

Estimating a Structural Model of Herd Behavior in Financial Markets

Marco Cipriani and Antonio Guarino*

July 2010

Abstract

We develop a new methodology to estimate the importance of herd behavior in financial markets. We build a structural model of informational herding that can be estimated with financial transaction data. In the model, rational herd behavior arises because of information-event uncertainty. We estimate the model using transaction data on a NYSE stock (Ashland Inc.) during 1995. We then detect the periods in the trading days in which herd behavior occurs. Herd behavior often arises and is particularly pervasive in some days. The proportion of herd buyers (sellers) is 2% (4%) and is greater than 10% in 7% (11%) of information-event days. Herd behavior causes important informational inefficiencies: on average, the distance of the price from the one that would have prevailed had traders not herded is 4% of the expected asset value.

*Cipriani: Department of Economics, George Washington University (e-mail: marco.cipriani@gwu.edu) and IMF; Guarino: Department of Economics and ELSE, University College London (e-mail: aguarino@ucl.ac.uk). We thank Orazio Attanasio, Pedro Carneiro, Andrew Chesher, Robert Engle, Christopher Flinn, Ana Fostel, Douglas Gale, Toru Kitagawa, Lars Nesheim, Henri Pages, Nicola Pavoni, Jean Marc Robin and seminar participants in various institutions for very helpful comments and suggestions. We thank Silvia Camussi, Minyu Chen and Thibaut Lamadon for excellent research assistance. We gratefully acknowledge the financial support of the Fondation Banque de France, the ERC and the ESRC. This work was partially conducted while Guarino was a visiting scholar at the IMF; Guarino is grateful to the IMF for its hospitality. We are responsible for any error. The views expressed in this paper are only ours, and should not be attributed to the IMF, its Executive Board, or its management.

1 Introduction

In recent years there has been much interest in herd behavior in financial markets. This interest has led researchers to look for theoretical explanations and empirical evidence of herding. There has been, however, a substantial disconnect between the empirical and theoretical literatures: the theoretical work has identified motives of herding in abstract frameworks, developing models that cannot easily be brought to the data; the empirical literature has mainly looked for a-theoretical, statistical evidence of trade clustering, interpreted as herding.

The purpose of this paper is to propose a novel approach: we develop a theoretical model of herding in financial markets that can be brought to the data, and show how it can be estimated with financial markets transaction data. This methodology allows us to measure the quantitative importance of herding, to detect when it happens, and to assess the informational inefficiency that it generates.

The theoretical work on herd behavior started 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 environment, 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 disregard their own private information and imitate. More recently, a number of papers (see, among others, Avery and Zemsky, 1998; Lee, 1998; Cipriani and Guarino, 2008) have focused on herd behavior in financial markets. In particular, these studies analyze a market where informed and uninformed traders sequentially trade a security of unknown value. The price of the security is set by a market maker according to the order flow. The presence of a price mechanism makes herding more difficult to arise. Nevertheless, there are cases in which it occurs. In Avery and Zemsky (1998), for instance, herd behavior can occur when there is uncertainty not only on the value of the asset but also on the occurrence of an information event or on the model parameters.

As we mentioned above, whereas the theoretical research has tried to

¹We only study informational herding. Therefore, we do not discuss herd behavior due to reputational concerns or payoff externalities. For an early critical assessment of the literature on herd behavior see Gale (1996). For recent surveys of herding in financial markets see Bikhchandani and Sharma (2001), Vives (2008) and Hirshleifer and Teoh (2009).

identify the mechanisms through which herd behavior can arise, the empirical literature has followed a different track. The existing work (see, e.g., Lakonishok *et al.*, 1992; Grinblatt *et al.*, 1995; and Wermers, 1999) does not test the theoretical herding models directly, but analyzes the presence of herding in financial markets through statistical measures of clustering.² These papers find that in some markets fund managers tend to cluster their investment decisions more than if they acted independently. This empirical research on herding is important, as it sheds light on the behavior of financial market participants and in particular on whether they act in a coordinated fashion. As the authors themselves emphasize, however, decision 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, that is, the decision to disregard one's private information to follow the behavior of others (see Bikhchandani and Sharma, 2001; and Hirshleifer and Teoh, 2009).

Testing models of informational herd behavior is difficult. In these models, a trader herds if he trades against his own private information. The problem that empiricists face is that there are no data on the private information available to the traders and, therefore, it is difficult to understand whether traders decide not to follow it. Our purpose is to present a methodology to overcome this problem. We develop a theoretical model of herding and estimate it using financial market transaction data. We are able to identify the periods in the trading day in which traders act as herders and to measure the informational inefficiency that this generates. This is the first paper on informational herding that, instead of using a statistical, a-theoretical approach, brings a theoretical social learning model to the field data.³

Our theoretical analysis builds on the work of Avery and Zemsky (1998), who use a sequential trading model à la Glosten and Milgrom (1985) to show the conditions under which herding can arise in financial markets. They show that, if the only source of uncertainty is the asset's fundamental value, 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

²See also the recent paper by Dasgupta *et al.* (forthcoming), who study the effect of institutional herding on long term returns, and the literature cited therein.

³Whereas there are no direct empirical tests of herding models, there is experimental work that tests these models in the laboratory (see, e.g., Cipriani and Guarino 2005 and 2009; and Drehman *et al.*, 2005).

that agents neglect their information to imitate previous traders' decisions (and herd). In contrast, when there are multiple sources of asymmetric information between the traders and the market maker (e.g., asymmetric information on the asset's volatility) herd behavior may arise.⁴

In our model, herding arises for a mechanism similar to that exposed by Avery and Zemsky (1998). However, whereas they were interested in providing theoretical examples of herding, our aim is to provide an empirical methodology to gauge the importance of herding in actual financial markets. For this purpose, we build a model of herding that can be estimated with financial market transaction data. In the model, an asset is traded over many days; at the beginning of each day, an informational event may occur, in which case the fundamental asset value changes with respect to the previous day. If an informational event has occurred, some traders receive private information on the new asset value.⁵ These traders trade the asset to exploit their informational advantage on the market maker. In contrast, if no event has occurred, all traders in the market are noise, that is, they trade for non-information reasons only (liquidity or hedging motives). Whereas the informed traders know that they are in a market with private information (since they themselves are informed), the market maker does not. This asymmetry of information between the traders and the market maker implies that the market maker moves the price too "slowly" in order to take into account the possibility that the asset value may have not changed (in which case all trading activity is due to non-informational motives). As a result, after, for instance, a history of buys, a trader, even with a bad signal, may value the asset more than the market maker does. He will, therefore, trade against his own private information and herd buy.

We estimate the model with stock market transaction data via maximum likelihood, using a strategy first proposed by Easley *et al.* (1997). They show how to use transaction data to estimate the parameters of the Glosten and Milgrom (1985) model. There is an important difference, however, between their methodology and ours. In their set up, informed traders are perfectly informed about the value of the asset; as a result, the traders' decisions are never affected by the decisions of the previous traders, and they never herd. Therefore, only the total *number* of buys, sells and no trades in each

⁴A similar mechanism is also in Gervais (1997).

⁵The event is called informational precisely because some traders in the market receive private information on it.

day matters; *the sequence* in which these trades arrives is irrelevant. In contrast, in our framework, the precision of private information is one of the parameters that we estimate. This opens the possibility that informed traders may receive noisy signals, and that they may find it optimal to ignore them and engage in herd behavior. In this circumstance, the sequence by which trades arrive in the market does matter: in contrast to Easley *et. al.* (1997), we cannot estimate our model only from the number of buy or sell orders in a given day, but we must consider the whole history of trading activity in each day of trading.

As an illustration of the methodology, we estimate the model using transaction data on a NYSE stock (Ashland Inc.) during 1995. The restriction that private signals are perfectly precise is rejected by the data, which implies both that herd behavior arises in equilibrium and that there is information content in the sequence of trades. In particular, we find that informed traders receive incorrect information 40% of the time.

This has important consequences on the estimates of trading informativeness. A large literature has studied the information content of the trading activity using a measure (i.e., the PIN) based on the Easley *et. al.* (1997) methodology. Using that methodology, the measure of information-based activity in our sample would be 9%. Using our methodology, instead, we obtain 19%. The difference is due to the fact that in the previous literature incorrect trades (e.g., sell in a good-event day) can only be due to exogenous, non-informative (e.g., liquidity) reasons, whereas in our setup we do not exclude that they may be due to informed traders who either receive the incorrect information or herd.

Given our estimated parameters, we study how the traders' beliefs evolve during each day of trading. By comparing these beliefs to the prices, we are able to detect the periods of the trading day in which traders herd, and find that in most of the trading periods, a positive (albeit small) measure of informed traders herd. In an information-event day, on average, between 2% (4%) of informed traders herd buy (sell).

Herd behavior generates serial dependence in trading pattern, a phenomenon documented in the empirical literature. Herding also causes informational inefficiencies in the market. On average, the misalignment between the price we observe and that we would observe in the absence of herding is equal to 4% of the asset unconditional fundamental value.

The paper is organized as follows. Section 2 describes the theoretical model. Section 3 presents the likelihood function. Section 4 describes the

data. Section 5 presents the parameter estimates. Section 6 describes herding behavior and discusses informational inefficiency, and Section 7 concludes. An Appendix contains the proofs.

2 The Model

Following Easley and O’Hara (1987), we generalize the original Glosten and Milgrom (1985) model to an economy where trading happens over many days.

An asset is traded by a sequence of traders who interact with a market maker. Trade occurs over many trading days, indexed by $d = 1, 2, 3, \dots$. Time within each day is discrete and indexed by $t = 1, 2, 3, \dots$

The asset

We denote the fundamental value of the asset in day d by V_d . The asset value does not change during the day, but can change from one day to the other. At the beginning of the day, with probability $1 - \alpha$ the asset value remains the same as in the previous day ($V_d = v_{d-1}$), and with probability α it changes.⁶ In this case, since, as we will see, there are informed traders in the market, we say that an information event has occurred. If an information event occurs, with probability $1 - \delta$ the asset value decreases to $v_{d-1} - \lambda^L$ (“bad informational event”), and with probability δ it increases to $v_{d-1} + \lambda^H$ (“good informational event”), where $\lambda^L > 0$ and $\lambda^H > 0$. Informational events are independently distributed over the days of trading. To simplify the notation, we define $v_d^H := v_{d-1} + \lambda^H$ and $v_d^L := v_{d-1} - \lambda^L$. Finally, we assume that $(1 - \delta)\lambda^L = \delta\lambda^H$, which, as will become clear in the next pages, guarantees that the closing price is a martingale.

The market

The asset is exchanged in a specialist market. Its price is set by a market maker who interacts with a sequence of traders. At any time $t = 1, 2, 3, \dots$ 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. The trader’s action space is, therefore, $\mathcal{A} = \{buy, sell, no\ trade\}$. We denote the action of the trader at time t in day d by X_t^d . Moreover, we denote the history of trades and prices until time $t - 1$ of day d by H_t^d .

⁶Note that v_{d-1} is the realization of the random variable V_{d-1} . Throughout the text we will denote random variables with capital letters and their realizations with lower case letters.

The market maker

At any time t of day d , the market maker sets the prices at which a trader can buy or sell the asset. When posting these prices, he must take into account the possibility of trading with traders who (as we shall see) have some private information on the asset value. He will set different prices for buying and for selling, that is, there will be a bid-ask spread (Glosten and Milgrom, 1985). We denote the ask price (the price at which a trader can buy) at time t by a_t^d and the bid price (the price at which a trader can sell) by b_t^d .

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

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

The traders

There are a countable number of traders. Traders act in an exogenous sequential order. Each trader is chosen to take an action only once, at time t of day d . Traders are of two types, informed and noise. The trader's own type is private information.

In no-event days, all traders in the market are noise. In contrast, in information-event days, at any time t an informed trader is chosen to trade with probability μ and a noise trader with probability $1 - \mu$, with $\mu \in (0, 1)$.

Noise traders trade for unmodeled (e.g., liquidity) reasons: they buy with probability $\frac{\varepsilon}{2}$, sell with probability $\frac{\varepsilon}{2}$ and do not trade with probability $1 - \varepsilon$ (with $0 < \varepsilon < 1$). Informed traders have private information on the asset value. They receive a private signal on the new asset value and observe the previous history of trades and prices, and the current prices. The private signal S_t^d has the following value-contingent densities:

$$g^H(s_t^d | v_d^H) = 1 + \tau(2s_t^d - 1),$$

$$g^L(s_t^d | v_d^L) = 1 - \tau(2s_t^d - 1),$$

with $\tau \in (0, \infty)$. (See Figure 1.)

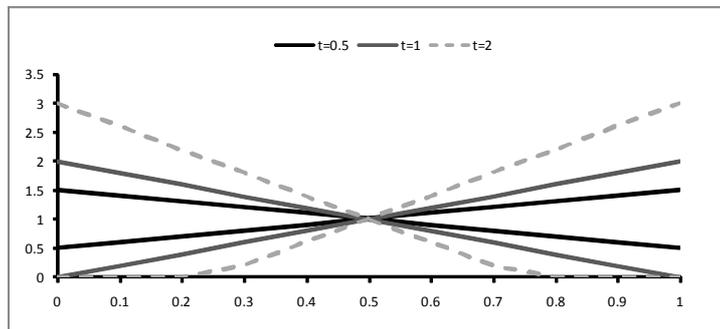


Figure 1: Signal state-contingent density functions for different values of τ .

For $\tau \in (0, 1]$, the support of the densities is $[0, 1]$. In contrast, for $\tau > 1$, the support shrinks to $[\frac{\tau-1}{2\tau}, \frac{\tau-1+2\sqrt{\tau}}{2\tau}]$ for g^H and to $[\frac{\tau+1-2\sqrt{\tau}}{2\tau}, \frac{\tau+1}{2\tau}]$ for g^L (in order for the density functions to integrate to one). Note that, given the value of the asset, the signals S_t^d are i.i.d. The signals satisfies the monotone likelihood ratio property. At each time t , the likelihood ratio after receiving the signal $\frac{\Pr(V_d=v_d^H|h_t^d, s_t^d)}{\Pr(V_d=v_d^L|h_t^d, s_t^d)} = \frac{g^H(s_t^d|v_d^H)}{g^L(s_t^d|v_d^L)} \frac{\Pr(V_d=v_d^H|h_t^d)}{\Pr(V_d=v_d^L|h_t^d)}$ is higher than that before receiving the signal if $s_t^d > 0.5$, and lower if $s_t^d < 0.5$. For this reason we refer to a signal larger than 0.5 as a “good signal” and to a signal smaller than 0.5 as a “bad signal.”

The parameter τ measures the informativeness of the signals. When $\tau \rightarrow 0$, the densities are uniform and the signals are completely uninformative. As τ increases, the signals become more and more informative. For $\tau \in [0, 1)$, the support of the distribution of the likelihood ratio is bounded away from 0 and infinity, while for $\tau \geq 1$ it is not. Following Smith and Sørensen (2000), in the first case we say that beliefs are bounded, and in the second case, that they are unbounded. With bounded beliefs, no signal realizations (even the most extreme ones) reveal the asset value with probability one. With unbounded beliefs, in contrast, some high (low) signal realizations are only possible when the asset value is high (low), and therefore, the signal can be perfectly informative.⁷ As τ tends to infinity, the measure of perfectly informative signals tends to one.

An informed trader knows that an information event has occurred, and

⁷In particular, any signal greater than or equal to $\frac{\tau+1}{2\tau}$ reveals that the asset value is v_d^H , whereas a signal lower than or equal to $\frac{\tau-1}{2\tau}$ reveals that the asset value is v_d^L .

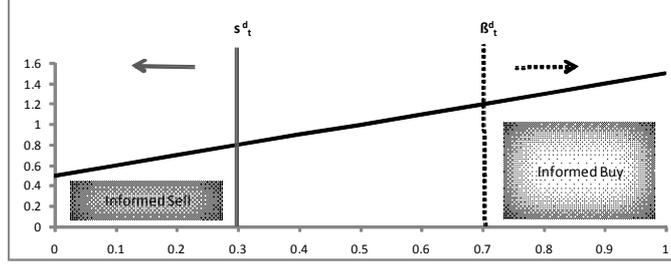


Figure 2: Informed trader's decision. The figure illustrates the signal realizations for which an informed trader decides to buy or sell when $V_d = v_d^H$ (the signal density function is conditional on v_d^H).

that, as a result the asset value has changed with respect to the previous day. Moreover, his signal is informative on whether the event is good or bad. Nevertheless, according to the signal realization that he receives and to the precision τ , he may not be completely sure of the effect of the event on the asset value. For instance, he may know that there has been a change in the investment strategy of a company, but not be sure of whether this change will affect the asset value in a positive or negative way. The parameter τ can be interpreted as measuring the precision of the information that the trader receives, or the ability of the trader to process such private information. Finally, note that, given our signal structure, informed traders are heterogenous, since they receive signal realizations with different degrees of informativeness about the asset fundamental value.

An informed trader's payoff function, $U : \{v_d^L, v_d^H\} \times \mathcal{A} \times [v_d^L, v_d^H]^2 \rightarrow \mathbf{R}^+$, is defined as

$$U(v_d, X_t^d, a_t^d, b_t^d) = \begin{cases} v_d - a_t^d & \text{if } X_t^d = \text{buy}, \\ 0 & \text{if } X_t^d = \text{no trade}, \\ b_t^d - v_d & \text{if } X_t^d = \text{sell}. \end{cases}$$

An informed trader chooses X_t^d to maximize $E(U(V_d, X_t^d, a_t^d, b_t^d) | h_t^d, s_t^d)$ (i.e., he is risk neutral). Therefore, he finds it optimal to buy whenever $E(V_d | h_t^d, s_t^d) \geq a_t^d$, and sell whenever $E(V_d | h_t^d, s_t^d) \leq b_t^d$. He chooses not to trade when $b_t^d < E(V_d | h_t^d, s_t^d) < a_t^d$.

It is important to remark that at each time t , the trading decision of an informed trader can be simply characterized by two thresholds, σ_t^d and β_t^d , satisfying the equalities

$$E [V_t^d | h_t^d, \sigma_t^d] = b_t^d$$

and

$$E [V_t^d | h_t^d, \beta_t^d] = a_t^d.$$

An informed trader will sell for any signal realization smaller than σ_t^d and buy for any signal realization greater than β_t^d . Obviously, the thresholds at each time t depend on the history of trades until that time and on the parameter values.⁸

In Figure 2 (drawn for the case of a good informational event) we illustrate the decision of informed traders. An informed trader buys the asset with a signal higher than the threshold value β_t^d , sells it with a signal lower than σ_t^d , and does not trade otherwise. The measure of informed traders buying or selling is equal to the areas (labelled as “Informed Buy” and “Informed Sell”) below the line representing the signal density function.

2.1 Herd Behavior

To discuss herd behavior, let us start by introducing some formal definitions.

Definition An informed trader engages in herd buying at time t of day d when he buys upon receiving a bad signal, that is,

$$E(V_d | h_t^d, s_t^d) > a_t^d \text{ for } s_t^d < 0.5.$$

Similarly, an informed trader engages in herd selling at time t of day d when he sells upon receiving a good signal, that is,

$$E(V_d | h_t^d, s_t^d) < b_t^d \text{ for } s_t^d > 0.5.$$

In simple words, a trader herds when he trades against his own private information. Since traders in our model receive different signals, it may well be (and typically it will be the case) that, at a given point in time, traders with less informative signals (i.e., close to 0.5) will herd, whereas traders

⁸Since noise traders buy and sell with probabilities bounded away from zero, standard arguments prove that both the bid and ask prices, and the informed traders’ signal thresholds exist and are unique.

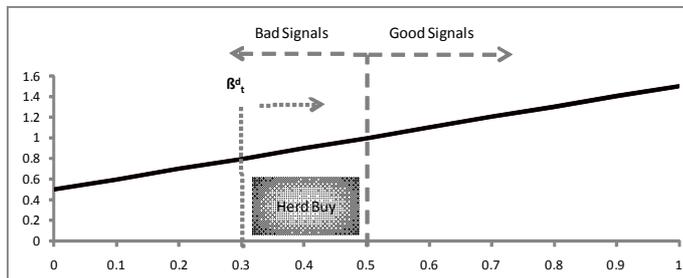


Figure 3: Herd buy. In the figure, an informed trader buys even upon receiving a bad signal (higher than 0.3).

with more informative signals (close to the extremes of the support) will not. We are interested in periods of the trading day in which traders engage in herd behavior for at least some signal realizations. At any given time t , we can detect whether an informed trader herds for a positive measure of signals by comparing the two thresholds σ_t^d and β_t^d to 0.5. Since a trader engages in herd buying behavior if he buys despite a bad signal ($s_t^d < 0.5$), there is a positive measure of herd buyers whenever $\beta_t^d < 0.5$.⁹ Similarly for herd sellers. Here is a formal definition:

Definition There is herd behavior at time t of day d when there is a positive measure of signal realizations for which an informed trader either herd buys or herd sells, that is, when

$$\beta_t^d < 0.5 \text{ or } \sigma_t^d > 0.5.$$

Figures 3 and 4 show an example of herd buy (3) and herd sell (4) in a day with a good information event. The areas below the signal density function and between the thresholds and 0.5 represent the measures of informed traders who herd buy and herd sell.

The reason why herd behavior arise is that prices move “too slowly” as buy and sell orders arrive in the market. Suppose that, at the beginning of an information-event day, there is a sequence of buy orders. Informed traders, knowing that there has been an information event, attach a certain

⁹We identify an informed trader with the signal he receives: we say “a positive measure of herd buyers” to mean “a positive measure of signal realizations for which an informed trader herd buys.”

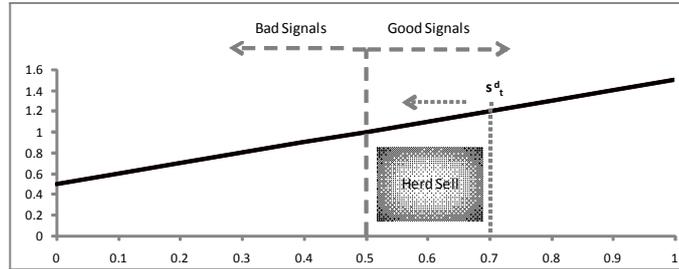


Figure 4: Herd sell. In the figure, an informed trader buys even upon receiving a good signal lower than 0.3.

probability to the fact that these orders come from informed traders with good signals. The market maker, however, attaches a lower probability to this event, as he takes into account the possibility that there was no event, and that all the buys came from noise traders. Therefore, after a sequence of buys, he will update the prices up, but by less than the movement in traders' expectations. Because traders and market maker interpret the history of trades differently, the expectation of a trader with a bad signal may be higher than the ask price, in which case he herd buys. Obviously, traders who receive signals close to 0.5 will be more likely to herd, since the history of trades has more weight in forming their beliefs.

We state this result in the next proposition:

Proposition For any finite τ , herd behavior arises with positive probability.

Furthermore, herd behavior can be misdirected, that is, an informed trader can engage in herd buy (sell) in a day of bad (good) information event.

Proof See the Appendix.

Avery and Zemsky (1998) have shown how herding can arise because of uncertainty on whether an information event has occurred (see their IS2 information setup). In our model herding arises for the same reason. Our contribution is to embed this theoretical reason to herd in a model that is suitable to empirical analysis.

When $\tau > 1$, extreme signals reveal the true value of the asset, and traders receiving them never herd. In the limit case of τ tending to infinity,

all signal realizations become perfectly informative, with the result that no informed trader herds. Therefore, while our model allows for herd behavior, it also allows for the possibility that some traders (when $\tau > 1$) or all traders (when $\tau \rightarrow \infty$) only rely on their private information and never herd.

The probability of herding depends on the parameter values. To take an extreme example, when α is arbitrarily close to zero, the market maker has a very strong prior that there is no information event. He barely updates the prices as trades arrive in the market, and herding arises as soon there is an imbalance in the order flow, as it happens in the seminal model of Bikhchandani *et al.* (1992). In contrast, if α is close to 1, the market maker and the informed traders update their beliefs in very similar manners, and herding rarely occurs.

Herding is important also for the informational efficiency of the market. During periods of herd behavior, private information is aggregated less efficiently by the price as informed traders with good and bad signals may take the same action. The most extreme case is when traders herd for all signal realizations (e.g., traders herd buy even for $s_t^d = 0$). In such a case, the market maker is unable to make any inference on the signal realization from the trades. The market maker, however, updates his belief on the asset value, since the action remains informative on whether an information event has occurred or not.¹⁰ Since the market maker never stops learning, he gradually starts interpreting the history of past trades more and more similarly to the traders and, as a result, the measure of herders shrinks.

During an information-event day, the measure of herders changes with the sequence of trades, and can become positive more than once at different times of the day. Given that information always flows to the market, however, the bid and ask prices converge to the asset value almost surely.¹¹ Eventually the market maker learns whether a good, bad or no event occurred.¹²

¹⁰The market maker learns since in periods of herding, the proportion of buys and sells is different from that in an uninformed day. Essentially, whereas in our model there is herd behavior, there is never an informational cascade.

¹¹The proof of convergence is standard and we omit it.

¹²Recall that we have assumed that $(1 - \delta)\lambda^L = \delta\lambda^H$. This implies that $E(V_{d+1}|V_d = v_d) = v_d$. Since the price converges to the fundamental value almost surely, this guarantees that the martingale property of prices is satisfied.

3 The Likelihood Function

To estimate the herding model presented above, we have to specify its likelihood function. Let us denote the history of trades at the end of a trading day by $h^d := h_{T_d}^d$, where T_d is the number of trading dates in day d . We denote the likelihood function by

$$\mathcal{L}(\Phi; \{h^d\}_{d=1}^D) = \Pr(\{h^d\}_{d=1}^D | \Phi),$$

where $\Phi := \{\alpha, \delta, \mu, \tau, \varepsilon\}$ is the vector of parameters.

Note that we write the likelihood function for the history of trades only, disregarding bid and ask prices. In our model there is no public information: for this reason, there is a one-to-one mapping from trades to prices, and adding prices would be redundant.

The one-to-one mapping from trades to prices breaks down in the presence of public information, since price changes may be the result of public information arrival (as opposed to being only determined by the order flow). Nevertheless, our likelihood function for the history of trades would still be correctly specified. The reason is that the probability of any given trade only depends on whether the trader is informed, and, in such a case, on whether his belief is higher or lower than the market maker's; neither event is affected by the arrival of public information (since this would affect traders' and market maker's beliefs in the same way, shifting all beliefs by the same amount).

Remember that information events are assumed to be independent. Moreover, the probability of the sequence of trades in a day only depends on the value of the asset that day. For this reason, the likelihood of a history of trades over multiple days can be written as the product of the likelihoods of the histories of trades for each day:

$$\mathcal{L}(\Phi; \{h^d\}_{d=1}^D) = \Pr(\{h^d\}_{d=1}^D | \Phi) = \prod_{d=1}^D \Pr(h^d | \Phi).$$

Let us focus on the probability of a history of trades in a single day. As we have written, the *sequence* of trades, and not just the *number* of trades, conveys information. Having many buy orders at the beginning of the day is not equivalent to having the same number of buy orders spread out during the day. In fact, a particular sequence of buy or sell orders may create herd behavior: in periods of herding, the probability of a trade depends on the measure of informed traders who herd and is different from the probability

in the absence of herding. Therefore, we have to compute the probability of a history of trades recursively, that is,

$$\Pr(h_t^d|\Phi) = \prod_{s=1}^t \Pr(x_s^d|h_s^d, \Phi),$$

where the probability of an action at time t of day d , $\Pr(x_t^d|h_t^d, \Phi)$, depends on the measure of informed traders who buy, sell or do not trade after a given history of trades h_t^d .

Using the law of total probability, at each time t , we compute $\Pr(x_t^d|h_t^d, \Phi)$ in the following way:

$$\begin{aligned} \Pr(x_t^d|h_t^d, \Phi) &= \Pr(x_t^d|h_t^d, V_d = v_d^H, \Phi) \Pr(V_d = v_d^H|h_t^d, \Phi) + \\ &\Pr(x_t^d|h_t^d, V_d = v_d^L, \Phi) \Pr(V_d = v_d^L|h_t^d, \Phi) + \Pr(x_t^d|h_t^d, V_d = v_{d-1}, \Phi) \Pr(V_d = v_{d-1}|h_t^d, \Phi). \end{aligned}$$

Let us consider, first, the probability of an action conditional on a good-event day. As illustrated above, at each time t , in equilibrium there is a signal threshold β_t^d such that an informed trader buys for any signal realization greater than β_t^d , that is,

$$E(V_d|h_t^d, \beta_t^d) = a_t^d = E(V_d|h_t^d, X_t^d = buy, a_t^d, b_t^d),$$

which can be written as

$$\begin{aligned} v_{d-1} + \lambda^H \Pr(v_d^H|h_t^d, \beta_t^d) - \lambda^L \Pr(v_d^L|h_t^d, \beta_t^d) = \\ v_{d-1} + \lambda^H \Pr(v_d^H|h_t^d, buy_t^d) - \lambda^L \Pr(v_d^L|h_t^d, buy_t^d), \end{aligned}$$

or, after some manipulations, as¹³

$$\Pr(v_d^H|h_t^d, \beta_t^d) - \Pr(v_d^H|h_t^d, buy_t^d) = \frac{\delta}{1 - \delta} (\Pr(v_d^L|h_t^d, \beta_t^d) - \Pr(v_d^L|h_t^d, buy_t^d)).$$

The probabilities in this equation can easily be expressed as a function of the traders' and market maker's beliefs a time $t - 1$ and of the parameters. (See the Appendix for further details.)

¹³Note that for simplicity's sake in the probabilities to compute the ask we have omitted a_t^d and b_t^d in the conditioning. More importantly, note that the magnitude of the shocks that buffet the asset's value (λ^L and λ^H) do not appear in this equation (they cancel out since $(1 - \delta)\lambda^L = \delta\lambda^H$). This is important, since it implies that we do not need to estimate the magnitude of the shocks that hit the asset value.

Note that at time $t = 1$, the prior beliefs of the traders and of the market maker are a function of the parameters only. Therefore, we can easily compute β_1^d as the solution to the equation above. After observing x_1^d , we can update the market maker's and traders' beliefs, repeat the same procedure for time 2 and compute β_2^d .¹⁴ We do so recursively for each time t , always conditioning on the previous history of trades.

From β_t^d , we can compute the probability of a buy order in a good-event day. For the sake of exposition, let us focus on the case in which $\tau \in [0, 1]$, that is, let us concentrate on the case of bounded beliefs. In this case:

$$\begin{aligned} \Pr(\text{buy}_t^d | h_t^d, V_d = v_d^H, \Phi) = \\ \mu \int_{\beta_t^d}^1 (1 + \tau(2s_t^d - 1)) ds_t^d + (1 - \mu) \left(\frac{\varepsilon}{2}\right) = \\ \left(\left(\tau(1 - \beta_t^{d2}) + (1 - \tau)(1 - \beta_t^d) \right) \mu + (1 - \mu) \left(\frac{\varepsilon}{2}\right) \right). \end{aligned}$$

We use a similar procedure to compute the probability of a sell, that is,

$$\begin{aligned} \Pr(\text{sell}_t^d | h_t^d, V_d = v_d^H, \Phi) = \\ \mu \int_0^{\sigma_t^d} (1 + \tau(2s_t^d - 1)) ds_t^d + \left(\frac{\varepsilon}{2}\right) (1 - \mu) = \\ \left(\left((1 - \tau)\sigma_t^d + \tau\sigma_t^{d2} \right) \mu + \left(\frac{\varepsilon}{2}\right) (1 - \mu) \right). \end{aligned}$$

Finally, the probability of a no trade is simply the complementary to the probabilities of a buy and a sell.

The analysis for the case of a bad information event ($V_d = v_d^L$) follows the same steps. Also the case of unbounded beliefs, where $\tau \geq 1$, can be dealt with in a similar manner. The only change are the extremes of integration when computing the probability of a trade.

The case of a no-event day ($V_d = v_{d-1}$) is easy, since the probabilities of a buy or sell is $\frac{\varepsilon}{2}$ and the probability of a no-trade is $1 - \varepsilon$.

Finally, to compute $\Pr(x_t^d | h_t^d, \Phi)$, we need the conditional probabilities of V_d given the history until time t , that is, $\Pr(V_d = v | h_t^d, \Phi)$ for $v = v_d^L, v_{d-1}, v_d^H$. These can also be computed recursively by using Bayes's rule. This completes the description of the likelihood function.

¹⁴To update the belief of traders and market maker after x_1^d we also need to know how to compute the probability of a sell and of a no trade, which we discuss below.

To conclude, let us give an intuition of how the model is identified. For simplicity’s sake let us consider only the number of buys, sells and no trades in each day.¹⁵ Similarly to analogous structural models of market microstructure, our model classifies days into high-volume days with a prevalence of buys (“good event” days), high-volume days with a prevalence of sells (“bad event” days) and low-volume days (“no event” days). The parameter α defines the probability that there is an event at the beginning of a trading day. We use data over many days of trading to identify it. The direction of the imbalance in the event days identifies δ . No-event days allow us to identify ε , since in no-event days only noise traders trade. Finally, in good event days, the ratio between buys and sells is determined by the proportion of traders who trade in the right direction (i.e., buy when there is a good event), which depends on μ and τ . An analogous argument holds for bad-event days. To any given estimate of μ and τ corresponds only one predicted ratio between buys and sells in the two types of days.¹⁶

4 Data

The aim of our study is mainly methodological. We are not trying to gauge the importance of herding across different stocks or different markets. Our goal is to develop a methodology to carry out a structural estimation of herding based on a market microstructure model. For this reason, we performed our empirical analysis on a single stock that had been studied before. Our choice fell on Ashland Inc., a stock traded in the New York Stock Exchange and used in Easley *et al.*, 1997.¹⁷ We obtained the data from the TAQ (Trades and Quotes) dataset.¹⁸ The dataset reports a complete listing of the posted bid and ask prices (the “quotes”), the prices at which the transactions occurred (the “trades”), and the time when the quotes were posted or the transactions occurred. We use transactions data on Ashland Inc. in 1995, for a total of 252 trading days. The data refer to trading in the New York Stock Exchange, the American Stock Exchanges, and the consolidated

¹⁵In our estimation we use much more information than that, since we take into account the entire sequence of trades when constructing the likelihood function.

¹⁶For a further argument for identification, see footnote 21.

¹⁷The name of the stock is slightly different, since the company changed name in 1995, and Easley *et al.* (1997) use 1990 data.

¹⁸Hasbrouck (2004) provides a detailed description of this dataset.

regional exchanges.

The TAQ dataset does not sign the trades, that is, does not report whether a transaction was a sale or a purchase. To classify a trade as a sell or a buy order, we used the standard algorithm proposed by Lee and Ready (1991). We compared the transaction price with the quotes that were posted just before a trade occurred.¹⁹ Every trade above the midpoint was classified as a buy order, and every trade below the midpoint was classified as a sell order; trades at the midpoint were classified as a buy or a sell order according to whether the transaction price had increased (uptick) or decreased (downtick) with respect to the previous one. If there was no change in the transaction price, we looked at the previous price movement and so on.²⁰

Moreover, TAQ data do not contain any direct information on no trades. We used the established convention of inserting no-trades between two transactions if the time elapsed between them was longer than a particular time interval (see, e.g., Easley *et al.*, 1997). Specifically, we obtained this interval by computing the ratio between the total trading time in a day and the average number of buy and sell trades over the 252 days (see, e.g., Chung *et al.*, 2004). In our 252 trading day window, the average number of trades per day was 90.2. We divided the total daily trading time (390 minutes) by 90.2, and obtained a unit-time interval of 259 seconds (i.e., on average, a trade occurred every 259 seconds). If there was not any trading activity for 259 seconds or more, we added one or more no trades to the sequence of buy and sell orders. The number of no-trades that we added between two consecutive transactions was equal to the number of 259-seconds time intervals elapsing between them. Finally, to check the robustness of our results, we also replicated the analysis for other no-trade time intervals (2, 3, 4, 5, 6 and 7 minutes).

Our sample of 252 trading days contained on average 149 decisions (either a buy, or a sell or a no-trade) per day. The sample was balanced, with 30% of buys, 31% of sells and 40% of no-trades.

¹⁹Given that transaction prices are reported with a delay, we followed Lee and Ready (1991)'s suggestion of moving each quote ahead in time of five seconds. Moreover, following Hasbrouck (1991, p. 581), we ignore quotes posted by the regional exchanges.

²⁰We classified all trades, with the exception of the opening trades, since these trades result from a trade mechanism (an auction) substantially different from the mechanism of trade during the day (which is the focus of our analysis).

5 Results

In presenting the results, we first show the estimates of the model parameters, and then illustrate the importance of herd behavior in the trading activity.

5.1 Estimates

We estimated the parameters through maximum likelihood. We used both a direct search method (Nelder-Mead simplex) and the Genetic Algorithm.²¹ The two methods converged to the same parameter values. Table 1 presents the estimates and the standard deviations for the five parameters of the model.²²

Parameter	Estimate	S.D.
α	0.28	0.03
δ	0.62	0.06
μ	0.42	0.01
τ	0.45	0.02
ε	0.57	0.00

Table 1: Estimation Results.

The table shows the estimates for the five parameters of the model and their standard deviations.

Information events are relatively frequent: from the estimate of α , we infer that the probability of an information event is 28%, that is, in almost a third of trading days trading activity is motivated by private information. There is a small imbalance between good and bad-event days: the probability of a good information event is 62% (although the parameter has a relatively high standard deviation).²³ During event days, the proportion of traders with

²¹We also simulated the theoretical model and verified that we could recover the model's parameters. Both methods converged to the true parameter values, which provides further evidence in favor of identification.

²²The standard deviation for ε is 0.00172 (rounded to 0.00 in the table). The same comment also applies to Table 2 and Table 3.

²³Note that δ is greater than 0.5, although in the sample the number of buys and sells is essentially balanced. This happens because among the days with high trading volume (classified as event days), a higher number of days have a positive trade imbalance than a negative one. To see this, consider the posterior beliefs of δ and α at the end of each day. In 22% of days the posterior belief of both alpha and delta is above 0.5 (i.e., we are

private information is 42%. The remaining trading activity comes from noise traders, who trade 57% of the time. Moreover, private information is noisy (that is, it is not perfectly informative). The estimate for τ is 0.45, which means that the probability of receiving an “incorrect signal” — i.e., a signal below 0.5 when we are in a good-information event day or a signal above 0.5 when we are in a bad-information event day — is 39%.²⁴

As we explained above, we constructed our dataset adding a no-trade each 259 seconds of trading inactivity. As a robustness check, we repeated the estimation on several other datasets, where we added a no trade for different intervals of trading inactivity. We report these estimates in Table 2.

	NT=120	S.D.	NT=180	S.D.	NT=240	S.D.
α	0.20	0.02	0.25	0.02	0.30	0.03
δ	0.72	0.02	0.67	0.04	0.61	0.03
μ	0.26	0.01	0.36	0.01	0.41	0.01
τ	0.44	0.01	0.40	0.03	0.44	0.03
ε	0.33	0.00	0.45	0.00	0.54	0.00
β	0.64		0.63		0.62	
Γ	0.32		0.33		0.34	
	NT=300	S.D.	NT=360	S.D.	NT=420	S.D.
α	0.30	0.03	0.27	0.03	0.27	0.04
δ	0.60	0.03	0.61	0.03	0.55	0.07
μ	0.40	0.01	0.39	0.01	0.37	0.01
τ	0.52	0.03	0.59	0.03	0.67	0.01
ε	0.62	0.00	0.69	0.00	0.74	0.00
β	0.62		0.62		0.62	
Γ	0.32		0.31		0.30	

Table 2: Robustness Checks for Different No-trade Intervals.

The table shows the estimates for various no-trade intervals, from 2 to 7 minutes.

The last two rows report two more statistics derived from the estimated parameters and explained in the text.

The estimates of the probability of an information event (α) and of a good in an good event day), whereas in only 12% of days the posterior belief of alpha is above 0.5 and that of delta is below 0.5 (i.e., we are in a bad event day).

²⁴Given the signal density functions, the probability of an incorrect signal is given by $0.5 - 0.25\tau$.

event (δ) are fairly similar over the different numbers of seconds defining a no-trade interval. The estimate of ε increases with the the no-trade interval: this is expected since the number of no-trades in the sample (and also in the no-event days) become smaller and smaller. To have a better description of the trading activity in no-event days, we computed the probability of observing at least one trade during a 5-minute interval in a no-event day: $\beta = 1 - (1 - \varepsilon)^{\frac{300}{\text{Seconds}}}$ (where “Seconds” is the no-trade interval). Table 2 shows this probability to be independent of the choice of the no-trade interval.

The parameter μ is quite stable across samples, whereas τ increases. To understand this, it is useful to observe that if both τ and μ were constant, as ε increases the estimated proportion of trading activity due to traders not having correct information (either because they are noise or because their signal is incorrect) would increase. In contrast, this proportion should obviously be independent of our choice of no-trade interval. This is indeed the case. To show this we computed the parameter

$$\Gamma = \frac{\mu(0.5 + 0.25\tau)}{(1 - \mu)\varepsilon + \mu},$$

which represents the proportion of correctly informed traders (e.g., informed traders with a signal greater than 0.5 in a good-event day) over the sum of all informed traders and the noise traders who trade. In other words, Γ is approximately equal to the fraction of trades coming from informed traders with the correct signal.²⁵ It is remarkable that Γ , which equals 0.34 when the no-trade interval is 259, is constant across all the different dataset that we used to estimate the model’s parameters. This shows the robustness of our results to the choice of the no-trade interval.

Let us now discuss how our results compare to different specifications of the model. A natural comparison is with a model in which the signal precision is not estimated, but is restricted to be perfectly informative (i.e., $\tau \rightarrow \infty$). This is the case studied by Easley *et al.* (1997). In this case, all informed traders follow their own private information, the sequence of trades has no informational content beyond the aggregate numbers of buys, sells and no trades, and herding never arises. As a result, the likelihood function does not need to be computed recursively (see Easley *et al.* (1997) for a detailed description). Table 3 presents the estimated parameters.

²⁵The approximation is due to the fact that we are ignoring that, because of the bid-ask spread, a small measure of informed traders may not trade. Easley *et al.* (1997) report a similar composite parameter when analyzing their results for different no-trade intervals.

Parameter	Estimate	S.D.
α	0.33	0.04
δ	0.60	0.06
μ	0.17	0.01
ε	0.58	0.00

Table 3: Parameter Estimates for the Easley et al. (1997) Model. The table shows the estimates for the four parameter model of Easley et al. (1997), in which informed traders know the true asset value. The no-trade interval is 259 seconds.

The estimates for α and δ are very close to those we obtain for our model. This shows that the classification of days is not affected by the specification of the signal structure. Similarly, the estimates for ε in the two models are almost the same. This is not surprising since ε captures the trading activity of noise traders, and is not affected by assumptions on the structure of private information. The parameter μ is instead quite smaller in the restricted model, which is intuitive, since in this model it is imposed that informed traders all receive the correct signal (i.e., they know whether a good or a bad information event occurred).

We can test whether the restriction in Easley *et al.* (1997) is rejected by the data. The likelihood Ratio test overwhelmingly rejects the hypothesis of perfectly informative signals, with a LR statistics of 272.15 (and a p-value of zero).²⁶ This is important for our aims, since the fact that signals are not perfectly informative implies that the sequence in the order flow matters. In other words, the number of buys, sells and no trades at the end of the day is not a sufficient statistics for the pattern of trading activity. Depending on the sequence, herd behavior by informed traders may occur in equilibrium.

In the market microstructure literature, a great deal of attention has been given to the PIN, a measure of the probability that a trade comes from an informed trader (see, among others, Easley *et al.* (1996) and the literature cited in Chung *et al.* (2005)). This measure is given by $\text{PIN} = \frac{\alpha\mu}{\alpha\mu + \varepsilon(1-\alpha\mu)}$, where the numerator is the beginning-of-the-day probability that a trade is information based and the denominator is the probability that a trade occurs. With the estimated parameters of our model, the PIN equals 19%,

²⁶A note of warning on the result of the test is needed here, since the null hypothesis is on the boundary of the parameter space (see Andrews, 2001).

whereas if we compute it for the Easley *et al.* (1997) model, it is only 9%.²⁷ If we adjust for the fact that in our model the information may not be correct (i.e., we multiply the numerator by the probability of a correct signal $0.5 + 0.25\tau$), the measures of information-based trading given by the two models become almost the same. Since the null that the signal is perfectly precise is rejected by the data, our results suggest that the PIN, as usually computed, measures the proportion of informed-based trading coming from traders receiving the correct information and not the overall proportion of information-based trading. Essentially, the difference is due to the fact that in the previous literature incorrect trades (e.g., sell in a good-event day) can only be due to exogenous, non-informative (e.g., liquidity) reasons to trade, whereas in our setup we do not exclude that they may come from informed traders who either receive the incorrect information or herd.

To conclude, note that a 99% confidence interval for τ does not include 1.²⁸ This means that there is evidence in our sample that there are no realizations of the signal that reveal the true asset value with probability one. In the jargon of the social learning literature, signals are bounded.

5.2 Herd Behavior

The estimates of the parameters α and τ imply that herd behavior can occur in our sample. Since the estimate of α is clearly lower than 1,²⁹ there is information uncertainty in the market, which is a necessary condition for the mechanism of herd behavior highlighted in Section 2 to work. Moreover, the estimate $\tau = 0.44$ means that traders receive a signal that is noisy (i.e., not perfectly informative) and may decide to act against it (i.e., buy upon receiving a bad signal or sell upon receiving a good one).

²⁷We compute the PIN for our model using the same formula as Easley *et al.* (1996). They interpret the PIN as the probability of a trade coming from an informed trader at the beginning of the day. In our model, since the signal is continuous, the interpretation is correct only if we ignore the bid-ask spread (otherwise, some informed traders may decide not to trade because their expectations fall inside the bid-ask spread.) We use this approximation for simplicity's sake and to keep comparability with the existing work on the PIN.

²⁸Since the parameter's standard deviation is 0.02, this is the case for any reasonable confidence interval.

²⁹The parameter's standard deviation is 0.03. See the argument in the previous footnote.

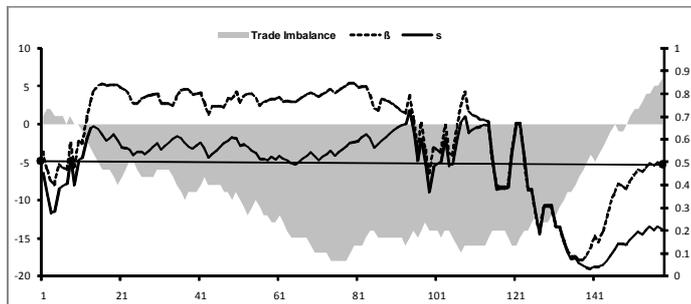


Figure 5: A day of Trading. The figure reports the evolution of trade imbalance (shaded line), buy threshold (dashed line) and sell threshold (solid line) in one day of trading. The thresholds are measure on the vertical axis on the right, and the trade imbalance on the vertical axis on the left. Herd buy occurs when the solid line is above 0.5 (indicated by an horizontal line) and herd sell when the dashed line falls below 0.5.

5.2.1 The Frequency of Herding

Recall that there is herd behavior at time t of day d when there is a positive measure of signal realizations for which an informed trader either herd buys or herd sells, that is, when, in equilibrium, either $\beta_t^d < 0.5$ (herd buy) or $\sigma_t^d > 0.5$ (herd sell). To gauge the frequency of herd behavior in our sample, for each trading day we computed the buy thresholds (β_t^d) and the sell thresholds (σ_t^d) given our parameter estimates. As an illustration, Figure 5 shows the thresholds for one day out of the 252 days in the sample. Whenever the buy threshold (dotted line) drops below 0.5 or the sell threshold (solid line) goes above 0.5, there is herd behavior. The shaded area represents the trade imbalance, that is, the number of buys minus the number of sells until the period before (measured on the vertical axis on the left). As one can see, herd buying occurs at the beginning of trading activity, as the trade imbalance is positive, that is, as more buy orders arrive in the market. This is followed by a long stretch of herd sells, as sell orders arrive and the trade imbalance becomes negative. At the very end of the day, herd behavior effectively disappears.

In this section we look at the frequency of herding, asking how often the buy threshold is below 0.5, or the sell threshold is above 0.5. It should be clear that this analysis is relevant both for event and no-event days. In both

cases, the existence of signal realizations for which informed traders herd (if an information event has occurred — which neither the market maker nor an external observer knows at the moment of the trade) modifies the way the market maker updates the price, which, as we shall see, affects the market informational efficiency.

In our sample, herding happens quite frequently: over the 252 days of trading, in 30% of trading periods β_t^d is below 0.5, and in 37% of trading periods σ_t^d is above 0.5. Moreover, there are some days where herding is very pronounced. Table 4 reports the proportion of days in which the buy or sell thresholds β_t^d or σ_t^d cross 0.5 at least 10, 30, and 50 percent of the time. Herd buy was observed in more than 50% of the trading times in 23% of the 252 days in our sample. Similarly, herd sell was observed in more than 50% of the trading times in 35% of the 252 days in our sample.

	> 10%	> 30%	> 50%
Herd Buy	0.79	0.47	0.23
Herd Sell	0.83	0.58	0.35

Table 4: Days with High Frequency of Herd Behavior.

The table shows the proportion of days in which the percentage of trading periods with herd behavior was higher than 10, 30 or 50. For instance, in 23 percent of days, herd buy periods were at least 50 percent of the total.

	Herd Buy	Herd Sell
Mean Length	9	10
Length S.D.	13	15
Max Length	99	120
Mean Time	60	61

Table 5: Length and Time of a Herd.

The table shows the mean, the standard deviation and the maximum for the number of consecutive trading periods in which there was herd behavior. The last row shows the mean trading period in which herd behavior occurred.

Table 5 reports the mean period of the day when we observed herding, that is, approximately the 60th trading period for both herd buy and herd sell

(which, in clock time, is roughly after 2 hours and 36 minutes of trading).³⁰ It is also interesting to ask how long periods of herding last, that is, for how many trading periods after a herd starts we observe a positive measure of herders. Herd buys last on average 9 trading periods (corresponding to about 24 minutes) and herd sells last on average for 10 periods. There is, however, pronounced variability in the length of herds, with a standard deviation of 13 and 15 trading periods. The longest herd buy lasted for 99 periods (about 257 minutes) and the longest herd sell lasted 120 periods (312 minutes).

5.2.2 Proportion of Herders

The previous analysis helps us gauge how often herding occurs in our sample. The fact that at a given time the buy (sell) threshold is lower (higher) than 0.5, however, does not tell us how likely it is for an informed traders to herd at that time. This is captured by the measure of signal realizations for which an informed trader herds. As we will discuss in detail in the next subsection, this measure is also very important for the informational efficiency of the market. The higher the measure of signal realizations for which traders herd, the lower the informational efficiency.

To compute the measure of herd-signal realizations, however, we need to know the distribution of signals in any given day. This, in turn, depends on whether the day of trading was a good-event or a bad-event day. To this purpose, we classified a day as a good-event day (bad-event day) if two conditions hold: a) $\Pr(V_d \neq v_{d-1} | h_{T_d}) > 0.9$ and b) either $\Pr(V_d = v^H | h_{T_d}, V_d \neq v_{d-1}) > 0.9$ (good-event day) or $\Pr(V_d = v^L | h_{T_d}, V_d \neq v_{d-1}) > 0.9$ (bad-event day). That is, we classified a day as a good (bad) event day if at the end of the day, the posterior probabilities of an information event occurring, and of the event being good (bad) were both higher than 0.9.³¹

We concentrated our analysis on the days classified as good-event or bad-event days. For each trading period, we computed the proportion of bad signals for which informed traders would herd buy, and, similarly, the proportion of good signals for which informed traders would herd sell. Let us consider, for instance, a good-event day. In such a day, the signal is distributed according to $g^H(s_t^d | v_d^H) = 1 + 0.44(2s_t^d - 1)$. For each trading period

³⁰On average there are 149 decisions in a day (390 minutes). Therefore, the average length of a trading period is 2.6 minutes.

³¹Of course, the 0.9 threshold is arbitrary. As a robustness check we repeated the calculations for 0.75, 0.8 and 0.85 and we obtained very similar results (see the Appendix).

in which $\beta_t^d < 0.5$, we computed the measure of signals between β_t^d and 0.5 (i.e., the measure of signals for which an informed trader herd buys). We then divided this measure by the measure of signals between 0 and 0.5 (the measure of all bad signals, i.e., all signals for which the informed trader could potentially herd buy). We refer to this ratio as the “proportion of herders,” since it shows the proportion of informed traders who would herd buy or herd sell were they to trade in that period. We report the average results in Tables 6 and 7.

	Average	S.D.	Max
Herd Buyers	2%	3%	11%
Herd Sellers	4%	6%	29%

Table 6: Percentage of Herders.

The table shows the mean, the standard deviation and the maximum of the percentage of herd buyers and herd sellers.

	> 1%	> 5%	> 10%
Herd Buyers	0.53	0.16	0.07
Herd Sellers	0.55	0.22	0.11

Table 7: Event Days with a High Percentage of Herders.

The table shows the proportion of days in which the percentage of buy or sell herders was higher than 1, 5 or 10. For instance, in 11 percent of days, the percentage of herd sellers was at least 10.

On average, across all event days, the proportion of buy herders was 2% and that of sell herders 4%. This proportion is, however, quite variable across days, reaching a maximum of 11% for herd buy and 29% for herd sell. Misdirected herding (i.e., herd buying in a bad-event day and herd selling in a good-event day) does occur: on average, in a bad-event day the proportion of herd buyers was 1%; in a good-event day, the proportion of herd sellers was 2%.

As Table 7 shows, there are a substantial number of days where the percentage of herd buyers or sellers is significant: for instance, in 7% of event days, the proportion of informed traders who herd buy was higher than 10%; similarly, in 11% of event days, the proportion of informed traders who herd

sell was higher than 10%. This confirms the result of the previous section that herding behavior seems to be particularly concentrated in some days of trading.

An important question is whether herding usually happens after buy or sell orders have accumulated in the market. Table 8 shows that this is actually the case. The first row of Table 8 reports the average level of the trade imbalance in periods of herd buy (i.e., when $\beta_t^d < 0.5$). The trade imbalance is on average positive both in good and in bad event days. This means that herd buy usually happens when there has been a preponderance of buys. Similarly, the second row of Table 8 shows that herd selling usually occurs when there has been a preponderance of sells.

	Good-Event Days	Bad-Event Days
$\beta_t^d < 0.5$	9	3.5
$\sigma_t^d > 0.5$	-8.2	-12.9

Table 8: Trade Imbalance in Periods of Herding.

The table shows the average level of the trade imbalance in periods of herd buy and herd sell.

	Good-Event Days	Bad-Event Days
$\Pr(buy \beta_t^d < 0.5)$	0.48	0.31
$\Pr(buy)$	0.43	0.26
$\Pr(sell \sigma_t^d > 0.5)$	0.39	0.48
$\Pr(sell)$	0.27	0.45

Table 9: Proportion of Buys and Sells in Periods of Herding.

The table shows the proportion of buys and sells in periods of herd buy, herd sell and in the whole sample.

Obviously, by definition, herd buy increases the proportion of buys and herd sell increases the proportion of sells. Table 9 illustrates this point, by showing the frequency of buy and sell orders that we observe in periods of herd buy and herd sell and contrasting them with the overall frequencies. In good-event days, for instance, the overall frequency of buy orders is 43%. This frequency raises up to 48% when there is herd buying. It is important to remark that the results of Tables 8 and 9 taken together imply that higher

positive (negative) levels of the trade imbalance increase the probability of a buy (sell) order. That is, herd behavior generates serial dependence in the trading pattern during the day.³²

5.2.3 Informational Inefficiency

In periods of herd behavior, a proportion of informed traders do not trade according to their private information; as a result, information is aggregated less efficiently by the price. It is easy to show, in fact, that trades convey the maximum amount of information when informed traders buy upon receiving a good signal and sell upon receiving a bad one.³³ In periods of herding, in contrast, traders may buy even with a bad signal or sell even with a good one.

To quantify the informational inefficiency caused by herding, we proceeded in the following way. We simulated the history of trades and prices over many days for our theoretical model, using our estimates as parameter values. We then compared the simulation results with two benchmarks that capture the price behavior in an informationally efficient market. In a first benchmark, we simulated the model forcing informed traders to buy (sell) upon a good (bad) signal. In other words, informed traders (irrationally) never herded and always followed their private information. As a second benchmark, we considered the case in which there is no information uncertainty, that is, the market maker knows whether there has been an informational event. As a result, he updates his beliefs (and prices) exactly as informed traders do, and, because of this, informed traders never herd. Essentially, in the first benchmark, the market is efficient because traders (irrationally) follow their signals; in the second, the market is efficient because the informational asymmetry between traders and market makers due to event uncertainty is eliminated. The difference between the two scenarios

³²In a related strand of literature, Hasbrouck (1991, 1991a) and Chung *et al.* (2005), among others, provide evidence of autocorrelation in trades. Easley *et al.* (1997a) and Easley *et al.* (2008) recognize the importance of the sequence of trades in conveying information. To capture it, they allow for path dependence in noise trading which is due to unmodelled reasons. In contrast, the sequence becomes important in our setup for reasons dictated by economic theory.

³³By this, we mean that the strategy $X_t^d = -1$ for $s_t^d < 0.5$ and $X_t^d = 1$ for $s_t^d > 0.5$ minimizes $E[(E(V_d|h_t^d, x_t^d) - V_d)^2]$. Note that, for simplicity's sake, we are assuming that traders can only buy or sell, since in equilibrium they abstain from trading only when their expectation is in the bid-ask spread.

is caused by the bid and ask spread. When, in the first benchmark, we force informed traders to buy (sell) upon a good (bad) signal, we disregard not only the incentive to herd, but also that to abstain from trading (because a trader’s expectation may fall within the bid-ask spread).³⁴

We simulated the price paths for 100,000 days of trading (with 149 trading periods per day) for our theoretical model and for the two benchmarks. Then, at each time t of any day d , we computed the distance (i.e., the absolute value of the difference) between the public belief $E(V_d|h_t^d)$ in our model and that in the benchmark.³⁵

In Table 10, we present the average distance taken over all trading periods as a percentage of the expected value of the asset. For the first benchmark, we present the average distance both over all days and over informed days only. For the second benchmark, by construction, only the average distance over event days is meaningful.

No-Herd Benchmark		No Event Uncertainty Benchmark
All Days	Event Days	Event Days
4%	10%	7%

Table 10: Informational Inefficiency

The table shows the average distance between the public belief in the model and that in the benchmarks as a percentage of the expected value of the asset.

Av. Distance	No-Herd Benchmark		No Event Uncertainty Benchmark
	All Days	Event Days	Event Days
$\geq 10\%$	0.13	0.37	0.21
$\geq 15\%$	0.07	0.20	0.11
$\geq 20\%$	0.04	0.11	0.06

Table 11: Days of High Informational Inefficiency

The table shows the proportion of days in which the average distance between the public belief in the model and that in the benchmarks is higher than 10, 15, or 20 percent of the expected asset value.

³⁴See previous footnote.

³⁵The public belief $E(V_d|h_t^d)$ is always between the bid and ask prices. It is common in the literature to interpret it as the price (abstracting from the bid-ask spread).

The average distance between the prices in our model and in the first benchmark over all days amounts to 4% of the asset expected value. If we focus our attention on event days, the distance is higher (10%), since the imbalance between buys and sells causes herding to arise more often. The average distance between the public beliefs using the second benchmark is similar, equal to 7%. In Table 11 we repeated these computations for days in which herding was more pronounced. In 7% of days (20% of event days) the distance between the price and the non-herding price is greater 15%. This suggests that there are days when intraday herding affects the informational properties of the price in a very significant manner.

6 Conclusion

We have developed a theoretical model of herd behavior in financial markets amenable to structural estimation with transaction data. We have estimated the model using data for a NYSE stock (Ashland Inc.) in 1995. Using our estimates, we have detected the periods in each trading day in which informed traders herd. We have found that herding is present in the market and fairly pervasive in some trading days. Moreover, it generates important informational inefficiencies.

The main contribution of the paper is methodological: it provides an empirical strategy to analyze herding within a structural estimation framework. This contrasts with the existing empirical studies of herding, based on a-theoretical, statistical measures of trade clustering.

In future research, we plan to use our methodology to investigate the importance of herding for a large number of stocks, by analyzing how herding changes with the stock characteristics (e.g., large stocks versus small ones) and with the macroeconomic environment (e.g., crises versus tranquil periods). We also plan to contribute to the existing literature on information and asset pricing, both by studying whether our measure of market informativeness (which takes into account that information may not be perfectly precise) improves the performance of the information factor, and by seeing whether herding itself is a risk factor priced in the market.

Finally, we observe that whereas our interest was in learning in financial markets, in the future our methodology could be fruitfully used in fields other than financial economics. The voluminous and growing theoretical literature on social learning has been intensively tested in laboratory exper-

iments. There is, however, a general acknowledgement that these models cannot be easily studied with field data, because we lack of data on private information. Our study shows how this problem can be overcome.

References

- [1] Avery, Christopher and Zemsky, Peter. “Multidimensional Uncertainty and Herd Behavior in Financial Markets.” *American Economic Review*, 1998, 88, 724-748.
- [2] Andrews, Donald. “Testing When a Parameter is on the Boundary of the Maintained Hypothesis.” *Econometrica*, 2001, 69, 683-734.
- [3] Banerjee, Abhijit. “A Simple Model of Herd Behavior.” *Quarterly Journal of Economics*, 1992, 107, 787-818.
- [4] Bikhchandani, Sushil, Hirshleifer, David and Welch Ivo. “A Theory of Fads, Fashion, Custom and Cultural Change As Informational Cascades.” *Journal of Political Economy*, 1992, 100, 992-1027.
- [5] Bikhchandani, Sushil and Sharma Sunil. “Herd Behavior in Financial Markets: A Review.” IMF Working paper, 2001, 47, 3.
- [6] Chung, Kee, Li, Mingsheng and McInish Thomas H. “Information-based Trading, Price Impact of Trades, and Trade Autocorrelation.” *Journal of Banking and Finance*, 2005, 29, 1645–1669.
- [7] Cipriani, Marco and Guarino Antonio. “Herd Behavior in a Laboratory Financial Market.” *American Economic Review*, 2005, 95, 1427-1443.
- [8] Cipriani, Marco and Guarino Antonio. “Herd Behavior and Contagion in Financial Markets.” *B.E. Journal of Theoretical Economics*, 2008, 8, Article 24.
- [9] Cipriani, Marco and Guarino Antonio. “Herd Behavior in Financial Markets: An Experiment with Financial Market Professionals.” *Journal of the European Economic Association*, 2009, 7, 206-233.
- [10] Dasgupta, Amil, Prat, Andrea and Verardo, Michela. “Institutional Trade Persistence and Long-Term Equity Returns.” *Journal of Finance*, forthcoming.
- [11] Drehmann, Mathias, Oechssler, Jörg and Rider, Andreas. “Herding and Contrarian Behavior in Financial Markets - An Internet Experiment.” *American Economic Review*, 2005, 95, 1403-1426.

- [12] Easley David and O'Hara, Maureen. "Price, Trade Size, and Information in Securities Markets." *Journal of Financial Economics*, 1987, 19, 69-90.
- [13] Easley David, Kiefer, Nicholas, O'Hara, Maureen and Paperman, Joseph. "Liquidity, Information and Infrequently Traded Stocks." *Journal of Finance*, 1996, 10, 805-835.
- [14] Easley David, Kiefer, Nicholas and O'Hara, Maureen. "One Day in the Life of a Very Common Stock." *Review of Financial Studies*, 1997, 10, 805-835.
- [15] Easley David, Kiefer, Nicholas and O'Hara, Maureen. "The Information Content of the Trading Process." *Journal of Empirical Finance*, 1997a, 4, 159-186.
- [16] Easley David, O'Hara, Maureen, Engle, Robert and Wo, Liuren. "Time-Varying Arrival Rates of Informed and Uninformed Trades." *Journal of Financial Econometrics*, 2008, 6, 171-207.
- [17] Gale, Douglas. "What Have We Learned from Social Learning." *European Economic Review*, 1996, 40, 617-628.
- [18] Gervais, Simon. "Market Microstructure with Uncertain Information Precision: A Multi-Period Analysis." Mimeo, Wharton School, 1997.
- [19] 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, 71-100.
- [20] 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, 1088-1105.
- [21] Hasbrouck, Joel. "The Summary Informativeness of Stock Trades: An Econometric Analysis." *Review of Financial Studies* 1991, 4, 571-595.
- [22] Hasbrouck, Joel. "Measuring the Information Content of Stock Trades." *Journal of Finance* 1991a, 46, 179-207.
- [23] Hasbrouck, Joel. "Working with the WRDS (TAQ and CRSP) Data," mimeo, NYU Stern School of Business, 2004.

- [24] Hirshleifer, David, and Teoh, Siew. “Thought and Behavior Contagion in Capital Markets,” in Reiner Schenk-Hoppé K., and Hens T., ed. *Handbook of Financial Markets: Dynamics and Evolution* Elsevier, 2009, 1-56.
- [25] Lakonishok, Joseph, Shleifer Andrei and Vishny Robert W. “The Impact of Institutional Trading on Stock Prices.” *Journal of Financial Economics*, 1992, 32, 23-43.
- [26] Lee, Charles and Ready, Marc. “Inferring Trade Direction from Intraday Data.” *Journal of Finance*, 1991, 46, 733-746.
- [27] Milgrom, Paul and Stokey, Nancy. “Information, Trade and Common Knowledge.” *Journal of Economic Theory*, February 1982, 26, 17-27.
- [28] Smith, Lones and Sørensen, Peter. “Pathological Outcomes of Observational Learning.” *Econometrica*, 2000, 68, 371-398.
- [29] Vives, Xavier. *Information and Learning in Markets: The Impact of Market Microstructure*, Princeton University Press, Princeton, 2008.
- [30] Welch, Ivo. “Sequential Sales, Learning, and Cascades.” *Journal of Finance*, 1992, 47, 695-732.

7 Appendix

7.1 Proof of Proposition 1

Consider the belief $E(V_d|h_2^d = \{buy_1\})$ of a trader who observes a buy order at time 1. Let $E(V_d|h_2^d = \{buy_1\}) - v_{d-1} = \kappa$, where $\kappa > 0$. Consider now the belief $E(V_d|h_2^d = \{buy_1\}, s_2^d = 0.5 - \vartheta)$ of a trader with a bad signal $s_2^d = 0.5 - \vartheta$ (with $\vartheta > 0$) after he observes a buy order at time 1. For any κ there is a ϑ such that $E(V_d|h_2^d = \{buy_1\}, s_2^d = 0.5 - \vartheta) - v_{d-1} > 0$. Let us denote this difference by ϕ .

Consider now a history of trades in which in each odd period there is a buy order and in each even period there is a sell order. Given the presence of noise traders, such a history has positive probability. For any ϕ , there exists an (odd) t such that at time $t + 1$, in equilibrium, $a_{t+1}^d - v_{d-1} < \phi$ (since the market maker after such a history attaches a lower and lower

probability to the day being informed). Finally, note that at $t + 1$ the belief of an informed trader is exactly as at time 2, since he knows that an event has occurred and the equal number $\frac{t-1}{2}$ of buy and sell orders before time t offset each other in the updating of the belief about the event being good or bad. Therefore, a trader with signal $s_{t+1}^d \in [0.5 - \vartheta, 0.5)$ has an expectation $E(V_d | h_{t+1}^d, s_{t+1}^d) > v_{d-1} + \phi > a_{t+1}^d$ and herds. Since this history occurs with positive probability when $V_d = v_d^L$, this shows that incorrect herd buy occurs with positive probability. The proof for herd sell is analogous.

7.2 Further computations for the likelihood function

The probabilities in the equation in the text can easily be expressed as a function of β_t^d and of the parameters. For instance, the probability that an informed trader with signal β_t^d attaches to the good informational event can be computed as

$$\Pr(v_d^H | h_t^d, \beta_t^d) = \frac{2\beta_t^d \Pr(v_d^H | h_t^d, V_d \neq v_{d-1})}{2\beta_t^d \Pr(v_d^H | h_t^d, V_d \neq v_{d-1}) + 2(1-2\beta_t^d)(1-\Pr(v_d^H | h_t^d, V_d \neq v_{d-1}))}.$$

(Note that $\Pr(v_d^H | h_t^d, V_d \neq v_{d-1})$ is conditioned on $V_d \neq v_{d-1}$ since the informed trader knows that an event has occurred.) The probability that the market maker attaches to the good event can, instead, be computed as

$$\Pr(v_d^H | h_t^d, buy_t^d) = \frac{\left(\left(1-\beta_t^{d^2}\right)\mu + \left(\frac{\varepsilon}{2}\right)(1-\mu)\right) \Pr(v_d^H | h_t^d, V_d \neq v_{d-1}) \Pr(V_d \neq v_{d-1} | h_t^d)}{D},$$

where

$$D = \left(\left(1-\beta_t^{d^2}\right)\mu + \left(\frac{\varepsilon}{2}\right)(1-\mu)\right) \Pr(v_d^H | h_t^d, V_d \neq v_{d-1}) \Pr(V_d \neq v_{d-1} | h_t^d) + \left(\left(\beta_t^{d^2} - 2\beta_t^d + 1\right)\mu + \left(\frac{\varepsilon}{2}\right)(1-\mu)\right) (1 - \Pr(v_d^H | h_t^d, V_d \neq v_{d-1})) \Pr(V_d \neq v_{d-1} | h_t^d) + \left(\frac{\varepsilon}{2}\right) \Pr(V_d = v_{d-1} | h_t^d).$$

7.3 Robustness checks for Table 6 for 0.75, 0.8 and 0.85

	Average	S.D.	Max
Herd Buyers	3%	4%	20%
Herd Sellers	3%	5%	29%

	Average	S.D.	Max
Herd Buyers	3%	4%	20%
Herd Sellers	3%	5%	28%

	Average	S.D.	Max
Herd Buyers	3%	4%	20%
Herd Sellers	3%	5%	29%