Estimating Herd Behavior in Financial Markets: A Structural Approach

Marco Cipriani and Antonio Guarino*

September, 1st, 2005
PRELIMINARY AND INCOMPLETE. PLEASE DO NOT QUOTE

Abstract

We estimate a structural model of herd behavior in financial markets. A sequence of traders exchanges an asset with a market maker. Trade occurs over many days. Herd behavior can arise despite the fact that the price is efficiently set by the market maker. The price is updated too slowly by the market maker and there are periods in which traders choose the same action independently of their private signal. We estimate the model by maximum likelihood using transaction data on one NYSE stock in the first quarter 1995. We detect the periods of the trading day in which there could be herd behavior. We find that such periods account for 15% of the total number of trading periods.

*Cipriani: Department of Economics, George Washington University (e-mail: mciprian@gwu.edu); Guarino: Department of Economics and ELSE, University College London (e-mail: aguarino@ucl.ac.uk). We thank Andrew Chesher, Robert Engle, Christopher Flinn, Douglas Gale, Lars Nesheim and Jean Marc Robin for very helpful comments and suggestions. We gratefully acknowledge the financial support of the Fondation Banque de France. We are responsible for any error.

1 Introduction

In recent years there has been much 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. This interest has led researchers to look for both theoretical explanations and empirical evidence (for a recent survey, see, e.g., Hirshleifer and Teoh, 2003). The main questions addressed in this literature are: How pervasive is herd behavior in financial markets? Can it arise even if traders act rationally?

The theoretical research started at the beginning of the 90's with the seminal papers of Banerjee (1992), Bikhchandani et al. (1992), and Welch (1992).[1] These papers do not discuss herd behavior in financial markets, but in an abstract setting, in which agents with private information make their decisions in sequence. They show that, after a finite number of agents have chosen the same action, all following agents will disregard their own private information and imitate that action. More recently, a number of papers (see, e.g., Avery and Zemsky, 1998, Lee, 1998, Cipriani and Guarino, 2001, Chari and Kehoe, 2002) have focused on herd behavior in financial markets. In particular, all these studies analyze a market where agents sequentially trade a security of unknown value. The price of the security is efficiently set by a market maker according to the order flow. The presence of a price mechanism makes herding more difficult to arise. Still there are cases in which it occurs. In Avery and Zemsky (1998) people can herd when there is uncertainty not only on the value of the asset but also on other parameters of the model. In Cipriani and Guarino (2001) agents herd because they have other reasons to trade, in addition to informational motives. In this model, not only agents herd, but, in contrast with Avery and Zemsky, a complete informational cascade arises. In Chari and Kehoe (2002) agents choose when to enter the market (i.e., timing is endogenous) and can decide to herd on early traders. While theoretical work has tried to explain herding in a world of rational agents, the empirical literature has followed a different track. The existing work (see, e.g., Lakonishok et al., 1992, Grinblatt et al. , 1995, Welch, 2000 and the other papers cited in the survey of Hirshleifer and Teoh, 2003) does not test these models directly. Instead, it analyzes

[1] In this paper we only study informational herding. Therefore, we do not discuss herd behavior due to reputational concerns or payoff externalities. For a taxonomy of herding models, see Hirshleifer and Teoh (2003).

the presence of herding in financial markets through statistical measures of clustering. These papers find that in some cases fund managers (or financial analysts) tend to cluster their investment decisions more than if they acted independently. This work is important as it sheds light on the behavior of financial market participants. It suffers, however, of a major and well known drawback. The decisions clustering may or may not be due to herding. For instance, it may be the result of a common reaction to public announcements. These papers cannot distinguish spurious herding from true herd behavior (see Bikhchandani and Sharma, 2000, and Welch, 2000). The problem that empiricists face in the task of detecting herding is that there are no data on the private information available to the traders and, therefore, it is difficult to understand whether traders make similar decisions because they disregard their own information and imitate or because of other reasons.

The purpose of this paper is to overcome this problem and offer an empirical analysis of herd behavior which is not purely statistical. We present a theoretical model of herding and estimate it. We are able to identify the periods in the trading day in which traders act as herders according to the model. We do not have data on private information, but we are able to estimate our model of herding by using transactions data only. This is the first empirical paper on informational herding that, instead of using a statistical, a-theoretical approach, directly estimates a herding model.[2]

The model that we present is inspired by the work of Avery and Zemsky (1998). They use a sequential trading model *à la* Glosten and Milgrom (1985). Traders trade an asset of unknown value with a market maker. Traders receive private information on it. The market maker is uninformed and sets the price of the asset on the basis of the buy and sell orders that he receives. Avery and Zemsky show that traders will 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 only). Therefore, it will never be the case that agents neglect their information and imitate previous traders' decisions. Avery and Zemsky, however, also show some examples in which there are multiple sources of uncertainty, i.e., uncertainty not only on the asset value but also on other parameters of the model. In these examples, although the price aggregates private infor-

[2]While there are no direct empirical tests of herding models, there is experimental work that tests these models in the laboratory. Herding in a laboratory financial market is studied by Cipriani and Guarino (2005).

mation, herd behavior arises. The price moves “too slowly” in response to the order flow. Traders have multiple sources of informational advantage on the market maker and, therefore, they update their beliefs on the asset value more quickly that the price. When the price is not very responsive to the order flow, we are in a situation similar to that analyzed by Banerjee (1992) and Bikhchandani et al. (1992) in which the price (i.e., the cost of an action) is fixed. Hence, herding arises here as it does there.

In our model, based on Easley and O’Hara (1992), there is uncertainty not only on the value of the asset but also on the proportion of informed traders in a trading day. The asset is traded over many days. At the beginning of each day, an informational event can occur or not. In the former case, the fundamental asset value changes with respect to the value of the previous day. It can be higher (in the case of a good informational event) or lower (in the case of a bad informational event) than the value in the previous day. In the latter case, instead, it remains unchanged. If an event has occurred, some agents will receive private information on the new asset value. These agents will trade the asset to exploit their informational advantage on the market maker. On the contrary, if no event has occurred, all traders in the market will be uninformed: they will trade for other reasons only, like liquidity or hedging motives. While the informed traders know that they are in a market in which there is private information (since they themselves are informed) the market maker does not know whether he is in an informed or uninformed market for that day. This asymmetry in information determines a different way of updating the beliefs on the asset value by traders and market maker. The market maker will move the price “slowly” since he has to take into account the possibility that the asset value has not changed, the market is uninformed and all orders are coming from uninformed (noise) traders. His interpretation of the history of trades will be different from the traders’. There can be times in which, irrespective of their private signals, all informed agents will value the asset more than the market maker and, therefore, find it optimal to buy; or they will all value the asset more than the market maker and find it optimal to sell. These are periods in which herd behavior arises.

To estimate our model, we will use a strategy proposed by Easley, Kiefer and O’Hara (1997). They show how to use transaction data to estimate the parameters of the Glosten and Milgrom model by means of maximum likelihood. We will construct the likelihood function for the trading of an asset over multiple days. Our function takes into account that for some histories of trade agents will herd. This means that in these histories, the probability of a buy or a sell order will be different from that when each agent follows his own private information. Our task will be more complicated than Easley, Kiefer and O'Hara's. In their set up, informed traders are perfectly informed on the value of the asset. Given that their signal is perfectly informative, these traders' decisions will never be affected by the previous decisions and they will never herd. Therefore, the only thing that matters is the *number* of buys, sells and no trades during the day. The *sequence* in which orders arrive is irrelevant. In contrast, in our framework, history matters. Sequences with the same number of buys and sells may occur with different probabilities depending on the order in which buy and sell decisions arrive in the market. Therefore, we cannot limit our analysis to the number of orders but we have to consider the entire sequence of trades.

We applied our methodology to the trading activity of Ashland Oil stock, a stock traded in the NYSE, in the first quarter of 1995. We find that 7.3% of the trading periods are characterized by herd buy and 8% by herd sell. In 23 (out 63) days of trade, herding periods were higher than 15% of the total periods of trade. Herding behavior generated a significant deviation of the price path from the full information level.

The paper is organized as follows. Section 2 describes the theoretical model. Section 3 characterizes herd behavior and shows how it can arise. Section 4 describes the estimation strategy. Section 5 describes the data. Sections 6 presents the results, and Section 7 concludes.

2 The model

Our model is based on Easley and O'Hara (1987), which is a derivation of Glosten and Milgrom (1985). The Glosten and Milgrom model can be interpreted as describing the process of trading during a single day. The present model, by contrast, describes trading over multiple days. A sequence of agents exchanges an asset of unknown value with a market maker. At the beginning of each day, nature chooses the value of the asset for that day. There are two possibilities: either the value remains the same as in the previous day, or it changes, going up or down. In the latter case, we say that there has been a (good or bad) informational event. The reason we call this event "informational" is that some traders in the market will receive private information on it. Therefore, the decisions of the traders during the day will contain information on the event and, hence, on the value of the asset.

Let us now move to a formal description of the model.

We consider an asset whose value at the end of a trading day $d = 1, 2, 3...$is a random variable V_d. In day 1 the asset value V_1 is distributed on the support $\{V_1^L, V, V_1^H\}$ with the following probabilities: $\Pr(V_1 = V) = (1-\alpha)$, $\Pr(V_1 = V_1^L) = \alpha\delta$, and $\Pr(V_1 = V_1^H) = \alpha(1-\delta)$. In any future day $d \geq 2$, the asset value can remain the same as the realized value of the previous day, or change. In particular, V_d is equal to V_{d-1} with probability $(1-\alpha)$ and changes with probability α. In the latter case, it is equal to $V_{d-1} - \Delta_d^L$ with probability $\alpha\delta$ and to $V_{d-1} + \Delta_d^H$ with probability $\alpha(1-\delta)$, where $\Delta_d^L > 0$ and $\Delta_d^H > 0$. We assume that these events are independent over time. To alleviate notation, we define $V_d^H := V_{d-1} + \Delta_d^H =$ and $V_d^L := V_{d-1} - \Delta_d^L$. Moreover, we assume that $\delta\Delta_d^H = (1-\delta)\Delta_d^L$.

The asset is exchanged in a specialist market. Its price is set by a competitive market maker (the specialist) who interacts with a sequence of traders. There is a continuum of traders. At any time $t = 1, 2, 3...$during the day a trader is randomly chosen to act and can buy, sell or decide not to trade. Each trade consists of the exchange of one unit of the asset for cash.

In the market there can be two types of traders: informed and uninformed traders. Uninformed traders buy, sell or do not trade with fixed probability, for unmodelled reasons. They buy with probability $\frac{\varepsilon}{2}$, sell with probability $\frac{\varepsilon}{2}$ and do not trade with probability $(1-\varepsilon)$. Moreover, we assume that they decide independently of one another. Informed traders are risk neutral. They receive a private signal on the asset and maximize their expected profit based on that signal. The signal x is distributed on the support $\{x^L, x^H\}$ (with $x^L < x^H$) according to the conditional probability function $\Pr(x = x^L | V_d^L) = \Pr(x = x^H | V_d^H) = q$, with $\frac{1}{2} < q < 1$. Informed agents know that an event has occurred and their signal is informative on whether the event is good or bad. Nevertheless, they are not completely sure of the effect of the event on the asset value. For instance, they know that there has been a change in the investment strategy of a company, but they cannot be completely sure that this change will affect the asset value in a positive or negative way. The precision of the signal (q) can also be interpreted as measuring the ability of traders to process the private information that they receive.

When no event occurs, there is nothing agents can learn in that day and, therefore, all agents in the market are uninformed. On the contrary, when there is an event, a positive proportion of agents receives a private signal on the new asset value. In this case, the proportion of informed traders is μ.

Of course, the idea that an event may occur only at the beginning of a trading day is a pure abstraction. What is really relevant, from a theoretical point of view, is that new information on changes in the asset value arrives at discrete times. We will comment more on this point when we will discuss the mechanism through which herd behavior arises.

At time t of day d, the market maker and the traders know the history of trades (H_t^d) until time $t-1$ of the same day. We denote the expected value of the asset for an informed trader in day d at time t by $E(V_d|H_t^d, x)$.

On the other hand, the expected value for the market maker will be conditioned only on the public information available at time t, i.e., it will be

$$P_t^d = E(V_d|H_t^d). \tag{1}$$

We refer to the market maker's expectation as the price of the asset, although it is not the posted price. Indeed, the market maker must take into account the possibility that he may trade with agents more informed than he is. Therefore, he will set a bid-ask spread between the price at which he is willing to sell and to buy (Glosten and Milgrom, 1985). Given that the market maker behaves competitively by assumption, he will make zero expected profits. Hence, the bid and ask price for the asset will be

$$Bid_t^d = E(V_d|H_t^d, h_t^d = sell) \text{ and } Ask_t^d = E(V_d|H_t^d, h_t^d = buy), \tag{2}$$

where h_t^d is the action taken by the trader who arrives in day d at time t. At any time t, there exists a unique bid and ask price for the asset, which satisfies $Bid_t^d \leq P_t^d \leq Ask_t^d$.[3] The market maker takes into account that buying or selling orders contain private information and sets a spread between the price at which he is willing to sell and buy. Equilibrium prices always exist because noise traders are willing to accept any loss and, therefore, the market will never shut down.

3 Herd Behavior

Very preliminary...to be completed.

[3] For a formal proof see the Appendix.

In our model agents can rationally decide to neglect their signal and follow the decisions of previous traders. We will, first, introduce the formal definition of herding and then illustrate the result.

Definition 1 *There is herd behavior at time t of day d if any informed trader called to trade chooses the same action, independently of his private information.*

Our definition of herding is standard in the literature (see, e.g., Gale, 1996, and Smith and Sørensen, 2000). Agents herd when they act independently of their own private information. Because of this, there is conformity in the market. Avery and Zemsky (1998) have analyzed cases in which herd behavior can arise that is similar in the spirit to that presented here (see their models IS2 and IS3).

It is clear that traders who herd can make the wrong decision. For instance, it is possible that in a trading day characterized by a positive informational event traders in a period of herding neglect their positive information on the true asset value and decide to sell it. In the cases in which agents neglect the correct signal and take the opposite action, we say that herding is misdirected.

We can prove the following Proposition:

Proposition 1 *During an informed trading day, herd behavior arises with positive probability. Furthermore, herd behavior can be misdirected.*

Intuitively, the reason for the occurrence of herding is that the price, although efficiently set by the market maker, moves "too slowly." When informed traders and market maker look at the past history of trades, they interpret it in different ways. Suppose that a sequence of buy orders arrives on the floor. Informed traders, knowing that there has been an informational event, attach a certain probability to the fact that these orders come from agents with positive private information. The market maker attaches a lower probability to this event, as he has to take into account the possibility that there was no event and that all traders in the market trade for non-speculative reasons only. Therefore, after this sequence of buys, he will update the price up less than if he knew that an event had occurred. It can be that even a trader with negative information has an expected value of the asset higher than the (ask) price and, therefore, ignores his signal to herd buy. The same logic explains herd selling: after a sequence of sells, the bid price may be high enough that also traders with positive information prefer to sell.

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 incorporated in the price, as it normally occurs in the absence of herding. In a situation of herding, actions do not reflect private information, as agents are ignoring their signal. Therefore, the price cannot aggregate such information.

While traders herd, the market maker knows that he is in one of two situations: either he is in a day with informational event, where traders are herding; or he is in a day where agents trade for non speculative reasons only. While the price does not aggregate private information efficiently, even during a period of herding, the market maker does learn something on the true asset value. Indeed, even in a period of herding, he can update his belief on the fact that there has been an informational event or not, i.e., that the asset value has changed or not. Therefore, the bid and ask prices are updated even in a period of herding. To continue on the same example above, if the market maker sees other traders buying the asset, he will give more and more weight to the fact that these traders are informed (liquidity traders would indeed buy or sell with the same probability). Hence, he will post higher prices. Because of this movement in prices, herd behavior will eventually disappear. Agents will no longer find it optimal to neglect their signal and private information will be aggregated in the usual manner. The price will converge to the true asset value.[4]

To fix ideas, we present now a simple example in which herd behavior arises and then disappears.

Example 1: How Herd Behavior Arises

Consider an asset whose value in day $d-1$ was $\frac{1}{2}$. In the case of an informational event, which occurs with probability $\alpha = 0.01$, the value will be either 0 or 1, with equal probability $\delta = 0.5$. In the case of an informational event, $\mu = 0.5$, i.e., 50% of traders in the market are informed and 50% are

[4] As we have seen, there is never a complete blockage of information in this model. This is equivalent to say that there is never an informational cascade. For a model that illustrates the possibility of cascades in financial markets, see Cipriani and Guarino (2001).

noise. The private signal has precision $q = 0.65$. Finally, noise traders buy, sell or do not trade with equal probability $\varepsilon = \frac{1}{3}$.

Suppose that an informational event occurs and the first two traders in the market decide to buy.

At time 1, the price is equal to the unconditional expectation, i.e., $P_1 = \frac{1}{2}$. The market maker computes the bid and ask prices assuming that an informed trader with a negative signal wants to sell and an informed trader with a positive signal wants to buy. The ask and bid prices will be $A_1 = 0.502$ and $B_1 = 0.498$. At these prices, an informed trader wants indeed to buy if he receives a positive signal and sell if he receives a negative one. Therefore, these are equilibrium ask and bid. Conditional on receiving a buy or a sell order, the market maker will update upward the probability that an event occurred. In the absence of an event, a buy and a sell happen each with probability 0.33, while after an event they happen with a higher probability, given that at time 1 informed traders never decide not to trade.

At time 2, after receiving a buy order, the market maker will update his belief and the bid and ask prices in the same fashion. The equilibrium ask price will be $A_2 = 0.509$. Note that the ask price barely moves. This is because the market maker has a strong prior that no event occurred. Moreover, the probability of being in an uninformed state of the world declines slowly, as only the lower probability of no trades in an informed than in an uninformed market makes it decline.

At time 3, if the market maker computed the ask as before, i.e., assuming that informed traders choose different actions depending on their signal, an informed trader with a *negative* signal facing such an ask price, would decide to buy. Observing the previous two buys, this trader raises his valuation of the asset more than the market maker, since he knows that an event occurred and there are informed traders in the market. Indeed, his valuation of the asset after two buy orders and upon receiving a low signal is 0.574. The market maker realizes that both types of informed traders would buy and that the buy order reveals nothing on the value of the asset. We are a situation of herd buying. Therefore, the market maker sets the ask assuming that all informed traders would buy, which occurs in equilibrium. The new ask is $A_3 = 0.528$. Note that this value differs from A_2 despite the fact that now we are in a herding period. The market maker updates the probability of an informational event conditional on receiving a new buy order. Indeed, he knows that a new buy order is more likely to come from an informed trader (as all informed traders buy) than from an uninformed trader. Note also that, despite the action of the traders does not reveal anything on the value of the asset, still the bid and ask spread does not collapse to zero. In fact, at time 3 the bid price is $B_3 = 0.504$. In an informed market after two buy orders the probability of a sell is much lower than that of a buy. Therefore, the specialist attaches different probabilities to the informational event after observing a sell or a buy.

Example 2: How herd behavior ends.

As we have pointed out in the last part of the previous section, the market maker keeps updating his beliefs even during periods of herding. Indeed, the probability of being in an uninformed market is revised downwards conditional on a new buy order, because the market maker knows that in an informed market buys are very likely, given that all informed traders would now buy. In contrast, informed traders do not update their beliefs since they know that each action in the market either comes from a noise trader or from another informed trader who is neglecting his signal. In the example above, traders keep buying from period 4 to period 10. After all these buy orders, the bid and ask become $B_{10} = 0.932$ and $A_{10} = 0.969$. At these quotes, an informed trader with a bad signal wants to sell (since his belief is 0.930) and a trader with a good signal wants to buy (since his belief is 0.9786).

This simple example makes clear a more general point. The possibility of informational events makes the market maker interpret the past history of trades in a different manner from the traders. The market maker updates his beliefs too slowly and, therefore, traders may find it optimal to disregard their signals. After a sequence of buys (sells), all informed agents may be willing to take advantage of the slow movement of the ask (bid) price to buy (sell) the asset. The actions taken by the previous traders affect the decisions of the following traders. During the herding period, the market maker still learns something, namely the informational structure of the market. When the market maker's belief on μ approaches the traders' belief, traders will resume trading according to their own private information as they do not interpret the history of trades different from the market maker any longer.

Given that information always flows to the market, despite herd behavior, the price is able to aggregate the information that traders receive, although at a slower rate. Given that this information is on average correct, the price will converge to the true asset value. We prove this result formally in the next proposition:

Proposition 2 *In any day d the asset price converges almost surely to the realized value of the asset.*

Proof. See the Appendix.

The new dimension of uncertainty does not prevent the market from learning the true value of the asset. Although herding arises, there is not a complete blockage of information. Eventually, the market maker will learn if he is in a day characterized by an informational event or not. The market maker and the traders have converging beliefs on the true state of the world, herding will be broken and prices will converge to the fundamentals.

4 Estimating the model

The main objective of our work is to estimate the structural model of herd behavior that we have just illustrated. Once we have recovered the parameters of the model, we can track down the beliefs of the traders and compare them to the prices set by the market maker during each day of trading: this will allow us to detect periods of herding. The parameters that we want to estimate are: the probability of an informational event (α), the probability that the information event is good (δ), the proportion of informed traders if an information event occurs (μ), the precision of the signal that informed traders receive (q), and the probability that a noise trader decides to trade (ε).

We now write the likelihood function for the history of trades, i.e., for the sequence of buys, sells and no trades over many trading days. Note that the histories of bid and ask prices and of transaction prices do not enter the likelihood function as they do not convey any additional information besides that contained by the sequence of transactions. Furthermore, in contrast to Easley, Kiefer and O'Hara (1997), our likelihood function cannot be written as a simple function of the number of buys, sells and no trades in each day. According to our model, the *sequence* of trades, not just the *number* of transactions, conveys information.

Clearly, to estimate all our parameters we need data on more than one trading day. The parameters α and δ define the probability that there is an event at the beginning of a trading day and that the event is good or bad. Therefore, if we used data on one day only, this would be equivalent to observing just one draw from a joint distribution with parameters α and δ. Clearly we would be unable to estimate these two parameters. For this reason, we will use multiple days data.

While α and δ govern the probability of having a (good or bad) informational event at the beginning of a day, the remaining parameters μ, ε, and q, define the probability of observing a particular history of trades during a trading day, *given* that the day is good, bad or with no event. In each day there will be a number of buys, a number of sells and a number of no trades. In our set up, this three-dimensional statistics is not sufficient for the probability distribution of the trading sequences. Indeed, as we showed in the previous section, in our framework history matters. Having many buy orders at the beginning of the day is not necessarily equivalent to having the same number of buy orders spread during the day. In fact, if there is a concentration of buys, this may generate a period of herding. The market maker in this period will have to update his quotes (beliefs) in a different way than in the absence of herding. Furthermore, the probability of a particular order of trades in such a period of herding is different from the same order in the absence of herding. Therefore, to estimate our three parameters, we will not only use the number of buys, sells, and no trades. Rather, we will use the entire sequence of trades during each day.

From the description of the model and the discussion above, it is clear that histories of trade in days when there is an informational event will be quite different from histories in days when an event occurred. In a no-event day, we may expect liquidity traders to buy or sell in a balanced way. In contrast, given that informed traders follow an informative signal, when there is an informational event there will be either a prevalence of buys or a prevalence of sells, so that the price converges to the new value.

Informational events at the beginning of each day are independent of each other. Furthermore, trades in a day are independent of the history of trades in previous days. Therefore, a particular history of trades over multiple days can be written as the product of the probability of the history of each single day. Let us define the complete history of trades in a single day as $H^d := H^d_{T_d}$, where T_d is the number of trading periods in day d. Then,

$$\Pr\left(\{H^d\}_{d=1}^D|\alpha,\delta,q,\mu,\varepsilon\right) = \Pi_{d=1}^D \Pr(H^d|\alpha,\delta,q,\mu,\varepsilon).$$

The probability of the history during a trading day, can be written as

$$\Pr(H^d|\alpha,\delta,q,\mu,\varepsilon) =$$
$$(1-\alpha)\Pr(H^d|V_d = V_{d-1}) + \alpha(1-\delta)\Pr(H^d|V_d = V_d^L) + \alpha\delta\Pr(H^d|V_d = V_d^H)$$

(note that to simplify the notation we have omitted the parameters in the conditional probabilities in the last line).

Now we describe how to compute these conditional probabilities, starting with the case of no information event. During a non informed day, the probability of each buy or sell is $\frac{\varepsilon}{2}$, while the probability of each no trade is $(1-\varepsilon)$. Given that noise traders act independently of one another, the probability of all sequences of B_d buys, S_d sells and N_d no trades in day d is simply

$$\begin{aligned}\Pr(H^d|\alpha,\delta,q,\mu,\varepsilon,V_d &= V_{d-1}) = \Pr(B_d,S_d,N_d|\alpha,\delta,q,\mu,\varepsilon,V_d = V_{d-1}) \\ &= K\left(\frac{\varepsilon}{2}\right)^{B_d+S_d}(1-\varepsilon)^{N_d},\end{aligned}$$

where K is the number of permutations of B_d buys, S_d sells and N_d no trades.

When there is an informational event, things are significantly more complicated as we have to distinguish times in which agents follow their own signal from times in which they herd. As we illustrated above, this depends on the expectations of the market maker and of the traders. Let us start from the case in which the market maker sets the bid and the ask assuming that informed traders will follow their signal. If indeed, $E(V|H_t^d, x^H) \geq Ask_t^d$ and $E(V|H_t^d, x^L) \leq Bid_t^d$, then these are equilibrium prices. In this case, the probability of observing a buy, a sell or a no trade at time t are

$$\Pr(Buy_t^d|H_t^d, V_d^H) = \mu q + (1-\mu)\frac{\varepsilon}{2},\ \Pr(Buy_t^d|H_t^d, V_d^L) = \mu(1-q) + (1-\mu)\frac{\varepsilon}{2},$$

$$\Pr(Sell_t^d|H_t^d, V_d^H) = \mu(1-q) + (1-\mu)\frac{\varepsilon}{2},\ \Pr(Sell_t^d|H_t^d, V_d^L) = \mu q + (1-\mu)\frac{\varepsilon}{2},$$

$$\Pr(NT_t^d|H_t^d, V_d^H) = \Pr(NT_t^d|H_t^d, V_d^L) = (1-\mu)(1-\varepsilon).$$

As we illustrated above, after a prevalence of buys, the expectation of a trader with a negative signal can be higher than the ask price: in such a case, all informed traders will buy, independently of their signal. In equilibrium, the market maker will set the ask and the bid taking into account that all informed traders want to buy, i.e., that there is herd buying. In this case, the probability of observing a buy, a sell or a no trade at that time are

$$\begin{aligned}
\Pr(Buy_t^d|H_t^d, V_d^H) &= (Buy_t^d|H_t^d, V_d^L) = \mu + (1-\mu)\frac{\varepsilon}{2}, \\
\Pr(Sell_t^d|H_t^d, V_d^H) &= \Pr(Sell_t^d|H_t^d, V_d^L) = (1-\mu)\frac{\varepsilon}{2}, \\
\Pr(NT_t^d|H_t^d, V_d^H) &= \Pr(NT_t^d|H_t^d, V_d^L) = (1-\mu)(1-\varepsilon).
\end{aligned}$$

Similarly, after a prevalence of sells, we will have $E(V|H_t, x^H) < Bid_t^d$: in such a circumstance there will be herd selling and the probabilities of each action will be

$$\begin{aligned}
\Pr(Buy_t^d|H_t^d, V_d^H) &= \Pr(Buy_t^d|H_t^d, V_d^L) = (1-\mu)\frac{\varepsilon}{2}, \\
\Pr(Sell_t^d|H_t^d, V_d^H) &= \Pr(Sell_t^d|H_t^d, V_d^L) = \mu + (1-\mu)\frac{\varepsilon}{2}, \\
\Pr(NT_t^d|H_t^d, V_d^H) &= \Pr(NT_t^d|H_t^d, V_d^L) = (1-\mu)(1-\varepsilon).
\end{aligned}$$

Finally, there are two intermediate cases. The first occurs when, after a positive trade imbalance, $Bid_t^d < E(V|H_t^d, x^L) < Ask_t^d$. In this case, in equilibrium, the market maker computes the bid assuming that only noise traders sell while agents receiving negative signal do not trade, and this is indeed what happens. In this instance, the probabilities of trades will be

$$\begin{gathered}
\Pr(Buy_t^d|H_t^d, V_d^H) = \mu q + (1-\mu)\frac{\varepsilon}{2},\ \Pr(Sell_t^d|H_t^d, V_d^H) = (1-\mu)\frac{\varepsilon}{2}, \\
\Pr(Buy_t^d|H_t^d, V_d^L) = \mu(1-q) + (1-\mu)\frac{\varepsilon}{2},\ \Pr(Sell_t^d|H_t^d, V_d^L) = (1-\mu)\frac{\varepsilon}{2}, \\
\Pr(NT_t^d|H_t, V_d^H) = (1-\mu)(1-\varepsilon) + \mu(1-q), \\
\Pr(NT_t^d|H_t, V_d^L) = (1-\mu)(1-\varepsilon) + \mu q.
\end{gathered}$$

Analogously, when after a negative trade imbalance, $P_t < E(V|H_t^d, x^H) < Ask_t^d$, in equilibrium the market maker sets the ask assuming that only noise traders buy, as it occurs. Therefore, we have

$$\Pr(Buy_t^d|H_t, V_d^H) = (1-\mu)\frac{\varepsilon}{2}, \Pr(Sell_t^d|H_t, V_d^H) = \mu(1-q) + (1-\mu)\frac{\varepsilon}{2}$$
$$\Pr(Buy_t^d|H_t, V_d^L) = (1-\mu)\frac{\varepsilon}{2}, \Pr(Sell_t^d|H_t, V_d^L) = \mu q + (1-\mu)\frac{\varepsilon}{2},$$
$$\Pr(NT_t^d|H_t, V_d^H) = (1-\mu)(1-\varepsilon) + \mu q,$$
$$\Pr(NT_t^d|H_t, V_d^L) = (1-\mu)(1-\varepsilon) + \mu(1-q).$$

If we define five indicator functions $I, I_{hb}, I_{hs}, I_{mb}, I_{ms}$ for each case above (i.e., I_{hb} takes value 1 if there is herd buy and 0 otherwise), we can write the probability of, for instance, a buy order in the case of a good event day as

$$\Pr(Buy_t^d|H_{t-1}^d, V_d^H) =$$
$$\left[\mu + (1-\mu)\frac{\varepsilon}{2}\right] I_{hb} + \left[(1-\mu)\frac{\varepsilon}{2}\right](I_{hs} + I_{mb}) + \left[\mu q + (1-\mu)\right](I + I_{ms})$$

We can write the probability of any other action for either a good event or a bad event day in a similar way. Given that the probability of a trade at any time depends on the evolution of beliefs until then, the probability of a history of trades must be computed as $Pr(H_t^d|V_d) = \Pi_{s=1}^t \Pr(h_s^d|H_s^d, V_d)$ (where h_s^d is an action at time t of day d), i.e., the probability of each trade depends on the previous history of trades.

To conclude our description, it is useful to remark the following point. Herd buying arises when all informed traders value the asset more than the ask price posted by the market maker, independently of the private information that they receive. Similarly, herd selling arises when they value it less than the bid. Therefore, one could believe that in order to compare the traders' and the market maker's beliefs and decide in which of the five cases illustrated above we are at any time t, we would need data on the magnitude of the good or bad news at the beginning of the day. This would definitely complicate our analysis, as we would need to estimate also Δ_d^H and Δ_d^L. However, we can easily show that this is not the case. The expected value of the asset for a trader at time t is

$$E(V_d^L|H_t^d, x) = V_d^L \Pr(V_d^L|H_t^d, x) + V_d^H \Pr(V_d^H|H_t^d, x) =$$
$$V_{t-1} - \Delta_d^L \Pr(V_d^L|H_t^d, x) + \Delta_d^H \Pr(V_d^H|H_t^d, x).$$

On the other hand, the expected value of the market maker is

$$E(V|H_t^d) = V_d^L \Pr(V_d^L|H_t^d) + V_{t-1} \Pr(V_{t-1}|H_t^d) + V_d^H \Pr(V_d^H|H_t^d) =$$
$$(V_{t-1} - \Delta_d^L) \Pr(V_d^L|H_t^d) + V_{t-1} \Pr(V_{t-1}|H_t^d) + (V_{t-1} + \Delta_d^H) \Pr(V_d^H|H_t^d) =$$
$$-\Delta_d^L \Pr(V_d^L|H_t^d) + V_{t-1} + \Delta_d^H \Pr(V_d^H|H_t^d),$$

Therefore, the difference between the two expectations is

$$\left[-\Delta_d^L \Pr(V_d^L|H_t^d) + V_{t-1} + \Delta_d^H \Pr(V_d^H|H_t^d)\right] -$$
$$\left[V_{t-1} - \Delta_d^L \Pr(V_d^L|H_t^d, x) + \Delta_d^H \Pr(V_d^H|H_t^d, x)\right]$$
$$= \left[-\Delta_d^L \Pr(V_d^L|H_t^d) + \Delta_d^H \Pr(V_d^H|H_t^d)\right] -$$
$$\left[-\Delta_d^L \Pr(V_d^L|H_t^d, x) + \Delta_d^H \Pr(V_d^H|H_t^d, x)\right]$$
$$= \left[-\frac{1-\delta}{\delta}\Delta_d^H \Pr(V_d^L|H_t^d) + \Delta_d^H \Pr(V_d^H|H_t^d)\right] -$$
$$\left[-\frac{1-\delta}{\delta}\Delta_d^H \Pr(V_d^L|H_t^d, x) + \Delta_d^H \Pr(V_d^H|H_t^d, x)\right] =$$
$$\Delta_d^H \left[\frac{1-\delta}{\delta}\left(\Pr(V_d^L|H_t^d, x) - \Pr(V_d^L|H_t^d)\right) - \Pr(V_d^H|H_t^d, x) + \Pr(V_d^H|H_t^d)\right]$$

whose sign is independent of how big the positive or negative news is.[5]

5 Data

The main interest of this paper is methodological. We are not interested in analyzing herding in a particular market or stock, rather in establishing an estimation strategy to study herding on the basis of a theoretical model. We use data for one stock, Ashland Oil. Our choice is exclusively motivated by convenience: this stock has already been used for similar analyses (see, e.g., Easley, Kiefer and 0'Hara, 1997, and Easley et al., 2001). We decided to be conservative in our study. We do not want to study herding in the trading of a peculiar stock, but just study the trading process of a standard stock.

We use transactions data on this stock for the first quarter of 1995. In this period there were 63 trading days. The probabilistic structure of our

[5] For simplicity we have considered just the "price" of the asset. It is straightforward to repeat the argument for the bid and the ask prices.

model is quite complicated. Potentially there are three different types of trading days. For a given type of day, there are many different sequences (not only numbers) of transactions. Each sequence can lead to herding in different moments of the day. Therefore, we needed to consider a fairly large trading-day window.

We took data for this stock from the TAQ dataset. This dataset reports a complete list of the posted prices (quotes), the price at which the transactions occurred (trade), the size of the transactions and, of course, the time when the quotes were posted and the transaction occurred.

In order to extract the relevant information from these data, we had to make several transformations. First, these data do not say who initiated the trade, i.e., whether the transaction was a sell or a buy. In order to classify a trade as a sell or buy, we had to compare the trade with the quotes that were posted at the time of the transaction. For this purpose, we used the standard algorithm proposed by Lee and Ready (1991). Given that quotes are reported with a delay, we used their suggestion of moving each quote behind in time of five seconds. For instance, a quote that results to have been posted exactly at 10am is recorded in our data as occurring at 5 seconds to 10am. We then compared the transaction price with the quotes that were posted just before the transaction occurred. In some cases, to decide whether a transaction was a buy or sell was easy: if it occurred at the ask it was unambiguously a buy; if it occurred at the bid, it was unambiguously a sell. Some transactions, however, are finalized at a price that is within the bid-ask spread. In this case, we classified a transaction as a buy if it occurred above the midpoint and as a sell if it occurred below it. If the price was exactly equal to the midpoint, we looked at whether the midpoint was higher or lower than the previous one. If there was an uptick, the order was classified as a buy, otherwise as a sell. If there was no change, then we looked at the previous price movement and so on.

Second, clearly these data do not contain any direct information on no trades. But, of course, there is a significant difference between an hour of trading in which there are 100 transactions and an hour in which there is none. We used the convention of classifying as a no trade any period of five minutes in which no transaction occurred. For instance, we considered a period of 20 minutes between two transactions as four no trades. The choice of five minutes is, of course, arbitrary, as it would be for any other integer. As a robustness check, we also used other alternative intervals.

As we said, we considered a period with 63 trading days. On average, we observed 100 transactions (either buys or sells) and 33 no trades a day.

6 Results

Table 1 shows the results of our estimation.

Table 1: Estimation Results

Parameter	Estimate	Standard deviation
α	0.79	0.057
δ	0.68	0.097
μ	0.24	0.018
ε	0.66	0.027
q	0.69	0.006

If the specification we adopted is correct, informational events are quite frequent: 79% of trading days are in fact classified as days in which some trading activity was motivated by private information. There is a certain imbalance between good and bad news. Good informational events account for almost 70% of the informed days. During informed days, the proportion of traders with private information is, on average, 24%. The remaining trading activity is explained by noise traders. Such noise traders almost equally split their decisions among buying, selling, and not trading (34%). The precision of private information is just below 70%. Such a precision of the signal strictly lower than 1 opens the door to herd behavior. On the basis of these parameters, we tracked down the beliefs of the traders (with a good or bad signal) and the beliefs of the market maker (i.e., the bid and ask prices) during each trading day. By comparing such beliefs we can detect periods in which, according to our model, there was herd behavior in the market. These are periods in which an informed trader would have made the same decision independently of the signal he received. For instance, when the belief of the trader is higher than the equilibrium ask even if he received a bad signal, we classify this period as herd buying. Similarly, when the belief of the trader is lower than the equilibrium bid even if he received a good signal, we classify this period as herd selling. We found that 7% of trading periods are of herd buying and 8% are of herd selling. It should be noted that these are periods in which there was potential herding. We cannot detect whether the trader really had information at odds with the public information conveyed by the history of trades. To check this would require data on private information that we do not have. Nevertheless, it is also important to remark that herding periods are relevant for the informational (in)efficiency of the market. Indeed, during herding periods, the market is unable to learn whether the traders received a good or a bad signal. Although the market maker still learns something from the trading activity (namely whether he is in an informed or uninformed day), information is aggregated less efficiently and the price converges more slowly to the fundamental asset value. For the period under analysis, we found that herding was pronounced in 23 days (out of 63). In such days, at least 15% of the trading periods were characterized by herding behavior. In 7 days, in particular, herding was very pronounced, since it characterized more than half of the trading periods.

Table 2: Number of days in which herding periods were frequent.

$> 50\%$	$> 30\%$	$> 15\%$
7	15	23

To have a better understanding of the inefficiency produced by herding behavior, we simulated $10,000$ days of trading in a financial market with our estimated parameters. We then simulated the same days of trading with the same parameters but assuming that traders, instead of behaving rationally as in our model, always followed private information. We took this as a benchmark case, since in this case all private information would be revealed by the trading decisions. We compared the price paths under the two scenarios. We considered the absolute difference at each time of very trading day. We found that the average distance between the two prices is 0.4%. In other words, the presence of herding determined a deviation of the price from the full information level of 0.4% on average during each day. While this is a fairly high distortion from the full information level, it should also be noted that such an information inefficiency is temporary only, since the price eventually converges to the fundamental.

7 Conclusion

We estimated a model of herding behavior in financial markets. We used transaction data for a stock traded in the NYSE. We estimated the parameters of the structural model and detected periods in each trading day in which, according to our model, informed traders chose the same action independently of whether they had good or bad private information on the value of the stock. We found that herding is present in the market. In some days of trade it is fairly pervasive. In our future research we will apply our methodology to verify whether herding is more pronounced in particular markets, or in particular times (like during financial crises).

References

[1] Avery, Christopher and Zemsky, Peter. "Multidimensional Uncertainty and Herd Behavior in Financial Markets." *American Economic Review*, September 1998, 88 (4), pp.724-748.

[2] Banerjee, Abhijit, "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, August 1992, 107 (3), pp. 787-818.

[3] Bikhchandani, Sushil; Hirshleifer, David and Welch Ivo. "A Theory of Fads, Fashion, Custom and Cultural Change As Informational Cascades." *Journal of Political Economy,* October 1992, 100 (5), pp. 992-1027.

[4] Bikhchandani, Sushil and Sharma Sunil. "Herd Behavior in Financial Markets: A Review." 2000, IMF Working paper.

[5] Brunnermeier, Markus. *Asset Pricing Under Asymmetric Information. Bubble, Crashes, Technical Analysis and Herding.* Oxford University Press, Oxford, 2001.

[6] Cipriani, Marco and Guarino Antonio. "Herd Behavior and Contagion in Financial Markets." mimeo, NYU, 2001.

[7] Cipriani, Marco and Guarino Antonio. "Herd Behavior in a Laboratory Financial Market." *American Economic Review*, December 2005.

[8] Easley David and Maureen O'Hara. "Time and the Process of Security Price Adjustment." *Journal of Finance*, 1992

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

[10] Easley D. and M. O'Hara, 1987, "Price, Trade Size, and Information in Securities Markets", *Journal of Financial Economics*, 19, 69-90.

[11] Easley David, Maureen O'Hara, Robert Engle and Liuren Wo, 2001, "Time-Varying Arrival Rates of Informed and Uninformed Trades." Working Paper.

[12] Gale, D., 1996, "What Have We Learned from Social Learning", *European Economic Review*, 40, 617-628.

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

[14] Grinblatt Mark, Sheridan Titman and Russ Wermers, "Momentum Investment Strategies, Portfolio Performance, and Herding: A Study of Mutual Fund Behavior." *American Economic Review*, December 1995

[15] Hirshleifer D. and S.H. Teoh, 2003, "Herd Behavior and Cascading in Capital Markets: a Review and Synthesis," *European Financial Management*, 9, 25-66.

[16] Lakonishok, Joseph, Shleifer Andrei and Vishny Robert W. "The Impact of Institutional Trading on Stock Prices," *Journal of Financial Economics*, 1992, 32, pp. 23-43.

[17] Lee, C., M. Ready, 1991. "Inferring Trade Direction from Intraday Data." *Journal of Finance 46*, 733-746.

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

[19] Smith, L. and P. Sørensen. "Pathological Outcomes of Observational Learning." *Econometrica, 2000, 68,* pp. *371-398.*

[20] Welch, Ivo. "Sequential Sales, Learning, and Cascades." *Journal of Finance*, 1992, 47, pp. 695-732.

[21] Welch, Ivo. "Herding Among Security Analysts." *Journal of Financial Economics* 58-3, December 2000, 369-396.