

How does the market react to your order flow?

B. Tóth^{a,b}, Z. Eisler^b, F. Lillo^{a,c,d},
J.-P. Bouchaud^b, J. Kockelkoren^b, J.D. Farmer^a

April 5, 2011

^a *Santa Fe Institute, 1399 Hyde Park Rd., Santa Fe NM 87501*

^b *Capital Fund Management, 6, blvd Haussmann 75009 Paris, France*

^c *Scuola Normale Superiore di Pisa, Piazza dei Cavalieri 7, I-56126 Pisa, Italy*

^d *Dipartimento di Fisica, Università di Palermo, Viale delle Scienze, I-90128 Palermo, Italy*

Abstract

We present an empirical study of the intertwined behaviour of members in a financial market. Exploiting a database where the broker that initiates an order book event can be identified, we decompose the correlation and response functions into contributions coming from different market participants and study how their behaviour is interconnected. We find evidence that (1) brokers are very heterogeneous in liquidity provision – some are consistently liquidity providers while others are consistently liquidity takers. (2) The behaviour of brokers is strongly conditioned on the actions of *other* brokers. In contrast brokers are only weakly influenced by the impact of their own previous orders. (3) The total impact of market orders is the result of a subtle compensation between the same broker pushing the price in one direction and the liquidity provision of other brokers pushing it in the opposite direction. These results enforce the picture of market dynamics being the result of the competition between heterogeneous participants interacting to form a complicated market ecology.

Keywords: Financial markets, market microstructure, limit order book, order flow, behavioural economics

1 Introduction

Empirical studies of order flow and market impact have recently boomed due to the availability of high frequency data, where all market events (trades, limit orders, cancellations) are recorded. These data sets allow one to investigate many interesting statistical regularities at the order book level, and shed light on the price formation mechanisms. One of the most interesting results established in the recent literature concerns the long-ranged correlated nature of the order flow, and a detailed understanding of the impact of individual transactions on prices (see [1] for a recent review). However, most empirical studies to date rely on a purely anonymous order flow: trades, limit orders and cancellations cannot be traced to a particular agent in the market. In order to access this data, some special agreement must be reached

with exchanges, which regulators allow or even promote in certain conditions (for example in order to investigate the role of “high-frequency traders” in the market, see [2, 3]).

The data that we exploit here is unfortunately not that detailed, but allows us to identify the activity of market members of the LSE (London Stock Exchange) in the period May, 2000 – December, 2002. Since members are often brokers who act on the behalf of many final clients, the granularity of the order flow is still quite coarse, but some interesting conclusions can be still drawn from this data, as has been shown in [4, 5, 6]. Note that since most members also act as brokers, we will throughout use the word “broker” as synonymous with “member” and we will use these terms interchangeably. Here, we want to adapt a formalism introduced by some of us in Ref. [7] to investigate the correlation and impact of various types of order book events. In that paper, events were broken down into six categories: market orders, limit orders and cancellations, and for each type whether the event immediately changes the midpoint price or not. In principle, further categories can be envisaged, and here we add the brokerage code as an additional tag.

This decomposition allows us to separate the total impact of a given type of order book event into a contribution from the same broker and a contribution from all other brokers. We find that these two contributions very nearly offset each other, leading to a total impact that is nearly constant in time and *much smaller than either of these contributions*. This is the central result of this paper, which confirms the dynamical liquidity picture put forth in [8, 9, 10, 11, 7], according to which the highly persistent sign of market orders must be buffered by a fine-tuned counteracting limit order flow in order to maintain statistical efficiency (i.e. that the price changes are close to unpredictable, in spite of the long-ranged correlation of the order flow). We believe that the quantitative result presented here is a very important ingredient to understand the dynamics of markets, since it explicitly demonstrates that the stability of markets relies on a rather precise balance between liquidity taking and liquidity providing, and that small fluctuations of one or the other can lead to micro-liquidity crises and price jumps [12, 13].

We also report some additional statistical regularities in order placement that result from interaction/imitation between different brokers, depending on the type of order book event, and on clear signatures of heterogeneity between different order placement strategies. These raw empirical results would certainly deserve an additional layer of quantitative interpretation, which we, however, leave quite open at this stage, partly for lack of more precise data.

2 Data and notation

In this paper we analyze data on 7 of the most liquid stocks traded at the LSE during the period May, 2000 – December, 2002*. Our data set contains all on-book events for the stocks. We only consider the usual trading time between 8:00–16:30, all other periods are discarded. The unique feature of our data set is that each order is characterized by a membership code identifying the initiator of the order. These codes uniquely represent the member firms of the LSE even if we are not able to identify the firms by name[‡]. The activity level of brokerages is very heterogeneous. For example, in a typical month for AZN the 5 most active market members are responsible for 40-50% of transactions and the 15 most active ones are responsible

*The studied stocks are AstraZeneca (AZN), BHP Billiton (BLT), Lloyds Banking Group (LLOY), Prudential (PRU), Rentokil Initial (RTO), Tesco (TSCO), Vodafone Group (VOD).

[‡]Actually the brokerage codes in our data are reshuffled at the beginning of each month. However, since most of the results we show happen at the intraday scale, we can ignore the effect of the reshuffling.

ticker	number of events	tick/avg price [$\times 10^{-4}$]	typical number of brokers	Gini coeff.	std($\log_{10} \alpha_i$)
AZN	1,846,922	3.49	95	0.80	1.12
BLT	958,573	7.69	69	0.75	1.03
LLOY	1,761,548	7.84	104	0.79	1.13
PRU	1,312,220	7.44	88	0.77	1.10
RTO	714,665	11.07	65	0.75	1.02
TSCO	1,175,353	10.68	87	0.78	1.10
VOD	2,712,084	15.23	134	0.82	1.13

Table 1: Summary statistics for the data we study here. The first three columns are the ticker, the total number of events and the ratio of the tick size to the average price. Letting α_i be the fraction of the trades during a given month initiated by broker i , the other fourth column shows the typical (average) number of active brokerages (defined to be those with $\alpha_i > 0.01$). The fourth column is the Gini coefficient of α_i , and the sixth is the standard deviation of $\log_{10} \alpha_i$.

for 80-90% of transactions. Thus the trading activity is strongly concentrated in a relatively small number of member firms. In Table 1 we show summary statistics of the stocks and the brokerages. We present the number of events for each stock, the ratio of the tick size to the average price, the typical (average) number of active brokers. In a given month we define α_i as the fraction of trades initiated by broker i . The typical Gini coefficient of α_i is about 0.8, and the typical standard deviation of $\log_{10} \alpha_i$ is a little more than one, indicating that the typical difference between the activity of two brokers chosen at random is more than an order of magnitude.

Following Ref. [7] we will analyse time series of order book events. We use the name “event” for any change in the order book that modifies the bid or ask price or the volume quoted at these prices. Events will be used as the unit of time. Since there can be many events between two transactions this notion of “event time” is similar but more fine-grained than the notion of transaction time used in many papers. The price just before the t^{th} event is defined as the midpoint price p_t , i.e. the average of the best ask and best bid quote. We use ticks as the units of price. The type of event at time t will be denoted by π_t . The upper index ‘ (‘prime’) denotes that an event changed the price p_t , and the upper index 0 that it did not. The possible events are:

- MO^0 is an effective market order¹ that does not change the price.
- MO' is an effective market order that does change the price.
- LO^0 is a limit order at the current bid or ask (which does not change the price).
- LO' is a limit order inside the spread (which does change the price).
- CA^0 is cancellation at the bid or ask that does not remove all the volume quoted (and thus does not change the price).
- CA' is a cancellation at the bid or the ask that does remove all the volume quoted (and thus does change the price).

¹ By *effective market order* we mean any event that generates immediate transactions with existing orders in the limit order book.

Abbreviations without the upper index (MO, CA, LO) refer to events whether or not they change the price. We will not explicitly consider limit orders and cancellations inside the order book, because they do not have an immediate effect on the best quotes. The event at time t is given a sign ϵ_t according to its expected long-term effect on the price. For market orders and limit orders this corresponds to order signs, i.e., $\epsilon_t = 1$ for buy orders and -1 for sell orders. Cancellation of a sell limit order has $\epsilon_t = 1$ while cancellation of a buy limit order has $\epsilon_t = -1$ (the sign is reversed because the effect on the price is in the opposite direction).

We will use the indicator function $I(A)$ which is defined as $I(A) = 1$ if the condition A is true and $I(A) = 0$ otherwise. For example, the indicator variable $I(\pi_t = \pi)$ is 1 if the event at time t is of type π and zero otherwise. The unconditional probability of the event type π is by definition $P(\pi) = \langle I(\pi_t = \pi) \rangle$. The market member acting at time t will be denoted by b_t , and the indicator function $I(b_t = b)$ is 1 if the broker acting at time t is b and zero otherwise.

The average behaviour of the price ℓ time steps after an event of a particular type π_1 defines the corresponding response function (or average impact function) [14]

$$\mathcal{R}_{\pi_1}(\ell) = \frac{\langle (p_{t+\ell} - p_t) I(\pi_t = \pi_1) \epsilon_t \rangle}{P(\pi_1)}. \quad (1)$$

This response can be divided into a part that is the response due to the actions of the same broker as the one active at time t ($\mathcal{R}_{\pi_1}^{\text{same}}(\ell)$), and the response due to other brokers ($\mathcal{R}_{\pi_1}^{\text{diff}}(\ell)$). The response due to the same broker can be written

$$\mathcal{R}_{\pi_1}^{\text{same}}(\ell) = \frac{\left\langle \sum_{t'=t}^{t+\ell-1} (p_{t'+1} - p_{t'}) I(b_{t'} = b_t) I(\pi_{t'} = \pi_1) \epsilon_{t'} \right\rangle}{P(\pi_1)}. \quad (2)$$

$\mathcal{R}_{\pi_1}^{\text{same}}(\ell)$ is the expected price change between time t and $t + \ell$ caused by the further actions of the same broker that acted at time t and ignoring all other brokers (since $I(b_{t'} = b_t)$ picks out only events from the same broker). Conversely the response that is only due to other brokers can be written

$$\mathcal{R}_{\pi_1}^{\text{diff}}(\ell) = \frac{\left\langle \sum_{t'=t}^{t+\ell-1} (p_{t'+1} - p_{t'}) I(b_{t'} \neq b_t) I(\pi_{t'} = \pi_1) \epsilon_{t'} \right\rangle}{P(\pi_1)}. \quad (3)$$

Trivially $\mathcal{R}_{\pi_1}^{\text{same}}(\ell) + \mathcal{R}_{\pi_1}^{\text{diff}}(\ell) = \mathcal{R}_{\pi_1}(\ell)$. Furthermore, we can define the different contributions to the response $\mathcal{R}_{\pi_1}^{\text{same}}(\ell)$ coming from the possible π_2 types of events occurring at t' as

$$\mathcal{R}_{\pi_1, \pi_2}^{\text{same}}(\ell) = \frac{\left\langle \sum_{t'=t}^{t+\ell-1} (p_{t'+1} - p_{t'}) I(\pi_{t'} = \pi_2) I(b_{t'} = b_t) I(\pi_t = \pi_1) \epsilon_{t'} \right\rangle}{P(\pi_1)}, \quad (4)$$

and similarly for $\mathcal{R}_{\pi_1, \pi_2}^{\text{diff}}(\ell)$.

In a similar way one can define the signed-event correlation function

$$C_{\pi_1, \pi_2}(\ell) = \frac{\langle I(\pi_{t+\ell} = \pi_2) \epsilon_{t+\ell} I(\pi_t = \pi_1) \epsilon_t \rangle}{P(\pi_1) P(\pi_2)}, \quad (5)$$

which again can be divided as the sum of the correlation between events initiated by the same broker ($C_{\pi_1, \pi_2}^{\text{same}}(\ell)$) and the correlation between events initiated by different brokers ($C_{\pi_1, \pi_2}^{\text{diff}}(\ell)$),

$$C_{\pi_1, \pi_2}^{\text{same}}(\ell) = \frac{\langle I(\pi_{t+\ell} = \pi_2) \epsilon_{t+\ell} I(\pi_t = \pi_1) \epsilon_t I(b_{t+\ell} = b_t) \rangle}{P(\pi_1)P(\pi_2)}, \quad (6)$$

$$C_{\pi_1, \pi_2}^{\text{diff}}(\ell) = \frac{\langle I(\pi_{t+\ell} = \pi_2) \epsilon_{t+\ell} I(\pi_t = \pi_1) \epsilon_t I(b_{t+\ell} \neq b_t) \rangle}{P(\pi_1)P(\pi_2)}. \quad (7)$$

For simplicity, when talking about the response functions and the correlations, we will simply refer to $\mathcal{R}_{\pi_1}^{\text{same}}(\ell)$ and $C_{\pi_1, \pi_2}^{\text{same}}(\ell)$ as the contribution of the *same broker*, while we will refer to $\mathcal{R}_{\pi_1}^{\text{diff}}(\ell)$ and $C_{\pi_1, \pi_2}^{\text{diff}}(\ell)$ as the contribution of *other brokers*.

When showing our results, we averaged over all the 7 stocks studied. However, we checked all the results for each stock individually and found that the results are very similar.

3 Heterogeneity of broker liquidity provision

Are brokers homogeneous in the sense that an event from a given broker has the same statistical properties as an event from any other broker? Or are brokers heterogeneous in the sense that their events have different statistical properties? In this section we show that when it comes to liquidity provision brokers are very heterogeneous.

We first study liquidity provision. Let $\#MO'_b$ be the number of price changing market orders placed by broker b in a given month, and similarly $\#LO'_b$ be the number of price changing limit orders placed by broker b in that month. The fraction

$$f_b^{\text{MO}'} = \frac{\#MO'_b}{\#MO'_b + \#LO'_b} \quad (8)$$

is the fraction of price changing market orders placed by broker b in a given month (we use a month because the brokerage codes are shuffled every month). A high value of $f_b^{\text{MO}'}$ thus implies that a broker tends to be a liquidity taker and a low value of $f_b^{\text{MO}'}$ implies that she tends to be a liquidity provider². The distribution³ of the average values of $f_b^{\text{MO}'}$ is shown in the left panel of Fig. (1). This distribution is extremely broad, with average values of $f_b^{\text{MO}'}$ ranging from $f_b^{\text{MO}'} \approx 0$, indicating a broker that acts as a market maker, using limit orders almost exclusively, to $f_b^{\text{MO}'} \approx 0.8$, indicating a broker that predominately uses market orders.

Statistical testing shows that the wide variation in $f_b^{\text{MO}'}$ is due to real heterogeneity among the brokers rather than statistical fluctuations. If we define a series with values of one or zero depending on whether a price changing order is an effective market order or a limit order, the characteristic time for the autocorrelation to decay into the noise level is less than ten price changing orders. This enables us to estimate the standard deviation of the $f_b^{\text{MO}'}$ value for each broker by drawing block bootstrap samples for all brokers in all periods (with block sizes of 10 orders). The typical standard deviation for the $f_b^{\text{MO}'}$ values is less than 0.04. Thus the broad

²In studying $f_b^{\text{MO}'}$ we look only at price changing market and limit orders (and not all MO and LO), since these are what have an immediate effect on the price. This we do in order to include less bias in the statistics due to orders being placed and cancelled almost immediately.

³Note that the distribution is not weighted by the size of the broker.

distribution of values of $f_b^{MO'}$ shown in Fig. (1) is almost entirely real variation, corresponding to heterogeneous behaviour across brokers.

We find that $f_b^{MO'}$ does not show significant dependence on the size of the broker (defined by his total number of orders of any type), so there is no simple relation between transaction volumes and the use of market vs. limit orders that would reflect, for example, some systematic difference of information between large and small brokers.

Nonetheless, the results summarized in the right panel of Fig. (1) suggest that brokers with different levels of liquidity provision specialize in different types of execution. We show the average immediate impact $\mathcal{R}_\pi(\ell = 1)$ of an order of type $\pi = LO'$ and $\pi = MO'$ as a function of $f_b^{MO'}$. To reduce the statistical variation we bin the brokers into five groups according to $f_b^{MO'}$, and plot the average value of $\mathcal{R}_\pi(\ell = 1)$ in each bin against the average value of $f_b^{MO'}$ for that bin. The impact for market orders decreases slightly with $f_b^{MO'}$, indicating that brokers who use market orders more frequently get better execution. For limit orders, in contrast, the lowest impact is for low values of $f_b^{MO'}$, indicating that brokers who use limit orders more frequently get better execution. For the top two quintiles of $f_b^{MO'}$ there is essentially no difference between the impact of limit orders and market orders, while for the lowest quintile the difference is dramatic – limit orders have an impact of about a tick, while market orders have more than 1.4 ticks of impact.

There are two possible interpretations of this result. One is that brokers who use a given type of order more often are more skillful at using this type of order (assuming that moving the price less means better execution, this might not always be true). Another is that brokers who are predominantly doing market making tend to transact passively with less information, so that their limit orders have less impact than those of directional traders. However, in the rare cases where market makers use market orders, e.g. to flatten their positions at the end of the day, they tend to execute under unfavourable circumstances, e.g. under more time pressure to execute a large quantity.




Figure 1: (left) The distribution of the ratio $f_b^{MO'}$, defined in Eq. (8), which takes on low values for brokers who tend to be liquidity providers and higher values for brokers that tend to be liquidity takers. We only include brokers whose number of trades is at least 1% of the total in a given month. (right) The immediate impact $\mathcal{R}_\pi(\ell = 1)$ for price changing limit orders $\pi = LO'$ and $\pi = MO'$ orders is plotted as a function of $f_b^{MO'}$, the fraction of price changing market orders. To reduce statistical fluctuations the data are binned into five groups according to $f_b^{MO'}$.

4 Regularities in order placement

This analysis can also be used to gain insight into the origin of long-ranged correlation in the sign of orders. Using the above formalism, we can break up this correlation into different contributions, depending on whether or not the event is price changing and whether the broker is the same or different. As we will see, the response of brokers to their own price changes is quite different than their response to the price changes of others.

In the left panel of Fig. (2) we plot some relevant correlations for market orders, in particular $C_{MO^0,MO^0}(\ell)$ and $C_{MO',MO^0}(\ell)$, conditioning on the same broker and on different brokers. It can be seen that:

- i) The autocorrelations $C_{MO^0,MO^0}^{same}(\ell)$, $C_{MO',MO^0}^{same}(\ell)$, and $C_{MO^0,MO^0}^{diff}(\ell)$ all behave similarly. They are positive for lags up to more than 500, and they all decay roughly as power laws, $\ell^{-\gamma}$ with $\gamma \approx 0.5$. This is very close to the exponent found in [8, 14] for the unconditional autocorrelation of market order sign. The three autocorrelations differ in amplitude: $C_{MO^0,MO^0}^{same}(\ell)$ is largest, with a value of roughly 0.8 for lag one, then $C_{MO',MO^0}^{diff}(\ell)$, with a value of roughly 0.2 for lag one, and finally $C_{MO',MO^0}^{same}(\ell)$, with a value less than 0.1 at lag one. They all decay roughly in the same way across the entire range (see also [16]).
- ii) In contrast, the autocorrelation $C_{MO',MO^0}^{diff}(\ell)$, which measures the response of another broker to a price changing market order, is weakly but consistently negative and shows no clear pattern of decay. Furthermore its behaviour is completely different to that of $C_{MO^0,MO^0}^{diff}(\ell)$.

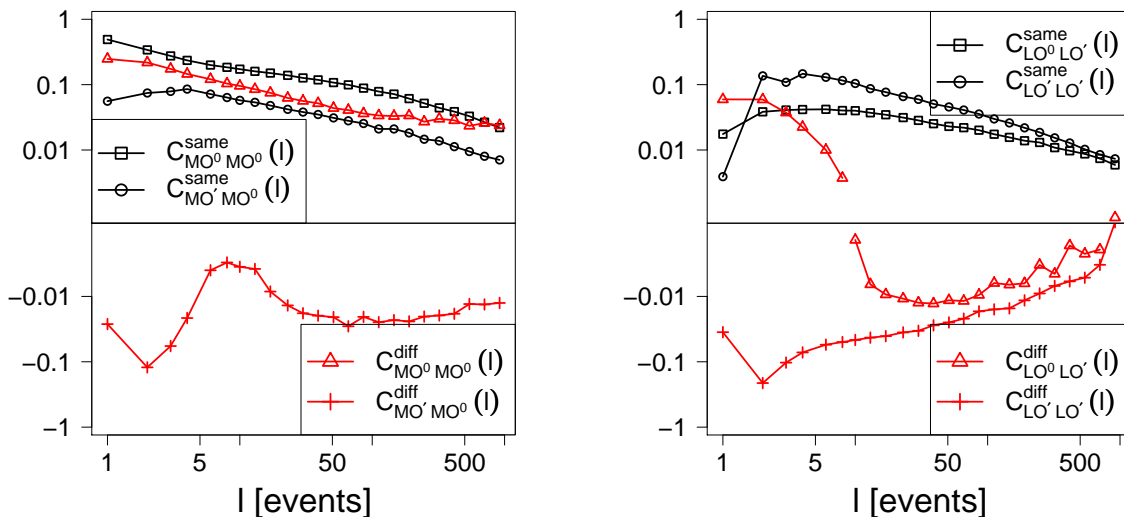


Figure 2: Correlation functions between events. (left) The correlation $C_{MO^0,MO^0}(\ell)$ between non-price-changing market orders at two different times, and the correlation $C_{MO',MO^0}(\ell)$ between a price-changing market order and a non-price-changing market order. In each case we present results conditioned on the same broker for the two events vs. a different broker for the two events. (right) The correlation $C_{LO^0,LO'}(\ell)$ for a non-price changing and a price changing limit order, and the correlation $C_{LO',LO'}(\ell)$ for two price changing limit orders.

Stated in different terms, the first set of results imply that the response to a non-price-changing market order is always the same: Whether subsequent orders come from the same

broker or a different broker, subsequent order placements tend to be of the same sign. This is in contrast to the response to a price-changing market order. In this case the same broker tends to keep placing orders of the same sign as her original order, while other brokers tend to place orders of opposite signs to the original order.

The observation of a long-memory autocorrelation functions for orders placed by the same broker provides additional evidence that at brokerage level the long-memory of order flow is primarily driven by the splitting of large metaorders into small pieces, as posited in [17] (see also [11, 16] for additional evidence supporting this hypothesis). This happens whether or not the original market order causes a price change. Given that a price change in the same direction is unfavourable (a buyer does not want the price to rise), this is a bit surprising at first sight. However, note that although there is long-memory in either case, the magnitude of the effect is roughly a factor of eight smaller when the original market order causes a price change.

The fact that order flow by the same broker continues in the same direction suggests that parent orders (or ‘metaorders’) are executed to a large extent independently of the price change, even when this change is caused by trading. From a behavioural finance perspective, agents doing splitting appear to act like “noise traders”, i.e. they adapt their own order flow very little to the effect of their own recent trading, as assumed in [17]. This could be because agents have already factored in the impact of their trades in their estimates of trading costs and thus are rationally following their plan to let their order run until executed, having already taken this impact for granted.

The fact that the autocorrelation in response to a non-price-changing market order is also positive, even for different brokers, suggests herding behaviour. This could happen in two ways: One is that after a non-aggressive market order MO^0 , other brokers jump on the bandwagon, thinking there might be some information in the initial trade, or alternatively, the other brokers might be responding to the same information signals as the original broker, but with a slight lag.

Surprisingly, though, if the original market order is price-changing, the sign of this effect is reversed. In this case the original market order triggers the activity of “other brokers” with the opposite sign. The observation of a price rise converts the other brokers (or at least a majority of them) from buyers to sellers, or from sellers to buyers. This is compatible with the idea of a large liquidity buffer that reveals itself as soon as the price changes, which seems to be enough to overwhelm the herding effect that is seen when there is no price change.

Similar behaviour is seen for limit orders as well when the broker is the same, but the situation is somewhat altered when the broker is different. The right panel of Fig. 2 shows $C_{LO^0,LO'}(\ell)$ and $C_{LO',LO'}(\ell)$. Again, we find that $C_{LO^0,LO'}^{\text{same}}(\ell)$ is similar to $C_{LO',LO'}^{\text{same}}(\ell)$. In this case, however, there is no herding on the part of other brokers since both $C_{LO^0,LO'}^{\text{diff}}(\ell)$, $C_{LO',LO'}^{\text{diff}}(\ell)$ are negative (except at very short lags $\ell < 5$ for $C_{LO^0,LO'}^{\text{diff}}(\ell)$). Thus at long time lags other brokers respond to limit orders by placing orders of the opposite sign, whether or not the original order was price-changing.

Even restricting to LO’s and MO’s and discarding CA’s, one can define a total of 32 different correlation functions, whereas Fig. 2 only shows 8 of these. It is not worth showing all 24 others, but Fig. 3 shows three that are interesting. The question we ask is the following: conditioned on a non-price-changing order by broker b , what is the excess probability that the same broker b issues a price-changing order with the same sign after ℓ trades? We ask this question when the original order is either a market order MO^0 or a limit order LO^0 . The four probabilities are given by $P(\pi_2)C_{\pi_1,\pi_2}^{\text{same}}(\ell)$ with $\pi_1 = MO^0$ or LO^0 and $\pi_2 = MO'$ or LO' , as

shown in Fig. 3. (Note that $C_{LO',LO'}^{\text{same}}(\ell)$ already appears in Fig. 2, right, so that we only add three new quantities to the above.)




Figure 3: The probability for a given broker to stick with non-price-changing limit orders or market orders (conditioned on starting with non-price-changing orders). In each case we show the excess probability $P(\pi_2)C_{\pi_1, \pi_2}^{\text{same}}(\ell)$ that the same broker b issues a price-changing order with the same sign after ℓ trades, with $\pi_2 = LO'$ or MO' . The left panel shows $\pi_1 = MO^0$ and the right panel $\pi_1 = LO^0$.

What transpires from these plots is that conditioned on the fact that a broker decided to execute using a non-aggressive market order, he will make roughly equal use of MO' and LO' in the future, whereas after deciding to place a non-aggressive limit order, the probability to continue using limit orders in the future is much larger than switching to market orders. This result, however, mostly comes from the contribution of the brokers with the smallest $f_b^{MO'}$, i.e. those brokers who act as market makers, for which such a strategy is indeed expected.

5 Balance between liquidity taking and liquidity providing

Let us investigate another aspect of the intertwined liquidity dynamics. It is known from previous results [18, 14, 7] that the average impact of market orders, $\mathcal{R}_{MO}(\ell)$, first increases rapidly with ℓ and then becomes flat for large ℓ . To understand the reason for the flattening of the response function, we studied the contributions coming from the actions of the same broker and of other brokers. When disaggregating these two effects, we find that the flat response function comes from a nearly exact cancellation between $\mathcal{R}_{MO}^{\text{same}}(\ell)$, which increases as roughly $\sqrt{\ell}$, and $\mathcal{R}_{MO}^{\text{diff}}(\ell)$, which decreases as roughly $\sqrt{\ell}$. The growth of $\mathcal{R}_{MO}^{\text{same}}(\ell)$ can be directly understood from the self-correlation $C^{\text{same}}(\ell)$ described above[†]. The impact functions illustrating this behaviour for event types MO^0 and MO' can be seen in Figure 4. To better show the power law increase of the absolute values, in Figure 5 we plot $\mathcal{R}_{MO}^{\text{same}}(\ell)$ on a log-log scale, together with $-\mathcal{R}_{MO}^{\text{diff}}(\ell) + \text{const.}$, where a constant term was just added in order to better visualise the similarity of the curves.

[†]More precisely, for large tick stocks the response is related to the integral of the correlation functions [7]. Therefore, since the correlations $C_{MO, \pi_2}^{\text{same}}(t')$ all have a power law decay with exponents close to $\gamma = 0.5$, we expect the response to increase with an exponent close to $1 - \gamma = 0.5$. In fact $\mathcal{R}_{MO^0}^{\text{same}}(\ell)$ increases as a power law with an exponent between 0.5 and 0.63 for all the studied stocks, while $\mathcal{R}_{MO'}^{\text{same}}(\ell)$ show slightly lower exponents between 0.38 and 0.51.

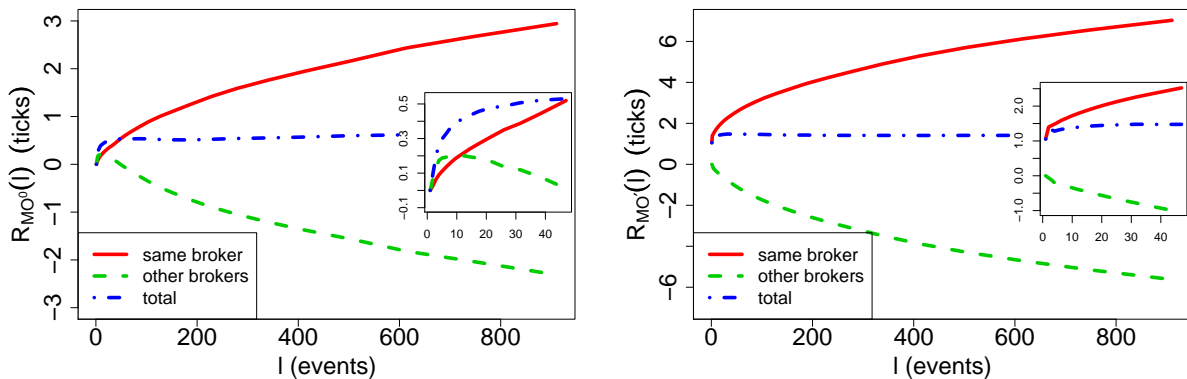


Figure 4: The response function $\mathcal{R}_{\pi_1}(\ell)$ and its contributions coming from the orders of the same broker ($\mathcal{R}_{\pi_1}^{\text{same}}(\ell)$) and of different brokers ($\mathcal{R}_{\pi_1}^{\text{diff}}(\ell)$). (left) The case of $\pi_1 = \text{MO}^0$. (right) The case of $\pi_1 = \text{MO}'$. The insets show a zoom for small ℓ .

Interestingly, the sum of these two responses give a total response that is much weaker in absolute value and is flat for time lags $\ell \gtrsim 100$. This is the central result of our study. The short time behaviour ($\ell \lesssim 10 - 20$) is different when (i) the initial order left the price unchanged (MO^0) and (ii) when it changed the price (MO'). In case (i) the contribution coming from different brokers is initially positive and then becomes negative, while in case (ii) it is immediately negative. Therefore, upon an aggressive buy market order (MO') from one broker, the other brokers (probably those with small $f_b^{\text{MO}'}$'s) react by immediately providing liquidity to the market, and continue to do so during the whole ‘buying spree’, thereby limiting the total upward price shift. After a non-aggressive market order, on the other hand, the herding mechanism described in the previous section explains the initial positive contribution of $\mathcal{R}_{\text{MO}^0}^{\text{diff}}(\ell)$ seen in Figure 4, left (inset).

The above findings extend to the decomposition of all types of impact functions $\mathcal{R}_\pi(\ell)$. After any type of event, the response due to the same broker’s actions is monotonically increasing. In contrast, the response due to the rest of the market is always monotonically decreasing for $\ell \gtrsim 10 - 20^*$.

6 Conclusions

The aim of this paper was to use the formalism introduced in [7] to exploit a database where the broker initiating an order book event can be identified, and decompose the correlations and response functions into different contributions. This allowed us to identify several interesting regularities in order placement, as summarized below:

- In Section 3 we present clear evidence that brokers are heterogeneous in their liquidity provision. Different brokers use different type of strategies, from market making strategies with a small fraction of market orders to liquidity taking strategies with a very large fraction of market orders.

*The frequency of events varies among stocks and changes in time. On average the frequency of events in the period studied was 0.28 events/sec.




Figure 5: The contributions to the response function $\mathcal{R}_{\pi_1}(\ell)$ from the same broker and from different brokers on a log-log scale. We show $\mathcal{R}_{MO^0}^{\text{same}}(\ell)$, together with $-\mathcal{R}_{MO^0}^{\text{diff}}(\ell) + \text{const.}$, where a constant term was just added in order to better visualise the similarity of the curves. (left) The case of $\pi_1 = MO^0$. (right) The case of $\pi_1 = MO'$.

- In Section 4 we confirm that the long-range correlation in the sign of market orders comes mostly from order splitting from a unique broker. There is, however, a certain degree of herding behaviour from other brokers which is evident as long as the price does not change. After a price changing market order, however, the broker responsible for it continues trading in the same direction regardless of his own impact, whereas other brokers start firing market orders in the opposite direction.
- The total impact of market orders is the result of a nearly perfect compensation between two opposite contributions: one resulting from the accumulation of orders in the same direction coming from the same broker, while the reaction of brokers who provide liquidity results in an impact of roughly equal magnitude but opposite sign. (The contribution by other brokers is always a bit smaller, so that the total impact has the correct sign).

These results suggest a picture of the ecology of markets anticipated in several papers (see e.g. [19, 20, 9, 4, 5]), where agents are both broadly heterogeneous in their expectations and strategies, and strongly interacting, with a complex intertwined dynamics between liquidity providers and liquidity takers. It is tempting to conjecture that these ingredients are crucial to understand the incipient instabilities of financial markets, epitomized by price jumps and volatility clustering [12, 13, 2].

References

- [1] J.-P. Bouchaud, J. D. Farmer, F. Lillo, *How markets slowly digest changes in supply and demand*, in: Handbook of Financial Markets: Dynamics and Evolution, North-Holland, Elsevier, 2009.
- [2] A. A. Kirilenko, A. S. Kyle, M. Samadi, T. Tuzun, *The Flash Crash: The Impact of High Frequency Trading on an Electronic Market*, Working Paper, (2010)
- [3] A. J. Menkveld *High Frequency Trading and The New-Market Makers*, Working Paper (2011)

- [4] F. Lillo, E. Moro, G. Vaglica, R. N. Mantegna, *Specialization and herding behavior of trading firms in a financial market*, New Journal of Physics 10, 043019 (2008)
- [5] E. Moro, J. Vicente, L. G. Moyano, A. Gerig, J. D. Farmer, G. Vaglica, F. Lillo, R. N. Mantegna, *Market impact and trading profile of hidden orders in stock markets*, Physical Review E 80, 066102 (2009)
- [6] B. Tóth, F. Lillo, J. D. Farmer, *Segmentation algorithm for non-stationary compound Poisson processes*, Eur. Phys. J. B **78**, 235-243 (2010)
- [7] Z. Eisler, J.-P. Bouchaud, J. Kockelkoren, *The price impact of order book events: market orders, limit orders and cancellations*, to appear in Quantitative Finance, <http://arxiv.org/abs/0904.0900>
- [8] F. Lillo, J. D. Farmer, *The long memory of the efficient market*, Studies in Nonlinear Dynamics & Econometrics, 8, 1 (2004).
- [9] J.-P. Bouchaud, J. Kockelkoren, M. Potters, *Random walks, liquidity molasses and critical response in financial markets*, Quantitative Finance, **6**, 115 (2006).
- [10] J. D. Farmer, A. Gerig, F. Lillo, and S. Mike. *Market Efficiency and the Long-Memory of Supply and Demand: Is Price Impact Variable and Permanent or Fixed and Temporary*. Quantitative Finance. **6** (2), 107 (2006).
- [11] A. Gerig, Ph.D. thesis, University of Illinois, (2007)
- [12] A. Joulin, A. Lefevre, D. Grunberg, J.-P. Bouchaud, *Stock price jumps: news and volume play a minor role*, Wilmott Mag., Sept/Oct 2008
- [13] J.-P. Bouchaud, *The Endogenous Dynamics of Markets: Price Impact, Feedback Loops and Instabilities*, in 'Lessons from the Credit Crisis', A. Berd Edt., Risk Publications (2011).
- [14] J.-P. Bouchaud, Y. Gefen, M. Potters, M. Wyart, *Fluctuations and response in financial markets: The subtle nature of random price changes*, Quantitative Finance, **4**, 176, (2004)
- [15] B. Efron, *Bootstrap methods: Another look at the Jackknife*, The Annals of Statistics, Vol. 7., No. 1., 1, (1979)
- [16] B. Tóth, I. Palit, F. Lillo, J. D. Farmer, *in preparation*
- [17] F. Lillo, S. Mike, J. D. Farmer, *Theory for long memory in supply and demand*, Physical Review E 71, 066122 (2005)
- [18] J. Hasbrouck, *Measuring the information content of stock trades*, Journal of Finance, **XLVI**, 179, (1991)
- [19] P. Handa, R. A. Schwartz, A. Tiwari, *The Ecology of an Order-Driven Market*, Journal of Portfolio Management, Winter, pp. 4756. (1998)
- [20] J. D. Farmer, *Market Force, Ecology and Evolution*. Industrial and Corporate Change **11** (5) 895-953 (2002).