

Why is order flow so persistent?

Bence Tóth^{a,b}, Imon Palit^c, Fabrizio Lillo^{b,d,e}, J. Dooyne Farmer^{b,f}

^aCapital Fund Management, 6, blvd Haussmann 75009 Paris, France

^bSanta Fe Institute, 1399 Hyde Park Rd., Santa Fe NM 87501, USA

^cCentre of Computational Finance and Economic Agents (CCFEA), Univeristy of Essex, UK

^dDipartimento di Fisica, Università di Palermo, Palermo, Italy

^eScuola Normale Superiore di Pisa, Piazza dei Cavalieri 7, 56126 Pisa, Italy

^fLUISS Guido Carli, Viale Pola 12, 00198, Roma, Italy

Abstract

Equity order flow is persistent in the sense that buy orders tend to be followed by buy orders and sell orders tend to be followed by sell orders. For equity order flow this persistence is extremely long-ranged, with positive correlations spanning thousands of orders, over time intervals of up to several days. Such persistence in supply and demand is economically important because it influences the market impact as a function of both time and size and because it indicates that the market is in a sense out of equilibrium. Persistence can be caused by two types of behavior: (1) Order splitting, in which a single investor repeatedly places an order of the same sign, or (2) herding, in which different investors place orders of the same sign. We develop a method to decompose the autocorrelation function into splitting and herding components and apply this to order flow data from the London Stock Exchange containing exchange membership identifiers. Members typically act as brokers for other investors, so that it is not clear whether patterns we observe in brokerage data also reflect patterns in the behavior of single investors. To address this problem we develop models for the distortion caused by brokerage and demonstrate that persistence in order flow is overwhelmingly due to order splitting by single investors. At longer time scales we observe that different investors' behavior is anti-correlated. We show that this is due to differences in the response to price-changing vs. non-price-changing market orders.

JEL codes: G12, D44, D61, C62.

Keywords: Financial markets, behavioral economics

Contents

1	Introduction	2
2	Data	5
3	Decomposing order flow	6
4	Main empirical results	8
4.1	Splitting dominates herding	8
4.2	Heterogeneity of individual members	10
5	Do these results apply to investors?	12
5.1	Three hypotheses about persistence in order flow	12
5.1.1	Order splitting	13
5.1.2	Herding model I: Public information	13
5.1.3	Herding model II: Imitation	14

5.2	Models of brokerage	15
5.2.1	Fixed random brokerage	15
5.2.2	Dynamically random brokerage	16
5.2.3	Correlation of social network and brokerage	17
5.3	Comparison with real data	17
5.3.1	Public information/fixed random brokerage	17
5.3.2	Generalized imitation model/fixed random brokerage	17
5.3.3	Any investor model/dynamic random brokerage	18
5.3.4	Order splitting/partially random brokerage	18
6	Clues about the cause of anti-herding	19
6.1	Anti-herding is statistically significant	19
6.2	What underlies “anti-herding”?	19
7	Conclusions	22
8	Acknowledgments	23

1. Introduction

Recent studies have shown that order flow is remarkably persistent in the sense that orders to buy tend to be followed by more orders to buy and orders to sell tend to be followed by more orders to sell. Positive serial autocorrelation for the first autocorrelation of signed order flow has been observed by many authors in many different markets, including the Paris Bourse by Biais, Hillion and Spratt (1995), in foreign exchange markets by Danielsson and Payne (2001), and the NYSE by Ellul et al. (2007) and Yeo (2008).

In fact, *all* the coefficients of the autocorrelation function are positive out to large lags, corresponding to thousands of transactions (Bouchaud et al., 2004; Lillo and Farmer, 2004). An example is shown in Figure 1. Although most of the autocorrelations are quite small, the fact that they are all positive implies a substantial degree of predictability. This has now been studied in many different equity markets, including the London, New York, Paris, and Madrid stock markets, on a variety of different stocks¹. The behavior is remarkably consistent.

In an efficient market the autocorrelation of price returns must decay to zero quickly. How can strong persistence in order flow co-exist with efficient price returns? Two theories have been offered. Lillo and Farmer (2004) show that the problem can be solved if market impact is permanent but liquidity is asymmetric and time varying: When buy orders are more likely, the price response to buy orders is smaller than the response to sell orders, and when sell orders are more likely, the price response to sell orders is smaller than that for buy orders. Bouchaud et al. (2004, 2006) have shown that it can also be solved if market impact is constant but temporary, slowly decaying to zero. Farmer et al. (2006) demonstrate time varying liquidity explicitly, and show that for models based on market order impact the two hypotheses are in a sense equivalent².

The persistence of order flow thus has important economic consequences. It indicates that supply and demand have non-trivial dynamical properties that constrain their functional forms. For example, Farmer et al. (2011) argue that the persistence of order flow determines the shape

¹ Lillo and Farmer (2004) originally studied New York and London, Bouchaud et al. (2004) Paris, and Vaglica et al. (2008) and Moro et al. (2009) Madrid.

² See also Gerig (2007) and Bouchaud et al. (2009) for a detailed discussion of the equivalence. More recent work by Eisler et al. (2011a) shows that when limit orders and cancellations are taken into account these models are no longer equivalent.

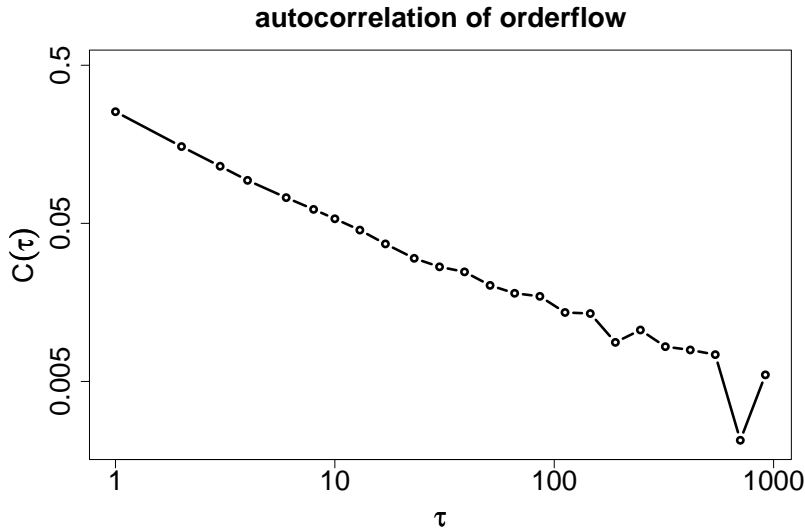


Figure 1: Autocorrelation function of market order signs for the stock AZN in the first half of 2009, plotted on double logarithmic scale. The time lag τ is measured in terms of number of market order placements. The estimated autocorrelations are all positive and statistically significant out to lags of more than 100.

of the market impact of large institutional orders, as well as its permanent and temporary components. Persistence of order flow implies that the market deviates substantially from the equilibrium that would prevail if all participants were forced to reveal their true intentions at the outset of trading.

Although nothing in our analysis here depends on this, in many equity markets order flow has been observed to obey a long-memory process³. This means that the autocorrelation function $C(\tau)$ of the order signs ($\epsilon_t = +1$ for buy and $\epsilon_t = -1$ for sell) asymptotically decays in time⁴ for large τ as $C(\tau) \sim \tau^{-\gamma}$, where $0 < \gamma < 1$. The observation of long-memory in stock markets is very robust; in London, for example, every stock examined shows long-memory under strict statistical tests (Lillo and Farmer, 2004).

The persistence associated with long-memory is remarkable in many ways. The slow decay with τ implies that the autocorrelation function is not integrable. This means that order flow is highly predictable, with events in the distant past having non-negligible influence on the present. This predictability is persistent, in the sense that prediction errors asymptotically decay as a power law, implying one can predict further into the future than is possible if errors decay exponentially. We wish to emphasize, however, that nothing in this paper depends on whether or not order flow is actually long-memory; our work here is focused on the origin of the correlations, which are undeniably extremely long-range.

We distinguish two types of behavior that can give rise to persistent order flow:

1. *Order splitting*, corresponding to a sequence of orders of a given sign originating from a single investor. This has a natural strategic cause: To reduce impact investors split their orders into smaller pieces, and execute them gradually, as originally hypothesized by Kyle

³ A standard example of a long-memory process is a fractional Brownian motion. We use the term in its more general sense to mean any process whose autocorrelation function is non-integrable (Beran, 1994). This can include processes with structure breaks, such as those studied by Ding, Engle and Granger (1993). Long-memory was observed in the London and New York stock exchanges by Lillo and Farmer (2004), in the Paris stock exchange by Bouchaud et al. (2004) and in the Spanish stock exchange by Vaglica et al. (2008). The question of whether or not order flow in other asset classes has long-memory remains open.

⁴ We are using the notation $A(x) \sim B(x)$ to mean $\lim_{x \rightarrow \infty} A(x)/B(x) = C$ with $C \neq 0$.

(1985)⁵. If we call the underlying package of orders corresponding to a given decision a *metaorder*, heavy tails in the size distribution of metaorders will induce persistence in order flow⁶.

2. *Herding*, corresponding to a sequence of roughly contemporaneous orders of a given sign originating from different investors⁷. Herds can be directed exogenously or endogenously. We investigate an example of each:

- (a) *Public information*. Information is public and investors all respond in more or less the same way at more or less the same time. Persistence in information arrival, or alternatively heavy tails in the variability of the response to information, can generate persistence in order flow.
- (b) *Imitation*. Information is private but investors influence each other. If one investor buys, other investors are more likely to buy. Even if private information arrival is IID, the time needed for information to propagate can make order flow persistent.

Although one might not normally think of a common response to public information as herding, as we show here, the two can produce indistinguishable order flow. In the public information model the sheep are directed by an external shepherd, whereas in the imitation model they are directed by each other. We develop simple examples of processes of type (1), (2a), and (2b) and use this to formulate null hypotheses for order splitting and herding.

Our empirical analysis here rests on a method that we develop to decompose the autocorrelation of order flow into herding and splitting components. We apply this method to data from the London Stock Exchange (LSE) with codes indicating the member who executed the order. Our results show that the persistence of order flow is overwhelmingly caused by order splitting. In contrast the contribution of herding is negligible. In fact, over longer time scales we observe anti-herding, in which investors trade in the opposite direction of other investors.

Members of the exchange may trade for their own accounts, but they may also act as brokers for investors who are not members⁸. We thus have the problem that while we have useful information about identity, this information is not fine-grained enough to allow us to separate different investors. By *investor* we mean any entity with a coherent trading strategy. This could correspond to a specific trading account within an institution such as an investment bank or hedge fund, or a private individual trading for his or her own account. The lack of identifiers at the level of single accounts potentially makes our results difficult to interpret.

To deal with this problem we introduce the concept of a brokerage map associating an investor with one or more brokers. We make several different hypotheses about brokerage maps, and study their effect on the decomposition of the autocorrelation function under different hypotheses for investor behavior. This is done through a combination of closed form calculations, simulations, and bootstrapping with the real data. This allows us to get a good understanding

⁵See also Farmer et al. (2011); Toth et al. (2011b).

⁶ Assume a metaorder is split into pieces and executed incrementally. Given that we have observed N executions, what is the probability of observing more? It is possible to show that if the distribution of metaorder sizes is exponential, past executions give no information about whether execution will continue, i.e. they give no information about the size of the metaorder. If the distribution is heavier-tailed than exponential, continuation becomes more likely with each execution, and if it is thinner-tailed than exponential, continuation becomes less likely with each execution. See Lillo et al. (2005), LeBaron and Yamamoto (2010) and Farmer et al. (2011)

⁷ Most studies of herding have focused on imitation. See for example (Banerjee, 1992, 1993; Bikhchandani et al., 1992; Orlean, 1995; Raafat et al., 2009). LeBaron and Yamamoto 2007; 2008 constructed an agent-based model of imitation and demonstrated that it could produce long-memory in order flow. Several previous studies have focused on the effect of communication network structure on price fluctuations; see (Kirman and Teyssiere, 2002; Iori, 2002; Cont and Bouchaud, 2000; Tedeschi et al., 2009). Barber et al. (2009) provides an example in which herding is externally directed. Lakonishok et al. (1992) and Wermers (1999) find only weak evidence for herding.

⁸ We will use the terms “member” and “broker” interchangeably.

of how brokerage influences the results. We show that, barring unrealistic brokerage maps, our main qualitative conclusions can be reliably interpreted as reflecting the behavior of investors.

Our results here add to earlier evidence that order splitting is dominant. Early studies showing that order splitting is widespread were done by Chan and Lakonishok (1993, 1995). Lillo, Mike and Farmer (2005) introduced a model in which information is completely private, and participants cannot observe each others' order flow (or in any case do not respond to it). This model assumes that investors draw the total quantity V they wish to trade at random from a pre-defined distribution $P(V)$, divide V into smaller, equal sized pieces, and then incrementally execute each piece at a constant rate until the entire order of size V is filled. Under the assumption that in the limit $V \rightarrow \infty$ the distribution $P(V)$ has a Pareto (power law) tail, of the form $P(V > v) \sim v^{-\beta}$, then in the limit $\tau \rightarrow \infty$ the resulting autocorrelation $C(\tau)$ of the order flow is a power law $C(\tau) \sim \tau^{-\gamma}$, with

$$\beta = 1 + \gamma. \quad (1)$$

This hypothesis was tested by comparing empirical measurements of the tail exponents β for trading volume in block markets to empirically measured correlation scaling exponents γ using data for 20 stocks from the London Stock Exchange. They showed that Eq. (1) is well-satisfied by the data. Another supporting piece of evidence is due to Gerig (2007), whose results (based on only one stock) suggested that trades coming from the same brokerage have long-memory, whereas trades from different brokerages do not (see also Bouchaud, Farmer, and Lillo (2009)). Finally, Vaglica et al. (2008) reconstructed the volumes V of the split orders from broker data (similar to the one used here), and found that their size is asymptotically Pareto distributed with a tail exponent close to the one predicted by Eq. (1).

Thus there is already preliminary evidence in favor of order splitting. This paper makes the evidence much stronger by developing precise null hypotheses to test for splitting v.s. herding, and systematically analyzing six stocks over a period of ten years, a sample of more than 39 million orders.

In Section 2 we describe our data. In Section 3 we develop a decomposition of the autocorrelation function into splitting and herding components. In Section 4 we apply this to order flow data identifying member firms and show that it overwhelmingly favors splitting over herding. We study the heterogeneity of individual members and show that most of them are extremely persistent in terms of both trade direction and the clustering in time of their trades. In Section 5 we investigate the question of whether our analysis reliably reflects the behavior of investors. We develop simple models of herding and splitting and study how brokerage distorts the decomposition of Section 3, to formulate null hypotheses for investor behavior and brokerage. Comparing the data to the null hypotheses shows that extreme assumptions are required to contradict our qualitative conclusions. In Section 6 we investigate anti-herding more carefully, and show that it is driven by heterogeneity in the response of investors to market orders that change the price, vs. those that don't. In the conclusions we summarize and reflect on the economic implications of our results.

2. Data

This study is performed using data from the London Stock Exchange (LSE). There are two parallel markets, the on-book market (SETS) and the off-book market (SEAQ). In the on-book market trades take place via a fully automatic electronic order matching system, while in the off-book market trades are arranged bilaterally via phone calls. We restrict our study to the on-book market. Note that there are no official market makers, though it is possible for any member firm to act as a market maker by posting bids and offers simultaneously.

We study six stocks in the period from June 2000 to June 2009, with the exception of a six month period from January to May in 2003. We divide the data for each stock into 17 subperiods of six months each, for a total of $17 \times 6 = 102$ samples. The six stocks we study are AstraZeneca (AZN), BHP Billiton (BLT), British Sky Broadcasting Group (BSY), Lloyds Banking Group (LLOY), Prudential (PRU), Vodafone Group (VOD)⁹. In cases where we present data for only one period we will use AZN in the first half of 2009. In this period the typical number of market orders per day is between 4,000 and 5,000. The number of market orders in each sample ranges from 91,710 to 1,416,000, with an average of 382,500.

Our data set contains all orders placed in the on-book market. Each order is characterized by a membership code of the member who placed the order¹⁰. The number of members varies throughout the sample, but there are typically about 100 members. To avoid data with poor statistics, in each six month period we remove members who make less than 100 transactions across the full period, which typically leaves about 80 active members. The activity level is very heterogeneous. For example, the 5 most active market members are responsible for 40-50% of transactions and the 15 most active ones are responsible for 80-90% of transactions. The value of the Gini coefficient of member activity averaged across the 102 samples is 0.87. This is quite consistent across stocks and time periods, with a standard deviation of only 0.02. Thus the trading activity is strongly concentrated in a relatively small number of member firms.

We have performed the analysis given here in three different ways: (1) market order signs ϵ_t , where a market order is defined as any event that results in an immediate transaction; (2) signs of all orders, including both market and limit orders; (3) signed volume v_t of transactions. The results are similar in all three cases, except that for signed volume they are somewhat noisier. We present only the results for case (1), market order signs. The number of market orders is used to measure time, i.e. $t' = t + 1$ whenever there is a market order placement.

3. Decomposing order flow

In this section we derive a decomposition of the autocorrelation of order flow. This decomposes it in two ways: (1) Contributions due to trading direction (buying vs. selling) vs. activity clustering (when investors make their trades), and (2) contributions due to herding vs. order splitting.

Consider a time series of orders of sign ϵ_t , where $\epsilon_t = +1$ for buy orders and $\epsilon_t = -1$ for sell orders. For convenience we measure time in terms of order arrivals, so that time advances by one unit every time a new order arrives. We define the signed order flow ϵ_t^i to be zero if the order at time t was not placed by investor i and to be the order sign otherwise, i.e.

$$\begin{aligned}\epsilon_t^i &= 1 \Rightarrow \text{buy order placed by investor } i \\ \epsilon_t^i &= 0 \Rightarrow \text{order placed by another investor} \\ \epsilon_t^i &= -1 \Rightarrow \text{sell order placed by investor } i.\end{aligned}$$

As already mentioned the analysis can be performed using market orders (as we have done here), all orders, or with signed order volumes, with similar results.

⁹ We have also investigated other stocks and find similar results.

¹⁰ We do not have the actual names of the members, but rather anonymized codes uniquely identifying each member. For the period 2000-2002 we are only able to track members for one month, because in this part of the data codes are reshuffled each month. Because we are operating on timescales of at most a few hours this does not affect the results.

The sample autocorrelation is

$$\begin{aligned} C(\tau) &= \frac{1}{N} \sum_t \epsilon_t \epsilon_{t+\tau} - \left(\frac{1}{N} \sum_t \epsilon_t \right)^2, \\ &= \frac{1}{N} \sum_t \sum_{i,j} \epsilon_t^i \epsilon_{t+\tau}^j - \left(\frac{1}{N} \sum_t \sum_i \epsilon_t^i \right)^2, \end{aligned} \quad (2)$$

where N is the length of the time series.

To decompose this further assume stationarity, let N^i be the number of orders due to investor i and let $N^{ij}(\tau)$ be the number of times that an order from investor i at time t is followed by an order from investor j at time $t + \tau$. Similarly let $P^i = N^i/N$ be the fraction of orders placed by investor i , and $P^{ij}(\tau) = N^{ij}(\tau)/N$ be the fraction of times that an order from investor i at time t is followed by an order from investor j at time $t + \tau$. We can then write

$$\begin{aligned} C(\tau) &= \sum_{i,j} P^{ij}(\tau) \left[\frac{1}{N^{ij}(\tau)} \sum_t \epsilon_t^i \epsilon_{t+\tau}^j \right] - \left(\sum_i P^i \frac{1}{N^i} \sum_t \epsilon_t^i \right)^2, \\ &= \sum_{i,j} \left(P^{ij}(\tau) \left[\frac{1}{N^{ij}(\tau)} \sum_t \epsilon_t^i \epsilon_{t+\tau}^j \right] - P^i P^j \left[\frac{1}{N^i} \sum_t \epsilon_t^i \right] \left[\frac{1}{N^j} \sum_t \epsilon_t^j \right] \right). \end{aligned} \quad (3)$$

To decompose and interpret the autocorrelation it is useful to define some simplifying notation. Let the sample mean of the order sign for investor i be

$$\mu^i \equiv \frac{1}{N^i} \sum_t \epsilon_t^i.$$

The sample autocorrelation in the order signs of agents i and j (including the zeroes when they are not trading) is

$$C^{ij}(\tau) \equiv \frac{1}{N^{ij}(\tau)} \sum_t \epsilon_t^i \epsilon_{t+\tau}^j - \mu^i \mu^j. \quad (4)$$

The weighting function $P^{ij}(\tau)$ captures the extent to which investors i and j tend to be active with a given time lag τ , regardless of the sign of their orders. If investors act independently then in the large N limit

$$P^{ij}(\tau) = P^i P^j \quad \forall i \text{ and } \forall j, \quad (5)$$

and $P^{ij}(\tau)$ is independent of τ . Deviations from independence can be characterized by

$$\tilde{P}^{ij}(\tau) \equiv P^{ij}(\tau) - P^i P^j. \quad (6)$$

With these definitions the autocorrelation function of order flow can be written

$$C(\tau) = \sum_{i,j} P^{ij}(\tau) C^{ij}(\tau) + \sum_{i,j} \tilde{P}^{ij}(\tau) \mu^i \mu^j. \quad (7)$$

This is naturally further decomposed as

$$C(\tau) = C_{split}(\tau) + C_{herd}(\tau), \quad (8)$$

where

$$\begin{aligned} C_{split}(\tau) &\equiv \sum_i P^{ii}(\tau) C^{ii}(\tau) + \sum_i \tilde{P}^{ii}(\tau) (\mu^i)^2, \\ C_{herd}(\tau) &\equiv \sum_{i \neq j} P^{ij}(\tau) C^{ij}(\tau) + \sum_{i \neq j} \tilde{P}^{ij}(\tau) \mu^i \mu^j. \end{aligned} \quad (9)$$

This decomposition will be the main analysis tool used in this paper.

$C_{split}(\tau)$ is the autocorrelation of orders coming from the same investor, and $C_{herd}(\tau)$ is the autocorrelation of orders coming from different investors. There might be many reasons why the orders from the same investor might be positively correlated; we have chosen to label this as C_{split} because the economic motivations for doing this are clear and well-known. Another obvious interpretation might be that the private information arriving to investors is autocorrelated; we cannot distinguish these explanations, so the reader should bear this in mind when we use the term “splitting”. Similarly we refer to any interaction that induces autocorrelations among different investors as “herding”.

It is useful to compare the first and second terms of the decomposition in Eq. (7). The first term weights the correlation of order signs by the activity $P^{ij}(\tau)$, and the second term weights deviations in activity by the mean order signs μ^i and μ^j . Since buying and selling roughly balance, for large samples μ^i is small and the second term is negligible¹¹. Typical values for our dataset are $|\mu^i| = 0.03$ and $\sum_{i,j} \tilde{P}^{ij}(1)\mu^i\mu^j = 10^{-4}$. Thus the first term of Eq. (7) is three orders of magnitude larger than the second one, and it is a very good approximation to take

$$C(\tau) \simeq \sum_{i,j} P^{ij}(\tau)C^{ij}(\tau). \quad (10)$$

Similar approximations can be used for $C_{split}(\tau)$ and $C_{herd}(\tau)$. Thus one can think of the autocorrelation as a product of two terms, one that depends on activity, as reflected by $P^{ij}(\tau)$, and the other that depends on trading direction, as reflected by $C^{ij}(\tau)$.

4. Main empirical results

We now compute the decomposition of the order flow of the data. To do this we make the bold leap that members of the exchange can be treated as if they were investors, i.e. we associate the indices i and j of the previous section with exchange membership. As we have already commented, this has the very serious problem that members often act as brokers, so that orders associated with a given member may correspond to a large number of different investors, and similarly, a given investor may use more than one broker. In the next section we will explicitly model the distortion this introduces under carefully constructed null hypotheses and show that, barring extreme behavior, the qualitative conclusions of this section apply to investors as well.

4.1. Splitting dominates herding

Figure 2 shows an example of the decomposition of the autocorrelation function into $C(\tau) = C_{split}(\tau) + C_{herd}(\tau)$ according to Eq. 8. This illustrates the main result of this paper: *Splitting dominates herding*. This is evident in several ways. The splitting term is always positive and is larger than the herding term at all lags. This is particularly true for $\tau \gtrsim 10$, since in this case the herding term is negative. We term such behavior *anti-herding*.

The behavior shown in Figure 2 is remarkably consistent across all 102 samples. To illustrate this, in the left panel of Figure 3 we plot the decomposition $C(\tau) = C_{split}(\tau) + C_{herd}(\tau)$ averaged across all of the 102 samples, and also plot the standard deviation across the samples for each time lag. The average behavior is very close to that of AZN. Furthermore, the standard deviations are small compared to the difference between C_{split} and C_{herd} . For $\tau \leq 100$, where

¹¹ It is not surprising that this term is small – in the long-run one expects μ^i to be close to zero. If ϵ_t^i represents signed trading volume, for example, then $\sum_t \epsilon_t^i$ is the total inventory change during the period of the study. If the inventory remains bounded, then $\mu^i \rightarrow 0$ in the limit $N \rightarrow \infty$. Even if ϵ_t^i represents order signs, then if investors buy as often as they sell, μ^i is small.




Figure 2: A decomposition of the autocorrelation into $C(\tau) = C_{split}(\tau) + C_{herd}(\tau)$ according to Eq. (8); the horizontal axis is plotted on logarithmic scale and the vertical axis on linear scale. Splitting dominates at all lags. Note the negativity of the herding component for $\tau > 10$. Here we show AZN in the first half of 2009, but all of the 102 samples look similar.

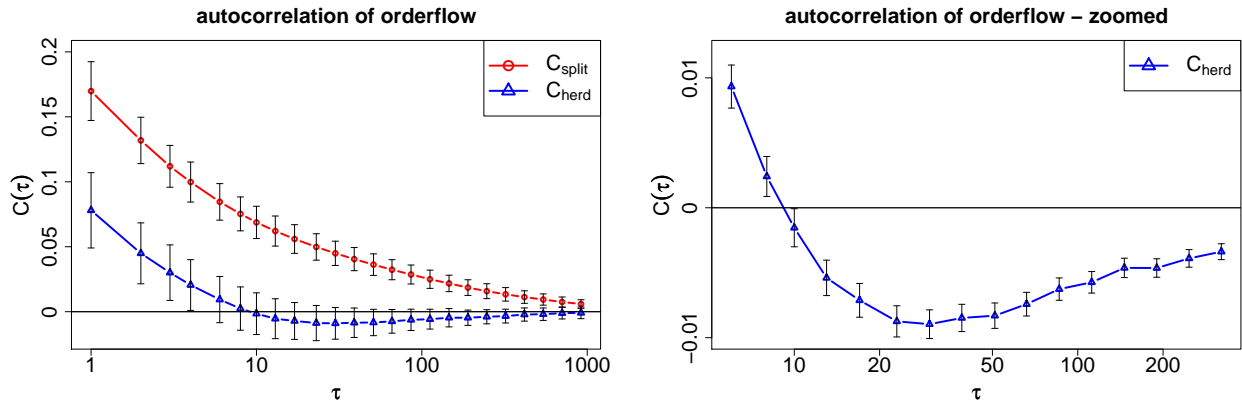


Figure 3: An illustration of the consistency with which splitting dominates herding. *Left panel:* The autocorrelations of market order signs averaged across all 102 samples (spanning nine years and six stocks). The bars are standard deviations. *Right panel:* The herding component in the left panel is magnified to better observe the anti-herding effect. The bars in this panel are standard errors. For both plots we use logarithmic scale on the horizontal axis and linear scale on the vertical axis.




Figure 4: Splitting becomes steadily more dominant over herding as time increases. *Left panel:* The splitting ratio $S(\tau) = C_{split}(\tau)/C(\tau)$ for the first half of 2009 for AZN. *Right panel:* The splitting ratio averaged over all 102 samples. The bars represent standard deviations. The red dashed line corresponds to the herding null hypothesis of Eq. (15) with parameters matched to the data.

we have the best statistical reliability, there is not a single case in which the herding term is larger than the splitting term.

The anti-herding effect observed for AZN is not special to this stock or this time period: Almost all stocks show similar behavior. To examine this in more detail, in the right panel of Fig. 3 we enlarge the scale and plot only C_{herd} . Instead of showing the standard deviation across the samples, we show the standard error. The fact that the average value of C_{herd} is consistently negative for $10 \leq \tau \leq 250$, at many lags by more than three times the standard error, suggests that this effect is real. We will return in Section 6 to examine anti-herding in more detail.

To have a quantitative comparison of the size of the two terms and to see how this changes with τ , in the left panel of Figure 4 we plot the fraction of the autocorrelation due to splitting, i.e.

$$S(\tau) = \frac{C_{split}(\tau)}{C(\tau)}. \quad (11)$$

Note that $S(\tau)$ is larger than one if C_{herd} is negative. In the left panel we show results for AZN for the first half of 2009, and in the right panel we show the average result for all 102 samples. For $\tau = 1$ the splitting accounts for about 75% of the autocorrelation and herding explains about 25%. For larger lags splitting becomes relatively even more important. For lags larger than roughly 10, $S(\tau)$ becomes larger than one because the herding part becomes negative, and $S(\tau)$ rises to about 1.5. This means that the splitting term is about 2.5 times the absolute value of the (anti)herding term. For large τ the standard deviation across the samples becomes large because C_{split} and particularly C_{herd} are small, so we are taking ratios of small numbers.

4.2. Heterogeneity of individual members

In this section we study the behavior of individual members of the exchange and show that their behavior is remarkably consistent, with a few exceptions. The decomposition of Eq. (7) makes it clear that persistence in order flow is driven by two factors: Persistence in order sign, measured by $C^{ij}(\tau)$, and persistence in activity, measured by $P^{ij}(\tau)$. Because they appear multiplicatively, both factors are needed to get a large autocorrelation. The fact that splitting is dominant over herding indicates that the bulk of the persistence is due to the diagonal terms $C^{ii}(\tau)$ and $P^{ii}(\tau)$.

To understand how these vary across the members of the exchange, in Fig. 5 we show $C^{ii}(\tau)$ and $\tilde{P}^{ii}(\tau)$ for the 15 most active members. We have chosen to look at $\tilde{P}^{ii}(\tau)$ rather than

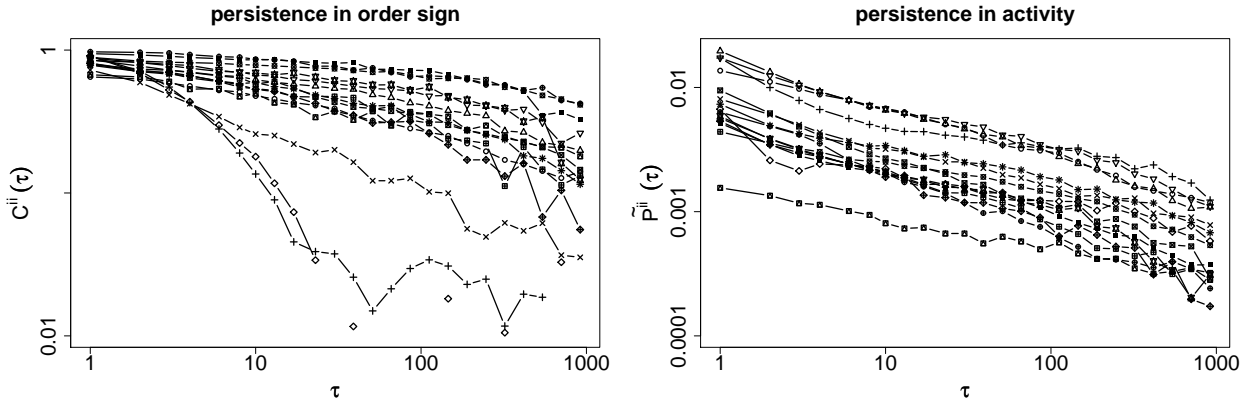


Figure 5: Heterogeneity of the contribution of individual exchange members (labeled by i) in the diagonal component of the autocorrelation of order flow. Following the decomposition of Eq. (7), we plot $C^{ii}(\tau)$, which measures persistence in trading direction (left) and $\tilde{P}^{ii}(\tau)$, which measures persistence in activity (right). We show the 15 most active members for AZN for the first half of 2009, using distinctive symbols to indicate each member (consistently in both plots). For one the members (labeled by diamonds) $C^{ii}(\tau)$ decays more quickly than the others and for some values of $\tau > 10$ is negative, hence the diamonds cannot be plotted on logarithmic scale and cannot be connected. This member is the fifth most active.

$P^{ii}(\tau)$ in order to highlight the approximate power law behavior¹² with τ . For $\tilde{P}^{ii}(\tau)$ all of the members show strikingly similar behavior; the individual curves are surprisingly straight and parallel, indicating a good correspondence to the same power law, though with differences in the overall scale. $C^{ii}(\tau)$ also decays very slowly, though the curves tend to steepen slightly with increasing τ , and the relative dispersion also increases. There are three members for which $C^{ii}(\tau)$ decays significantly faster than the others, suggesting that these members follow a different business model.

The correlation in trading direction C^{ii} and the trading frequency P^i are essentially uncorrelated: The Spearman rho of P^i and $C^{ii}(1)$ is 0.06. In contrast the Spearman rho of P^i and $\tilde{P}^{ii}(1)$ is typically very high, for example 0.92 for AZN in the first half of 2009. Because $(P^i)^2$ is subtracted in computing \tilde{P}^{ii} , such a large correlation is not automatic. This indicates that the most active members also tend to trade in intermittent bursts – if they become active, they tend to remain active, and vice versa. Not surprisingly the Spearman rho between $C^{ii}(1)$ and $\tilde{P}^{ii}(1)$ is small, indicating that correlation in trading direction does not imply correlation in trading activity.

The relationship between conditional and unconditional activity is illustrated differently in Figure 6. We plot $P^{ij}(\tau = 1)$ against $P^i P^j$ for all pairs of members i and j . The off-diagonal elements $i \neq j$ tend to be very close to the identity line, indeed somewhat below it, while the diagonal elements are well above the identity. This shows that there is strong persistence in the activity of brokers, but little coordination between the activities of different brokers. Larger values of τ show similar behavior though with somewhat smaller amplitude.

These results indicate that the persistence of order flow stems from the persistence in both trading direction and activity of individual members of the exchange. They also show that the dominance of splitting over herding is also apparent in the activity level P^{ij} , where we see that the diagonal elements P^{ii} dominate over the off-diagonal elements P^{ij} with $i \neq j$.

¹² $P^{ii}(\tau)$ converges to a constant value when τ is large. In contrast, the deviation in activity level, $\tilde{P}^{ii}(\tau) = P^{ii}(\tau) - (P^i)^2$, converges to zero. Thus on double logarithmic scale one is able to see that $\tilde{P}^{ii}(\tau)$ is asymptotically roughly a straight line, in contrast to P^{ii} .




Figure 6: The joint probability of activity $P^{ij}(\tau = 1)$ as a function of $P^i P^j$, the hypothetical joint probability under the assumption of independence. Each symbol indicates a given pair of members i and j for AZN during the first half of 2009. Diagonal elements $i = j$ are indicated by blue crosses and off-diagonal elements $i \neq j$ by black circles. The plot is on double logarithmic scale. For comparison the identity line $P^{ij}(\tau) = P^i P^j$ is shown as a solid black line. The fact that the diagonal elements $P^{ii}(\tau)$ consistently lie well above the identity line is consistent with the large contribution of splitting. In contrast the off-diagonal elements P^{ij} (with $i \neq j$) tend to cluster along or even slightly below the identity line.

5. Do these results apply to investors?

If we assume a one-to-one mapping between investors and member firms, then the results of the previous section show that splitting strongly dominates herding. However, most of the member firms are at least partially acting as brokers, lumping together orders from many investors, while investors may be using multiple brokers, or varying their choice of broker. To what extent do the results of the previous section also apply to investors?

In this section we confront this problem by developing a set of null hypotheses. This involves formulating more explicit models of investor behavior for which we can compute the decomposition of order flow. We then explicitly model brokerage, formalizing the notion of a brokerage map assigning investors to brokers, and studying how brokerage distorts the decomposition of order flow. It is then possible to compare the behavior observed in the data to a set of null hypotheses combining different scenarios for investor behavior and brokerage. While brokerage can substantially distort the decomposition, we show that the dominance of splitting over herding for member firms is so strong that under reasonable assumptions it must be true for investors as well.

5.1. Three hypotheses about persistence in order flow

We now formulate models of splitting and herding. At this stage we leave the definitions as general as possible, specifying only what is needed to derive the decomposition, and make them more specific later as needed.

As we have defined it here, “herding” refers to any phenomenon that induces synchronized trading behavior between different individuals, inducing positive correlations in order flow. There are several quite different causal mechanisms capable of generating such behavior. To illustrate the notion of a common cause we introduce a model of herding based on a common response to public information. To illustrate the notion of mutual influence we introduce a model of herding based on imitation. Order splitting is more straightforward, so one model

suffices. There are of course many variations in how one can define such models; we make some arbitrary choices, aiming for simplicity.

5.1.1. Order splitting

Under order splitting the autocorrelation of order flow is generated entirely by investors who split large trades into small pieces and execute them incrementally. In a “pure” order splitting model information is private and IID. Since by assumption the autocorrelation between different investors is zero, the persistence is entirely due to splitting,

$$\begin{aligned} C_{split}(\tau) &= C(\tau), \\ C_{herd}(\tau) &= 0. \end{aligned} \tag{12}$$

5.1.2. Herding model I: Public information

Our first herding model produces herding behavior via a similar response to a common external cause – the herd is directed from outside. Assume an exogenously given public information signal I_t with homogeneous investors who respond to the signal similarly except for possible variations in factors such as speed. I_t is random real-valued variable. Let the trading frequency of investor i be $P^i > 0$, normalized so that $\sum_{i=1}^M P^i = 1$, where M is the number of investors. Assume an information signal I_t arrives at time t and generates n_t orders with the same sign as I_t , where $n_t = f(|I_t|)$ and f is a non-decreasing function. The investor placing each order is drawn according to that investor’s trading frequency P^i . We assume sampling with replacement, so that a given investor may be drawn more than once. At time $t + n_t$ a new information signal I_{t+n_t} arrives and the process is repeated¹³.

Persistence in order flow can result from several different effects: (1) strong autocorrelations in I_t , such as those of a long-memory process. (2) heavy tails in the distribution $P(I_t)$; (3) nondeterministic behavior in f , generating or contributing to heavy tails in n_t ; (4) any combination of the previous three effects. All of these produce similar results. For definiteness we assume (1), but this choice makes no difference.

The decomposition of the autocorrelation can be computed in closed form. Since the investors are chosen randomly the timing and sign of the trade are independent of the investor, which implies

$$\begin{aligned} C^{ij}(\tau) &= C(\tau), \\ P^{ij}(\tau) &= P^i P^j. \end{aligned}$$

The second of these relations implies that $\tilde{P}^{ij}(\tau) = 0, \forall i, j$. Eq. (9) then implies that

$$\begin{aligned} C_{split}(\tau) &= \sum_i P^{ii}(\tau) C^{ii}(\tau) = C(\tau) \sum_i (P^i)^2 \\ C_{herd}(\tau) &= \sum_{i \neq j} P^{ij}(\tau) C^{ij}(\tau) = C(\tau) \sum_{i \neq j} P^i P^j. \end{aligned} \tag{13}$$

The decomposition thus depends only on the sum of the square of the trading frequency P^i of the investors. This is closely related to the variance

$$\text{Var}[P] = \frac{1}{M} \sum_i (P^i)^2 - \left(\frac{1}{M} \sum_i P^i \right)^2 = \frac{1}{M} \sum_i (P^i)^2 - \frac{1}{M^2}. \tag{14}$$

¹³ As stated this model generates n_t strings of sequential orders all of the same sign. It is possible to make the order flow look more realistic by injecting orders with a random sign, but the only effect is to decrease the prefactor of the autocorrelation without otherwise affecting its time dependence.

Substituting into Eq. (13) yields

$$\begin{aligned} C_{split}(\tau) &= C(\tau) \left(\frac{1}{M} + M \text{Var}[P] \right) \\ C_{herd}(\tau) &= C(\tau) \left(\frac{M-1}{M} - M \text{Var}[P] \right). \end{aligned} \quad (15)$$

Thus the relative size of C_{split} and C_{herd} only depends on the number of investors M and the variance of their trading frequencies P^i .

The variance of P lies in the range $0 \leq \text{Var}[P] \leq \frac{1}{M} \left(1 - \frac{1}{M}\right)$. The bounds occur when:

- *All investors are equally active, $\text{Var}[P] = 0$.*

$$\begin{aligned} C_{split} &= \frac{C(\tau)}{M} \\ C_{herd} &= \frac{M-1}{M} C(\tau). \end{aligned}$$

- *All trading is concentrated in one investor, $\text{Var}[P] = \frac{1}{M} \left(1 - \frac{1}{M}\right)$.*

$$\begin{aligned} C_{split} &= C(\tau) \\ C_{herd} &= 0. \end{aligned}$$

If $C(\tau) > 0$ then C_{split} is always positive and C_{herd} is always non-negative. As the trading goes from uniformly distributed to concentrated in a single member, C_{split}/C grows from 0 to 1 while C_{herd}/C decreases from 1 to 0. Thus the heterogeneity in investor's trading frequencies controls the relative importance of the herding vs. splitting components of the autocorrelation function.

5.1.3. Herding model II: Imitation

Our second herding model produces herding via investors influence on each other, so that herding is generated endogenously without need for an external coordination mechanism. We assume that investors exist within a social network where information is transmitted between neighbors. The transmission of information depends on the topology of the network, which also determines the autocorrelation of order flow. In this model the investors herd by following each other.

The nodes of the social network correspond to investors and links correspond to influence. The network is connected, so that all nodes are joined to the same graph by at least one link. Investor i has a binary state $s_t^i = \pm 1$ that indicates whether in the absence of influence this investor prefers to buy or sell. The dynamics are characterized by a parameter $p \in [0, 1]$ that describes the degree of imitation between agents.

The order flow is generated using the following algorithm: At time t an investor i is chosen randomly and an order with sign s_t^i is submitted. Then $t \rightarrow t + 1$ and each neighbor j of investor i is considered. With probability p investor j submits an order $\epsilon_t^j = s_t^i$ and changes state to $s_{t+1}^j = s_t^i$, and with probability $1 - p$ investor j submits an order s_t^j and her state is left unaltered. The time is once again incremented and the next neighbor, if any, is considered. Once all the neighbors of investor i have been considered a new node i' is chosen at random and the process continues.

There is no need for exogenous information in this model. Providing the system is initially placed in a sufficient diverse initial configuration, it will generate time correlated but otherwise random order flow. In the noiseless version this model has two fixed points where all the

agents are found in the same state (buy or sell). If the network is large enough this state cannot be reached in a reasonable time. It is also possible to generalize the model to inject external information by occasionally randomly altering the sign of the states of one of the nodes. Providing this is not done too often, however, it does not substantially alter the characteristics of the order flow. The degree distribution $p(\ell)$ and the parameter p determine the autocorrelation of the order flow¹⁴.

We cannot compute the splitting vs. herding decomposition for this model in closed form. The simulations presented in the next section show that the imitation model has essentially the same decomposition as the public information model, i.e. Eq. (15) is a good approximation of its decomposition.

5.2. Models of brokerage

In the real data we cannot observe investors – we only know the member of the exchange who executes the trade. As already mentioned, members often act as brokers for other investors. The distortion this introduces has the potential to alter the decomposition. In this section we address this problem by introducing the concept of a brokerage map and studying how the decomposition is altered under various hypotheses about brokerage. We will consider several different types of brokerage maps, as described below.

5.2.1. Fixed random brokerage

This model assumes that each investor randomly chooses a single broker and thereafter always executes through that same broker and that broker only. To match the real data we have to constrain the random choices to match the brokerage trading frequency P'^i . This is important, because as we will show the amount of distortion is strongly dependent on P'^i . Mirroring the same notation that we used for investors, assume there are M' brokers labeled by an index i and that P'^i is the trading frequency of broker i .

Under this brokerage map the decomposition is distorted as follows:

- *Order splitting.* The decomposition is unchanged, i.e. $C_{split}(\tau) = C(\tau)$ and $C_{herd} = 0$. This is true regardless of P'^i .
- *Public information.* The decomposition is given by Eq. (15), except that $M = M'$ and $\text{Var}[P] = \text{Var}[P']$. Thus, depending on $\text{Var}[P']$, the result can either be dominated by herding (when there are many brokerages with the same trading volume) or dominated by splitting (when one brokerage dominates the trading volume). While it might seem surprising that this no longer depends on M or $\text{Var}[P]$, one should bear in mind that under fixed random brokerage we are assuming the correlation for investors within a given broker is no different than for investors with different brokers, and $\text{Var}[P']$ and hence the decomposition at the level of brokerage data are determined by M , M' and $\text{Var}[P]$, which have to be consistent with P'^i .
- *Imitation.* We are not able to compute the distortion in closed form. Instead we simulate the imitation model and compare it to the public information model under the assumption that both have the same autocorrelation function $C(\tau)$. We construct a social network using preferential attachment (Simon, 1955; Barabasi et al., 1999). An initial node is created, and then new nodes are incrementally generated and connected to a randomly chosen pre-existing node with a probability of attachment proportional to its degree. Because we connect to only one pre-existing node at a time the resulting graph is a tree¹⁵.

¹⁴ The degree of a node of a graph is the number of links connected to that node

¹⁵ We could more realistically attach to multiple pre-existing nodes, but this does not affect the results.

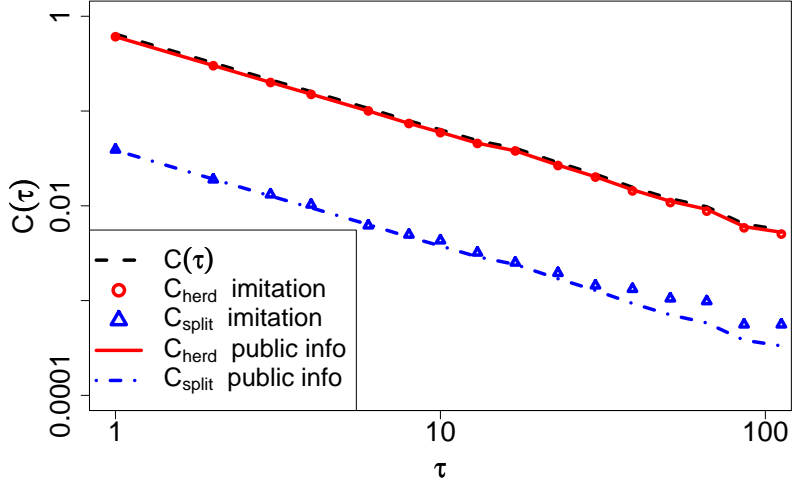


Figure 7: Splitting vs. herding decomposition for the public information and imitation models with a fixed random brokerage map. For the imitation model order flow is generated as described in the text, and the splitting (blue triangles) and herding (red circles) contributions are computed numerically using Eq. (9). For the public information model we assume the same total autocorrelation $C(\tau)$ (black dashed line) as the imitation model, and compute the splitting (blue dots) and herding (red line) contributions using Eq. (15). Broker trading frequencies P'^i are chosen to match AZN in the first half of 2009. Herding strongly dominates splitting and the decomposition for the two models is nearly identical.

Once the investor social network is built we construct a fixed random brokerage map. We assume $M' = 50$ brokers and match their trading frequency to the 50 largest exchange members for *AZN* during the first half of 2009. The imitation frequency $p = 0.9$ with $M = 10,000$ investors. Simulations are run for 10^6 time steps.

The resulting autocorrelation and its decomposition under both the imitation model and the public information model are shown in Fig. 7. The result is completely dominated by the herding component. While there is a non-zero splitting component, it is more than an order of magnitude smaller. Furthermore the decomposition for the two models is nearly identical.

For the imitation model the persistence of order flow derives from the heavy tailed degree distribution of the social network. The process of imitation converts the scale free power law degree distribution of the social network, which by construction is of the form $p(\ell) \sim \ell^{-\eta+1}$, into order flow with power law decay of autocorrelations. In fact, in the limit $p \rightarrow 1$ it can be shown that the autocorrelation function decays asymptotically as $C(\tau) \sim \tau^{-(\eta-1)}$.

Thus for order splitting the fixed random brokerage map introduces no distortion, whereas for herding models the distortion is strongly dependent on the brokerage trading frequencies. When the trading frequencies are uniform these models strongly favor herding, but as it becomes concentrated in a few brokerages the splitting component increases.

5.2.2. Dynamically random brokerage

To conceal order flow some investors may vary their choice of brokers. Dynamically random brokerage captures this strategy in its extreme form: On every trade each investor randomly chooses a broker according to the broker trading frequencies P'^i . There is thus no memory of the past, and no allegiance of any investor to any broker. Random choice implies that $P'^{ij} = P'^i P'^j$, and therefore the decomposition is given by Eq. (15). The resulting decomposition is thus identical to that of the public information model under a fixed random broker map. This is true regardless of the investor model.

As we saw in the previous section, when we use realistic broker trading frequencies the herding component is much larger than the splitting component. Thus, we have the interesting result that the distortion caused by dynamically random brokerage creates the appearance of herding, regardless of what the investors are actually doing. Given that we observe a strong dominance of splitting this is serendipitous, as it implies that dynamically random brokerage is a small effect, a result that is interesting in and of itself.

5.2.3. Correlation of social network and brokerage

Perhaps the previous examples are too random. What if brokerage choice is correlated with investor behavior? To study this question we have generalized the imitation model to allow for the possibility that neighbors in the influence network also tend to use the same broker.

In the generalized imitation model the social network and the brokerage assignments are made in tandem. As we construct the social network, when we attach each new investor to a previous investor, we assign the same broker with probability Φ or a random broker with probability $1 - \Phi$. As a result there is a correlation between neighborhood relationships in the trading graph and brokerage assignment¹⁶.

5.3. Comparison with real data

We now have all the tools needed to address the main question: Given that brokerage can distort the splitting/herding decomposition, can we nonetheless say with confidence that the dominance of splitting over herding observed in the exchange member trading data reflects the behavior of investors? We address this question by examining different scenarios for investor/brokerage to provide null hypotheses for comparison to the data. A summary of the comparisons is given in Figure 8 and discussed in detail below.

5.3.1. Public information/fixed random brokerage

This scenario combines the public information model and fixed random brokerage, as described in Section 5.2.1. The decomposition is given by Eq. (15) and depends only on the number of brokerages M' and their variance $\text{Var}[P']$. In the right panel of Figure 4 we compare the predictions to real data. For the data $S(\tau) = C_{split}(\tau)/C(\tau)$ increases from about 0.7 to 2 with increasing τ , in contrast to the predicted value, which is roughly 0.06 independent of τ . Thus the order splitting component is a factor of 10 – 30 higher than predicted.

A further comparison is given in Figure 8, where we plot the splitting ratio $S(\tau) = C_{split}(\tau)/C(\tau)$ averaged for lags $1 \leq \tau \leq 50$ as a function of the variance of the brokerage trading frequencies, assuming $M' = 50$. The values from the real data samples are very well-separated from the prediction for all 102 cases, indicating that this scenario cannot explain the data.

5.3.2. Generalized imitation model/fixed random brokerage

This scenario combines the generalized imitation model (including correlations between neighbors and brokerage) and fixed but random brokerage, as discussed in Section 5.2.3. The predicted decomposition is compared to the 102 data samples in Figure 8. The figure shows the splitting ratio $S(\tau) = C_{split}/C(\tau)$ and the broker-neighbor correlation Φ as a function of the brokerage trading frequency variance $\text{Var}[P']$. Because $S(\tau)$ and Φ are not linearly related we use different scales on the right and left. While a few of the data sets approach the prediction, almost all of them are well-separated from the predictions under this scenario. Realistic values

¹⁶ An inconvenient feature of this model is that Φ and P'^i are not independent, so it is harder to construct a network with given values of P'^i . As Φ varies from zero to one the trading frequencies of the brokers necessarily become more concentrated.




Figure 8: The splitting decomposition for the 102 real data samples (indicated by red x's) compared to the predictions of three different null hypotheses. The variance of the brokerage trading frequencies varies along the horizontal axis. On the left the trading frequencies are equal, on the right trading is concentrated in a few brokerages. The left vertical axis indicates the splitting ratio $S(\tau) = C_{split}(\tau)/C(\tau)$ averaged over $1 \leq \tau \leq 50$ and the right vertical axis indicates the broker-neighbor correlation Φ (for the generalized imitation model only). The null hypotheses are generated by combining a model for investor herding with a model of brokerage. The blue circles represent simulations of the generalized imitation model with Fixed Random Brokerage (FRB). The black curve is the public information model (with $\Phi = 0$) and Fixed Random Brokerage; alternatively it represents the prediction of the Dynamically Random Brokerage (DRB) model under any investor model. All 102 samples are outside the range predicted by any of the three null hypotheses.

of S in the center of the distribution of real values require $\Phi \approx 1$, which seems absurdly high, and more than half of the values of S cannot be reached at all, indicating that this scenario cannot explain the data.

5.3.3. Any investor model/dynamic random brokerage

As already discussed, if investors randomly pick their brokers, with no allegiance to any broker, the investor's behavior is irrelevant, and the properties of the decomposition are determined entirely by the brokerage trading frequencies. The prediction under this scenario is represented by the solid black curve in Figure 8. This is the same as the public information/fixed random brokerage scenario. The data are very well-separated from the predictions, indicating that this scenario cannot explain the data.

5.3.4. Order splitting/partially random brokerage

What about order splitting? As already discussed in the introduction, the order splitting model of Lillo et al. (2005) is capable of matching the empirical autocorrelation $C(\tau)$ quite well. The splitting ratio of this model is $S(\tau) = 1$, independent of τ and independent of the brokerage map. This ratio is too high for $\tau \leq 5$, where $S \approx 0.7$. This could be due to one of two reasons, that we cannot distinguish: (1) Some short term herding behavior or (2) partially dynamic brokerage choice. By this we mean that some investors (but not others) might randomly vary their brokers, or investors might have partial allegiance to a small set of brokers, randomly choosing between them but avoiding the others. As the magnitude of the random choice increases so does the herding component, and as we know from the previous scenario, unless a single brokerage dominates the trading (which is not the case for the real data), this strongly distorts the decomposition in favor of herding.

For $\tau \geq 10$ the splitting ratio $S(\tau) = 1$ predicted under the order splitting scenario is actually too low, due to anti-herding. We present a few clues about anti-herding in the next section.

For the combinations of herding and brokerage models considered here the distortion caused by broker maps cannot explain the empirical dominance of splitting over herding. We can thus be confident that there are no reasonable circumstances in which the distortion introduced by the use of brokerage data could plausibly be large enough to be consistent with the empirical results.

6. Clues about the cause of anti-herding

As we have shown in Section 4.1, the herding component of the order flow autocorrelation is often negative for $\tau > 10$. This implies that buying by one investor tends to invoke selling by other investors, a phenomenon that gives the appearance of “anti-herding”. In this section we first test the statistical significance of anti-herding and then perform some empirical investigations that give a clue as to its origin.

6.1. Anti-herding is statistically significant

To assess whether anti-herding is statistically significant we perform a one-sided significance test by comparing the real data to the null hypothesis that both the signs and brokerage codes are assigned randomly. Realizations of this null hypothesis are obtained by randomly shuffling both the signs and brokerage codes. For each of the 102 datasets we produced 10^6 realizations of the null hypothesis, and for each realization we computed the splitting and herding components of the autocorrelation. Then for each lag τ we estimated the fraction of random realizations having a herding component smaller than the value observed in the corresponding real sample. If this fraction is smaller than 5%, we reject the one-sided null hypothesis. Figure 9 shows the fraction of sets for which we reject the null hypothesis. For small values of τ we never reject the hypothesis because the real herding component of the autocorrelation function is significantly positive (as we have seen above). For values of τ between roughly 15 and 80 we reject the null hypothesis in approximately 80% of the sets. This means that for these lag values the real herding component of the autocorrelation is statistically significantly negative and thus anti-herding is statistical significant. We now explore the possible origin of this phenomenon.

6.2. What underlies “anti-herding”?

We now show that anti-herding is associated with a difference in the response of brokers to market orders, depending on whether or not the order changes the price and whether or not it is from the same or a different broker. If a market order placed by broker i changes the price, broker j is less likely to place market orders in the same direction, while the behavior of broker i is unchanged¹⁷.

We use the notation MO_t^0 for a market order at time t that does not change the price and MO_t' for a market order that changes the price. Conditioned on either of these events, the probabilities for subsequent market orders to have the same sign are

$$\begin{aligned} P(\epsilon_t = \epsilon_{t+\tau} \mid MO_t^0) \\ P(\epsilon_t = \epsilon_{t+\tau} \mid MO_t') \end{aligned} \tag{16}$$

Assuming that on average buy and sell trades have the same size and that brokers’ inventories are bounded, these probabilities for large τ should converge to the unconditional probability

¹⁷ Our guess that whether or not a market order changes the price is an important determinant of the behavior was inspired in part by a previous study by Eisler et al. (2011b); see also Toth et al. (2011a).




Figure 9: The fraction of the 102 data sets for which the herding component of the autocorrelation C_{herd} has a p value less than 5% under the IID null hypothesis.

$P(\epsilon_t) = 1/2$. In Figure 10 these are shown for AZN, plotted as a function of τ on a double logarithmic scale. To make the τ dependence clearer, we also plot the excess probabilities

$$\begin{aligned}\tilde{P}^0(\tau) &\equiv P(\epsilon_t = \epsilon_{t+\tau} \mid MO_t^0) - 1/2 \\ \tilde{P}'(\tau) &\equiv P(\epsilon_t = \epsilon_{t+\tau} \mid MO_t') - 1/2.\end{aligned}$$

The decay of both \tilde{P}^0 and \tilde{P}' is approximately a power law. However, when the original market order changes the price, the decay is faster. This means that, all else being equal, when a market order does not change the price the persistence of the sign of subsequent orders is much stronger than when it does change the price.

We now study how the identity of the broker affects this behavior. We look at the probability that the signs of the orders at t and $t+\tau$ are the same, conditioned on whether they were placed by the same broker and also on whether the event at t changed the price or not. We study the following probabilities:

$$\begin{aligned}P(\epsilon_t^i = \epsilon_{t+\tau}^j \mid i = j ; MO_t^0) & \quad P(\epsilon_t^i = \epsilon_{t+\tau}^j \mid i \neq j ; MO_t^0) \\ P(\epsilon_t^i = \epsilon_{t+\tau}^j \mid i = j ; MO_t') & \quad P(\epsilon_t^i = \epsilon_{t+\tau}^j \mid i \neq j ; MO_t')\end{aligned}\tag{17}$$

In Figure 11 we plot the probabilities above on a double logarithmic scale. When the broker at time t and $t+\tau$ is the same there is no qualitative difference, regardless of whether the order at time t changed the price, other than a small shift in scale.

In contrast, when different brokers place the orders at time t and $t+\tau$, the behavior changes dramatically. If the market order at time t does not change the price, the probability that the signs of the two orders are the same is lower than before, but still positive. In contrast, if the market order at time t changes the price, for most lags τ the subsequent order is more likely to have the opposite sign. Thus, after a market order that does not change the price other brokers tend to place their orders in the same direction, indicating a slight herding on short time scales, but after a market order that changes the price, the opposite happens: they are more likely to put their orders in the opposite direction, corresponding to anti-herding¹⁸.

¹⁸ The interesting point is that splitting traders seem to act like “noise traders”, i.e. they do not adapt their

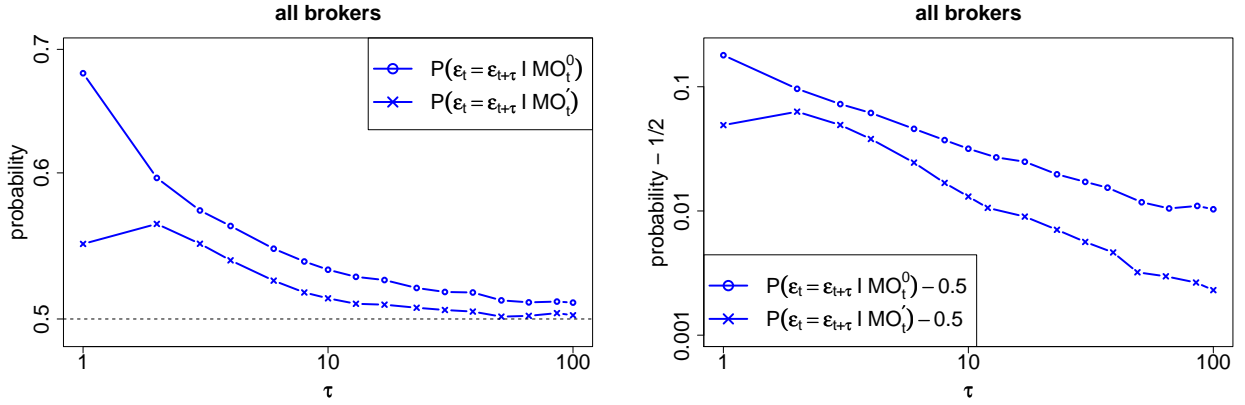


Figure 10: Probability that market orders at times t and $t + \tau$ have the same sign, conditioned on whether or not the first market order changed the price. The case of no price change is given by $P(\epsilon_t = \epsilon_{t+\tau} | MO_t^0)$ and the case of a price change by $P(\epsilon_t = \epsilon_{t+\tau} | MO_t^1)$. The plots on the left and right are the same, except that on the right we subtract the unconditional probability $P(\epsilon_t) = 1/2$ to make the functional dependence on time and the difference in the convergence rates clearer. When the original market order changes the price, the decay is much faster. Data is for AZN for 2000-2009.




Figure 11: The probability that market order signs separated by lag τ are the same, conditioned on whether the members placing the orders are the same and whether or not the first market order changes the price, as described in Eq. (17). *Left panel:* $i = j$: The same member places the order at t and $t + \tau$. Whether or not the market order at time t changed the price (MO_t^1 vs. MO_t^0) makes little difference in the behavior, other than a slight shift in scale. *Right panel:* $i \neq j$: Different members place the order at t and $t + \tau$. Whether or not the market order at time t changes the price now makes a big difference in the behavior. After an event that does not change the price, other brokers tend to put their order in the same direction with a probability higher than 0.5. After an event that changes the price, in contrast, at most lags it is slightly more probable that other brokers put their orders in the opposite direction. Data is for AZN for 2000-2009.




Figure 12: Splitting and herding component of the autocorrelation function of market order signs conditional on the event at time t . The left panel is conditioned on the case where the original market order does not change the price, and the right panel is the case where it does. The negative contribution to the herding term is entirely due to the latter case. Data is for AZN for 2000-2009.

This is confirmed by decomposing the autocorrelation conditional on whether or not the market order at time t changed the price, shown in Figure 12. More precisely, we apply the herding/splitting decomposition of Eq. (9) to the conditional autocorrelations $E[(\epsilon_t - \mu)(\epsilon_{t+\tau} - \mu)|MO_t^0]$ and $E[(\epsilon_t - \mu)(\epsilon_{t+\tau} - \mu)|MO_t^1]$, where μ is the mean market order sign. When the market order at t does not change the sign, the herding component decays to zero but never becomes very negative; in contrast, when the market order at t changes sign, it is negative for roughly $\tau > 5$. The shape of the herding term is qualitatively similar to the probabilities shown in Fig. 11.

The conclusion is that the observed anti-herding effect is related to the difference in the response of brokers to market orders placed by *others*, depending on whether or not they changed the price. When a market order does not cause a price change brokers continue being more likely to place orders of the same sign, regardless of who placed the original market order. In contrast, if a market order triggers a price change, other brokers place fewer market orders in the same direction than in the opposite direction.

7. Conclusions

We have shown that in the LSE the cause of long-memory on short timescales (corresponding to 500 transactions or less, typically about an hour) is overwhelmingly due to autocorrelated trading by investors (which we call splitting), rather than correlated trading by different investors (which we call herding). Our belief is that this is mainly due to single individuals who consistently execute their large orders through at most a few brokers, splitting them into small pieces. This is exactly what one expects from algorithmic brokerages, who take large orders, split them into pieces and execute them throughout the day. We observe only weak herding at short time scales, and for longer times we observe anti-herding, i.e. the trading of a given individual is anti-correlated with that of other individuals. We are only able to obtain statistically significant results for fairly short time scales of less than an hour. It is quite possible that herding is a stronger factor on longer timescales.

The analysis we have performed is based on data identifying the members of the exchange,

behavior to their own impact. The reason for this could be that they have already calculated their impact in their estimates of the trading cost, so it is expected. The same effect is also mentioned by Toth et al. (2011a).

who often act as brokers for other investors. In order to understand whether these results apply to investors, we developed a set of models for investor behavior and choice of broker. We then combined these into scenarios of investor behavior/brokerage, and used these as null hypotheses. The decomposition of order flow under these null hypotheses is then compared to the data. Even though our analysis was done at the level of brokers, we are able to show that the qualitative conclusion that splitting dominates herding applies at the level of investors as well.

This behavior is extremely consistent across different stocks and time periods. It is also fairly consistent across member firms. While a few member firms have less directional persistence than others, the vast majority are quite persistent, and in an almost identical way. Furthermore the persistence exhibits itself similarly in both trading direction and trading activity.

Under the interpretation that the strong positive autocorrelation of investors is due to order splitting, the fact that investors split their orders so strongly implies that the market is in a certain sense typically out of equilibrium. That is, if investors revealed their intentions when they made decisions, rather than concealing their intentions and trading gradually, prices at any given moment might be substantially different.

8. Acknowledgments

We thank Jean-Philippe Bouchaud and Terry Odean for useful comments, and National Science Foundation grant 0624351 for support. Any opinions, findings, and conclusions or recommendations expressed are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- Banerjee, A., 1992. A simple model of herd behavior. *Quarterly Journal of Economics* 107 (3), 797–817.
- Banerjee, A., 1993. The economics of rumours. *Review of economic studies* 60 (2), 309.
- Barabasi, A.-L., Albert, R., Jeong, H., 1999. Mean-field theory for scale-free random networks. *Physica A* 272, 173–189.
- Barber, B., Odean, T., Zhu, N., 2009. Systemic noise. *Journal of Financial Markets* 12, 547–569.
- Beran, J., 1994. *Statistics for Long-Memory Processes*. Chapman & Hall, New York.
- Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of the limit order book and order flow in the paris bourse. *The Journal of Finance* 50 (5), 1655–1689.
- Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *The Journal of Political Economy*, 100 (5), 992–1026.
- Bouchaud, J.-P., Farmer, J. D., Lillo, F., 2009. How markets slowly digest changes in supply and demand. In: Hens, T., Schenk-Hoppe, K. (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*. Elsevier, pp. 57–156.
- Bouchaud, J.-P., Gefen, Y., Potters, M., Wyart, M., 2004. Fluctuations and response in financial markets: The subtle nature of “random” price changes. *Quantitative Finance* 4 (2), 176–190.
- Bouchaud, J.-P., Kockelkoren, J., Potters, M., 2006. Random walks, liquidity molasses and critical response in financial markets. *Quantitative Finance* 6 (2), 115–123.

- Chan, L. K., Lakonishok, J., 1993. Institutional trades and intraday stock price behavior. *Journal of Financial Economics* 33, 173–199.
- Chan, L. K., Lakonishok, J., 1995. The behavior of stock prices around institutional trades. *The Journal of Finance* 50 (4), 1147–1174.
- Cont, R., Bouchaud, J.-P., 2000. Herd behavior and aggregate fluctuations in financial markets. *Macroeconomic Dynamics* 4 (2), 170–196.
- Danielsson, J., Payne, R., 2001. Measuring and explaining liquidity on the limit order book: Evidence from reuters d2000-2. Tech. rep., Financial Market Group, London School of Economics.
- Ding, Z., Granger, C. W. J., Engle, R. F., 1993. A long memory property of stock returns and a new model. *Journal of Empirical Finance* 1 (1), 83–106.
- Eisler, Z., Bouchaud, J.-P., Kockelkoren, J., 2011a. Models for the impact of all order book events. Tech. rep., arXiv:1107.3364.
- Eisler, Z., Bouchaud, J.-P., Kockelkoren, J., 2011b. The price impact of order book events: market orders, limit orders and cancellations. *Quantitative Finance*, 1–25arXiv:0904.0900.
- Ellul, A., Holden, C., Jain, P., Jennings, R., 2007. A comprehensive test of order choice theory: Recent evidence from the nyse. *Journal of Empirical Finance* 14 (5), 636–661.
- Farmer, J. D., Gerig, A., Lillo, F., Waelbroeck, H., 2011. How efficiency shapes market impact. Tech. rep., <http://arxiv.org/abs/1102.5457>.
- Gerig, A., 2007. A theory for market impact: How order flow affects stock price. Ph.D. thesis, University of Illinois.
- Iori, G., 2002. A microsimulation of traders activity in the stock market: the role of heterogeneity, agents' interactions and trade frictions. *Journal of Economic Behavior and Organization* 49 (2), 269–285.
- Kirman, A., Teyssiere, G., 2002. Microeconomic models for long memory in the volatility of financial time series. *Studies in Nonlinear Dynamics and Econometrics* 5 (4), 281–302.
- Kyle, A. S., 1985. Continuous auctions and insider trading. *Econometrica* 53, 1315–1335.
- Lakonishok, J., Shleifer, A., Vishny, R., 1992. The impact of institutional trading on stock prices. *Journal of Financial Economics* 32 (1), 23–43.
- LeBaron, B., Yamamoto, R., 2007. Long-memory in an order-driven market. *Physica A* 383, 85–89.
- LeBaron, B., Yamamoto, R., 2008. The impact of imitation on long-memory in an order-driven market. *Eastern Economic Journal of Econometrics* 34 (4), 504–517.
- LeBaron, B., Yamamoto, R., 2010. Order-splitting and long-memory in an order-driven market. *European Physical Journal B*. 73 (1), 51–57.
- Lillo, F., Farmer, J. D., 2004. The long memory of the efficient market. *Studies in Nonlinear Dynamics & Econometrics* 8 (3).
- Lillo, F., Mike, S., Farmer, J. D., 2005. Theory for long memory in supply and demand. *Physical Review E* 7106 (6), 066122.

- Moro, E., Moyano, L. G., Vicente, J., Gerig, A., Farmer, J. D., Vaglica, G., Lillo, F., Mantegna, R., 2009. Market impact and trading protocols of hidden orders in stock markets. *Physical Review E*. 80 (6), 066102.
- Orlean, A., 1995. Bayesian interactions and collective dynamics of opinion: Herd behavior and mimetic contagion. *Journal of Economic Behavior and Organization* 28 (2), 257–274.
- Raafat, R., Chater, N., Frith, C., 2009. Herding in humans. *Trends in Cognitive Sciences* 13 (10), 420–428.
- Simon, H. A., 1955. On a class of skew distribution functions. *Biometrika* 42 (3/4), 425–440.
- Tedeschi, G., Iori, G., Gallegati, M., 2009. The role of communication and imitation in limit order markets. *The European Physical Journal B - Condensed Matter and Complex Systems* 71 (4), 489–497.
- Toth, B., Eisler, Z., Lillo, F., Bouchaud, J.-P., Kockelkoren, J., Farmer, J., 2011a. How does the market react to your order flow? Tech. rep., <http://arxiv.org/abs/1104.0587>.
- Toth, B., Lemperiere, Y., Deremble, C., de Lataillade, J., Kockelkoren, J., Bouchaud, J.-P., 2011b. Anomalous price impact and the critical nature of liquidity in financial markets. Tech. rep., <http://arxiv.org/abs/1105.1694>.
- Vaglica, G., Lillo, F., Moro, E., Mantegna, R., 2008. Scaling laws of strategic behavior and size heterogeneity in agent dynamics. *Physical Review E*. 77, 0036110.
- Wermers, R., 1999. Mutual fund herding and the impact on stock prices. *Journal of Finance* 55 (2), 581–622.
- Yeo, W. Y., 2008. Serial correlation in the limit order flow: Causes and impact. In: FMA European Conference (Prague). <Http://69.175.2.130/~finman/Prague/Papers/SCorLOFCL.pdf>.