Herd Behavior in a Laboratory Financial Market

By MARCO CIPRIANI AND ANTONIO GUARINO*

We study herd behavior in a laboratory financial market. Subjects receive private information on the fundamental value of an asset and trade it in sequence with a market maker. The market maker updates the asset price according to the history of trades. Theory predicts that agents should never herd. Our experimental results are in line with this prediction. Nevertheless, we observe a phenomenon not accounted for by the theory. In some cases, subjects decide not to use their private information and choose not to trade. In other cases, they ignore their private information to trade against the market (contrarian behavior). (*JEL* C92, D8, G14)

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 may be a reason for excess price volatility and financial systems fragility.

The theoretical research on herd behavior starts with the seminal papers by Abhijit Banerjee (1992), Sushil Bikhchandani et al. (1992), and Ivo 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. This is an important result, because it provides a rationale for the imitating behavior that we observe in consumers' and investors' decisions. In these first models of herding, however, the cost of taking an action (e.g., investing in a new project) 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, Christopher Avery and Peter Zemsky (1998) have studied herd behavior in a financial market where the price is efficiently set by a market maker according to the order flow. They show that the presence of an efficient price mechanism makes an informational cascade (i.e., a situation in which an agent does not use his own information and herds) impossible. Agents always find it optimal to trade on the difference between their own information (the history of trades and the private signal) and the commonly available information (the history of trades only). For this reason, the price aggregates the information contained in the history of past trades correctly.[2]

It is difficult to test these theoretical models of herding empirically. The existing literature

* Cipriani: Department of Economics, George Washington University, Old Main, 1922 F Street, NW, Washington, DC 20052 (e-mail: marco.cipriani@gwu.edu); Guarino: Department of Economics and ELSE, University College London, Gower Street, London, WC1E 6BT, UK (e-mail: a.guarino@ucl.ac.uk). We are indebted to Andrew Schotter for invaluable guidance. We thank Franklin Allen, Alberto Bisin, William Baumol, Ken Binmore, Andrew Caplin, Douglas Gale, Philipp Hartmann, Steffen Huck, Graciela Kaminsky, Marco Ottaviani, H. O. Stekler, Joseph Swierzbinski, and Antonella Trigari for useful comments. The revision of the paper owes very much to the comments of the co-editor and two anonymous referees. We also thank seminar participants at several institutions. 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, the ELSE Center at University College London, the ESRC, and the Leverhulme Trust. All errors are ours.

[1] In this paper we study only informational herding. We do not discuss herd behavior arising because of reputational concerns, as in David Scharfstein and Jeremy Stein (1990), or payoff externalities.

[2] For other theoretical contributions on informational herding in financial markets, see In Ho Lee (1998), Cipriani and Guarino (2001), and V. V. Chari and Patrick Kehoe (2004). For recent surveys of herding in financial markets, see David Hirshleifer and Siew H. Teoh (2003) and Christophe Chamley (2004).

(see, e.g., Joseph Lakonishok et al., 1992; Mark Grinblatt et al., 1995; Russ Wermers, 2000; and the other papers cited in the survey of Hirshleifer and Teoh, 2003) does not test these models directly, but analyzes only the presence of herding in financial markets through statistical measures of clustering. This literature finds that fund managers tend to cluster their investment decisions. Such clustering, however, may or may not be due to informational herding: for instance, it may be the result of a common reaction to public announcements. The problem for the empirical research on herd behavior 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. In an experiment, we can observe variables not available for actual markets, in particular, the private information that agents have when making their decisions. In our laboratory market, subjects receive private information on the value of a security and observe the history of past trades. Given these two pieces of information, they choose, sequentially, if they want to sell, to buy, or not to trade one unit of the asset. By observing the way in which they use their private information and react to the decisions of the previous traders, we can detect the occurrence of herding.[3] By testing directly the prediction of the theoretical work, we create a bridge between the existing empirical and theoretical literatures.

Our results on herd behavior are 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 flexible. We implement the flexible-price case in two ways: in one the price is updated according to a deterministic rule based on the order flow, and in the other it is set by experimental participants. We find that, with either price-updating mechanism, when the price is flexible, subjects disregard their private information and herd much less frequently than when the price is held constant.

Herd behavior prevents the price from aggregating private information dispersed across market participants and from converging to the fundamental value. In our laboratory market, herding rarely occurs and is not a serious source of informational inefficiency and price misalignment. Early trades do not have the disproportionate effect on later decisions implied by the first models of herding, and do not affect the ability of the market price to aggregate private information. Our results show that herd behavior should not be a concern in a financial market in which informed traders trade for informational reasons only.

Although the theory is able to predict the effect of a flexible price on herd behavior, it is unable to account for two phenomena that we observe in the laboratory. First, when the price is flexible, subjects sometimes choose not to trade. Second, sometimes (although less frequently) they choose to trade against their private information. In some cases, they do so because they engage in "contrarian behavior," i.e., they buy when the price is low or sell when it is high. The occurrence of contrarian behavior and of no-trade decisions reduces the ability of the price to aggregate private information dispersed across market participants. In particular, the relatively high incidence of no trades suggests that limited market participation may be an important source of financial markets' informational inefficiency.

The structure of the paper is as follows. Section I describes the theoretical model and its predictions. Section II presents the experimental design. Section III illustrates the results of the first three treatments. Section IV discusses the fourth treatment with endogenous price setting. Section V concludes.

I. The Theoretical Model

A. *The Model Structure*

Our experimental analysis is based on the model by Lawrence Glosten and Paul Milgrom (1985). In our economy, 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 dates indexed by $t = 1, 2, 3, \ldots$

The Market.—The fundamental value of the asset, V, is a random variable distributed on $\{0,$

[3] For previous experimental analyses of herd behavior based on the Banerjee (1992) and Bikchandani et al. (1992) models, see, e.g., Lisa Anderson and Charles Holt (1997), Steffen Huck and Jörg Oechssler (2000), and Dorothea Kübler and Georg Weizsäcker (2004).

100} with the same probability ½. At each time t, a trader can exchange the asset with a specialist (market maker). The trader can buy, sell, or decide not to trade. Each trade consists of the exchange of one unit of the asset for cash. The trader's action space is, therefore, $\mathcal{A} = \{buy, sell, no\ trade\}$. We denote the action of the trader at time t by $h_t \in \mathcal{A}$. Moreover, we denote the history of trades and prices until time $t - 1$ by H_t.

The Market Maker.—At any time t, the market maker sets the price at which a trader can buy or sell the asset. The market maker is allowed to set one price only (i.e., we do not allow for a bid-ask spread). We consider two setups, a fixed-price setup and a flexible-price setup. In the first setup, the market maker sets the price equal to the asset's unconditional expected value, i.e.,

$$P_t = E(V) = 50 \qquad \text{for all } t.$$

In the second setup, the market maker sets the price equal to the expected value, conditional on the information available at time t, i.e.,[4]

$$P_t = E(V|H_t).$$

The Traders.—There is a countably infinite number of traders. Traders act in an exogenously determined sequential order. Each trader, indexed by t, is chosen to take an action only once, at time t. Traders have private information on the asset value.[5] If at time t a trader is chosen to trade, he observes a private signal on the realization of V. The signal is a random variable X_t distributed on {0, 100}. We denote the conditional probability function of X_t, given a realization of V by $q(x_t|v)$ where x_t is a realization of X_t and v is a realization of V. We assume that the random variables X_t are independently and identically distributed across time. In particular, we assume that

$$q(0|0) = q(100|100) = 0.7.$$

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

Traders are endowed with an amount $K > 0$ of cash. Their payoff function $U : \{0, 100\} \times \mathcal{A} \times [0, 100] \rightarrow \mathbb{R}^+$ is defined as

$$U(v, h_t, P_t) = \begin{cases} v - P_t + K & \text{if } h_t = buy, \\ K & \text{if } h_t = no\ trade, \\ P_t - v + K & \text{if } h_t = sell. \end{cases}$$

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

B. *Prediction when the Price Is Fixed*

In order to study the theoretical predictions of our model, we need to introduce the concepts of trade imbalance and informational cascade. For any trading sequence, we define the trade imbalance at time t as the number of buy orders minus the number of sell orders until time $t - 1$. If a no trade can be the outcome of a rational decision and reveals a trader's private signal, we also take it into account in the computation of the trade imbalance.[6] Moreover, following

[4] 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. We avoid the presence of two prices (the bid and the ask) and assume that the market maker sets only one price equal to the expected value of the asset. The presence of only one price makes our experiment easier to run.

[5] In the original Glosten and Milgrom (1985) model, a proportion of traders are uninformed and trade for exogenous reasons, without regard for their profits. The presence of these noise traders is necessary for the market not to break down. 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 Nancy Stokey, 1982). In our setup, for simplicity, all traders are informed and we assume that the market maker is willing to trade, even if he makes negative profits in expected value.

[6] When the price is fixed, the decision of no trading can be rational. For instance, it is rational not to trade when there has been a buy order at time 1 and, at time 2, an agent receives a negative signal. From the first buy order, the agent can induce that the first agent had a positive signal. Therefore, his expected value of the asset, given one positive and one negative signal, is 50, equal to the asset price.

In the computation of the trade imbalance, a rational no trade is considered as a sell order if it reveals a negative signal, and as a buy order if it reveals a positive one.

Anderson and Holt (1997), we define an informational cascade as a situation in which it is optimal for a rational agent to ignore his own private information and conform to the established pattern of trade. During a cascade, agents herd, i.e., they all choose the same action.[7]

When the price is fixed, as in the seminal papers on herd behavior and information cascades, the following result holds:[8]

RESULT 1 (Bikhchandani et al., 1992): *When the asset price is fixed at the unconditional expected value* $E(V) = 50$, *an informational cascade occurs after a trade imbalance higher than or equal to* 2, *or lower than or equal to* -2.

To understand the intuition behind this result, consider the following example. Suppose that at time 3 there is a trade imbalance of 2, i.e., $H_3 = \{buy, buy\}$. Suppose also that the third trader receives the signal $x_3 = 0$. From the first two buy orders, he can infer that x_1 and x_2 equal 100. Therefore, by Bayes's rule, his expected value of the asset is 70. Given that the price is 50, he will ignore his signal and buy. This starts a cascade.

C. *Prediction when the Price Is Flexible*

Let us now discuss the case in which the price is flexible. The market maker updates the price in a Bayesian way on the basis of the order flow. In this setup, informational cascades cannot arise.

RESULT 2 (Avery and Zemsky, 1998): *When the market maker sets the price* P_t *equal to* $E(V|H_t)$, *agents always trade according to their private signal. An informational cascade cannot occur.*

To decide whether he wants to buy or sell the asset, an agent computes his expected value and compares it to the price. If at time t a trader receives a signal of 100, his expected value will be

$$\begin{aligned} &E(V|H_t, x_t = 100) \\ &\quad = 100 \Pr(V = 100|H_t, x_t = 100) \\ &\quad = 100 \frac{(.7)\Pr(V = 100|H_t)}{(.7)\Pr(V = 100|H_t) + (.3)(1 - \Pr(V = 100|H_t))} \\ &\quad > 100 \Pr(V = 100|H_t) = E(V|H_t). \end{aligned}$$

Similarly, if he receives a signal of 0, his expected value will be $E(V|H_t, x_t = 0) < E(V|H_t)$. This shows that an agent will always find it optimal to trade according to his private information, and an informational cascade cannot arise.

Note that, when a trader has the opportunity to trade at a certain price, knowing the history of trades does not give him additional information on the asset value beyond what is already contained in the asset price. Therefore, a rational agent should act according to his private signal, irrespective of whether he is able to observe the history of trades or not.

II. The Experiment and the Experimental Design

A. *The Experiment*

This was a paper and pencil experiment. We recruited subjects from undergraduate courses in all disciplines at New York University and University College London. They had no previous experience with this experiment. In total, we recruited 216 students to run 16 sessions (four sessions for each treatment).[9] We now describe the procedure for the first three treatments and postpone the discussion of the fourth to Section IV. In each session of these three treatments, we used 13 participants, one acting

[7] In the literature, it has been pointed out (see Lones Smith and Peter Sørensen, 2000) that rational herding and informational cascades are not identical concepts. For an experimental analysis of the difference between herd behavior and informational cascades, see Boğaçhan Çelen and Shachar Kariv (2004). In our setup with discrete signal space, however, an informational cascade implies rational herding, and vice versa. In the following pages, herding will indicate conformity of behavior, which can be rational (as in a cascade) or irrational.

[8] We assume that if an agent is indifferent he follows his private information. The result would also hold if we assumed that the agent randomizes among the choices.

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

as subject administrator and 12 acting as traders. The procedure was the following:

1. At the beginning of the sessions, we gave written instructions (available from the authors on request) to all subjects. We read the instructions aloud in an attempt to make the structure of the game common knowledge to all subjects. Then, we asked for clarifying questions, which we answered privately.
2. Each session consisted of ten rounds. In each round, we asked all subjects to trade one after the other.
3. The sequence of traders for each round was chosen randomly. At the beginning of the session, each subject picked a card from a deck of 13 numbered cards. The number that a subject picked was assigned to him for the entire session. The card number 0 indicated the subject administrator. In each round, the subject administrator called the subjects in sequence by randomly drawing cards (without replacement) from this same deck.
4. Before each round, an experimenter, outside the room, tossed a coin: if the coin landed tails, the value of the asset for that round was 100; otherwise it was 0. Traders were not told the outcome of the coin flip. During the round, the same experimenter stayed outside the room with two bags, one containing 30 blue and 70 white chips and the other 30 white and 70 blue chips. The two bags were identical. Each subject, after his number was called, had to go outside the room and draw a chip from one bag. If the coin landed tails, the experimenter used the first bag; otherwise he used the second. Therefore, the chip color was a signal for the value of the asset. After looking at the color, the subject put the chip back into the bag. Note that the subject could not reveal the chip color to anyone.
5. In the room, another experimenter acted as market maker, setting the price at which people could trade. After observing the chip color, the subject entered the room. He read the trading price on the blackboard and then declared aloud whether he wanted to buy, to sell, or not to trade. The subject administrator recorded all subjects' decisions and all trading prices on the blackboard.[10] Hence, each subject knew not only his own signal, but also the history of trades and prices.[11]
6. At the end of each round, i.e., after all 12 participants had traded once, the realization of the asset value was revealed and subjects were asked to compute their payoffs. All values were in a fictitious currency called lira. Their payoffs were computed as follows. In the event of a buy, the subject obtained $100 + v - P_t$ lire; in the event of a sell, he obtained $100 + P_t - v$ lire; finally, if he decided not to trade, he earned 100 lire. This is equivalent to giving each subject 100 lire each round, which he could use 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, subjects could never lose money.
7. After the tenth round, we summed up the per-round payoffs and converted them into dollars at the rate of 1⁄65. In addition, we gave $7 to subjects just for participating in the experiment. Subjects were paid in private immediately after the experiment and, on average, earned $25 for a 1.5-hour experiment.

B. *The Experimental Design*

Let us now describe the differences among these first three treatments. In the 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 a trade imbalance of two, an informational cascade should arise, i.e., subjects should buy despite a negative signal or sell despite a positive signal.

In the second treatment ("flexible-price" treatment), the price was updated after each trade decision in a Bayesian fashion. Rational subjects should always follow their signal, i.e., they should buy after seeing a positive signal and sell after seeing a negative one. No one

[10] This is true for all but one treatment, the "no-history" treatment, in which the history of trades and prices was not made public. We discuss this treatment in the next section.

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

should decide not to trade, as private information allows the traders to make money by trading with the market maker. Therefore, in this treatment, when a subject decided to buy, the price was updated assuming that he had seen a positive signal. Similarly, when a subject decided to sell, the price was updated assuming the subject had observed a negative signal. Finally, in the case of a no trade, the price was kept constant.

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

The third treatment ("no-history" treatment) was a control treatment: in order to understand the effect of history on the behavior of subjects, we ran an experiment where subjects could not observe the decisions of those who traded before them. When a subject had to make a decision, he could see the trading price on a piece of paper, but did not have the history of trades and of prices written on the blackboard. Although we did not want the subjects to know the past prices and decisions, we wanted them to know the mechanism of price formation. In order to make sure that the subjects understood this mechanism, not only did we describe 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. Starting with the fourth round, subjects were not allowed to see the history of trades and prices.

In the next section, we describe the results of these three treatments. The results refer to the last seven rounds of each session only.[12] We do not take the first three rounds into account for two reasons. First, although the experiment was very easy and subjects did not have problems in understanding the instructions, we believe that some rounds were needed to acquaint subjects with the procedures. We wanted to distinguish the decisions that subjects made in the learning stage from the decisions taken afterward. Second, considering only the last seven rounds helps to make the results comparable across the different treatments. (As explained above, in the no-history design, we did not allow the subjects to observe the history of trades and prices starting with the fourth round.)

III. Results

A. *Informational Cascades and Contrarian Behavior*

We start the presentation of our results by discussing informational cascades. Let us consider, first, the fixed-price treatment. In this case, theory predicts that an informational cascade occurs whenever a trade imbalance of at least two (in absolute value) arises. In this treatment, there were 58 periods of potential informational cascade, i.e., periods when the trade imbalance was at least 2 or not higher than -2 and, moreover, the subject received a signal that was against the trade imbalance.[13] In these periods, subjects engaged in cascade behavior in 52 percent of cases; in 26 percent of cases they decided not to trade, and in 22 percent of cases they decided to follow their signal.[14]

[12] In each round, the 12 subjects were asked to trade in sequence. Therefore, the results for each treatment refer to 336 decisions.

[13] Recall that in each round there were 12 decision-making opportunities. We refer to each decision-making opportunity as a "period."

[14] An important issue in the computation of the trade imbalance is how to handle the role of deviators. For instance, consider the fixed-price treatment, and suppose that there have been four buy orders. Suppose that the next subject decides to sell. In this case, his action is certainly irrational, since, no matter what his signal is, he should buy. 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 in Anderson and Holt, 1997) that, although irrational, this person must have received a negative signal; otherwise he would not have had 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 additional information) and now is destroyed by the sell order. The trade imbalance would go from four to one. Now consider the flexible-price treatment. In this treatment, informational cascades never happen; therefore, the trade imbalance should clearly be computed

TABLE 1—INFORMATIONAL CASCADES AND HERD-LIKE BEHAVIOR

	Fixed price	Flexible price	No history	Endogenous pricing
Trading against the signal	52 percent	12 percent	24 percent	21 percent
No trading	26 percent	42 percent	33 percent	34 percent
Trading following the signal	22 percent	46 percent	43 percent	45 percent
Relevant periods	58	66	70	58

Note: The relevant periods are those when the trade imbalance was at least 2 (or at most -2) and subjects received a negative (positive) signal.

What happens when, as in the flexible-price treatment, we allow the price to react to the order flow? Do subjects still neglect their signal? Table 1 shows the results of this treatment and contrasts them with those of the fixed-price treatment. In the flexible-price case there were 66 periods in which the trade imbalance was at least two (or at most -2) and the subject received a signal against it. In these periods, subjects decided to neglect their private information and engage in irrational herd-like behavior only in 12 percent of cases. In 42 percent of cases they decided not to trade, and in 46 percent they followed their signal even if it was at odds with the history of trades. These results show that subjects rarely decided to follow what other subjects had done. The experimental evidence supports the theoretical prediction that, with flexible prices, herd behavior should not arise. As theory suggests, the price movement reduced the incentive to imitate previous decisions. We ran a Mann-Whitney test for the hypothesis that the proportion of herding decisions was the same under the fixed-price and flexible-price treatments, and the null was rejected at the 5-percent significance level.[15]

It is also interesting to note that, in the fixed-price treatment, when the trade imbalance was zero (and, therefore, there was no scope for imitation), only 5 percent of decisions were against private information. This percentage climbed to 52 in periods when the absolute value of the trade imbalance was at least 2 and subjects received a signal against it. In contrast, in the flexible-price treatment, when the trade imbalance was 0, the percentage of decisions against private information was 10. This number barely increased (to 12 percent), when the absolute level of the trade imbalance was equal to or higher than 2 and subjects received a signal against it. Therefore, in this treatment, when there was scope for imitation, subjects did not disregard private information more often than when there was not. This suggests that imitation was not an important factor in determining the subjects' behavior.

To understand better the effect of past trades on subjects' decisions, let us consider the no-history treatment, in which subjects could not observe previous decisions. If, as theory suggests, in the flexible-price treatment the choices of the predecessors affected a subject's decision only through the prices, the results of the flexible-price and no-history treatments should be identical. If, in contrast, there were other effects beyond those reflected in the prices, the results of the two treatments could be different. In particular, if agents had an irrational taste for conformity, one should expect the percentage of irrational herding to be much higher when

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 measure of trade imbalance for the two treatments. We have decided to compute the trade imbalance in both treatments considering all past actions, irrespective of whether they could come from a rational decision. Our choice does not affect the results in a significant manner, since in the fixed-price treatment we observed very few deviations from a cascade. If we compute the trade imbalance by assuming that a deviator breaks the cascade (as in Anderson and Holt, 1997), in the fixed-price treatment, cascade behavior arose in 54 percent of cases, no trade in 24 percent, and following the signal in 22 percent.

[15] The test is carried out by taking the proportion of herding decisions in each of the four sessions and using the Mann-Whitney procedure.

subjects do observe previous decisions. This is not what happened (see Table 1). In the no-history treatment, out of the 70 periods in which the trade imbalance was at least 2 (or at most −2) and the subject received a signal against it, subjects irrationally traded against the signal in 24 percent of cases, versus the 12 percent of the flexible-price treatment.[16] As in the previous treatment, some decisions (33 percent of cases) were no trades.

Overall, our results show that, with flexible prices, subjects rarely decided to imitate their predecessors and, as a result, herding was not a significant source of informational inefficiency. In contrast, informational cascades arose much more frequently when the price was held constant. The behavior observed in the laboratory is in line with the theoretical predictions of Bikhchandani et al. (1992) and Avery and Zemsky (1998): informational cascades occur in a setup with fixed prices, but not in one with flexible prices.

So far, we have focused on subjects' behavior when they had private information at odds with that conveyed by the history of trades. It is also interesting to see how subjects acted on average during the experiment. Table 2 reports the proportion of decisions that were rational, i.e., consistent with the theory. In order to classify a decision as rational or irrational, we need to make some assumptions on each subject's belief on the choices of his predecessors. From now on, we assume that each subject believes that all his predecessors are rational, that all his predecessors believe that their predecessors are rational, and so on.[17] Under this assumption, a rational subject should always behave as predicted by the theoretical model.[18] In Section III C, we will discuss the possibility that subjects could hold different beliefs on their predecessors' rationality.

TABLE 2—RATIONAL AND IRRATIONAL DECISIONS

Fixed price	
Rational decisions	84 percent
No trading	3 percent
Buying or selling	81 percent
Irrational decisions	16 percent
No trading	11 percent
Buying or selling	5 percent
Flexible price	
Rational decisions	65 percent
Irrational decisions	35 percent
No trading	22 percent
Buying or selling	13 percent
No history	
Rational decisions	61 percent
Irrational decisions	39 percent
No trading	25 percent
Buying or selling	14 percent
Endogenous pricing	
Rational decisions	57 percent
Irrational decisions	43 percent
No trading	24 percent
Buying or selling	19 percent

Rational decisions amounted to 84 percent of the total in the fixed-price treatment and to 65 percent in the flexible-price treatment. Therefore, although the price mechanism seems able to discourage herding, it also seems to reduce the overall rationality of subjects' behavior. In particular, there are more irrational no trades (22 percent versus 11 percent) and there is a higher proportion of irrational buy and sell orders.[19] In the remain-

[16] Running a Mann-Whitney test, we cannot reject the hypothesis that the proportion of herds in the flexible and no-history treatments is the same (p-value = 0.19). In contrast, we can reject the hypothesis that the proportion of herds is the same in the no-history and fixed-price treatments.

[17] We do not need rationality to be common knowledge, since the behavior of a subject is unaffected by his belief in the rationality of subjects trading after him.

[18] In the experiment, sometimes subjects made decisions off the equilibrium path, i.e., decisions that could not be the outcome of a rational choice. An important issue is how to update subjects' beliefs after they observe such decisions. If the decision is a no trade, we assume that the following subjects do not update their beliefs (which is consistent with our price-updating rule). If the decision is a buy (sell), we assume that the following subjects update their beliefs as though it publicly revealed a positive (negative) signal (i.e., the signal that implies the lower expected loss). Note that this last assumption (which follows Anderson and Holt, 1997) is relevant only for the fixed-price treatment and is, in fact, quite innocuous. For instance, if we assumed that trades that cannot be the outcome of a rational choice do not convey any information, all the results in the fixed-price treatment would be almost identical (in particular, the level of rationality would be 81 percent).

[19] One may wonder whether these results depend on the behavior of some particular subjects or reflect homogeneous behavior. For instance, is the level of irrationality due to some people who behaved irrationally most of the time? Remember that each subject made ten decisions and we focused on the last seven. For each subject, we computed the number of times in which he acted rationally. In the flexible-price treatment, out of 48 participants, only one subject acted rationally fewer than three times. Moreover, only one subject decided not to trade more than three times.

TABLE 3—CONTRARIAN BEHAVIOR

	Fixed price	Flexible price	No history	Endogenous pricing
Contrarian behavior	0 percent	19 percent	10 percent	24 percent
No trading	1 percent	18 percent	25 percent	23 percent
Following private information	99 percent	63 percent	65 percent	53 percent
Relevant periods	121	132	134	116

Note: The relevant periods are those when the trade imbalance was at least 2 (or at most −2) and subjects received a positive (negative) signal.

der of this section, we discuss these two phenomena and try to offer possible explanations.

Let us start from the analysis of no-trade decisions. In the flexible-price treatment, the frequency of no trades increases with the absolute value of the trade imbalance: it is 16 percent for an absolute value of the trade imbalance of 0 or 1, 22 percent for an absolute value of 2 or 3, and 33 percent for an absolute value higher than 3. In other words, we observed more no-trade decisions when the price was closer to 0 or 100. Given the payoffs, it is difficult to reconcile this behavior with plausible levels of risk aversion.[20] A possible alternative explanation is that subjects preferred not to trade mostly when the trade imbalance was high because they faced a high maximum loss.[21]

As we commented above, in the flexible-price treatment, a percentage of trading decisions was taken against private information. To understand this result, we considered the possibility that subjects decided to trade against the signal for reasons other than imitation. We studied another form of irrational behavior (which we call contrarian behavior) in which a subject neglects his signal in order to buy at a low price or sell at a high price. In particular, we say that a subject is a contrarian when he buys, despite a negative signal, at a low price, i.e., at a price lower than 30, or sells, despite a positive signal, at a high price, i.e., a price higher than 70. Equivalently, we can say that a subject acts as a contrarian when he buys with a negative signal and there is a trade imbalance lower than or equal to −2, or sells with a positive signal and the trade imbalance is higher than or equal to 2. The definition of contrarian behavior is meant to capture the behavior of people who use the strategy of "going against the market." Contrarians disregard their positive signal to take advantage of the high price in the market, or disregard their negative signal to buy at a low price.

In the flexible-price treatment, there were 132 periods in which a subject could have acted as a contrarian. Of these 132 times, subjects behaved as contrarians in 19 percent of cases, whereas in 18 percent they decided not to trade, and in 63 percent they followed their private information. Similar behavior also arose when the history was not observable (see Table 3). In all these cases of contrarianism, the market was unable to aggregate private information correctly.[22] Table 3 also reports the percentage of contrarianism in the fixed-price treatment.[23] With a fixed price, subjects never adopted contrarian behavior. In the periods of potential contrarian behavior, they followed their signal in 99

Homogeneous behavior across subjects was also observed in the other treatments.

[20] Moreover, in a similar experiment, Mathias Drehmann et al. (2005) obtain a similar proportion of no trades by employing a binary lottery procedure, which should induce risk neutrality.

[21] The maximum loss is computed in the following way. Suppose that a subject faces a price p. If he buys and the value of the asset turns out to be 0, he loses p. If he sells and the value is 100, he loses $100 - p$. Therefore, the maximum loss is $\max\{p, 100 - p\}$.

[22] We have also computed the propensity to act as a contrarian for different levels of the trade imbalance. The results are very similar to those presented above. For example, when we consider a trade imbalance of at least 1, we find that 16 percent of the decisions are contrarian, 17 percent are no trade, and 67 percent are taken according to the signal. When we consider a trade imbalance of at least 3, these figures change only slightly to 21 percent, 19 percent, and 60 percent.

[23] In this treatment, given that the price is not updated, the behavior of a subject is defined as contrarian if he sells after a positive trade imbalance of at least 2, or if he buys after a negative trade imbalance of at least −2.

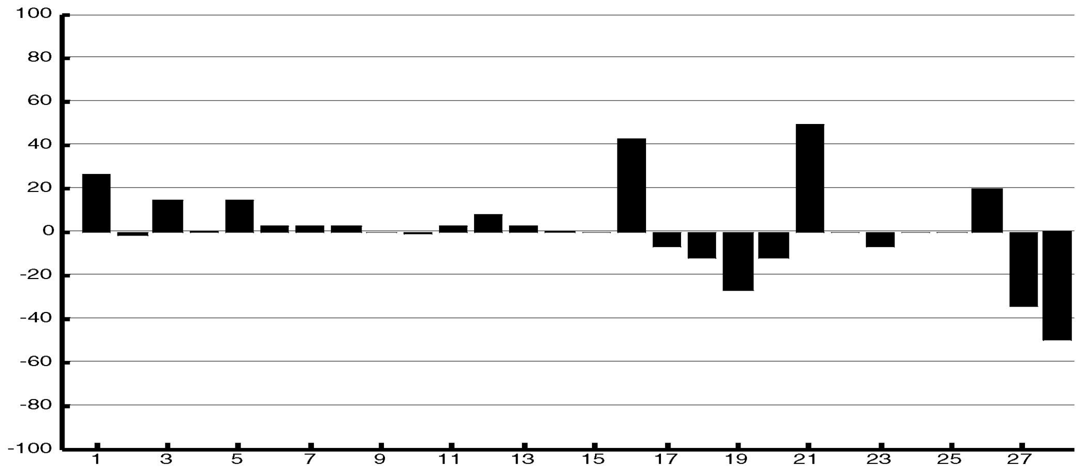

FIGURE 1. DISTANCE BETWEEN THE THEORETICAL AND THE ACTUAL LAST PRICES IN THE FLEXIBLE PRICE TREATMENT

percent of cases and decided not to trade in only 1 percent of cases.

The results reported in Tables 1 and 3 suggest that, while with flexible prices subjects seldom engage in irrational herding (and, therefore, the market is able to learn private information), they also have a lower incentive to use their information when this 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 and helps to explain why the percentage of irrational trades is higher in the flexible-price treatment than in the fixed-price treatment.

When people act as herders or as contrarians, they trade against their own signal and prevent private information from being correctly aggregated by the market price. It is important, however, to remark the different effect of herding and contrarian behavior on social learning. Herd behavior amplifies the importance of early decisions. If these decisions are incorrect, everyone makes the same mistake. In contrast, contrarian behavior goes against the previous history of trades and therefore reduces its importance. In terms of the price path, this means that herding can make the price converge to the wrong value (if, for instance, early traders sell and the value is 100), whereas contrarian behavior makes the price regress toward the mean (given that, for instance, if early traders sell, then contrarians buy).

B. *Actual and Theoretical Prices*

According to theory, the price should converge to the true value of the asset as the number of trading periods tends to infinity. In each period, by choosing to buy or to sell, subjects reveal their private information. Therefore, over time, by the law of large numbers, the price should reflect the fundamental value of the asset. In our treatment, with only 12 periods of trade, there is, of course, no guarantee that the price should always converge to the fundamental value, as private information may not be able (even if perfectly aggregated) to reveal the fundamental value of the asset. In order to have a clear idea of the price convergence in the laboratory, we proceeded in the following way. We studied the price level after all 12 subjects had traded and compared it to the levels that should have been observed theoretically. These theoretical prices were the prices implied by the Avery and Zemsky model, given the particular signal realizations (i.e., the prices that we would have observed if all subjects had followed their signals).

Figure 1 shows the difference between the theoretical and the actual last prices in the flexible-price treatment.[24] Note that 50 percent of the

[24] Remember that for each treatment we have observations for 28 rounds, i.e., for the last seven rounds of each of the four sessions.

time this difference was lower than 5, and 61 percent of the time it was lower than 10. The theoretical and actual prices never moved in the opposite direction (i.e., the distance was never greater than 50). However, 14 percent of the time the distance was greater than 30. The average difference (in absolute value) between the last actual and theoretical prices was 12.[25]

This relatively small distance between the last theoretical and actual prices is a direct consequence of the fact that, in the experiment, 65 percent of the time subjects followed their signal (as the theory suggests) and only 13 percent traded against it. Nevertheless, at least in some rounds, irrational decisions not to trade or to trade against the signal created a wedge between the theoretical and the actual prices.

C. *An Analysis of Errors*

As the previous analysis shows, during the experiment, subjects made errors, i.e., they did not behave rationally all the time. In computing their expected payoffs, the subjects who made a decision in the later periods could factor in the possibility of errors by their predecessors. This, in turn, could change their optimal trading decisions. To explore this issue, we performed an analysis of errors similar to that of Anderson and Holt (1997). We estimated the error rates, assuming that expected payoffs are subject to shocks distributed independently as a logistic random variable.[26] 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,[27] Π_t, i.e.,

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

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

The model implies that a subject may not choose the action that yields the highest payoff, i.e., that he may make a mistake. For each time t, we estimated the parameters of the model by regressing the trading decisions on Π_t. The analysis was recursive, i.e., we used the estimated parameters $\gamma_j^1, \dots, \gamma_j^{t-1}$ to compute the expected payoffs at time t.

Was herding or contrarian behavior ever rational in the flexible-price treatment when one recognizes that subjects make mistakes? In principle, this is a possibility. Theory rules out a decision at time t at odds with private information, assuming that the price is set efficiently by the market maker. In our experiment, however, whereas the price is set assuming that subjects do not make mistakes, people can indeed choose the wrong action. Suppose, for instance, that some subjects in the first periods made mistakes and bought the asset, although they had a negative signal. The market maker assumes that everyone is rational and updates the price after each buy on the basis of this assumption. If the next subject takes into account that previous subjects may have chosen the wrong action, he will realize that the price is too high, given the previous history of trades, and, therefore, will decide to sell, despite a negative signal. This would explain contrarian behavior. Similarly, think of the case where, after some buy orders, some subjects decide not to trade. If the next subject 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 explain the (modest) proportion of irrational herding that we found in the flexible-price treatment.

Table 4 reports the results of the analysis of errors for contrarian behavior. The table shows the percentage of time in which acting as a contrarian was "rational" for different levels of the trade imbalance (there is no column for the fixed-price treatment, since we observed no contrarianism in that treatment). While contrarianism cannot be reconciled with theory when it occurred at low levels of the trade imbalance, it can be rationalized when the trade imbalance was higher and, therefore, the price was more extreme.

Table 5 reports the results of the analysis of errors for herd behavior. The modest

[25] Drehmann et al. (2005) report a higher distance between the last theoretical and actual prices than we do. Their results and ours are not perfectly comparable, since they run different treatments with different parameter values, and they report the average difference between theoretical and actual prices across these treatments. Moreover, they report a lower level of rationality in the experiment than we do, which can also explain why last-period theoretical and actual prices are farther apart.

[26] See Richard McKelvey and Thomas Palfrey (1995, 1998).

[27] The expected payoff of a no trade does not enter the formula, since it is constant for all times t.

TABLE 4—ANALYSIS OF ERRORS FOR CONTRARIAN BEHAVIOR

Absolute value of the trade imbalance	Flexible price	No history	Endogenous pricing
2–3	0 percent	40 percent	69 percent
4–5	57 percent	100 percent	100 percent
> = 6	100 percent	100 percent	100 percent

Note: The table shows, for different absolute values of the trade imbalance, the percentage of "contrarian trades" that are rational under the assumption that subjects take into account the possibility that predecessors may have made mistakes.

TABLE 5—ANALYSIS OF ERRORS FOR INFORMATIONAL CASCADES AND HERD-LIKE BEHAVIOR

	Fixed price	Flexible price	No history	Endogenous pricing
Rational decisions	93 percent	0 percent	0 percent	0 percent

Note: The table shows the percentage of informational cascades (fixed price) or herd-like decisions (other treatments) that are rational under the assumption that subjects take into account the possibility that predecessors may have made mistakes.

proportion of irrational herd-like behavior that we observe in the treatments in which the price is flexible cannot be justified for any trade imbalance.

Finally, let us consider the fixed-price treatment. According to theory, informational cascades can arise as a subject's private information, after some trades, is 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 subjects buy the asset, it is not necessarily true that the third should buy as well if he receives a negative signal. In fact, if the probability of the first two subjects 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 shows that, even taking the possibility of errors into account, ignoring the signal and herding was rational in most of the cases that we classified as informational cascades: in fact, it was the rational decision in 93 percent of these cases (see Table 5).

IV. Endogenous Price Setting

In the flexible-price treatment, we tested how subjects used their own private information in a market in which the price is set as in Avery and Zemsky (1998). This treatment was meant to test a theoretical result, namely, that with such a price mechanism, agents act according to their private information and do not imitate their predecessors.

While the flexible-price treatment offers a clear benchmark to evaluate experimental traders' behavior, one may wonder what happens when, as in actual financial markets, the price is set by a person acting as market maker. To analyze these issues, we ran a treatment in which the price was set by two participants who acted as market makers.[28] This is important, since the Bayesian rule that we applied in the flexible-price treatment did not take into account subjects' actual behavior. In contrast, if market makers are also experimental subjects, they can modify the way they set the prices depending on how the other subjects trade. This, in turn, may have an effect on trading decisions.

The treatment, which we label "endogenous-pricing" treatment, was run according to the following procedures. The two experimental market makers chose the prices at which traders could trade, and updated them on the basis of the order flow. They chose their prices simultaneously and independently without observing each other's decision. The subject administrator recorded the prices on the blackboard, which

[28] We thank Douglas Bernheim (the co-editor) for proposing this treatment. For other experiments with endogenous market making, see, e.g., Robert Bloomfield and Maureen O'Hara (1998, 1999, 2000).

both market makers and traders could see. Then he called a subject to trade. After observing the subject's decision, the market makers chose the prices for the following trader. The procedures for the rest of the treatment were identical to those described in Section III.[29]

When a subject was called to trade, he could trade at the better of the two prices set by the market makers. He could buy at the lower price and sell at the higher price.[30] The traders' payoff was computed as in the other treatments. As with the traders, the market makers' payoffs also depended on the difference between the realized value of the asset and the price of the transaction.[31] Since the traders are informed, whereas the market makers are not, in expected value the market makers would always lose money by trading with the traders.[32] Therefore, we compensated them with a fixed amount of 110 lire for each trading period (i.e., 10 more lire than we gave to the traders).[33] Given this payoff structure, competition guarantees that the market makers should set the price equal to the expected value of the asset conditional on the order flow. Since the market makers' task was harder than that of the traders, they participated in training sessions before the experiment.[34] Only the market makers participated in these sessions. This assures that the behavior of subjects acting as traders is comparable to that observed in the other treatments.

A. *Price-Setting Decisions*

Now let us discuss the results. Prices were set at similar levels to those of the flexible-price treatment (the average distance was 5.5 lire). Figure 2 shows how the market makers updated the prices after a trading order for different absolute levels of the trade imbalance.[35] The figure contrasts these price changes with those of the flexible-price treatment, i.e., with the theoretical Bayesian updating. When the trade imbalance was 0, after a buy or a sell order, market makers updated the price less than we did in the flexible-price treatment. In other words, market makers did not give as much weight to the arrival of just one buy or sell order. In contrast, when the trade imbalance increased to 1 or 2 (in absolute value), the market makers updated the price slightly more than in the theoretical model. For higher trade imbalances, their updating was close to the theoretical one. It is also important to note that, after a no trade, market makers barely moved the price.

Given these price levels and the realizations of the signals, following private information would have been the rational decision 90 percent of the time (versus 100 percent of the time in the flexible-price treatment). In the remaining 10 percent of the time, traders should have traded against the signal (which was never the optimal action in the flexible-price treatment). Note that, in order to determine the optimal choice, we use the same maintained hypotheses about subjects' beliefs as in the analysis of the

[29] In this treatment, the sessions lasted approximately 2.5 hours.

[30] With such a mechanism, a rational agent finds it optimal to buy if his expected value is higher than the average price, and to sell otherwise. If a subject buys, his expected profit is $E(V|H_t, x_t) - \min(P_t^I, P_t^{II})$, where P_t^I and P_t^{II} are the two posted prices. If he sells, his expected profit is $\max(P_t^I, P_t^{II}) - E(V|H_t, x_t)$. Therefore, he will buy if $E(V|H_t, x_t) \geq [(\min(P_t^I, P_t^{II}) + \max(P_t^I, P_t^{II}))/2] = (P_t^I + P_t^{II})/2$ and sell if $E(V|H_t, x_t) \leq (P_t^I + P_t^{II})/2$.

[31] In particular, if a trader bought from a market maker, this market maker earned $110 + P_t - v$ lire; if a trader sold to a market maker, this market maker earned $110 + v - P_t$ lire; if there was a no trade, the market maker obtained 110 lire. Note that, in each trading period, at most one market maker traded. The other market maker's payoff was equal to a no-trade payoff. When the market makers chose the same prices, we randomly decided which market maker actually traded.

[32] Recall that there are no uninformed traders and the market makers are not allowed to post a bid and an ask price.

[33] At the end of the experiment, we divided the total amount of lire earned by the market makers by 12 and then exchanged it at the same exchange rate used for the traders. We divided the total payoff by 12, since the market makers received the fixed endowment 12 times as often as the traders. This guarantees that market makers and traders could earn a similar amount of money in the experiment.

[34] A training session for market makers is also used in Bloomfield and O'Hara (1998, 1999, 2000). In our training sessions, participants received the same instructions as in the experiment. They were told that they were participating in a training session aimed at making the rules of the experiment clear to them. In the training session, the role of traders was played by a computer software program, which simulated a sequence of trading orders. At the end of each round, each market maker was informed of the realized value of the asset and could compute his payoff. The training sessions lasted, on average, 2.5 hours.

[35] The averages were computed considering the price changes of both market makers in all four sessions.

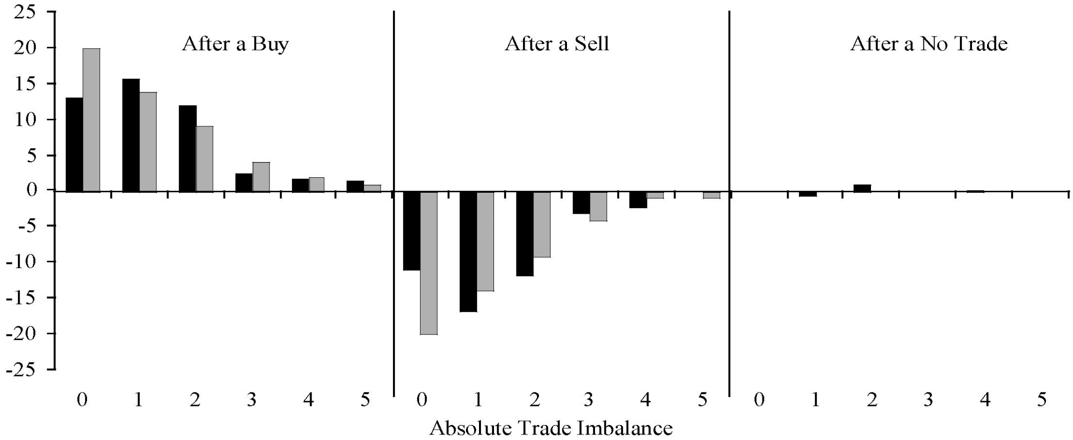

FIGURE 2. PRICE UPDATING IN THE ENDOGENOUS PRICING TREATMENT

Note: The black bar shows the average price updating of market makers for different absolute values of the trade imbalance in the endogenous pricing treatment. The grey bar shows the price updating that we used in the flexible-price and no-history treatments.

previous treatments. Note also that in this treatment the time-t optimal decision depends not only on the subjects' beliefs, but also on the prices posted at time t and in all the previous periods. Prices posted in previous periods are relevant, since, when computing his expectation, a subject must take into account that, given the posted prices, previous subjects may have found it optimal to act against their signal.[36]

B. *Herding and Contrarian Behavior*

Let us now turn to the traders' decisions. As just noted, a trader's rational choice depends on the prices set by the market makers (and does not necessarily coincide with following one's own signal). Rational decisions account for 57 percent of all decisions. No trades amount to 24 percent. These results are in line with those observed in the flexible-price treatment.

The fifth column of Table 1 reports the results on herd behavior for this treatment. There were 58 periods in which the trade imbalance was at least two (or at most -2) and the subject received a signal against it. In these periods, subjects decided to ignore private information and engage in irrational herding in 21 percent of cases; in 34 percent of cases they decided not to trade, and in 45 percent they followed their signal. These results, and in particular the low proportion of herd-like decisions, are similar to those of the flexible-price treatment.[37] Note that, in this treatment, herding could potentially be optimal if market makers updated prices incorrectly. For instance, if market makers updated the price too little, it could be optimal to imitate the predecessors. In our experiment, however, this was never the case.[38]

The observed herd-like behavior cannot be explained even by an analysis of errors (see Table 5, column 5). The level of prices was such that, in the experiment, there were no periods in which it would have been optimal to neglect the signal in order to herd, even recognizing that

[36] As with the fixed-price treatment, in this treatment it could be that a decision to buy or sell was out of the equilibrium path. For instance, it could be that, given the posted price, a rational subject should have bought independently of his signal. If in that period the subject sold, we treat such a decision as we did for the fixed-price treatment (i.e., as revealing a negative signal). Our results are unchanged under the alternative assumption that irrational decisions do not reveal any signal.

[37] A Mann-Whitney test for the hypothesis that the proportion of herding decisions was the same under this treatment and the flexible-price treatment cannot be rejected at the 5-percent significance level (p-value $= 0.15$). The hypothesis that the proportion of herding decisions was the same under this treatment and under the no-history treatment cannot be rejected either (p-value $= 0.67$).

[38] This is true under the same hypotheses on subjects' beliefs described in the last paragraph of Section IV A.

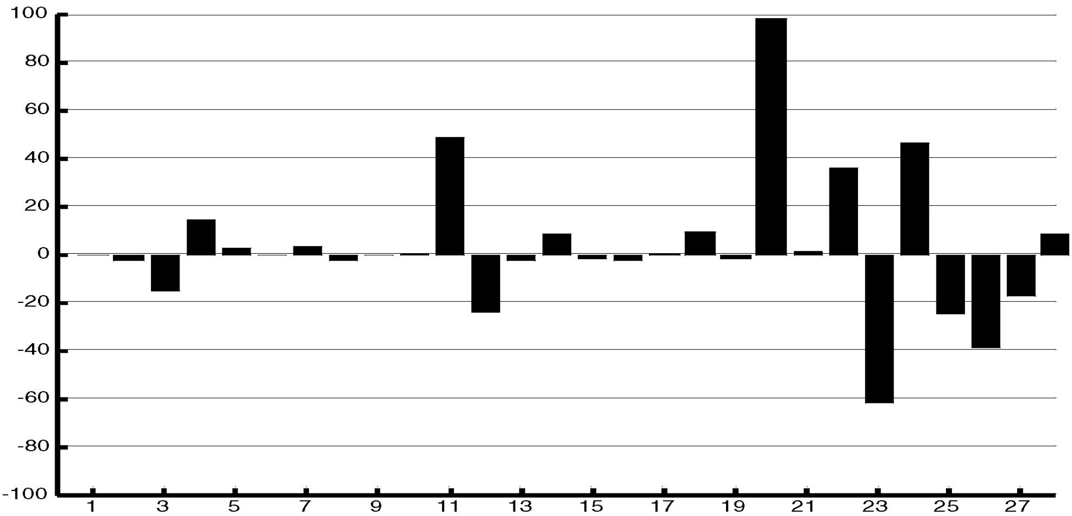

FIGURE 3. DISTANCE BETWEEN THE THEORETICAL AND THE ACTUAL LAST PRICES IN THE ENDOGENOUS PRICING TREATMENT

predecessors could have made mistakes. This means that experimental market makers were able to set prices at which, given the level of rationality in the experiment, no one should have herded.

Let us now discuss contrarian behavior. There were 116 periods in which subjects could have potentially acted as contrarians. They behaved as contrarians in 24 percent of these periods, whereas in 23 percent they decided not to trade and in 53 percent they followed their private information. The results are similar to those of the previous treatments with flexible prices.[39] Part of this contrarianism can be justified by looking at the price levels set by the market makers. Indeed, given the prices set by the market makers, contrarian behavior was the optimal choice in 46 percent of the cases in which it was observed.[40] Also, the analysis of errors helps to explain part of contrarian behavior (see Table 4, column 4): in particular, when the absolute level of the trade imbalance was equal or higher than 4, all contrarian decisions were rational if we assume that traders could correctly estimate previous subjects' level of rationality.

Now, let us study price convergence. Figure 3 shows the difference between the theoretical and the actual last prices. The theoretical prices are those implied by the Avery and Zemsky model given the particular realizations of the signals (i.e., the prices that we would have observed if all subjects—traders and market makers—had behaved as in the theoretical model). As in the flexible-price treatment, 61 percent of the time the difference between the actual and the theoretical last prices was lower than 10, and 50 percent of the time it was lower than 5. Once, however, in round 20, the actual price moved close to 100, while the theoretical price was close to 0. Moreover, 21 percent of the time the distance was greater than 30. As a result, the average distance between actual and theoretical last prices was 17, slightly higher than what was observed in the flexible-price treatment.

C. *A Comment on the Bid-Ask Spread*

In all the treatments discussed so far, the market makers (i.e., the experimenter in the previous treatments or two subjects in the endogenous-pricing treatment) could set only one price and, as a result, there was no bid-ask spread. One may wonder whether the presence of a bid-ask spread would affect our results. To answer this question, we ran a related treatment in which

[39] A Mann-Whitney test for the hypothesis that the proportion of contrarian decisions was the same under this treatment and the flexible-price (or no-history) treatment cannot be rejected at the 5-percent significance level (the p-value in both cases is 0.56).

[40] See footnote 38.

two participants acting as market makers set two prices, a bid and an ask price.[41] For a detailed description of this treatment, we refer the reader to an addendum posted on the *American Economic Review's* Web site (http://www.e-aer.org/data/dec05_app_cipriani.pdf).[42] Here we summarize only the main results concerning traders' behavior.

In this treatment, we observed 19 percent of herd-like behavior, a percentage very close to that observed in the other treatments with a flexible price discussed above. Therefore, the presence of a bid-ask spread does not change our result on the effect of flexible prices on herd behavior. Contrarian behavior arose in 10 percent of cases. This percentage is identical to that of the no-history treatment, but lower than what was observed in the flexible-price and in the endogenous-pricing treatments. On average, during the entire experiment, subjects behaved rationally 58 percent of the time, practically the same percentage observed in the endogenous-pricing treatment. The proportion of no trades, 35 percent, was higher than in the other treatments. It should be noted, however, that almost half of the no-trade decisions were rational, as the market makers set a bid-ask spread so large (i.e., larger than what theory predicts) that trading was not optimal.

In conclusion, the presence of a bid and ask spread did not modify traders' propensity to herd, or the overall level of rationality in the experiment. It reduced contrarian behavior and increased the proportion of no trades. Since in actual financial markets the size of the bid-ask spread varies significantly, the results of our treatment suggest that we may find more contrarianism in liquid markets, i.e., markets where the bid-ask spread is lower. In most financial markets, however, the average size of the bid-ask spread is much smaller than in our experiment. Therefore, one should be cautious in evaluating the quantitative importance of the presence of a bid-ask spread in reducing contrarianism and increasing no trading.

[41] For other experiments with subjects setting a bid-ask spread, see Bloomfield and O'Hara (1998, 1999, 2000). The bid-ask spread has also been studied experimentally in Bloomfield (1996).

[42] The addendum is also available on the authors' Web pages: http://home.gwu.edu/~mciprian/; http://www.homepages.ucl.ac.uk/~uctpagu/.

V. Conclusions

In this paper, we have reported and discussed the results of an experimental study on herd behavior in financial markets. We have shown that, in a frictionless laboratory market in which informed traders trade for informational reasons only, herd behavior seldom occurs. This result is consistent with the theoretical predictions of Avery and Zemsky (1998). The result suggests that in order to understand herd behavior in actual financial markets, we must look for other explanations, such as reputation concerns (Scharfstein and Stein, 1990), or noninformational motives to trade (Cipriani and Guarino, 2001).

Theory, however, does not completely capture the behavior observed in the laboratory market. Sometimes subjects decided not to follow their private information. In some cases, they did so because they engaged in "contrarian behavior," i.e., they bought when the price was low or sold when it was high. More frequently, subjects preferred to ignore their private information and abstain from trading. This indicates that limited market participation may be an important source of financial markets' informational inefficiency.

REFERENCES

Anderson, Lisa R. and Holt, Charles A. "Information Cascades in the Laboratory." *American Economic Review*, 1997, *87*(5), pp. 847–62.

Avery, Christopher and Zemsky, Peter. "Multidimensional Uncertainty and Herd Behavior in Financial Markets." *American Economic Review*, 1998, *88*(4), pp. 724–48.

Banerjee, Abhijit V. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 1992, *107*(3), pp. 797–817.

Bikhchandani, Sushil; Hirshleifer, David and Welch, Ivo. "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades." *Journal of Political Economy*, 1992, *100*(5), pp. 992–1026.

Bloomfield, Robert. "Quotes, Prices, and Estimates in a Laboratory Market." *Journal of Finance*, 1996, *51*(5), pp. 1791–1808.

Bloomfield, Robert and O'Hara, Maureen. "Does Order Preferencing Matter?" *Journal of Financial Economics*, 1998, *50*(1), pp. 3–37.

Bloomfield, Robert and O'Hara, Maureen. "Market Transparency: Who Wins and Who

Loses?" *Review of Financial Studies*, 1999, *12*(1), pp. 5–35.

Bloomfield, Robert and O'Hara, Maureen. "Can Transparent Markets Survive?" *Journal of Financial Economics*, 2000, *55*(3), pp. 425–59.

Brunnermeier, Markus K. *Asset pricing under asymmetric information: Bubbles, crashes, technical analysis, and herding*. Oxford: Oxford University Press, 2001.

Çelen, Boğaçhan and Kariv, Shachar. "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review*, 2004, *94*(3), pp. 484–98.

Chamley, Christophe P. *Rational herds: Economic models of social learning*. Cambridge: Cambridge University Press, 2004.

Chari, V. V. and Kehoe, Patrick J. "Financial Crises as Herds: Overturning the Critiques." *Journal of Economic Theory*, 2004, *119*(1), pp. 128–50.

Cipriani, Marco and Guarino, Antonio. "Herd Behavior and Contagion in Financial Markets." Unpublished Paper, 2001.

Drehmann, Mathias; Oechssler, Jörg and Rider, Andreas. "Herding and Contrarian Behavior in Financial Markets: An Internet Experiment." *American Economic Review*, 2005, *95*(5), pp. 1403–1426.

Glosten, Lawrence R. and Milgrom, Paul R. "Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics*, 1985, *14*(1), pp. 71–100.

Grinblatt, Mark; Titman, Sheridan and Wermers, Russ. "Momentum Investment Strategies, Portfolio Performance, and Herding: A Study of Mutual Fund Behavior." *American Economic Review*, 1995, *85*(5), pp. 1088–1105.

Hirshleifer, David and Teoh, Siew H. "Herd Behaviour and Cascading in Capital Markets: A Review and Synthesis." *European Financial Management*, 2003, *9*(1), pp. 25–66.

Huck, Steffen and Oechssler, Jörg. "Informational Cascades in the Laboratory: Do They Occur for the Right Reasons?" *Journal of Economic Psychology*, 2000, *21*(6), pp. 661–71.

Kübler, Dorothea and Weizsäcker, Georg. "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory." *Review of Economic Studies*, 2004, *71*(2), pp. 425–41.

Lakonishok, Josef; Shleifer, Andrei and Vishny, Robert W. "The Impact of Institutional Trading on Stock Prices." *Journal of Financial Economics*, 1992, *32*(1), pp. 23–43.

Lee, In Ho. "Market Crashes and Informational Avalanches." *Review of Economic Studies*, 1998, *65*(4), pp. 741–59.

McKelvey, Richard D. and Palfrey, Thomas R. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior*, 1995, *10*(1), pp. 6–38.

McKelvey, Richard D. and Palfrey, Thomas R. "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics*, 1998, *1*(1), pp. 9–41.

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

Scharfstein, David S. and Stein, Jeremy C. "Herd Behavior and Investment." *American Economic Review*, 1990, *80*(3), pp. 465–79.

Smith, Lones and Sørensen, Peter. "Pathological Outcomes of Observational Learning." *Econometrica*, 2000, *68*(2), pp. 371–98.

Welch, Ivo. "Sequential Sales, Learning, and Cascades." *Journal of Finance*, 1992, *47*(2), pp. 695–732.

Wermers, Russ. "Mutual Fund Performance: An Empirical Decomposition into Stock-Picking Talent, Style, Transactions Costs, and Expenses." *Journal of Finance*, 2000, *55*(4), pp. 1655–95.