

HERD BEHAVIOR IN FINANCIAL MARKETS: AN EXPERIMENT WITH FINANCIAL MARKET PROFESSIONALS

Marco Cipriani

George Washington University
International Monetary Fund

Antonio Guarino

University College London

Abstract

We study herd behavior in a laboratory financial market with financial market professionals. An important novelty of the experimental design is the use of a strategy-like method. This allows us to detect herd behavior directly by observing subjects' decisions for all realizations of their private signal. In the paper, we compare two treatments: one in which the price adjusts to the order flow in such a way that herding should never occur, and one in which the presence of event uncertainty makes herding possible. In the first treatment, subjects seldom herd, in accordance with both the theory and previous experimental evidence on student subjects. A proportion of subjects, however, engage in contrarianism, something not accounted for by the theory. In the second treatment, the proportion of herding decisions increases, but not as much as the theory would suggest. Moreover, contrarianism disappears altogether. In both treatments, in contrast with what theory predicts, subjects sometimes prefer to abstain from trading, which affects the process of price discovery negatively. (JEL: C92, D82, G14)

1. Introduction

In recent years, there has been much interest, both theoretical and empirical, in the extent to which trading in financial markets is characterized by herd behavior. Such an interest stems from the effects that herding may have both on financial markets' stability and on the markets' ability to achieve allocative and informational efficiency.

The editors in charge of this paper were Xavier Vives and Douglas Gale.

Acknowledgments: We thank Syngjoo Choi, Burkhard Drees, Douglas Gale, Giovanni Guazzarotti, Vincenzo Guzzo, Steffen Huck, Sjaak Hurkens, Brett Rayner, and the participants in the Workshop in Industrial Organization and Finance at IESE, in the WEF Conference at Warwick, and in seminars at GWU, NYU, and SMU for helpful comments and suggestions. The paper has greatly benefited from the comments of the editors and of the three referees. We also thank Brian Wallace, who wrote the experimental program, and Tom Rutter, who helped to run the experiment. Guarino gratefully acknowledges the financial support of the ESRC (World Economy and Finance program). The views expressed in this paper are those of the authors and do not necessarily represent those of the IMF or IMF policy. We are responsible for any errors.

E-mail addresses: Cipriani: marco.cipriani@gwu.edu; Guarino: a.guarino@ucl.ac.uk

The theoretical literature has tried to identify the mechanisms that lead traders to herd (for surveys, see, e.g., Gale 1996; Hirshleifer and Teoh 2003; Chamley 2004; Vives 2008). The theoretical contributions have emphasized that, in financial markets, the fact that prices adjust to the order flow makes it more difficult for herding to arise than in other setups, such as those studied in the social learning literature, where there is no price mechanism. Nevertheless, it is possible that rational traders herd, because there are different sources of uncertainty in the market, for example.

To test herding models directly with data from actual financial markets is difficult. In order to test for herd behavior one needs to detect whether agents choose the same action independently of their private information.¹ The problem for the empiricist is that there are no data on the private information available to the traders. As a result, it is difficult to determine whether traders make similar decisions because they disregard their own information and imitate other traders, or because they are reacting to the same piece of public information, for instance.²

To overcome this problem, some authors (Cipriani and Guarino 2005a; Drehmann, Oechssler, and Rider 2005) have tested herd behavior in a laboratory financial market. In the laboratory, participants receive private information on the value of a security and observe the decisions of other subjects. Given these two pieces of information, they choose sequentially if they want to sell, buy, or not trade a security with a market maker. In the laboratory one can observe the private information that subjects have when making their decisions, and therefore it is possible to test models of herding directly.

Our paper contributes to the existing experimental literature on herd behavior in financial markets by innovating in three significant aspects.

First, our sample consists of financial market professionals. The existing experimental studies on herding in financial markets use college undergraduates as subjects. As a result, one may wonder how representative these laboratory experiments are of the behavior of professionals operating in actual financial markets. The external validity of experimental studies is, indeed, a well-known concern in the literature. In our specific case, one may imagine that professional behavior in the field might differ from students' behavior in the laboratory because of age difference and different levels of education or training. Moreover, professional

1. Here we only discuss herding informally. In the next section we will give a formal definition.

2. The existing empirical literature on herd behavior in financial markets (see, e.g., Lakonishok, Shleifer, and Vishny 1992; Grinblatt, Titman, and Wermers 1995; Wermers 1999; Sias 2004) does not attempt to identify informational herding as is usually defined in the theoretical and experimental literature. These empirical studies present a statistical analysis that measures the extent of decision clustering by financial market participants (e.g., fund managers), independently of the underlying reasons for such clustering (see, e.g., the comments in Bikhchandani and Sharma 2001). An exception is a recent paper by Cipriani and Guarino (2006) that estimates a structural model of informational herding.

expertise, developed by working daily in financial markets, may lead to the development of trading heuristics that are different from those used by non-financial professionals.

Second, the existing literature has tested for the presence of herding in a market where, according to the theory, herding should never arise. In contrast, we compare two treatments: one (Treatment I) in which, as in the previous experimental work, subjects should always use their private information and never herd; the other (Treatment II) where, instead, herding becomes optimal because of event uncertainty, that is, uncertainty about the presence of informed traders in the market. The economy studied in Treatment II has never been analyzed experimentally (not even with a more standard pool of participants), although event uncertainty is recognized in the theoretical literature as one of the main channels of herding in financial markets (Avery and Zemsky 1998).

Third, we ran the experiment using a strategy method-like procedure that allowed us to detect herding behavior directly (whereas in previous work it could only be inferred indirectly). In particular, in previous experimental work subjects were asked to trade in sequence, one by one. Each subject received a private signal and then made a decision. In contrast, in our experiment, all subjects who have not yet traded make their decisions conditional on all signal realizations. Only after all subjects have chosen their strategies is one subject randomly chosen to trade and his strategy implemented for the realized signal value. This is a significant procedural novelty in the experimental literature on herding and informational cascades: Because each subject makes a decision for each signal realization, we can observe directly whether and when he chooses the same action irrespective of his private information. Moreover, because in each period of trading all subjects who have not yet traded are asked to make a decision for each signal, our data set is much larger than it would have been in the earlier experimental designs. This was particularly important given the difficulty of recruiting financial-market professionals in large numbers.

The results of our experiment show that, as theory suggests, the proportion of herding decisions is very low in Treatment I. Therefore, the theoretical prediction that price adjustment to the order flow reduces the scope for herding behavior is confirmed by the experimental data on financial market professionals. Moreover, also in accordance with the theory, herding increases in Treatment II, where the price adjustment rule is consistent with the presence of event uncertainty. Nevertheless, some important anomalies do occur in the laboratory. First, in Treatment I, some subjects engage in contrarianism, something not predicted by the theory. These subjects go against the market, selling (regardless of the private signal) when the price is high, and buying (regardless of the private signal) when it is low. Second, in Treatment II, herd behavior is lower than theory predicts. Third, in both treatments, subjects have a tendency to abstain from trading, which is not predicted by the theory. Abstention from trading implies that the market is unable

to infer the subjects' private signals, which lowers the informational efficiency of the market.

It is worth noting that our results in Treatment I are similar to those obtained by previous experimental work using student subjects. In both samples, the proportion of herding is low, as the theory predicts. Moreover, subjects in both samples share the propensity to act as contrarians and to abstain from trading more than is predicted by the theory. This reassures the reader of the validity of previous experimental work that relies on student subjects.

Before moving to the main analysis, we now provide a brief literature review.

Literature review. Our paper is related to the theoretical literature on herd behavior in financial markets. In particular, our experimental setup is based on the analysis of Avery and Zemsky (1998). They show that, in a sequential trading model such as Glosten and Milgrom (1985), where the price is correctly set by a market maker according to the order flow, traders never herd. Herding, however, arises if there is uncertainty not only regarding the fundamental value of the asset, but also regarding other characteristics of the market, such as the proportion of informed traders in the market (event uncertainty). Park and Sabourian (2006) have recently revisited Avery and Zemsky's model and provided different conditions on the signal structure under which herd behavior can arise. Other scholars have shown that informational cascades and herding in financial markets occur when traders have informational and non-informational (e.g., liquidity or hedging) motives to trade (Cipriani and Guarino 2008a), or when trading activity is affected by reputation concerns (Dasgupta and Prat, forthcoming).

Our work belongs to the experimental literature on herding in financial markets. We have already mentioned that Cipriani and Guarino (2005a) and Drehmann et al. (2005) have tested for herd behavior in financial markets using student subjects.³ One of the purposes of our paper is to compare the behavior of financial market professionals with that of students. In Section 5, we will discuss in detail how our results compare with those of these papers. Cipriani and Guarino (2008b) have shown that informational cascades form in a laboratory financial market in the presence of transaction costs. Because there are no transaction costs in the experiment described here, this type of cascade cannot arise.

Finally, our paper is also close in spirit to Alevy, Haigh, and List (2007). Like us, they use financial professionals in their experimental study. In contrast to our study, however, they test a standard informational cascade game based

3. Note that Drehmann et al. (2005) study herding behavior in an experimental financial market using a sample of both students and consultants. They use consultants as a control group, focusing on their level of rationality, which they find similar to that of student subjects; they do not present any result on the propensity of consultants to herd or act as contrarian. Consultants in their sample are not financial market professionals and therefore their behavior may differ from that of financial market actors because of differences in training, experience, and expertise.

on Bikhchandani, Hirshleifer, and Welch (1992) and not a model of trading in financial markets.⁴ They find that market professionals rely on their private information to a greater extent than student subjects do; as a result, fewer cascades (and especially fewer cascades on the wrong action) form in the laboratory.

The rest of the paper is organized as follows. Section 2 describes the theoretical model and its predictions. Section 3 presents the experimental design. Section 4 illustrates the main results. Section 5 compares them with the results in the existing experimental literature. Section 6 discusses individual behavior. Section 7 concludes.

2. The Theoretical Model

2.1. The Model Structure

As we mentioned in the Introduction, our experimental analysis is based on the theoretical model of Avery and Zemsky (1998), who analyze herd behavior in an economy similar to that of Glosten and Milgrom (1985) and Easley and O'Hara (1987). In contrast to these papers, however, we assume that the market maker can post only one price, that is, he is not allowed to post different prices at which traders can buy (the ask price) or sell (the bid price). We adopt this assumption to simplify the implementation of the trading game in the laboratory. All the results that we present in this theoretical section hold independently of whether the market maker is allowed to post a bid and an ask price.

In our market there is one asset traded by a sequence of traders who interact with a market maker. Time is represented by a countable set of trading periods, indexed by $t \in 1, 2, \dots$. Traders act in an exogenously determined sequential order. Each trader, indexed by t , is chosen to take an action only once, at time t .

The asset value. The fundamental value of the asset is a discrete random variable v . An information event occurs with probability p ;⁵ in this case, the asset value takes the values 0 and 100 with probability $1/2$. In contrast, with probability $(1 - p)$, there is no information event and v takes a value of 50. This assumption is meant to capture the idea that, during a day of trading, information may arrive in the market and push the fundamental value of the asset up or down. In contrast, in the case of no event, the asset value remains at its unconditional expected value.

4. Other experimental studies on non-financial herding and cascades, based on Bickchandani et al. (1992), include Anderson and Holt (1997), Çelen and Kariv (2004), Goeree et al. (2007), Huck and Oechssler (2000), and Kübler and Weiszäcker (2004).

5. The event is called "informational" because—as we shall see—when it occurs, some traders receive private information on it.

The market. At each time t , a trader can exchange the asset with a 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. We denote the action of the trader at time t by x_t and denote the history of trades and prices up to time $t - 1$ by h_t .

The market maker. At any time t , the market maker sets the price at which a trader can buy or sell the asset. He sets the price equal to the expected value conditional on the public information available at time t , that is,⁶

$$p_t = E(v | h_t).$$

The traders. Traders are of two types, noise traders and informed traders. If the value of the asset is 50 (i.e., there is no information event), there are only noise traders in the market. Noise traders act for “liquidity” or other exogenous reasons, buying, selling, or not trading with exogenously given probabilities. If, instead, an information event occurs and the value of the asset is either 0 or 100, then at each time t the trader acting in the market is an informed trader with probability μ and a noise trader with probability $1 - \mu$. Informed traders receive private information on the realization of the asset value. In particular, if at time t an informed trader is chosen to trade, he observes a symmetric binary signal on the realization of v with distribution

$$\Pr(s_t = 100 | v = 100) = \Pr(s_t = 0 | v = 0) = 0.7.$$

In addition to his signal, an informed trader at time t observes the history of trades and prices and the current price. Therefore, his expected value of the asset is $E(v | h_t, s_t)$. The informed traders’ payoff function is defined as

$$U(v, x_t, p_t) = \begin{cases} v - p_t & \text{if } x_t = \text{buy,} \\ 0 & \text{if } x_t = \text{no trade,} \\ p_t - v & \text{if } x_t = \text{sell.} \end{cases}$$

Informed traders are risk neutral and choose x_t to maximize $E(U(v, x_t, p_t) | h_t, s_t)$. Therefore, they find it optimal to buy whenever $E(v | h_t, s_t) > p_t$ and sell whenever $E(v | h_t, s_t) < p_t$. They are indifferent among buying, no trading, and selling when $E(v | h_t, s_t) = p_t$.

6. In the original Glosten and Milgrom (1985) model the market maker posts a bid price and an ask price and makes zero expected profits because of unmodeled potential competition. In our model, by setting one price only, the market maker earns negative expected profits. This is not a problem because in the experiment the market maker is not a subject, but an automaton.

2.2. Theoretical Predictions

We now illustrate the predictions of our model by analyzing two distinct parameterizations, each corresponding to one of the two treatments that we ran in the laboratory. In the first parameterization, we set $p = 1$, that is, we assume that an information event occurs with certainty. In this case we also assume that $\mu = 1$, that is, all traders in the market are informed. In the second parameterization, we set $p = 0.15$ and $\mu = 0.95$, that is, we assume that an information event occurs with probability strictly smaller than 1, and that, if the event occurs, there is a small proportion of noise traders in the market. Moreover, noise traders abstain from trading with probability 0.33 during an informed day and with probability 0.02 during an uninformed day and, if they trade, they buy and sell with equal probability.⁷

To discuss the theoretical predictions of the model, let us first introduce the formal definitions of cascade behavior, herd behavior, and contrarianism that we will use in our analysis.

DEFINITION 1. *An informed trader engages in cascade behavior if he chooses the same action independently of the private signal. If the chosen action conforms to the majority of past trades the trader engages in herd behavior. If the chosen action goes against the majority of past trades the trader engages in contrarian behavior.*

For instance, if a trader buys irrespective of whether he received a signal of 0 or 100, we say that he engages in *cascade behavior*. If the buy order follows a history in which there are more buy than sell orders, the trader herds.⁸ If instead the buy order follows a history with more sell than buy orders, the trader acts as a contrarian.

Herding and contrarianism are two particular types of cascade behavior. Cascade behavior, however, is a more general concept. For instance, a trader also engages in cascade behavior if he abstains from trading for any realization of his private signal. When describing the experimental results we will find it useful

7. This parameterization, with a strictly positive proportion of noise traders and a different probability of no trade by noise traders when there is no information event, makes the implementation of the model in the laboratory more natural. We will explain this in detail when we illustrate the experimental procedures.

8. It is worth clarifying the relation between the standard definition of herd behavior in the social learning literature and ours. In this literature (see, e.g., Gale 1996; Smith and Sørensen 2000), a herd is said to occur when a sequence of agents make the same decision (not necessarily ignoring their private information). Here, instead, we define herd behavior as a particular type of cascade behavior. Our departure from the standard definition is motivated by the fact that our definition is particularly convenient for the experimental analysis. In the analysis we elicit subjects' strategies conditional on the signal realizations, which is more informative than only observing the actions. Our definition of herding allows us to study when subjects ignore their private information to conform to the established pattern of trade.

to distinguish between “cascade trading behavior” (when a trader engaging in cascade behavior either buys or sells) and “cascade no-trading behavior” (when he instead decides to abstain from trading).

Following Avery and Zemsky (1998), it is easy to show that, in the first setup (i.e., when an informational event occurs with probability one), cascade behavior cannot arise; whereas in the second setup (with event uncertainty) cascade behavior (and, in particular, herd behavior) arises with positive probability. In contrast, contrarianism and the other type of cascade behavior mentioned previously never arise in equilibrium. We summarize this in the next two results.

RESULT 1. *If an informational event occurs with certainty ($p = 1$), in equilibrium traders always trade according to their private signal and never engage in cascade behavior.*

To explain the result, let us recall that, in order to decide whether to buy or to sell the asset, a trader computes its expected value and compares it to the price. If at time t a trader receives a signal of 100, his expected value is

$$\begin{aligned} E(v | h_t, s_t = 100) &= 100 \Pr(v = 100 | h_t, s_t = 100) \\ &= 100 \frac{(0.7) \Pr(v = 100 | h_t)}{(0.7) \Pr(v = 100 | h_t) + (0.3)(1 - \Pr(v = 100 | h_t))} \\ &> 100 \Pr(v = 100 | h_t) = E(v | h_t) = p_t, \end{aligned}$$

and, therefore, he buys. Similarly, if he receives a signal of 0, his expected value is lower than the market price and he sells. This shows that an agent always finds it optimal to trade according to his private information and cascade behavior cannot arise.

Let us turn now to the case in which $p = 0.15$, that is, in which there is uncertainty about whether or not the value of the asset changed from its unconditional expectation. In such a case, it can be optimal for agents to neglect their private information and herd.

RESULT 2. *In the presence of event uncertainty ($p < 1$), in equilibrium herd behavior occurs with positive probability.*

Here, we only discuss the intuition for this result and refer the reader to Cipriani and Guarino (2006) for a formal proof. When an informed trader receives a private signal, he learns that an event has occurred. Therefore, when he observes a sequence of trades, he knows that each buy or sell order comes from an informed trader with probability 0.95. He will update his belief on the asset value on the basis of this information. The market maker, by contrast, has a prior belief of




FIGURE 1. Prices and traders' expectations after a history of buys.

0.86 that a trade comes from a noise trader.⁹ Therefore, when he receives a buy or a sell order, he updates his belief (i.e., the price) by less than the traders. As a result, after a sequence of buy (sell) orders, the expectation of a trader may be higher (lower) than the price even if he receives a bad (good) signal.

In Figure 1, we show the sequence of expectations and prices after a series of buy orders. At time 3, the equilibrium price is lower than both the expectation of a trader receiving a good signal and the expectation of a trader receiving a bad signal. Therefore, the trader at time 3 will buy regardless of the signal he receives, that is, he will herd.

Note that, because the market maker updates his expectation (and the price) by less than the informed traders, it will never be the case that, after a history of buys, the expectation of a trader will be below the price for both signal realizations. Analogously, after a history of sell orders the expectation of an informed trader will never be above the price for both signal realizations. As a result, an informed trader will never engage in contrarian behavior.

The presence of herding in the market is, of course, important for the informational efficiency of prices. During periods of herd behavior, private information is not efficiently aggregated by the price. In these periods, traders do not make use of the private information they have and, as a result, the market cannot learn such information.

Even during a period of herding, although the price does not aggregate private information efficiently, the market maker does learn something on the true asset value. Indeed, even in a period of herding, he updates his belief on whether there has been an informational event.¹⁰ For this reason, in Figure 1 the price keeps moving even after time 3, even though traders are herding. The market

9. The value 0.86 is equal to $(1 - p) + p(1 - \mu)$.

10. Therefore, although in our model traders engage in herd behavior (and, hence, in cascade behavior), a blockage of information never occurs. In the social learning literature, such a blockage of information is called an *informational cascade*. In most setups, acting independently of the signal (i.e., engaging in cascade behavior) implies a blockage of information. This, however, is not true in our setup. For this reason, we prefer not to talk of informational cascades in the paper, and only use the concept of cascade behavior.

maker observes more and more traders buying the asset and gives more and more weight to the event that these traders are informed (noise traders would buy or sell with equal probabilities). Because of this price movement, herd behavior will eventually disappear. As shown in Figure 1, during a period of herding the traders' expectations do not move (because the traders already know that an event has occurred and they also know that informed traders are herding rather than using their signals). When the price becomes higher than the expectation conditional on a bad signal, agents will no longer find it optimal to herd. On the contrary, they will trade according to their private information. In our figure, this occurs at time 7. The model, therefore, explains temporary herd behavior. Clearly, Figure 1 is just an example because the occurrence and subsequent breaking of herd behavior depends on the specific sequence of trades.

3. The Experiment and the Experimental Design

3.1. *The Experiment*

We ran the experiment in the Experimental Laboratory of the ELSE Centre at the Department of Economics at UCL between December 2006 and February 2007. The participants were 32 financial professionals working for financial institutions operating in London. We ran four sessions and each subject participated in exactly one session.¹¹

The experiment was programmed and conducted with the software *z-Tree* (Fischbacher, 2007). The sessions started with written instructions given to all subjects.¹² We explained to the participants that they were all receiving the same instructions. Subjects could ask clarifying questions, which we answered privately. The experiment consisted of two treatments. The first treatment started with two practice rounds, followed by seven rounds in which subjects received monetary payments. After completing the first treatment, participants received the instructions for the second one. Then they took part in the second treatment, which consisted again of seven paid rounds.¹³

The participants acted as informed traders and could exchange an asset with a computerized market maker. In both treatments, we implemented our model conditioning on an information event having occurred. The two treatments differed with respect to the price-updating rule used by the market maker.

11. We also conducted a pilot session with eight more participants. In that session, we used a different payoff function to pay the subjects. For this reason, we do not include the data from the pilot session in the analysis of our results.

12. The instructions are available on the journal's and on the authors' Web pages: (<http://www.jeea.org>), (<http://home.gwu.edu/~mcprian>); (<http://www.homepages.ucl.ac.uk/~uctpagu>).

13. The seven rounds of the second treatment were not preceded by practice rounds because the two treatments were very similar.

Let us now explain the procedures for each round of the experiment in detail.

1. At the beginning of each round, the computer program randomly chose the asset value. The value was equal to 0 or 100 with probability $1/2$. Each random draw was independent.
2. Participants were not told the realization of the asset value. They knew, however, that they would receive information on the asset value in the form of a symmetric binary signal. If the asset value was equal to 100, a participant would get a “white signal” with probability 0.7 and a “blue signal” with probability 0.3. If the value was equal to 0, the probabilities would be inverted.¹⁴
3. Each round consisted of 8 trading periods. In the first trading period, all 8 subjects made two trading decisions, conditional on the two possible signal realizations. They had to choose whether they would like to buy or sell one unit of the asset (at the price of 50) or not to trade, both in the event of receiving a white signal and in the event of receiving a blue signal. After all eight participants made their decisions, the computer program randomly selected one of them (with equal probability) as the actual trader for that period. That subject received a signal (according to the rule indicated under point 2) and his decision conditional on the signal was executed.
4. The other subjects observed on their screens the executed trading decision, as well as the new price for period 2. The identity of the subject whose decision was executed, however, was not revealed.
5. In the second period, there were seven subjects whose decisions had not yet been executed. As in the first period, they indicated whether they wanted to buy, sell, or not to trade conditional on the white and the blue signal. Then, one of them was randomly selected, received a signal, and his decision conditional on that signal was executed.
6. The same procedures were repeated for eight periods, until all subjects had acted once. Note that all subjects (including those whose decision had already been executed) observed the trading decisions in each period and the corresponding price movement. Indeed, the computer program moved from one period to another only after all eight participants had observed the history of trades and prices, and had clicked on an “OK” button.
7. At the end of the round, after the decisions of all the eight subjects were executed, the realization of the asset value was revealed and each subject saw his own payoff for that round on the screen. The payoffs were computed as follows: if he had bought, the subject obtained $v - p_t$ of a fictitious experimental currency called “lira”; if he had sold, he obtained $p_t - v$ lire; finally, if he had decided not to trade, he earned (and lost) nothing. After participants had

14. That is, the white signal corresponded to $s_t = 100$ and the blue signal to $s_t = 0$.

observed their payoffs and clicked on an “OK” button, the software moved to the next round.

As should be clear from this description, compared to the existing experimental literature on informational cascades, we introduced the procedural novelty of a strategy-like method. This has the advantage that we could detect cascade behavior directly. A subject engages in cascade behavior when he makes the same decision independently of his signal realization. Because in our experiment a subject made a decision for each possible signal realization, we could directly observe whether he chose the same action for both signal realizations.¹⁵ Furthermore, with this method, we collect much more information on the subjects’ decision process than with the traditional procedures used in informational cascades experiments (in which a subject is first chosen to trade, then receives a signal and finally makes a decision). Indeed, in each treatment, we observed on average 36 decisions per subject, instead of just 7 (one per round). At the same time, our procedure was easy to implement and was quite natural for financial market professionals, because they are used to the idea of a conditional market order that is not necessarily executed.¹⁶

At the end of the experiment, we summed up the per-round payoffs of both treatments and converted them into pounds at the rate of 3 lire per pound. With this exchange rate the incentives were clearly much stronger than in most experiments. In addition, we gave subjects £70 just for participating in the experiment.¹⁷ On average, subjects earned £134 (approximately equal to £263 and €196 at the time of the experiment) for a 2.5-hour experiment. The minimum payment amounted to \$38 while the maximum was £268, with a standard deviation equal to £44.¹⁸

Finally, before leaving, subjects filled out a short questionnaire, in which they reported some personal characteristics (gender, age, education, work position, job tenure) and described their strategy and their beliefs on other subjects’ strategy

15. In the existing experimental literature, instead, cascade behavior is typically detected by focusing on the decisions of subjects when they receive a signal against the history of trades. The reason is that, in almost all the existing experiments, subjects first receive the signal and then are asked to make a decision. An important exception is Çelen and Kariv (2004), who employ continuous action and signal spaces to distinguish informational cascades from herd behavior in a non-market experiment.

16. Note that the procedure that we employ is not identical to the strategy method. With a strategy method, we should have asked each participant to make a decision for each possible contingency. Because there is a very large number of histories of trades, this would have been impossible to implement. In contrast, our method allowed us to collect a large dataset while keeping the process of trading simple.

17. The fixed payment was given to make sure that participants did not end up with losses.

18. We could have used the lottery method to pay our subjects in order to try to control for risk preferences. Because previous experimental work by Drehman et al. (2005) has found that using the lottery method does not produce significantly different results in this type of experiment, we preferred to use the more natural and simple way of computing payoffs.

in the experiment. Immediately after completing the questionnaire, subjects were paid in private and could leave the laboratory.¹⁹

3.2. *Experimental Design: The Two Treatments*

As we mentioned before, the difference between the two treatments is in the price-updating rule. In Treatment I, we implemented the model without event uncertainty described in Section 2 (i.e., the parameterization with $p = 1$ and $\mu = 1$). In Treatment II, we implemented instead the model with uncertainty about the informational event (i.e., with $p = 0.55$ and $\mu = 0.95$).

In Treatment I there is always an information event; whereas in Treatment II an information event occurs with probability 0.15. Nevertheless, in Treatment II we ran the experiment assuming that an information event had occurred. Therefore, from the participants' viewpoint, the main difference between the two treatments was how the price was updated for a given order flow.

Let us illustrate how we update the price. As explained in the previous section, the market maker sets only one price.²⁰ According to the theory, in Treatment I in equilibrium subjects should always follow their signal, that is, they should buy after seeing a white signal and sell after seeing a blue one. No one should decide not to trade, as private information allows the traders to make profits by trading with the market maker. Therefore, when a subject decides to buy, the price is updated assuming that he has seen a good signal. Similarly, when a subject decides to sell, the price is updated assuming that the subject has observed a bad signal. Finally, in the case of a no trade, the price is kept constant. As a result, in this treatment, the price moves through a grid. It starts at time 1 at the unconditional expected value of 50. After a sequence of buys, it moves, according to Bayesian updating, through a sequence of values 70, 84, 93, 97, 99, Similarly, after a sequence of sell orders, it moves through a sequence of values 30, 16, 7, 3, 1, The price at each time t only depends on the trade imbalance, that is, on the difference between the number of buy and sell orders observed up until the previous period $t - 1$.

In Treatment II, we change the price updating rule, following the theoretical model with event uncertainty. We implement the treatment in the laboratory by

19. In designing the experiment and the questionnaire we made sure to maintain subjects' anonymity. In particular, we made clear that the procedures were such that we would not be able to link each individual performance to a name or to a subject's institution. Moreover, in the experiment it was impossible to know the identity of other subjects in the sequence.

20. Allowing the market maker to set only one price makes the experiment easier to run. In their experiment with student subjects, Cipriani and Guarino (2005a) compare the results of a treatment with only one price set by an automaton (as in this paper) and a treatment where subjects acting as market makers were allowed to post bid and ask prices. They find that the results are not affected by the presence of the bid-ask spread.

explaining to the subjects that, in the second part of the experiment, the market maker will update the price as if, with high probability, it were trading not with informed traders, but with noise traders.²¹ As in the previous treatment, participants can observe the amount by which the computer updated the price before they made their decisions. Therefore, they have all the information needed to maximize their payoffs. Figures 1 and 2 show the price movement after a sequence of eight buy and eight sell orders. We have already commented on Figure 1 in the previous section. Let us focus on Figure 2 here. After the sell orders the price decreases, but by less than in Treatment I. As a result, subjects should follow the signal in the first two periods but then they should sell independently of the signal (herding on the previous actions) in periods 3 to 6. At time 7 the price is low enough that subjects should now sell only conditional on a blue signal (and buy conditional on a white one). Figure 3 offers another example of the price changes, following a sale at time 1 and a series of buy orders later on. In this case subjects should herd only starting at time 5, whereas they should follow their signals in the first 4 times. Note, that, as in Treatment I, the price is updated assuming that traders choose the optimal action, that is, they follow their private information when their expectation conditional on a white (blue) signal is above (below) the market price, and they buy (sell) irrespective of their signal when we are in a herd buy (herd sell) period.

3.3. *The Pool of Participants*

The study was conducted with 32 financial professionals employed in 13 different financial institutions, all operating in London. Out of the 32 participants, 28% were traders, 47% market analysts, 9% sale or investment management persons, 9% investment bankers, and 6% managers.²² Eighty-four percent of

21. Another difference between the parameterization of the first and the second treatment is that, in the second treatment, there were 5% of noise traders. We implemented this in the laboratory by having a 3.3% probability in each trading period of a wrongly executed trading order (e.g., with a 3.3% probability a sale or a no trade was executed, although the true order coming from the participant was a buy). This is equivalent to saying that there was a 5% probability that in each period the trade was coming from a “noise trader.” The presence of noise traders in the second treatment was necessary for the following reason. Suppose that at time t a rational subject should always buy (because we are in a herd buy period). If the subject chosen to trade decides to sell, in the absence of noise traders, the market maker would infer that the market is uninformed, that is, that all traders are noise traders. The market maker would, therefore, set the price equal to 50 for the entire round. Having a proportion of noise traders when there is an information event prevents this from happening. Also recall that, in the parameterization of the second treatment, the probability of a noise trader deciding not to trade differs according to whether an information event has occurred or not (33% and 2%, respectively). This is tantamount to imposing that no trades do not convey information on the likelihood of an information event to the market maker and, as a result, the Bayesian updating rule implies no change in the price after a no trade (as also happens in the first treatment), which is a natural and desirable feature.

22. We use “investment banking” in its stricter meaning, as one of financial institutions’ core functions. Moreover, “analyst” refers to the function within the institution and not to the rank.




FIGURE 2. Prices and traders' expectations after a history of sells.




FIGURE 3. Prices and traders' expectations after a sell followed by a history of buys.

subjects were men and 16% women. The participants' ages ranged between 21 and 40 years, with a mean equal to 28 years and a standard deviation equal to 4.9 years. The average job tenure was 4 years, with a range between 3 months and 16 years (standard deviation: 4.2 years). Finally, 8% of participants had a Ph.D., 61% an M.A./M.S., and 31% a B.A./B.S. Most participants (68%) with a B.A./B.S. degree had studied economics/finance/business; by contrast, the Masters degrees were split almost equally between economics/finance/business and scientific or technical disciplines such as physics, mathematics, or engineering; finally, the Ph.D. degrees were in physics or computer science.

4. Results: Rationality, Herding, and Contrarian Behavior

We now turn to a discussion of the results of the experiment. For expositional reasons, we find it convenient to present first the results of Treatment I and then (in Section 4.2) to illustrate those of Treatment II.

TABLE 1. Average behavior in Treatment I.

Decision	(%)
Following private information	45.7
Partially following private information	19.6
Cascade trading	19.0
Cascade no-trading	12.3
Errors	3.4
Total	100

4.1. Treatment I

Table 1 breaks down the participants' decisions in Treatment I according to how they used their own private information. In 45.7% of the cases, subjects just followed their private signal, buying on a white signal, and selling on a blue one. Recall that this is the rational behavior that theory predicts in equilibrium.²³ In 19.6% of the cases, instead, they followed one of the two signals, but preferred to abstain from trading conditional on the other. In 19% of the cases, they decided to disregard private information and buy or sell conditional on both signals—that is, they engaged in cascade trading behavior. In 12.3% of the cases, subjects preferred not to trade independently of their private information—that is, they engaged in cascade no-trading behavior. Finally, there are a few cases (3.4%) in which subjects made decisions that are self-contradictory for any possible belief.²⁴

This aggregate behavior clearly shows that, although the theory captures some of the trading rules that subjects used in the laboratory, there are some departures from the equilibrium predictions that must be explained.²⁵ First, we must understand why subjects sometimes decided to engage in cascade behavior and trade independently of the signal. One possibility is that a subject may neglect private information to herd. As we mentioned in Section 2, according to the theory, herding should not occur in this treatment. Subjects in the laboratory, however,

23. Following one's private information is rational only if each subject believes that all his predecessors are rational, that all his predecessors believe that their predecessors are rational, and so on. Furthermore, after a no-trade decision, which is always off the equilibrium path, subjects should not update their beliefs (which is consistent with our price updating rule), should believe that their predecessors did not update their beliefs, and so on.

24. For instance, we observed some decisions to sell conditional on a white signal, but not to trade conditional on a blue signal, which can only be interpreted as a mistake because a white signal always conveys more positive information about the asset value than a blue one.

25. Note that the results in Table 1 overweigh decisions taken in the first periods (when all subjects take a decision) with respect to those taken at later periods (when fewer subjects do so). This implies that the results overweigh decisions taken when the trade imbalance is 0 with respect to those taken when the trade imbalance is high. In the following analysis, we will control for this, by studying the decisions taken conditional on a given level of the trade imbalance.

TABLE 2. Cascade trading behavior in Treatment I.

Absolute value of the trade imbalance	Cascade trading (%)	Herd behavior (%)	Contrarian behavior (%)
0	5.8		
1	18.5	5.7	12.9
2	42.7	16.1	26.6
3	54.3	23.9	30.4
≥ 4	62.5	21.9	40.6

may give more weight to public information (i.e., the history of trades) than our price updating rule does, and believe that conditioning the trade on the private signal is not optimal when the order flow already shows evidence in favor of the asset value being high or low. A second possibility is that a subject may decide to act as a “contrarian” by going against the market. This behavior should not occur in equilibrium either, but a subject may use the strategy of going against the market to sell at a high price and buy at a low one.

Table 2 shows how cascade trading behavior evolved according to the absolute value of the trade imbalance, that is, the absolute value of the difference in the number of buy and sell orders. There is a monotonic increase in the proportion of cascade-trading decisions as the trade imbalance increases: When the trade imbalance is 0, cascade trading behavior accounts for less than 6% of decisions, whereas, for an absolute value of the trade imbalance of 3 or more, it accounts for more than 50% of decisions.

Note that, when the trade imbalance is 0, we cannot classify cascade behavior as herding or contrarianism. In such a case the number of buy and sell orders is identical, and the price is equal to the unconditional expected value of 50. Therefore, the subjects’ decisions to buy or sell independently of the signal cannot be explained either in terms of following the crowd or going against it. By contrast, when the absolute value of the trade imbalance is at least 1 we can distinguish between herd and contrarian behavior as explained in Section 2.

As Table 2 shows, the evolution of herding and contrarianism with the trade imbalance is quite different. When the absolute value of the trade imbalance increases, so does the evidence in favor of the asset value being 0 or 100. This could have induced subjects to follow more and more the predecessors’ decisions. As a matter of fact, herding almost triples when the imbalance goes from 1 to 2, but then it stabilizes at a level close to 20%. Contrarianism, instead, increases monotonically and by a substantial amount with the trade imbalance and accounts for a large percentage (40%) of all decisions when the trade imbalance is high (at least 4). Overall, our experiment seems to indicate that, with no event uncertainty,

subjects do not have a strong tendency to herd. In contrast, they do have a strong tendency to behave as contrarians.²⁶

One could wonder whether the observed deviations from the theory can be explained by the fact that a subject deciding in later periods may factor in the possibility of errors by their predecessors. In this case, his optimal trading decision may change. As is now standard in the experimental literature, we answered this question through an analysis of errors. We estimated the error rates assuming that expected payoffs are subject to shocks distributed independently as a logistic random variable (see McKelvey and Palfrey 1995). At each time t , the probability of an action is a function of the difference between the expected payoff of buying or selling the asset, Π_t , that is,

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

where $j = 0, 1, 2$ indicates a no trade, a buy, or a sell order, respectively.²⁷

The model implies that a subject may not choose the action that yields the highest payoff, that is, that he may make a mistake. For each period of trading, we estimated the parameters of the model by regressing all trading decisions up to that period across subjects and rounds on Π_t . The analysis was recursive, that is, we used the estimated parameters $\gamma_j^1, \dots, \gamma_j^{t-1}$ to compute the expected payoffs at time t . In other words, when computing a subject's expected payoff at time t , we assume that the subject incorporates in his expectations the likelihood that his predecessors have made mistakes.

By using this model, we can explain a proportion of contrarianism. In particular, when the absolute trade imbalance is equal to or higher than 4, all the contrarianism that we find in the data can be considered rational.²⁸ Contrarianism at lower levels of the absolute trade imbalance, however, remains a non-rational behavior even if one takes into account previous subjects' mistakes.²⁹

We will discuss individual behavior in detail in Section 5. Here, however, it is worth noting that there was significant heterogeneity in the decision to herd,

26. Our results on herding and contrarianism are further confirmed when one looks at the decisions to follow one of the two signals only (and not to trade conditional on the other). The figure reported in Table 1 (19.6%) results from two different types of behavior: the decision to follow the signal that agrees with the trade imbalance (e.g., the white signal after more buys than sells) and not to trade conditional on the signal at odds with it; and the decision to follow the signal that is at odds with the trade imbalance (e.g., the blue signal after more buys than sells) and not to trade conditional on the one that agrees with it. Interestingly, this latter type of behavior is more frequent (11.5%) than the former (6.6%), indicating, again, that subjects had a higher tendency to go against the market than to follow it.

27. The expected payoff of a no trade does not enter the model, because it is constant for all times t .

28. In a similar spirit, Bisière, Décamps, and Lovo (2008) show that contrarianism disappears when prices are set taking into account subjects' actual trading decisions.

29. Similarly, the modest proportion of herding remains non-national rational even taking into account the errors in the laboratory.

TABLE 3. No trade in Treatment I.

Absolute value of the trade imbalance	Cascade no-trading (%)
0	19.4
1	5.4
2	7.3
3	13.0
≥ 4	15.6

with the vast majority of subjects never herding. As a matter of fact, 24 out of the overall 39 decisions to herd for an absolute trade imbalance of at least 2 (i.e., 62% of these decisions) are due to two subjects only. If we exclude these two subjects, the percentage of herding becomes very low (only 8% of decisions taken for an absolute trade imbalance of at least 2). The results also show significant heterogeneity in the degree of contrarianism, with slightly more than half of the subjects never acting as contrarians. In contrast to herding, however, the overall proportion of contrarian decisions is not affected by the behavior of only few subjects.

Now, let us look at the decision of subjects not to participate in the market, that is, the decision not to trade independently of the signal (cascade no-trading).

Cascade no-trading occurred mainly under two circumstances: when the trade imbalance was 0 and when it was high (greater than or equal to 3) (see Table 3). A trade imbalance of 0 means that either no one has yet traded in the market or that the order flow has not taken any particular direction. In such circumstances, subjects have sometimes used the strategy of not taking a trading position, opting for trading only when the market has already taken a direction. For strictly positive levels of the absolute trade imbalance, the level of no trade is then monotonically increasing. It is worth recalling that a higher level of the trade imbalance is equivalent to a price farther away from the unconditional expected value. Therefore, a higher trade imbalance also meant that the possible loss (i.e., buying when the fundamental was 0 or selling when it was 100) was higher. The higher this potential loss, the lower was the participation in the market.

4.2. Treatment II

Let us now analyze subjects' decisions in Treatment II. Recall that the theoretical predictions for this treatment are different from those of Treatment I. In particular, in Treatment II, it is no longer the case that subjects should always follow their private information. After a given history of trades, it is possible that the optimal decision for a rational trader is to buy irrespective of the signal (herd buy periods)

TABLE 4. Average behavior in Treatment II.

Decision	(%)
Following private information	50.9
Partially following private information	20.1
Cascade trading	12.0
Cascade no-trading	16.5
Errors	0.05
Total	100.0

TABLE 5. Cascade trading behavior in Treatment II.

Absolute value of the trade imbalance	Cascade trading (%)	Herd behavior (%)	Contrarian behavior (%)
0	2.2		
1	8.2	4.4	3.8
2	23.0	18.4	4.6
3	34.3	30.3	4.0
≥ 4	40.4	40.4	0.0

or to sell irrespective of the signal (herd sell periods). Table 4 breaks down the participants' decisions in Treatment II according to how they used their own private information. In 51% of the cases, subjects followed their private signal, buying on a white signal, and selling on a blue one. Although in Treatment I the proportion of decisions in accordance with private information is also a measure of how the participants' strategies agreed with the theoretical predictions, this is no longer the case in Treatment II. For this reason, we also computed the percentage of times in which the participants' strategies agreed with the theoretical prediction, which was 48%, almost identical to that of Treatment I.³⁰ As for the other figures reported in Table 4, it is worth noting that there is slightly less cascade trading behavior than what was reported in Table 1 and slightly more cascade no-trading behavior. The strategy of following one of the two signals and not trading on the other was chosen almost the same percentage of times as in Treatment I.

The difference between the behavior in the two treatments becomes striking when one contrasts Table 5 with Table 2. In contrast with the previous treatment, contrarianism is now very modest. It does not increase at all with the trade imbalance and remains always at an almost negligible level. On the other hand, herd behavior is steadily increasing with the trade imbalance. For a trade imbalance of at least 4, herd behavior explains all cascade trading behavior; it amounts to 40%

30. In other words, this is the percentage of the time in which subjects followed the signals when theory prescribes to follow the signal and herded when the theory prescribes to do so. Of course, the same remark as in the first part of footnote 23 applies to this computation.

of all decisions taken for such levels of the imbalance. The different propensity to herd with respect to Treatment I can easily be appreciated by noting that in that treatment, even for the highest levels of the trade imbalance, herd behavior was around 20%, a relatively low increment from the 5.8% of cascade behavior when the trade imbalance was 0. In the present treatment, by contrast, cascade behavior is only 2.2% for a trade imbalance of 0 but jumps to 40.4% (all due to herding) for a trade imbalance higher than 3. Therefore, we can conclude that the theory correctly predicts the higher level of herding in this treatment with respect to the previous one.³¹

The level of herding observed in the laboratory, however, is lower than what the theory predicts. We computed the percentage of herd behavior in the periods in which herding is theoretically rational.³² The result is that herding occurred in only 23% of these periods.³³ Because this type of financial market has never been tested previously in the laboratory, we cannot compare our results to those of other studies, not even to experiments conducted with students. The closest study is the “fixed price treatment” presented in Cipriani and Guarino (2005a). In that treatment, subjects (undergraduate students) had three options, as in the present context, and the price was always set equal to the unconditional expected value of 50. Subjects engaged in herd behavior 50% of the time. The difference may well be due to the fact that here there is price movement, although it is less pronounced than in the previous treatment; this may have induced subjects to disregard the previous history of trades even in cases where doing so was not optimal. Our low level of herding, however, is also reminiscent of the results by Alevy et al. (2007), according to whom financial professionals put more weight on private information than students do and are less inclined to follow predecessors.³⁴

In summary, we can draw two conclusions on herding and contrarianism. First, whereas in Treatment I we observe a significant deviation from the theory because of contrarian behavior, this does not happen in Treatment II, where, as the theory predicts, contrarianism is not present. Second, the comparison between

31. Another significant difference with respect to Treatment I is that here, when participants followed only one signal and did not trade conditional on the other, they mainly followed the signal that agreed with the trade imbalance. In fact, this behavior accounts for 16.9% of subjects' decisions, out of the 19.6% of cases in which subjects followed only one signal (and decided not to trade for the other). This contrasts with what was observed in footnote 26 for Treatment I, namely, that when agents decided to follow only one signal, they mainly did so for the signal that did not agree with the trade imbalance.

32. The same remark as in the first part of footnote 23 applies to this computation, too.

33. Note, however, that if we take into account previous subjects' mistakes through an analysis of errors, the proportion of decisions in which traders correctly decided to herd increases to 31%.

34. As we have mentioned, Treatment II, in which the theory predicts herding because of event uncertainty, has never been run with student subjects. Nevertheless, given the difference in behavior between students and financial market professionals (see Section 5), one can speculate that, had we run the experiment with students, we might have observed a higher incidence of herd behavior, thus overestimating the importance of rational herding in the market.

TABLE 6. No trade in Treatment II.

Absolute value of the trade imbalance	Cascade no-trading (%)
0	25.3
1	16.0
2	8.0
3	4.0
≥ 4	2.1

the experimental results in the two treatments supports the theoretical prediction that informational uncertainty is a source of herding behavior. In particular, in Treatment II, herd behavior occurs, especially for high values of the trade imbalance, and occurs more often than what we observe in Treatment I. The level of herding observed in Treatment II, however, is lower than what theory predicts.

Another significant difference between the two treatments emerges when we look at the decisions not to trade. As Table 6 shows, in Treatment II, cascade no-trading is monotonically and sharply decreasing with the absolute value of the trade imbalance. Subjects decided not to participate in the market mainly for a trade imbalance of 0. To explain such a behavior it is worth recalling that in this treatment, even for a high level of the trade imbalance, the price never reached values close to the extremes (0 or 100) and, as a result, the maximum loss was never very high. A high trade imbalance revealed information on the asset value, without making the maximum loss too high. For a high value of the imbalance, when subjects wanted to use the option of not trading, they typically preferred to do so conditional on one signal only (the signal at odds with the trade imbalance) than to do so conditional on both.

5. Comparison with Previous Experimental Results

As we mentioned in the Introduction, Cipriani and Guarino (2005a) and Drehmann et al. (2005) have run experiments similar to our Treatment I with a different pool of subjects. It is useful to compare their results to ours.

Cipriani and Guarino (2005a) and Drehmann et al. (2005) reach similar conclusions: subjects have a modest propensity to herd, whereas, at the same time, there are deviations from the equilibrium predictions in terms of abstention from trading and of contrarian behavior. Our first treatment is very similar to Cipriani and Guarino's (2005a) "Flexible Price Treatment" (CG-FPT from now on) because the parameter values chosen to implement the experiment are the same. This makes the comparison with that study particularly easy. The main difference between Treatment I and CG-FPT is that, whereas we used a strategy-like method,

in CG-FPT each subject made only one decision per round, after observing the signal realization.³⁵ Therefore, comparing the statistics we have reported in the previous section with those reported in CG-FPT would not be correct.³⁶ In order to compare our experimental results with those of CG-FPT, we computed the same statistics as CG-FPT using our data set (e.g., we computed the proportion of rational decisions only considering those decisions that were actually executed, which is what we would have observed had we used the same procedures of that study).

In CG-FPT, the proportion of decisions that were rational (i.e., consistent with the theory) was 65%.³⁷ This is the same percentage that we obtain in our study. The average proportion of no trades was 22% in CG-FPT and is 24% in our experiment. Cipriani and Guarino (2005a) studied herd behavior by analyzing the subjects' decisions when they faced a trade imbalance of at least two (in absolute value) and received a signal against the imbalance. In CG-FPT subjects decided to neglect their private information and engage in herd behavior in 12% of the cases; in 42% of the cases they decided not to trade and in 46% they followed their signal. The corresponding numbers in our study are 5% for herding, 32% for no trade, and 63% for following the signal. Finally, contrarianism was studied in Cipriani and Guarino (2005a) by analyzing the case in which a subject observed a bad signal and a trade imbalance lower than or equal to -2 or a good signal and a trade imbalance greater than or equal to 2 . Using this criterion, we observed 28% of contrarianism versus 19% in CG-FPT.

It is clear from these numbers that the behavior of financial market professionals is not very dissimilar from that of the undergraduate students used in Cipriani and Guarino (2005a). The similarity of results is reassuring for previous experimental findings.³⁸ Our study confirms the low propensity to herd, and it shows an even more pronounced propensity to go against the market by financial professionals. Interestingly, it also shows that abstention from trading

35. Moreover, CG-FPT was run with paper and pencil and involved 12 subjects instead of 8. Cipriani and Guarino (2005a) also ran other treatments with slightly different procedures and found that their results are robust to these variations. Overall, their results are based on the observation of 192 subjects.

36. Indeed, the differences in procedures imply that even the definitions of rationality, herding, and contrarianism are different. For instance, we classified an action as rational when the subject made the correct decision (according to theory) conditional on both signals. In CG-FPT, instead, because subjects made a decision after observing the signal, rationality meant that the decision taken was correct given the observed signal. Clearly, the definition of rationality in this paper is stricter than that in CG-FPT.

37. For a structural approach to the analysis of rationality in the laboratory see also Cipriani and Guarino (2005b).

38. As we have mentioned, in contrast to CG-FPT, our experiment was run with a strategy-like method. The similarity between the results seems to suggest that using the strategy-like method *per se* does not induce a different behavior. In the experimental literature, other papers have reached the same conclusion (e.g., Brandts and Charness 2000; Oxoby and McLeish 2004), although there is no consensus on the issue (e.g., Güth, Huck, and Müller 2001; Brosig, Weimann, and Yang 2003).

TABLE 7. Percentage of decisions in accordance with the theoretical predictions at individual level.

Percentage of decisions in accordance with the theoretical predictions	Percentage of participants Treatment I	Percentage of participants Treatment II
0–20	18.7	25.0
21–40	21.9	12.5
41–60	31.3	18.8
61–80	18.8	34.4
81–100	9.4	9.4

remains an important deviation from the theoretical predictions, even for financial professionals.

6. Individual Behavior

In the previous section we characterized the aggregate choices of all the participants in the experiment. We now turn to discuss the degree of heterogeneity in individual behavior and its sources. Table 7 classifies individuals depending on the percentage of the time in which their decisions agreed with those predicted by the theory.

The table clearly shows that in both treatments participants behaved quite differently. For instance, in both treatments there are almost 10% of subjects who made the theoretically optimal decision more than 80% of the time; on the other hand, there are approximately 20% of subjects who made the theoretically optimal decision less than 20% of the time. It is worth studying whether such heterogenous behavior can be related to the participants' characteristics.

At the end of the experiment we collected information on the participants' age, gender, education, job tenure, and job position. Table 8 shows the results of regressing the proportion of decisions taken in accordance with theory for each participant against the participants' age, education, gender, and a dummy for traders.³⁹ Only the participants' age has a statistically significant and positive effect. The subjects' level of education, gender, and being an actual trader are not significant determinants of the level of rationality.⁴⁰

39. The variable education takes value 1 if the participant's highest degree of education is a B.A./B.S., 2 for an M.A./M.S., and 3 for a Ph.D. The dummy variable for trader takes value 1 if the participant is a trader and 0 otherwise.

40. We also used the job tenure as a regressor instead of age, and obtained similar results. Unfortunately, our data set does not allow us to disentangle which of these two (collinear) variables has an effect on rationality. If we include both age and job tenure as regressors, both coefficients become non-significant.

TABLE 8. Regressions of the level of rationality in the experiment on individual characteristics.

	(1)	(2)	(3)	(4)	(5)
Gender	0.039 (0.704)				0.019 (0.877)
Age		0.014 (0.013)			0.014 (0.028)
Degree			-0.134 (0.201)		-0.139 (0.215)
Trader				-0.114 (0.281)	-0.104 (0.386)
Observations	32	32	32	32	32
R^2	0.00	0.09	0.129	0.053	0.279

p-values are in parentheses.

Participants showed heterogeneity also in the specific trading strategies discussed in the previous sections (i.e., propensity to herd and act as contrarians). For instance, as already mentioned before, in Treatment I only very few subjects engaged often in herd behavior, whereas many never did.⁴¹ Through regression analysis, we have studied whether the participants' propensity to herd or act as a contrarian is affected by their personal characteristics. None of the characteristics that we analyzed had a significant impact, except gender: Women made significantly fewer contrarian decisions in the first treatment and more herd decisions in the second treatment.⁴²

Table 9 shows the relationship between a subject's payoff and his personal characteristics. Traders earned significantly more than the other participants. No other characteristics significantly affected the subjects' payoffs. The significantly higher payoff of traders was due to higher earnings in Treatment I, whereas no significant difference emerged in Treatment II.⁴³ It is, however, difficult to gauge from the data how traders achieved higher payoffs. Indeed, as Table 10 shows, being a professional trader did not change the tendency to act as a herder, or as a contrarian, or to abstain from trading or to behave rationally. It appears that professional traders had an ability to earn more money than the other participants, even though, with respect to herding, contrarianism, and no trading, their trading strategies do not look different.

41. For an absolute trade imbalance of at least 2, 66% of participants never herded.

42. In the interest of space, we do not report the regression results, because most of the coefficients are not significant.

43. The *p*-values of the per-treatment regressions (which, in the interest of space, we do not report) are 0.06 and 0.23, respectively.

TABLE 9. Regression of subjects' payoff at the end of the experiment on individual characteristics.

	(1)	(2)	(3)	(4)	(5)
Gender	-30.302 (0.640)				3.757 (0.962)
Age		3.779 (0.236)			4.142 (0.313)
Degree			-8.944 (0.759)		-10.230 (0.670)
Trader				67.403 (0.050)	69.831 (0.107)
Observations	32	32	32	32	32
R^2	0.00	0.02	0.001	0.054	0.080

p-values are in parentheses.

TABLE 10. Regressions of participants' proportion of herding, contrarianism, and no-trading on the trader's dummy. Herd 1 and Contrarian 1 refer to Treatment I.

	(1) Herd 1	(2) Herd 2	(3) Contr 1	(4) Contr 2	(5) No Trade
Trader	0.292 (0.204)	0.268 (0.206)	0.131 (0.199)	0.026 (0.744)	0.018 (0.855)
Observations	32	32	32	32	32
R^2	0.147	0.121	0.021	0.007	0.002

Herd 2 and Contrarian 2 refer to Treatment II. *p*-values are in parentheses.

7. Conclusions

In this paper we have analyzed cascade and herd behavior in a laboratory financial market. The main novelty of our experiment is the use of a sample of financial market professionals instead of college undergraduates. Moreover, with respect to the existing literature, we have introduced a new experimental design, akin to the strategy method. In the experiment, subjects make conditional orders, contingent on any signal realization, instead of choosing an action after observing a realization. With this method we are able to elicit beliefs and detect cascade behavior directly (whereas in previous experiments it could only be inferred indirectly).

We have run two treatments. In Treatment I, the price adjusts to the order flow in such a way that subjects should simply follow their own private information. In Treatment II, where the price adjustment rule is consistent with the presence of event uncertainty, it is sometimes rational for subjects to neglect their private information and imitate the predecessors. This is the first study where herding in financial markets with event uncertainty is tested in the laboratory.

In the first treatment, with no event uncertainty, agents rarely herded, which is consistent with the theory. At the same time, two anomalies emerged. First, subjects had a tendency to go against the market (contrarian behavior). Second,

they sometimes preferred to abstain from trading although they had an informational advantage over the market maker. The results of this first treatment are in line with previous experimental evidence on student subjects.

In the second treatment, with event uncertainty, the proportion of herd behavior increased with respect to the first treatment, as predicted by the theory. At the same time, such a proportion was lower than the equilibrium prediction. Contrarian behavior disappeared altogether in this treatment, whereas the decision not to trade was observed often, as in the case of no event uncertainty.

Our study combines the advantage of the controlled experiment with that of observing the behavior of professionals, who are engaged in the day-by-day activity of trading, pricing, and analyzing financial assets. We believe that the challenge for future research is twofold. On the one hand, the existing experimental results offer suggestions for research with field data, which should study whether the behaviors observed in the laboratory are present in actual financial markets. On the other hand, more theoretical work is needed to capture the behavior that the present model is unable to predict, such as contrarianism and abstention from trading activity.

References

- Alevy, Jonathan, Michael Haigh, and John List (2007). "Information Cascades: Evidence from a Field Experiment with Financial Market Professionals." *Journal of Finance*, 62, 151–180.
- Anderson, Lisa, and Charles Holt (1997). "Information Cascades in the Laboratory." *American Economic Review*, 87, 847–862.
- Avery, Christopher, and Peter Zemsky (1998). "Multidimensional Uncertainty and Herd Behavior in Financial Markets." *American Economic Review*, 88, 724–748.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992). "A Theory of Fads, Fashion, Custom and Cultural Change As Informational Cascades." *Journal of Political Economy*, 100, 992–1027.
- Bikhchandani, Sushil, and Sunil Sharma (2001). "Herd Behavior in Financial Markets." *IMF Staff Papers*, 47, 279–310.
- Bisière, Christophe, Jean-Paul Décamps, and Stefano Lovo (2008). "On the Information Content of the Order Flow: An Experiment." Working paper, Toulouse School of Economics and HEC-Paris.
- Brandts, Jordi, and Gary Charness (2000). "Hot vs. Cold: Sequential Responses and Preference Stability in Experimental Games." *Experimental Economics*, 2, 227–238.
- Brosig, Jeannette, Joachim Weimann, and Chun-Lei Yang (2003). "The Hot Versus Cold Effect in a Simple Bargaining Experiment." *Experimental Economics*, 6, 75–90.
- Çelen, Boğaçhan, and Shachar Kariv (2004). "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review*, 94, 484–497.
- Chamley, Christophe (2004). *Rational Herds*. Cambridge University Press.
- Cipriani, Marco, and Antonio Guarino (2005a). "Herd Behavior in a Laboratory Financial Market." *American Economic Review*, 95, 1427–1443.
- Cipriani, Marco, and Antonio Guarino (2005b). "Noise Trading in a Laboratory Financial Market: A Maximum Likelihood Approach." *Journal of the European Economic Association*, 3, 315–321.

- Cipriani, Marco, and Antonio Guarino (2006). "Estimating a Structural Model of Herd Behavior in Financial Markets." Working paper, GWU and UCL.
- Cipriani, Marco, and Antonio Guarino (2008a). "Herd Behavior and Contagion in Financial Markets." *The B.E. Journal of Theoretical Economics (Contributions)*, 8(1), article no. 24.
- Cipriani, Marco, and Antonio Guarino (2008b). "Transaction Costs and Informational Cascades in Financial Markets." *Journal of Economic Behavior & Organization*, 68(3–4), 581–592.
- Dasgupta, Amil, and Andrea Prat (forthcoming). "Information Aggregation in Financial Markets with Career Concerns." *Journal of Economic Theory*.
- Drehmann, Mathias, Jörg Oechssler, and Andreas Rider (2005). "Herding and Contrarian Behavior in Financial Markets—An Internet Experiment." *American Economic Review*, 95, 1403–1425.
- Easley, David, and Maureen O'Hara (1987). "Price, Trade Size, and Information in Securities Markets." *Journal of Financial Economics*, 19, 69–90.
- Fischbacher, Urs (2007). "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics*, 10, 171–178.
- Gale, Douglas (1996). "What Have We Learned from Social Learning?" *European Economic Review*, 40, 617–628.
- Glosten, Lawrence, and Paul Milgrom (1985). "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics*, 14, 71–100.
- Goeree, Jacob, Tom Palfrey, Brian Rogers, and Richard McKelvey (2007). "Self-Correcting Information Cascades." *Review of Economic Studies*, 74, 733–762.
- Grinblatt, Mark, Sheridan Titman, and Russ Wermers (1995). "Momentum Investment Strategies, Portfolio Performance, and Herding: A Study of Mutual Fund Behavior." *American Economic Review*, 85, 1088–1105.
- Güth, Werner, Steffen Huck, and Wieland Müller (2001). "The Relevance of Equal Splits in Ultimatum Games." *Games and Economic Behavior*, 37, 161–169.
- Hirshleifer, David, and Siew H. Teoh (2003). "Herd Behaviour and Cascading in Capital Markets: A Review and Synthesis." *European Financial Management*, 9, 25–66.
- Huck, Steffen, and Jörg Oechssler (2000). "Informational Cascades in the Laboratory: Do They Occur for the Right Reasons?" *Journal of Economic Psychology*, 21, 661–671.
- Kübler, Dorothea, and Georg Weizsäcker (2004). "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory." *Review of Economic Studies*, 71, 425–441.
- Lakonishok, Joseph, Andrei Shleifer, and Robert W. Vishny (1992). "The Impact of Institutional Trading on Stock Prices." *Journal of Financial Economics*, 32, 23–43.
- McKelvey, Richard, and Thomas Palfrey (1995). "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior*, 10, 6–38.
- Oxoby, Robert, and Kendra McLeish (2004). "Sequential Decision and Strategy Vector Methods in Ultimatum Bargaining: Evidence on the Strength of Other-regarding Behavior." *Economics Letters*, 84, 399–405.
- Park, Andreas, and Hamid Sabourian (2006). "Herd Behavior in Efficient Financial Markets." Working paper, University of Cambridge and University of Toronto.
- Sias, Richard (2004). "Institutional Herding." *Review of Financial Studies*, 17, 165–206.
- Smith, Lones, and Peter Sörensen (2000). "Pathological Outcomes of Observational Learning." *Econometrica*, 68, 371–398.
- Vives, Xavier (2008). *Information and Learning in Markets*. Princeton University Press.
- Wermers, Russ (1999). "Mutual Fund Herding and the Impact on Stock Prices." *Journal of Finance*, 54, 581–622.