven rounds only is also convenient in order to compare the results of the different treatments.

4 Results

4.1 Herding and contrarian behavior

We start the presentation of our results by discussing herd behavior. Let us consider first the fixed price treatment. In this case, theory predicts that all rational agents should neglect their private information and herd whenever a trade imbalance of at least two arises. In our experiment, there were 58 periods of potential herding, i.e., periods when the trade imbalance of at least two had occurred and, moreover, the trader had received a signal opposite to the trade imbalance. In these periods of potential herding, agents did herd in 52% of cases, decided not to trade in 26% of cases and decided to follow

their signal in 22% of cases.⁷

What happens when we allow the price to react to the order flow? Do people still herd? Table 1 shows the results of this treatment and contrasts them to those of the fixed price treatment. In the flexible price case there were 66 periods in which the trade imbalance was at least two and the trader received a signal against it. In these periods, agents decided to herd only in 12% of the cases. In 42% of cases they decided not to trade and in 46% of cases they followed their information even if it was at odds with the history of trades. These results show that agents rarely decided to follow what other agents had done. The price movement offset the incentive to imitate previous decisions: the propensity to herd is, indeed, considerably lower than in the case without a price mechanism. We ran a Mann-Whitney test for the hypothesis that the distribution of subjects' decisions was the same under the two treatments and the null is rejected at a 5% significance level (see Table 12).⁸

⁷An important aspect in the computation of the trade imbalance is how to handle the role of deviators. Consider the no price treatment and suppose that there have been four buy orders. Suppose that the next agent decides to sell. In this case, his actions is certainly irrational, since, no matter what is signal is, he should buy. Therefore, one can take into account this sell order in different ways. One could argue that the decision of this person should not be taken into account in the computation of the trade imbalance, as it is irrational and does not reveal anything about his signal. Therefore, after this sell, the trade imbalance should still be counted as four. On the other hand, one can argue (as Anderson and Holt, 1997) that, although irrational, this person is likely to have received a negative signal, otherwise he would not have any reason to sell. Therefore, his decision breaks the cascade. The cascade was created by the first two buy orders (the other two buys do not provide any information) and now is destroyed by the sell order. The trade imbalance would go from four to one. Consider now the flexible price treatment. In this treatment, informational cascades should never happen, therefore the trade imbalance should unarguably be computed by taking into account all past actions in the same way. Therefore, the trade imbalance in our example should go from four to three in this treatment. In this paper we want to understand the role of the price mechanism and compare the fixed price with the flexible price treatment. Therefore, we want to use the same definition of potential herding and the same measure of trade imbalance for the two treatments. We have decided to compute the trade imbalance in both treatments considering all past action, irrespective of whether they could have come from a rational agent or not. Our choice does not affect the results in a significant manner, as in the fixed price treatment we observed very few agents deviating from a cascade. If we compute the trade imbalance by assuming that a deviator breaks the cascade in the fixed price treatment, herd behavior arises 54% of cases, no trade 25%, and following the signal 21%.

⁸The test is carried out by taking the average proportion of trades due to herding in

It is important to note that, in the flexible price treatment, agents decided not to trade more often than in the fixed price treatment. For instance, the fact that after a series of buys the price went up did prevent more agents (with negative information) from buying the asset. In some cases, however, instead of following their signal, agents preferred not to trade. We will discuss more in detail this aspect in Section 6, when we move to an econometric analysis of data.

To understand better the effect of past trades on agents' decisions we ran a control treatment in which agents could not observe the past history of trades ("no history" treatment). If agents were significantly affected by previous agents' decisions, the results of the flexible price treatment and of the no history treatment should be very different. In particular, one could expect the percentage of herd behavior be much higher when agents do observe the previous decisions. This is not what happened (see Table 1). In the no history treatment, over the 70 periods of potential herding, agents herded in 24% of cases, versus the 12% of the flexible price treatment. As in the previous treatment, some decisions (33% of cases) taken in situations of potential herding were no trades. Running a Mann-Whitney test, we cannot reject the hypothesis that the two distributions of trading decisions are identical (see Table 12).

Our results of herd behavior in the no price treatment can be compared to what reported by Anderson and Holt (1997), who replicated in the laboratory the simple model of informational cascades of Bickchandani et al. (1992). The main difference between their setup and ours is that, while our agents had three options (buy, sell, and no trade), theirs had to make a binary decision.⁹ They found that agents decided to neglect their information and follow the previous agents' decisions in 73% cases of potential herding.¹⁰ A possible way of comparing our results with theirs is not to take into account the no trades, so that also in our case there are only two decisions, either follow the signal or not. If we compute the percentage of herds in this way, agents herded in 70% of cases, a result that is almost identical to their finding.

So far, we have illustrated the decisions taken in the experiment only in

each of the four sessions. Using the Mann-Whitney procedure, we tested whether the distribution of these four averages was significantly different between the two treatments.

⁹Moreover, Anderson and Holt (1997) computed the trade imbalance as indicated in the previous footnote.

¹⁰For a critique of Anderson and Holt (1997), see Huck and Oechssler (2000) and Kübler and Weizsäcker (2003).

situations where herding could arise. It is also interesting to see how agents acted on average during the experiment. Table 2 reports the proportion of agents who acted rationally, i.e., as the theory predict. For the no price treatment, this means that actions are classified as rational if agents herded in situations of potential herding and followed the signal otherwise. For the flexible price treatment, this means that decisions are classified rational if they were always based on the signal. In both treatments the proportion of irrational actions is not very high, 13% in the price treatment and 6% in the fixed price treatment. Interestingly, however, the presence of the price seems to negatively affect the behavior of agents in two ways: not only is the proportion of irrational decisions greater, but there is also more abstention from trading (22% versus 14%).¹¹

In order to understand better this result, we considered the possibility that agents could have incentives to trade against their signal (or not to trade) in circumstances other than potential herding. In particular, another possibility for irrational behavior is that agents may have acted as contrarians, i.e., that they may have neglected their signal to buy at a low price or sell at a high price. More specifically, we say that an agent is a contrarian when he buys, despite a negative signal, at a low price, i.e., at a price lower than 50, and sells, despite a positive signal, at a high price, i.e., a price higher than 50. Equivalently, we can say that an agent acts as a contrarian when he buys with a negative signal and there is prevalence of sell orders in the history of trades or sells with a positive signal and there is a prevalence of buys. The definition of a contrarian behavior is meant to capture the behavior of people who use the strategy of “going against the market.” These agents disregard their positive signal to take advantage of the high price in the market, or, vice versa, disregard their negative signal to buy the asset at a low price.

In the flexible price treatment, there were 165 periods in which agents could have potentially acted as a contrarian. That is, they had a high (low) signal when the price was high (low). Out of these 165 times, agents behaved

¹¹One may wonder whether these results depend on the behavior of some particular subjects or, on the contrary, reflect an average behavior. For instance, whether the level of irrationality is due to some people who behaved irrationally most of the times or whether it is the outcome of a common behavior. We computed for each subject the number of times in which he acted rationally, decided not to trade or chose the action opposite to what theory predicts. As Tables 5, 6 and 7 show that the behavior was quite homogeneous across subjects.

as contrarians in 16% of cases, whereas in 17% they decided not to trade and in 67% they followed their signal. A similar behavior arose also when the history was not observable (see Table 3). Therefore, when the price was flexible (with history observable or not) about 33% of times agents decided not to buy in the presence of a high price although they had a high signal, or decided not to sell in the presence of a low price although they had a negative signal. In all these cases the market was unable to aggregate private information correctly.¹²

Table 3 also reports the percentage of agents who acted as “contrarian” in the fixed price treatment.¹³ With a fixed price, agents adopted a contrarian behavior only in 1% of cases and decided not to trade only in 2% of cases. This suggests that, while the price mechanism significantly reduces the incentive of agents to herd, and in this way allows the market to aggregate private information, it also has a perverse effect, as it creates an incentive for agents not to reveal their information when it is consistent with the previous history of trades. This informational inefficiency is not present when the price is not allowed to respond to the history of trades.

Note that, when agents herd, the market does not learn private signals that are against the previous history of trades. When this information is trapped, we can have a low price in the market even if the fundamental value of the asset is high or, vice versa, a high price when the asset is worth zero. In contrast, when agents behave as contrarians, the private signals that are not revealed to the market are of the same type already revealed by the history of trades. This implies that the price will converge to the correct value, but more slowly than the theory predicts. In the next Section we discuss price convergence more in detail.

¹²We have computed the propensity to act as a contrarian also for different levels of the price (trade imbalance). The results are very similar to those presented above. For example, when we consider a trade imbalance of at least 2, we find that 19% of the decisions are contrarian, 18% are no trade and 63% are taken according to the signal. When we consider a trade imbalance of 3, these figures change only slightly to 21%, 19% and 60%.

¹³In this treatment, given that the price is not updated, the behavior of an agent is defined as contrarian if he sold after there were more buy orders than sells and he had a positive signal or if he bought after there were more sell orders than buys and he had a negative signal.

4.2 Price convergence

According to theory, in the flexible price treatment the price should always converge to the true value. In each period, by choosing to buy or to sell, agents reveal their private information. Over time, all private information is revealed to the market maker and the price reflects the fundamental value of the asset. Figure 1 shows the price path for a round in which the realized value of the asset was 0. The signals were all correct, i.e., all negative. Theoretically, the price should have decreased monotonically toward 0. In the experiment, the price did converge to the right value, but the speed of convergence was lowered by agents 2 and 5 who decided not to trade and agent 6 who behaved as a contrarian.

This price path was not an exception: indeed, we observed price convergence in many cases. Given the results reported in the previous section this is not surprising. We have already shown that in this treatment, even when the history of trades was contradicting the private information, only few decisions (12%) were against the signals. This implies that, at most, only in few cases the price might be misdirected, i.e., converge towards 0 when the value of the asset was 100 or vice versa. At the same time, however, the absence of trading in many rounds and the contrarian behavior should reduce the ability of the price to aggregate information efficiently and, therefore, reduce the speed of convergence.

We studied the price level after all 12 agents had traded and compared it to the levels that should be observed theoretically. Figures 2 and 3 show the histogram of the theoretical and actual last price when the true asset value happened to be 0 and when it happened to be 100 respectively. The theoretical price was computed assuming that all actions were based only on the signals. Remember that we considered the first three rounds of each session as a learning period and did not take them into account. Therefore, for each treatment we have observations for 28 rounds, 7 for each of the 4 sessions. In the flexible price treatment, the asset value happened to be 0 in 13 cases and 100 in 15 cases. When the value was 0 the price in the last period reached a level lower than 7 in 7 cases, in 3 cases it reached a level higher than 7 but lower than 30 and, finally, in 3 cases it was between 30 and 70. The price never converged to a value higher than 70. When the value was 100, on the other hand, the price in the last period was higher than 93 in 8 cases. In 4 cases it was between 70 and 93 and in 3 it was between 30 and 70. Therefore, in 54% of cases the price reached a level very close to

the fundamental one. In 25% of cases it moved in the right direction but was unable to become very close to the fundamental value and only in 21% of cases it failed to converge, remaining at a level close to the unconditional expectation of the asset value. On average, across the four sessions, the price in the last period reached a level of 14.9 when the fundamental value was 0 and a level of 93 when the value was 100. Figures 2 and 3 also show that the actual price behavior is very similar to what one could predict according to theory. We run a Mann-Whitney test to check whether the theoretical and actual distributions are different. At a 5% significance level, we cannot reject the null hypothesis that the two distributions are the same. By considering the average price level over the last three periods instead of the last period only, the results are not altered.

5 Informational Cascades and Trade Costs

In Section 2, we showed that the presence of a price mechanism destroys the possibility of an informational cascade. This result, however, hinges on the fact that the asset valuations of the market maker and of the traders differ only because of the private information that the traders have. When these valuations differ for other reasons, informational cascades are indeed possible, despite the fact that the price adjusts for the order flow. An asymmetry between the trader and the market maker arises, for instance, when trading is costly. If traders have to pay a trade cost to buy or sell the asset, their valuations will be equal to the expected value minus the trade cost. Therefore, trade costs induce a wedge between the traders and the market maker's expectations, making informational cascades possible (Lee, 1998). In this section of the paper, we study whether trade costs can generate informational cascades in the laboratory.

In order to see why trade costs make informational cascade possible, let us consider an economy similar to the one analyzed in Section 2, but where trade is costly, i.e., if an agent wants to buy or sell, he has to pay a trade cost $c \in (0, 100)$. In this economy, over time, the information contained in the history of trades will increase with respect to that contained in the private signal. As a consequence, when the history is long enough and there is a large enough imbalance between buys and sells, the difference between the price and the trader's expectations becomes smaller than the trade cost.

From that time on, every trader chooses not to trade and the asset price remains stuck for ever at the level that it had reached. Note that the price may remain stuck far away from the fundamental value. In particular, one can show that, irrespective of the fundamental, the price can converge to two different values, one close to 0 and one close to 100.¹⁴ Let us consider an economy where q equals 0.7, V is 100 and c is 9. In this economy, the price can end up at two different levels, 93 or 7. A price of 93 is reached when there is a trade imbalance of three buy orders. After such an imbalance, the difference between the expected value of the traders (with both a positive or a negative signal) and the price is smaller than 9 and, therefore, no one wants to trade. The price cannot move from that level as no more information is released to the market. The same happens at a price of 7, reached after there is a trade imbalance of three sells.¹⁵

In order to test whether in an experimental setup trade costs make informational cascades possible, we modified the flexible price experiment. If agents wanted to use the opportunity of trading, they had to pay a fixed cost of 9 lire. Therefore, in each round, the payoff was $100 + Value - Price - 9$ lire if the subject decided to buy; $100 + Price - Value - 9$ lire if he decided to sell; 100 lire if he decided not to trade. We refer to this experimental design as the “trade cost treatment.”

In the trade cost treatment agents decided not to trade in more than 51% of the trading periods, whereas they did so only 22% of times in the flexible price treatment (see Table 4). More importantly, these no trade decisions happened mostly when they should have occurred according to theory: indeed, 79% of no trades happened when the difference between the expected value of the trader and the price was not sufficient to compensate for the trading cost. In these cases, an informational cascade started, as people stopped trading according to their own private information. It is also worth mentioning that almost never did agents decide to trade when the trade cost was binding. Indeed, in the four sessions, we observed only 6 times when agents bought or sold the asset despite the fact that the expected gain from trade was lower than the trade cost. In other words, the cascade of no trades was almost never broken.

¹⁴Of course, the realized fundamental value will affect the probability with which these two levels are reached. If the fundamental is 0, it will be more likely that the price reaches the low value, if it is 100, it will be more likely that the price reaches the high value.

¹⁵We do not show all the computations behind these values, as they are a straightforward application of Bayes’s rule.

The fact that the trade costs induced people not to trade affected the price paths significantly. Figures 5 and 6 show the histograms of the final price (i.e., the price after all 12 subjects had traded) when the asset value was 0 and 100 respectively. The figures report the frequency of the prices actually observed and those that should have been observed, had agents acted according to the theory (i.e., following the signal when the trade cost was not binding and abstaining from trading otherwise).

Let us concentrate first on the observed prices. Out of the 28 rounds, the asset value happened to be 0 in thirteen cases and 100 in fifteen. When the value was 0, the price was lower than 7 in only one case, in eight cases it was between 7 and 30, it was twice at a level between 30 and 70, once at a level between 70 and 93 (included) and once higher than 93. When the value was 100, the price reached 12 times a level between 70 and 93, twice it remained at a level close to the unconditional expected value (i.e., between 30 and 70) and once it moved in the wrong direction converging to a value lower than 30.

This price behavior is significantly different from the one observed in the flexible price treatment. In particular, with trade costs the price was unable to converge close to the realized values. Furthermore, in some cases the price moved in the wrong direction. Instead, in the flexible price treatment, the price never converged to the wrong value. The revelation of the information dispersed across market participants and the price convergence was therefore significantly impaired by the presence of trade costs. As Figures 5 and 6 show, this was not the result of irrational behavior: indeed, the distribution of the last prices that we actually observed is very similar to the distribution that we would have obtained, had agents followed the theory exactly.

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 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 and reduce the ability of the price to reflect the fundamental value.

6 An Econometric Analysis

So far, we have reported and commented the main results of the experiment. Some questions naturally arise from the results that we reported. For instance, can we justify the contrarian behavior observed in the flexible price treatment? Was herding behavior in that treatment rational if one takes into account that agents could make errors? To answer these and other questions, we analyzed the data using a multinomial logit model.

6.1 A multinomial logit analysis

In the multinomial logit analysis, we defined our dependent variable “trade” as follows:

$$trade_t = \begin{cases} 0 & \text{to sell,} \\ 1 & \text{if the agent at time } t \text{ chose not to trade,} \\ 2 & \text{to buy.} \end{cases}$$

In the multinomial model, the probability of observing a particular choice $j = 0, 1, 2$ by the subject acting at time t is given by

$$\Pr(j) = \frac{e^{\beta'_j \mathbf{x}_t}}{\sum_{k=0}^2 e^{\beta'_k \mathbf{x}_t}}$$

where \mathbf{x}_t is the vector of independent variables.

We first ran a regression in which we use the trader’s private signal and the trade imbalance as independent variables. For the private signal, we constructed a variable called “signal,” set equal to -1 if the subject drew a blue chip and to 1 if he drew a white chip. We estimated the parameters β with maximum likelihood. We ran a logit regression for each treatment. For every treatment, we pooled the data of the four sessions, i.e., we had in each case 336 observations. The results of the four estimations are reported in Table 8, with p-values for the null that the coefficient is equal to 0 reported in parenthesis. In each cell we report first the coefficient relative to a sell order and, below, that relative to a buy order (the coefficients relative to a no trade are normalized to zero).

Let us consider, first, the estimation results for the fixed price treatment. Theory suggests that both the signal and the trade imbalance should play a role in determining the action of the traders, since agents should herd buy after a trade imbalance of 2 or more, and herd sell after a trade imbalance of -2 or less. This is indeed what we obtain in our estimation. The coefficients for both the signal and the trade imbalance are significant, with p-values of 0.00. There is statistical evidence that traders do follow the observed trade imbalance, i.e., they herd. The interpretation of the coefficients in a multinomial logit, however, is difficult, since their signs are not necessarily equal to the signs of the derivatives of the probabilities with respect to the dependent variables. For this reason, we also present, in Table 9, the estimated probabilities of a sell, a no trade and a buy for different values of the signal and of the trade imbalance. Subjects who observe a positive signal are clearly more likely to buy, for any level of the trade imbalance. Moreover, the trade imbalance modifies the probability of a buy or a sell significantly. For instance, the probability of a buy is about 62% when a subject has a positive signal and the trade imbalance is 0. With a trade imbalance of -2 this probability decreases to 25%; with a trade imbalance of 2, by contrast, it increases to 86%.

Let us look now at the results of the flexible price treatment. In this case, the probability of a trade should not be affected by the trade imbalance, only by the signal. This is what we find in the data: indeed, only the coefficients for the signal are significant, whereas the ones for the trade imbalance are not. Therefore, the analysis confirms that flexible prices destroy the tendency of agents to herd. For a given value of the signal, the probability of a buy or a sell is almost unaffected by different values of the trade imbalance (see Table 9).

We also ran a second regression where we used as regressors the signal and the maximum loss that the trader could experience when buying or selling the asset trading. The regressor maximum loss was built in the following way. Suppose that a subject has to make his decision when the price is 84. If he buys and the value of the asset turns out to be 0 he loses 84. If he sells and the values is 100, he loses 16. Therefore, the maximum loss in this case is 84. Clearly, we could not run a regression with the maximum loss as a regressor for the fixed price treatment, as the maximum loss would always be equal to 50.

In the flexible price treatment, we find that both the private signal and the

maximum loss affect the probability of a trade significantly.¹⁶ This suggests a possible explanation for the larger proportion (22%) of no trades that we found in the flexible price treatment relative to the fixed price treatment (14%). The logit analysis shows that agents preferred not to trade especially when the maximum loss was high, i.e., when the price was close to 0 or 100. As Table 9 shows, the probability of a no trade is strictly increasing in the maximum loss. When the maximum loss was 90 (i.e., the price was close to 90 or to 10) the probability of a no trade was almost the double than when it was 50 (i.e., the price was 50). Note that this behavior cannot be explained by risk aversion, since, when the price is close to 0 or 100, the variance of the payoff is lower than when the price is close to 50. Instead, what seems to induce agents not to trade is the perspective of experiencing a big loss compared to the status quo, a result reminiscent of Kahneman and Tversky (1979).

Finally, let us look at what happens in the trade cost treatment. In this case, theory suggests that the maximum loss should matter as well. A high value of the maximum loss means that the price has become close to 0 or 100. This, in turn, means that the difference between price and expected value of the asset is low. This should push traders to refrain from buying or selling, instead opting for no-trading. The econometric results confirm this. In fact, the maximum loss affects the probability of a sell or a buy significantly. When the price is 50 the probability of a no trade is 15% or 8%, depending on the signal. When the price is 10 or 90 this probability goes up to 67%.

6.2 An analysis of errors

During the experiment subjects may make errors. When errors occur in the early decisions, the following subjects may take these errors into account to compute their posteriors and choose their actions. Does the possibility of errors affect their behavior? We estimate the error rates assuming that the agents' expected payoffs are subject to shocks distributed independently as a logistic¹⁷. Moreover, our estimate is recursive, i.e., the estimate of the error rate in the first period is taken into account in the estimation of the error rate in the second period and so on.

¹⁶The results that we describe do not change if one runs a regression with signal, trade imbalance and maximum loss as independent variables (see Table 9).

¹⁷See McKelvey and Palfrey, 1995, 1997.

To focus ideas, let us consider first the choice problem that agent in period 1 faces. Suppose that he draws a white chip (high signal). His expected value of the asset is 70 lire, higher than the price of 50 lire. Therefore, if he buys, he has an expected payoff of $100 + (70 - 50) = 120$ lire. If he sells his expected payoff is $100 + (50 - 70) = 80$ lire. Finally, if he decides not to trade he has a payoff of 100 lire. Our logit analysis specifies that, at each time t , the probability of a buy, a sell or a no trade is a function of the difference between the expected payoff of buying or selling the asset (Π_t)¹⁸, i.e.,

$$\Pr(j) = \frac{e^{\gamma_j \Pi_t}}{\sum_{k=0}^2 e^{\gamma_k \Pi_t}}$$

Note that the model implies that the subject may not choose the action that yields a higher payoff (that is, that he may make a mistake). The γ 's used to calculate the probability of a particular action are estimated with a multinomial logit regression and allowed to change across time. At each time t , we ran a regression where the dependent variable were the trades of time $t - 1$ in all the rounds of all the sessions. The independent variable was the difference between the expected payoff of a buy and of a sell for the agents who made those trading decisions. These expected payoffs were calculated taking into account the past history of trades and the signal that the agents received. Note that, in estimating the γ 's for each period, we only used the information contained in the the trades of the previous period (i.e., we used a “rolling” regression).¹⁹

Let us move, now, to the second period. A person who has to make a decision in that round knows that the previous agent might have made a mistake. In particular, he can use the estimates of the γ 's to determine the probability that the previous agent had a white or blue chip depending on the decision he took. Given this probability (and his own private information), he will form his own expectation on the asset value. Therefore, the analysis allows us to calculate the expected value of the asset for a subject at time 2 who takes into account that the subject at time 1 may have made a mistake.

¹⁸The expected payoff of a no trade does not enter the model, since it is constant for all signals and for all times t .

¹⁹We also estimated the γ 's “recursively” (i.e., the γ 's for time T were estimated using all the data relative to all periods $t < T$). The results of the two estimation strategies are not very different.

Finally, we assume that the choice problem for the time 2 subject can be model in the same way as the choice problem for the time 1 subject. This allows us to calculate the expected payoff of buying, selling and not trading for the time 3 subject. This recursive structure is used to compute subjects' expectations in all the following periods.

We will use this analysis to interpret the data both in the flexible and in the fixed price treatment. In the flexible price treatment, there are two questions that we want to address. Was it ever rational for agents to herd when one takes into account that agents make mistakes? Was it ever rational to act as a contrarian? In principle, neither possibility can be excluded. Indeed, theory rules out herding and contrarian behavior of an agent at time t assuming that the price is set efficiently by the market maker and that all traders who traded before t acted correctly. In our experiment, however, while the price is set assuming that agents do not make mistakes, people could indeed choose the wrong action. To fix ideas, suppose, for instance, that some agents in the first periods made mistakes and bought the asset although they had a low signal. The market maker assumes that everyone is rational and therefore updates the price up after each buy on the basis of this assumption. When the next person has to make his decision, he will realize that the price is too high given the previous history of trades and, therefore, rationally decide to sell despite a negative signal. This would explain contrarian behavior. Similarly, think of the case where after some buy orders, some agents decide not to trade. If the next agent believes that these no trade decisions hide some positive signals, he will consider the price too low and therefore decide to herd buy, neglecting his negative signal. This would give a rational for the (admittedly modest) proportion of herding that we found in the flexible price treatment.

Tables 10 and 11 report the results of our estimation. The first row reports the actual number of herding (Table 10) and contrarian (Table 11) behavior in the first three treatments. The second row reports the percentage of times in which herding or acting as contrarians was “rational” if we assume that agents took into account the probability of a mistake by the previous agents.²⁰

As we have observed in Section 4, in the flexible price treatment agents decided to herd in 12% of cases of potential herding, for a total number of 8

²⁰The third row in both Tables refers to the case in which the γ 's are estimated with a “recursive” regression (see previous note).

herds. Even taking into account the possibility of errors, none of these herd decisions can be considered rational. Agents decided to act as contrarian 26 times, i.e., in 16% of cases of potential contrarian behavior. More than 70% of these decisions was indeed correct, taking into account the possibility of previous errors. The level of the price and the previous history of trades was such that acting as a contrarian was the appropriate (maximizing) choice.

Let us move now to the fixed price treatment.²¹ Theory suggests that herding is rational as the private information of the agent is, after some trades, overwhelmed by the public information contained in the history of trades. If people make mistakes, however, this is not necessarily the case. For instance, if two agents buy the asset, it is not necessarily true that the third should buy as well, having a negative signal. In fact, if the probability of the first two agents making mistakes is high, it may well be that the expected value of the asset, conditional on the first two buys and on the negative signal, is still lower than the price of 50. Our analysis suggests that even taking into account the possibility of errors, herd was rational in most of the cases in which traders decided to herd: in fact it was the rational decision in 67% of these cases.

7 Conclusions

In this paper we have reported and discussed the results of an experimental study on herd behavior and price convergence in financial markets. We have shown that, in a frictionless market where a market maker sets the price efficiently and agents trade only for pure informational reasons, herd behavior (informational cascades) occurs only rarely. Agents are not significantly affected by the decisions of the previous traders. In contrast, when the price of the asset is kept constant there are many situations when agents neglect their private information and rationally decide to herd. These results are consistent with the predictions of the theoretical models. Theory, however, does not capture another behavior observed in the laboratory market: the tendency of agents to act as contrarian, i.e., to trade against the market. Despite this tendency, the price on average aggregates private information and, therefore, converges to the true asset value. Finally, we have analyzed the effect of a trade cost in the market. When trade is costly, the price does not always fully aggregate private information. After some trades, agents

²¹A similar analysis of errors is performed by Anderson and Holt (1997).

prefer not to trade in order to save on the trade cost and the price is unable to aggregate private information efficiently.

8 Appendix

8.1 Instructions for the baseline experiment

8.1.1 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.

8.1.2 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 $\$1 = 65 \text{ lire}$. 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. 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*

$$100 + Value - p.$$

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

$$100 + p - Value.$$

- *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.

Therefore, when you buy, you will gain lire if the value of the good is higher than the price and lose them if it is lower. When you sell, you will gain lire if the price is higher than the value and lose them if it is lower. Always remember that the value of the good can be either 0 or 100. Notice also that you will never have a negative payoff, given that we give you 100 lire for trading. However, if you make the wrong decision, you will lose in part these 100 lire.

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 $100 + 0 - 70 = 30$. 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 $100 + 100 - 70 = 130$. 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 $100 + 70 - 0 = 170$. 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 $100 + 70 - 100 = 70$. 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 precede 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.

References

- [1] Anderson, Lisa and Holt Charles. "Information Cascades in the Laboratory." *American Economic Review*, December 1997, 87 (5), pp.847-862.
- [2] Avery, Christopher and Zemsky, Peter. "Multidimensional Uncertainty and Herd Behavior in Financial Markets." *American Economic Review*, September 1998, 88 (4), pp.724-748.
- [3] Banerjee, Abhijit, "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, August 1992, 107 (3), pp. 787-818.
- [4] Bikhchandani, Sushil; Hirshleifer, David and Welch Ivo. "A Theory of Fads, Fashion, Custom and Cultural Change As Informational Cascades." *Journal of Political Economy*, October 1992, 100 (5), pp. 992-1027.
- [5] Bikhchandani, Sushil and Sharma Sunil. "Herd Behavior in Financial Markets: A Review." 2000, IMF Working paper.
- [6] Bloomfield, Robert. "Quotes, Prices, and Estimates in a Laboratory Market." *Journal of Finance*, December 1996 (5), pp. 1791-1808.
- [7] Brunnermeier, Markus. *Asset Pricing Under Asymmetric Information. Bubble, Crashes, Technical Analysis and Herding*. Oxford University Press, Oxford, 2001.
- [8] Çelen, Boğaçan and Shachar Kariv. "Distinguishing Informational Cascades from Herd Behavior in the Laboratory," revised and resubmitted, *American Economic Review*, 2002.
- [9] Cipriani, Marco and Guarino Antonio. "Herd Behavior and Contagion in Financial Markets." mimeo, UCL and GWU, 2003.
- [10] Copeland, Thomas and Friedman Daniel. "The Effect of Sequential Information Arrival on Asset prices: An Experimental Study." *Journal of Finance*, July 1987 (3), pp. 763-797.
- [11] 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.

- [12] Huck, Steffen and Oechssler Jörg. “Informational Cascades in the Laboratory: Do They Occur for the Right Reasons?” *Journal of Economic Psychology*, 2000, 21, pp. 661-671.
- [13] Kahneman Daniel and Tversky Amos. “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica*, March 1979, 47, pp. 263-291.
- [14] Kübler Dorothea and Weizsäcker Georg, “Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory.” *Review of Economic Studies*, 2003, forthcoming.
- [15] Lakonishok, Joseph, Shleifer Andrei and Vishny Robert W. “The Impact of Institutional Trading on Stock Prices,” *Journal of Financial Economics*, 1992, 32, pp. 23-43.
- [16] McKelvey, Richard and Palfrey Thomas. “Quantal Response Equilibria for Normal Form Games.” *Games and Economic Behavior*, 1995, 10, pp. 6-38.
- [17] McKelvey, Richard and Palfrey Thomas. “Quantal Response Equilibria for Extensive Form Games.” *Experimental Economics*, 1997, 1, pp. 9-41.
- [18] Milgrom, Paul and Stokey, Nancy. “Information, Trade and Common Knowledge.” *Journal of Economic Theory*, February 1982, 26 (1), pp. 17-27.
- [19] Scharfstein, David and Stein Jeremy. “Herd Behavior and Investment.” *American Economic Review*, 80, 1990, pp. 465-479.
- [20] Smith, L. and P. Sørensen. “Pathological Outcomes of Observational Learning.” *Econometrica*, 2000, 68, pp. 371-398.
- [21] Welch, Ivo. “Sequential Sales, Learning, and Cascades.” *Journal of Finance*, 1992, 47, pp. 695-732.

Table 1: Trading decisions in periods of potential herd behavior

	Fixed Price	Flexible Price	No History
Herding	52%	12%	24%
No Trading	26%	42%	33%
Following Private Information	22%	46%	43%
Periods of Potential Herding	58	66	70

Table 2: Rational and irrational trading decisions during the experiment in the first three treatments

	Fixed Price	Flexible Price	No History
Rational Trades	80%	65%	61%
No Trading	14%	22%	25%
Irrational Trades	6%	13%	14%

Table 3: Trading decisions in periods of potential contrarian behavior

	Fixed Price	Flexible Price	No History
Contrarian Behavior	1%	16%	11%
No Trading	2%	17%	23%
Following Private Information	97%	67%	66%
Periods of Potential Contrarian Behavior	154	165	162

Table 4: Rational and irrational trading decisions during the experiment in the flexible price and the trade cost treatments

	Flexible Price	Trade cost	
Trade following the signal	65%	43%	
No Trading	22%	51%	79%
Irrational Trade	13%	6%	21%

Note: 79% of the no trade decisions in the trade cost treatment are rational.

Table 5: Individual behavior in the flexible price treatment (rational, no trading, and irrational decisions)

	Rational decision	No Trade	Irrational decision
$X \geq 5$	22	0	0
$3 \leq X < 5$	25	1	9
$X < 3$	1	47	39
Total	48	48	48

The table reports the number of subjects who traded rationally, irrationally or decided not to trade at least 5 times, between 3 and 5 times, or less than 3 times. For each individual we considered only the last 7 decisions that he made. For instance, the first number should be read as “22 subjects made the rational trade at least 5 times out of their 7 trading decisions.”

Table 6: Individual behavior in the fixed price treatment (rational, no trading, and irrational decisions)

	Rational decision	No Trade	Irrational decision
$X \geq 5$	39	0	1
$3 \leq X < 5$	8	0	3
$X < 3$	1	48	44
Total	48	48	48

The table reports the number of subjects who traded rationally, irrationally or decided not to trade at least 5 times, between 3 and 5 times, or less than 3 times. For each individual we considered only the last 7 decisions that he made. In this treatment, a trade was considered rational when the agent herded in situations of potential herding and followed the signal otherwise. For instance, the first number should be read as “39 subjects made the rational decision at least 5 times out of their 7 trading decisions.”

Table 7: Individual behavior in the fixed price and flexible price treatments (rational minus irrational decisions)

	Flexible Price	Fixed Price
$Z \geq 3$	71%	94%
$-3 < Z < 3$	29%	6%
$Z \leq -3$	0%	0%

The table reports the proportion of subjects for whom the difference between rational and irrational trades (excluding no trades) was at least 3, between -3 and 3, or less than or equal to -3. For each individual we considered only the last 7 decisions that he made. For instance, the first number should be read as “in the flexible price treatment, 71% of subjects made the rational trade at least 3 times more often than the irrational trade.”

Table 8: Estimated coefficients of the logit analysis for the four treatments

Regression 1 (independent variables: signal and trade imbalance)

	Flexible Price	No History	Fixed Price	Trade Cost
Signal	-0.3567 (0.0356)	-0.5732 (0.0005)	-1.1241 (0.0000)	-1.0214 (0.0000)
	1.2153 (0.0000)	0.8111 (0.0000)	1.7832 (0.0000)	1.0823 (0.0000)
Trade Imbalance	-0.0136 (0.7761)	-0.1387 (0.0031)	-0.4513 (0.0000)	0.0727 (0.0360)
	-0.0126 (0.7876)	-0.0300 (0.4894)	0.5453 (0.0007)	-0.1439 (0.3174)

Regression 2 (independent variables: signal and maximum loss)

	Flexible Price	No History	Fixed Price	Trade Cost
Signal	-0.3519 (0.0363)	-0.6003 (0.0003)	N/A	-0.8442 (0.0000)
	1.2201 (0.0000)	0.8511 (0.0000)	N/A	1.1624 (0.0000)
Max loss	-2.1236 (0.0177)	-1.2071 (0.1934)	N/A	-5.6736 (0.0000)
	-2.0226 (0.0424)	-1.9169 (0.0296)	N/A	-8.1823 (0.0000)

Regression 3 (independent variables: signal, trade imbalance and maximum loss)

	Flexible Price	No History	Fixed Price	Trade Cost
Signal	-0.347375 (0.0446)	-0.561762 (0.0008)	N/A	-0.882607 (0.0000)
	1.218688 (0.0000)	0.850004 (0.0000)	N/A	1.215613 (0.0000)
Imbalance	-0.006626 (0.8935)	-0.117018 (0.0147)	N/A	0.059174 (0.4713)
	0.002808 (0.9532)	0.018975 (0.6982)	N/A	-0.115662 (0.1512)
Max loss	-2.121041 (0.0185)	-0.783261 (0.4089)	N/A	-5.596668 (0.0000)
	-2.038680 (0.0430)	-2.080283 (0.0297)	N/A	-8.042086 (0.0000)

The table shows the estimated regression coefficients for each treatment. The first cell reports the estimated coefficient for the probability of a sell. The second reports the coefficient for the probability of a buy. The coefficient for the probability of a no trade is normalized to zero. The p-values for the null hypothesis that the coefficient is 0 are reported in parenthesis.

Table 9: Estimated probabilities of a sell, a no trade or a buy order in the four treatments

Regression 1 (independent variables: signal and trade imbalance)

Flexible Price Treatment

Trade Imbalance \ Signal	-2	0	2
-1	0.5916 0.2854 0.1229	0.5869 0.2909 0.1222	0.5820 0.2965 0.1214
1	0.1470 0.1447 0.7083	0.1462 0.1479 0.7059	0.1453 0.1511 0.7035

Fixed Price Treatment

Trade Imbalance \ Signal	-2	0	2
-1	0.9399 0.0586 0.0015	0.8579 0.1318 0.0102	0.6819 0.2584 0.0597
1	0.4684 0.2764 0.2553	0.1549 0.2254 0.6196	0.0295 0.1057 0.8648

No History Treatment

Trade Imbalance \ Signal	-2	0	2
-1	0.6457 0.2109 0.1434	0.5858 0.2525 0.1617	0.5231 0.2975 0.1794
1	0.1796 0.1846 0.6358	0.1481 0.2008 0.6512	0.1211 0.2168 0.6621

Trade Cost Treatment

Trade Imbalance \ Signal	-2	0	2
-1	0.3795 0.5348 0.0857	0.4228 0.5153 0.0619	0.4654 0.4904 0.0442
1	0.0370 0.4019 0.5611	0.0494 0.4644 0.4862	0.0645 0.5241 0.4115

Regression 2 (independent variables: signal and maximum loss)

Flexible Price Treatment

Max Loss \ Signal	50	70	90
-1	0.6825 0.1792 0.1383	0.6218 0.2497 0.1286	0.5480 0.3364 0.1156
1	0.1605 0.0852 0.7544	0.1513 0.1228 0.7259	0.1401 0.1739 0.6859

No History Treatment

Max loss \ Signal	50	70	90
-1	0.5693 0.2119 0.2187	0.5533 0.2622 0.1844	0.5284 0.3188 0.1528
1	0.1082 0.1339 0.7579	0.1156 0.1820 0.7024	0.1208 0.2422 0.6370

Trade Cost Treatment

Max loss \ Signal	50	70	90
-1	0.6931 0.1529 0.1540	0.5493 0.3768 0.0739	0.3110 0.6636 0.0253
1	0.0690 0.0824 0.8486	0.0823 0.3054 0.6123	0.0586 0.6771 0.2643

The table reports, for each treatment, the estimated probability that an agent sells the asset (first number in each cell), does not trade (second number) or buys the asset (third number), according to his own private information and to the trade imbalance (for regression 1) or according to his own private information and the maximum loss (for regression 2).

Table 10: Analysis of errors for the trading decisions in periods of potential herd behavior

	Fixed Price	Flexible Price	No History
Herds	30	8	17
Rational	20 (66.67%)	0 (0.00%)	0 (0.00%)
Rational II	16 (53.33)	0 (0.00%)	0 (0.00%)

The table shows the number and the percentage of herds that would have been rational even taking into account the probability that previous traders may have made mistakes. The first row indicates the number of herds actually observed. The second reports how many of those herds are rational according to the “rolling” regression. The third row reports how many of those herds are rational according to the “recursive” regression.

Table 11: Analysis of errors for the trading decisions in periods of potential contrarian behavior

	Fixed Price	Flexible Price	No History
Contrarian Behavior	1	26	17
Rational	0 (0.00%)	19 (73.08%)	13 (76.47%)
Rational II	0 (0.00%)	21 (80.77%)	13 (76.47%)

The table shows the number and the percentage of “contrarian trades” that would have been rational even taking into account the probability that previous traders may have made mistakes. The first row indicates the number of “contrarian trades” actually observed. The second reports how many of those trades are rational according to the “rolling” regression. The third row reports how many of those trades are rational according to the “recursive” regression.

Table 12: Results of the Mann-Whitney test

H⁰: The proportion of herding and contrarian behavior is the same in the two treatments

	Flexible price versus No History	Flexible Price versus Fixed Price
% of Herd Behavior	NO	YES
% of Contrarian Behavior	NO	YES

H⁰: The proportion of rational trades, no trades and irrational trades is the same in the two treatments


	Flexible price versus No History	Flexible price versus Fixed price	Flexible price versus Trade cost
Rational	NO	YES	YES
Irrational	NO	YES	YES
No Trade	NO	NO	YES

H⁰: The two distributions of the last price are the same

	Flexible price treatment	Trade cost treatment
Actual versus theoretical price	NO	NO


The table reports whether the above null hypotheses can be rejected (YES) or not (NO) at a 5% significance level. The tests are carried out by taking the averages of the relevant variable (e.g., the proportion of trades due to herding) for each of the four sessions. Using the Mann Whitney procedure, we tested whether the distribution of these four averages was significantly different between two treatments.

Figure 1: An example of theoretical and observed price paths when the realized asset value was 0 in the flexible price treatment




In this round all agents received the correct signal of 0. All should have sold the asset, thus driving the price monotonically to 0. The actual price converged fast, although agents 2 and 5 preferred not to trade and agent 6 acted as a contrarian, buying at a low price.

Figure 2: Histogram of the last price when the asset value was 0 in the flexible price treatment




The black histogram refers to the last prices (i.e., the prices at time 13) for all histories in which the fundamental was 0. The gray histogram refers to the last prices that would have occurred if all agents had behaved according to the predictions of the theoretical model.

Figure 3: Histogram of the last price when the asset value was 100 in the flexible price treatment




The black histogram refers to the last prices (i.e., the prices at time 13) for all histories in which the fundamental was 100. The gray histogram refers to the last prices that would have occurred if all agents had behaved according to the predictions of the theoretical model.

Figure 4: Histogram of the distance of the last price from the fundamental in the flexible price treatment




The black histogram represents the distance of the last prices (i.e., the prices at time 13) from the fundamental values. The gray histogram refers to the prices that would have occurred if all agents had behaved according to the predictions of the theoretical model.

Figure 5: Histogram of the last price when the asset value was 0 in the trade cost treatment




The black histogram refers to the last prices (i.e., the prices at time 13) for all histories in which the fundamental was 0. The gray histogram refers to the last prices that would have occurred if all agents had behaved according to the predictions of the theoretical model.

Figure 6: Histogram of the last price when the asset value was 100 in the trade cost treatment



The black histogram refers to the last prices (i.e., the prices at time 13) for all histories in which the fundamental was 100. The gray histogram refers to the last prices that would have occurred if all agents had behaved according to the predictions of the theoretical model.

Figure 7: Histogram of the distance of the last price from the fundamental when the asset value was 100 in the trade cost treatment



The black histogram represents the distance of the last prices (i.e., the prices at time 13) from the fundamental values. The gray histogram refers to the prices that would have occurred if all agents had behaved according to the predictions of the theoretical model.