

# Mechanical vs. informational components of price impact

J. Doyne Farmer<sup>1,\*</sup> and Neda Zamani<sup>2,1,†</sup>

<sup>1</sup>*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501*

<sup>2</sup>*School of IT, Faculty of Science, The University of Sydney, Australia*

We study the problem of what causes prices to change. It is well known that trading impacts prices – orders to buy drive the price up, and orders to sell drive it down. We introduce a means of decomposing the total impact of trading into two components, defining the *mechanical impact* of a trading order as the change in future prices in the absence of any future changes in decision making, and the *informational impact* as the remainder of the total impact once mechanical impact is removed. This decomposition is performed using order book data from the London Stock Exchange. The average mechanical impact of a market order decays to zero as a function of time, at an asymptotic rate that is consistent with a power law with an exponent of roughly 1.7. In contrast the average informational impact builds to approach a constant value. Initially the impact is entirely mechanical, and is about half as big as the asymptotic informational impact. The size of the informational impact is positively correlated to mechanical impact. For cases where the mechanical impact is zero for all times, we find that the informational impact is negative, i.e. buy market orders that have no mechanical impact at all generate strong negative price responses.

## I. INTRODUCTION

What causes prices to change? Despite a great deal of research on this subject the answer remains far from clear. On one hand, everyone agrees that prices respond to information – good news drives prices up and bad news drives it down. On the other hand, prices often change even when there is little information, sometimes by large amounts [7, 17]. How much of price changes are information driven, and how much are due to other factors?

At an immediate level it is clear that trading is an important cause of price change. When a trade is initiated by a buyer, the price tends to go up, and when it is initi-

ated by a seller, it tends to go down. The change in prices associated with a given trade is called price impact (or alternatively market impact), and has now been extensively studied [1, 2, 4–6, 8, 9, 11, 14–16, 18–20, 23–25, 28]. Price impact says nothing about why trades are made: They could be made because of new information or they could be made at random, for example, because a participant needs cash for reasons that are unrelated to anything going on in the market. This is further complicated by the fact that market participants regard trading and the price changes it produces as important signals about private information [12, 13]. When market participants observe a change in price, they may reason that others have private information, which may cause them to trade, setting off a cascading avalanche of price changes whose origin is difficult to as-

---

\*Electronic address: [jdf@santafe.edu](mailto:jdf@santafe.edu)

†Electronic address: [neda@cs.usyd.edu.au](mailto:neda@cs.usyd.edu.au)

certain.

Why do trades have price impact? One cause is mentioned above – they reflect information, and this information is incorporated in price changes. There is, however, an even more fundamental reason: When a trading order is placed it causes purely mechanical changes in prices. The word “mechanical” refers to the component of price changes that is deterministic, that occurs under the rules of the auction (dependent on the set of queued trading orders) even in the absence of any information. Although the mechanical impact of trading is often discussed (see, e.g. Hopman [16]), up until now no one has offered a precise definition of what it means, or made any suggestions as to how it can be measured.

In this paper we take advantage of the fact that in most modern financial markets prices are formed using a continuous double auction (see Section II A). Once the sequence of trading orders are given, the auction proceeds according to deterministic rules, acting as a deterministic dynamical system with exogenous inputs. While the placement of trading orders may depend on complicated factors, once the sequence of trading orders is given, price formation is purely mechanical.

We propose a definition for mechanical impact and introduce a practical method to compute it. The mechanical impact of a trading order can be defined as the change in future prices that occurs even if no other trading orders are changed in any way. This can be computed by introducing hypothetical alterations in the size of a trading order, and using the deterministic nature of the continuous double auction to simulate their effect on the future price sequence, holding all other aspects of the trading order sequence constant. In contrast, the informational impact of a trading order is what is left of the total impact after the mechanical impact is removed, i.e. it is the component that depends

on relationships between orders. As discussed in Section VII A, this can be either the causal effect that a given order has on future orders, or correlated effects between the placement of orders (e.g. due to a common cause). More precise definitions are given in Section III.

We find that mechanical impact behaves very differently than informational impact. The immediate effect of placing an order is entirely mechanical. The average mechanical impact decays to zero monotonically in time, although at a slow rate. The long time behavior is consistent with the hypothesis of a power law, with an exponent of about 1.7. In contrast, informational impact grows with time and approaches a constant value, at least over the time horizon where statistics remain reliable. On average the initial mechanical impact is about half the asymptotic informational impact. We find that the integrated mechanical impact and the informational impact are correlated, with the caveat that when the mechanical impact is identically zero, the informational impact remains nonzero but has the opposite sense that one normally expects (i.e. buy orders with zero mechanical impact generate negative price impact).

This paper is organized as follows: Section II gives a brief summary of the properties of our data set and gives some background information about the functioning of the continuous double auction and the London Stock Exchange in particular. In Section III we give more precise definitions of total price impact and its decomposition into mechanical and informational components. We then measure the average impacts and durations for real data in Section IV. In Section V we study the effects of long-memory in amplifying mechanical impact. In Section VI we study the correlations between mechanical and information impact. Finally in Section VII we discuss the implications and future

directions of this work.

## II. BACKGROUND

### A. Continuous double auction

We give a brief review of the continuous double auction, which is the most common mechanism used for trading in modern financial markets, and define some terminology that will be essential in the remainder of the paper. “Continuous” refers to the fact that the market is asynchronous, so that trading orders can be placed at any time, and orders are received one at a time, so there is unique ordering of events. “Double” refers to the fact that both buyers and sellers are allowed to update their orders at will. Orders contain both a trading quantity and a limit price, which is the worst price the trader is willing to accept<sup>1</sup>. The queue of unexecuted orders is called the *limit order book*. A transaction is generated whenever an order crosses the prices of orders of the opposite sign, e.g. if a buy order has a higher price than the lowest priced sell order. In this case we would say that the transaction is *buyer-initiated*; similarly, if a sell order crosses the best price, we say that the transaction is *seller-initiated*. If an order does not generate an immediate transaction, it is added to the limit order book without a transaction taking place. It is also possible to cancel any order sitting in the limit order book at any time.

Real markets have a variety of different types of possible orders that vary from mar-

ket to market, but for our purposes it is possible to categorize all orders into three types of *events*: *Effective market orders*, defined as any order or component of an order that generates a transaction; *Effective limit orders*, defined as any order or component of an order that does not generate an immediate transaction, and *cancellations*, defined as any when an order is removed from the limit order book without a transaction taking place. Our notion of effective events may not be in one-to-one correspondence with the actual orders that are placed. For example, in a situation where the lowest priced sell order on the book is to sell 1000 shares at 50 pounds, an order to buy 3000 shares at 50 pounds will result in a transaction for 1000 shares at 50 pounds and leave a buy order sitting in the book for 2000 shares. We will call such an order two effective orders, an effective market order followed by an effective limit order.

The LSE has two parallel markets, the on-book market and the off-book market. The on-book market operates via a continuous double auction as described above, in which the limit order book is transparently visible to everyone, but the identities of those placing the orders are concealed. The off-book market operates through a bilateral exchange in which agents contact each other via the telephone or message boards and know the identity of the agent they are trading with. Trades in the off-book market are revealed only after they occur, and the intention to trade beforehand is only communicated to a limited circle of contacts. For this reason the on-book market is widely regarded as the dominant force in price formation. It accounts for more than half the number of trades and about half the trading volume. Here we study only data from the on-book market.

---

<sup>1</sup> Sometimes the desired price is omitted, indicating a willingness to accept any price. This is called a *market order*. A buy market order is equivalent to a limit order with an infinite limit price, and a sell market order is equivalent to a limit price of zero.

## B. Data

For the purpose of this study we have used the TDS (Transaction Data Service) data set for the on-book market of the London Stock Exchange (LSE). The data set consists of records of orders placed or cancelled within a 3 year period during 2000-2002. We study three of the most liquid stocks, Astrazeneca (AZN), Vodafone (VOD), and Lloyds (LLOY). We will use AZN for all the figures in the paper, but we have repeated all the analyses for the other stocks as well. Orders placed during opening and closing auctions are excluded. The resulting reduced data set contains about 570K transactions and 3.7M events for AZN, 1.0M transactions and 4.0M events for VOD, and 600K transactions and 2.9M events for LLOY.

We have done some cleaning to reduce problems due to data errors. Because the data set contains both a record of transactions and a record of order placements we can test to be sure that both records are consistent. We find some problems that we do our best to correct. E.g. there are a few cases where orders are placed that are never removed. We remove such orders. A more serious problem is that the sequencing of the orders is not accurate for orders that are placed within the same second. This often results in nonsensical behavior, such as negative spreads. When this occurs we reorder the data to eliminate these problems. This reordering is not always unique. However, this is rare, so that the overall probability that an order is out of sequence is much less than one percent. Mis-orderings have a small effect at very short time scales (e.g. a few events), but essentially no effect on longer time scales.

In working with the data we have to deal with the problem of interday boundaries. Because the market closes at night and over weekends, there are gaps in the data, and the

behavior of the data across these gaps can be quite different than that within a trading day. To cope with this problem we have tried two different approaches. One is to reject any situations where we cross interday boundaries. The other is to include situations that cross interday boundaries, but to remove price changes that occur outside the period of our analysis. We use the latter approach for the results presented here, but we do not find that it makes a big difference in our results.

## III. DECOMPOSING PRICE IMPACT

### A. Total price impact

Throughout this paper we work with the logarithmic midprice  $p_t$ , which is defined as the average of the logarithm of the best price for a sell order and the logarithm of the best price for a buy order. The *total price impact* of an event at time  $t$  is defined as the difference between the price just before the event and the price  $\tau$  time units later, i.e.,  $\Delta p_\tau^T(t) = p_{t+\tau} - p_t$ . Letting  $s_t$  be the sign of an event at time  $t$ , we can merge buy and sell orders together and define the *average total price impact*<sup>2</sup> as

$$R^T(\tau) = \langle s_t(p_{t+\tau} - p_t) \rangle = \frac{1}{N} \sum_{t=1}^N s_t(p_{t+\tau} - p_t),$$

where  $N$  is the size of the sample and  $\langle \rangle_t$  indicates a time average over time  $t$ . This lumps together the price impact of buy orders (positive events) and sell orders (negative events).

---

<sup>2</sup> The multiplication by the sign avoids problems associated with asymmetries between buying and selling. This was introduced in reference [1], where it is called the response function.

In general these are different, but the difference is small and negligible for our purpose.

We have used two time units in this study: One is *event time*, where an event is an effective limit order, an effective market order, or a cancellation. The other is *transaction time*, where a transaction is defined as an effective market order<sup>3</sup>. The price corresponding to time  $t$  is defined to be just before the event or transaction occurs, and incremented by one immediately after that. For AZN about 17% of events in the data cause immediate transactions (i.e. market buy/sell events). Hence, for AZN every unit of transaction-time is the equivalent of roughly 6 event-time units. For AZN during this period on average one transaction time unit is about 40 seconds<sup>4</sup>.

## B. Mechanical impact

We define the mechanical impact in terms of the change in the midprice when an event is removed, but all other events are held constant. The idea is to measure the component of price change that is purely due to the presence of the event, excluding any effects that are caused by related changes in the sequence of events. For the purposes of this paper we base our definition on the removal of events, but this is not the only possibility – one can

<sup>3</sup> An effective market order may transact with several different orders in the limit order book. From the point of view of transaction time we consider this a single transaction.

<sup>4</sup> The definition of mechanical impact that we propose here relies on the use of natural time units, such as event time or transaction time, as opposed to real time units (seconds). The problem is that the number of events in any fixed interval of real time is highly variable, and so it is not clear how one would formulate Equation 2 in real time.

also use addition or partial additional or removal, as discussed in Section VII B.

Our definition of mechanical impact takes advantage of the fact that, under the rules of the continuous double auction, any initial limit order book and subsequent sequence of events generates a unique sequence of limit order books, which correspond to a unique sequence of midprices. The auction  $A$  can be regarded as a deterministic function

$$b_{t+1} = A(b_t, \omega_t)$$

that maps an event  $\omega_t$  and a limit order book  $b_t$  onto a new limit order book  $b_{t+1}$ . The event  $\omega_t$  can be an effective market order, limit order or cancellation, as described in Section II A. For a given sequence of events  $\Omega_{t+1}^{t+\tau} = \{\omega_{t+1}, \omega_{t+2}, \dots, \omega_{t+\tau}\}$  the auction  $A$  can be iterated to generate the limit order book  $b_{t+\tau}$  at time  $t + \tau$ ,

$$b_{t+\tau} = A(b_t, \Omega_{t+1}^{t+\tau}).$$

The continuous double auction can thus be thought of as a deterministic dynamical system with initial condition  $b_t$  and exogenous input  $\Omega_{t+1}$ .

Each limit order book  $b_t$  defines a unique logarithmic midprice  $p_t = p(b_t)$ . Simplifying the notation, we write the composition of the auction operator  $A$  and the midprice price operator  $p$  as  $\Pi = p \circ A$ . Thus, for any initial limit order book  $b_t$  and event sequence  $\Omega_{t+1}^{t+\tau}$  the logarithmic midprice at time  $t + \tau$  is

$$p_{t+\tau} = \Pi(b_t, \Omega_{t+1}^{t+\tau}). \quad (1)$$

There are a couple of situations that deserve more discussion. First, for the midprice to be well-defined it is necessary that the limit order book contain at least one order to buy and one order to sell. This is not always the case. For example for AZN, which is one of the more liquid stocks in the LSE, out

of roughly 570 thousand transactions we observe 17 cases where the midprice is not well-defined. Such situations become more common for less liquid stocks. In the case where this happens, we simply say that the price is undefined. In our empirical work we discard such cases. The second potentially problematic situation occurs when there is a cancellation of an order that does not exist. The most common cause of this is the lag between the receipt of information and the time needed to react to implement a cancellation, which can result in an order being executed before it can be cancelled. There are 231 cases of this for AZN during the period of our sample. When this occurs the exchange simply ignores the cancellation. We do the same, treating the cancellation as a null event, i.e. one that has leaves the limit order book unchanged.

For the purposes of this paper we define the mechanical impact of an event  $\omega_t$  in terms of its effect on the price under its hypothetical removal. More precisely, making use of Equation 1, we define the *mechanical impact*  $\Delta p_\tau^M(t)$  as

$$\Delta p_\tau^M(t) = \Pi(b_t, \Omega_{t+1}^{t+\tau}) - \Pi(b_{t-1}, \Omega_{t+1}^{t+\tau}). \quad (2)$$

In other words, the mechanical impact of event  $\omega_t$  at time  $t+\tau$  is the difference between the real price and the hypothetical price when  $\omega_t$  is removed, but the subsequent sequence of events is left the same. This isolates the part of the price impact that is “purely mechanical”, in the sense that it is generated solely by the effect of placing an order in the book and observing its effect under the deterministic operation of the continuous double auction. Although we have not been able to prove this, we conjecture that when  $\omega_t$  is a buy order ( $s_t = +1$ ) the continuous double auction guarantees that  $\Delta p_\tau^M \geq 0$  for all  $\tau$ , and for sell orders ( $s_t = -1$ )  $\Delta p_\tau^M \leq 0$ . In our empirical investigations we have not seen

any exceptions to this conjecture. For  $\tau \leq 0$  the mechanical impact is by definition zero.

Removing an order to compute the mechanical impact increases the number of cancellations of orders that do not exist. For example, suppose that in the real event sequence the buy market order  $\omega_t$  removes a sell limit order  $\hat{\omega}$ . When we generate the hypothetical sequence by removing  $\omega_t$ ,  $\hat{\omega}$  is left in the book, so that the best ask price remains lower than it did in the real sequence. This can cause a subsequent buy limit order to be executed that would otherwise have remained in the book and later been cancelled. When this occurs, as already stated, we treat this as a null event. For AZN, for example, we find that about 35% of the time removing  $\omega_t$  generates at least one cancellation of a non-existent order, i.e. 19% of the time there is one such event, 7% of the time two such events, etc., for an average of 0.73 such events per removal.

In the hypothetical series we also observe a slight increase in the number of undefined prices. For AZN, for example, this happens in 0.01% of the cases, in contrast to 0.003% for the real series. We omit these cases from our analysis.

Finally, we need to deal with the possibility that the assignment of effective orders might alter the number of events in the hypothetical series. This can happen, for example, because an order that is fully executed in the real series, and was therefore an effective market order, is now only partially executed, and so becomes an effective market order followed by an effective limit order. We deal with this by preserving the alignment of the hypothetical and real data series based on the real events. This prevents the possibility of a persistent misalignment which could create a persistent artificial price difference between the two series. In any case we find that such situations are rare.




FIG. 1: (a) Initial and modified price sequences for the removal of a buy market order and (b) the resulting mechanical impact (which is just the difference between the two price series in (a)). Note that this figure and all other figures in this paper are based on the stock AZN. Prices are in pence.

In this paper we study only the impact of effective market orders. In Figure 1 we show a typical example comparing a real price sequence to a hypothetical price sequence with a buy order removed<sup>5</sup>. It is clear that the mechanical impact is highly variable. To give

<sup>5</sup> The TDS data set comes with a series of bids and offers computed by the exchange. These do not

the reader a feeling for the variety of possibilities, in Figure 2 we show a set of four examples of the mechanical impact of effective market orders (in pence). The mechanical impact is highly variable. In some cases there is an initial burst of mechanical impact, which dies to zero and then remains there. In some cases there are long gaps in which the impact remains at zero and then takes on nonzero values after more than a thousand transactions. In other cases there is no mechanical impact at all.

### C. Informational impact

The *informational impact* is defined as the portion of total impact that cannot be explained by mechanical impact, i.e.

$$\Delta p_{\tau}^I = \Delta p_{\tau}^T - \Delta p_{\tau}^M.$$

Whatever components of the total impact not explained by mechanical impact must be due to correlations between the order  $\omega_t$  and other events. With the data we have it is impossible to say whether the placement of the order  $\omega_t$  causes changes in future events  $\Omega_{t+1}$ , or whether the properties of  $\Omega_{t+1}$  are correlated with those of  $\omega_t$  due to a common cause. In either case, changes in price that are not caused mechanically must be due to information – either the information contained in  $\omega_t$  affecting  $\Omega_{t+1}$ , or external information affecting both  $\omega_t$  and  $\Omega_{t+1}$ . See the discussion in Section VII A.

always match the prices that we compute by applying the continuous double auction algorithm to the event series. For consistency of comparison between the real and hypothetical series, we use the latter.




FIG. 2: Examples of mechanical price impact (in pence) in transaction time. The top two plots correspond to buy orders and the bottom two to sell orders.

## IV. EMPIRICAL STUDY OF IMPACT

In this section we perform statistical analysis of average properties of price impact. For the remainder of this paper we will only study the impacts of effective market orders, deferring the problem of studying effective limit orders and cancellations.

### A. Average impact

The first property we study is the average impact as a function of time. In Figure 3 we compare the average mechanical impact  $R^M(\tau) = \langle s_t \Delta p_\tau^M(t) \rangle_t$  and the average total impact<sup>6</sup>  $R^T(\tau) = \langle s_t \Delta p_\tau^T(t) \rangle_t$ . For event time the total impact and mechanical impact are by definition the same at  $\tau = 1$ . This is because in moving from  $\tau = 0$  to  $\tau = 1$  the only event that affects the price is the reference event  $\omega_t$  – alterations in  $\Omega_{t+1}$  cannot effect  $\Delta p_1^T$ . For larger values of  $\tau$  the mechanical impact decreases and the informational impact increases. Over the timescale shown here (100 events), when measured in units of the average spread, the mechanical impact is initially about 0.17, and then decays monotonically toward zero. In contrast the total impact increases toward what appears to be an asymptotically constant value slightly greater than 0.3. This is the source of our statement that the initial value of mechanical impact is about half the asymptotic value of the total impact. Similar results are observed for VOD and LLOY.

Figure 3(b) shows the same behavior in transaction time. In this case the mechanical

<sup>6</sup> In this and other figures we compute impacts in units of the average spread. For each stock and the full time period of each data set, the average spread is computed by sampling just before transactions.





FIG. 3: Average mechanical impact  $R^M(\tau) = \langle s_t \Delta p_\tau^M(t) \rangle_t$  (red squares) and total impact  $R^T(\tau) = \langle s_t \Delta p_\tau^T(t) \rangle_t$  (blue stars), in units of the average spread, plotted in (a) event time and (b) transaction time.

impact and total impact diverge immediately at  $\tau = 1$ . This is not surprising, since one increment of transaction time is equivalent to roughly six increments of event time. The initial gap between the total and the mechanical impacts in transaction time is roughly consistent with the results in event time at  $\tau = 6$ . Once again, except for statistical fluctuations the total impact appears to approach a constant value out to  $\tau = 100$ .




FIG. 4: Average mechanical impact in units of the average spread, plotted in double logarithmic scale. The x's are the data and the solid line shows a power law fit to the tail.

The total impact of market orders has received extensive study in several different markets [6, 8, 9, 14–16, 18, 20, 23–25]. Here we focus on the behavior as a function of time [1, 2, 4, 5, 11, 19, 28]. Our results on total impact are consistent with previous results. We find that total impact builds with time and appears to approach a constant, up to the point where statistical fluctuations make the results questionable.

In sharp contrast to the total impact, the mechanical impact decays toward zero. To get a better view in Figure 4 we plot the average mechanical impact for times up to 2000 transactions in log-log scale<sup>7</sup>. We fit a power law of the form  $K\tau^{-\lambda}$  to what we subjectively deem to be the asymptotic region of the tail. For AZN we find  $\lambda = 1.6$ , for VOD  $\lambda = 1.8$ ,

<sup>7</sup> To make the fluctuations in the data clearer in this and other figures, we have placed x's at diminishing intervals of time to get a better view in log-log scale.

and for LLOY  $\lambda = 1.7$ . Given that the scaling region is only over a little more than an order of magnitude, this is certainly not convincing evidence that the average mechanical impact scales as a power law. Nonetheless, plotting in semi-log coordinates makes it quite clear that the decay is slower than exponential. We have not attempted to put any error bars on these estimates because they are difficult to assign<sup>8</sup>. These should just be viewed as representative values.

It is clear from simple theoretical arguments that the mechanical impact must decay to zero. From the definition of Equation 2 the only difference between the real price series and the hypothetical series with the order removed is the initial order book  $B$ . As orders disappear from the order book the difference between the real and the hypothetical price series decreases. Once there are no orders that were in the original order book there is no longer any memory of the initial condition, and there will no longer be any difference between the real and the hypothetical price series. Previous studies suggest that the distribution of times for an order to remain in the book before being removed by a transaction decays as a power law with an exponent near 1.5, and that the distribution of times to cancellation decays as a power law with an exponent near 2 [3, 22]. These both suggest that we should expect an asymptotic power law decay of the average mechanical impact.

Under the rules of the LSE no order can persist for more than one month, which would seem to imply an upper bound on the per-

sistence of the mechanical impact. However, it is quite common for orders to be immediately replaced, so that from an effective point of view some orders can persist for very long times, e.g. six months or more [29]. For practical reasons related to limitations in computation time we have not measured the mechanical impact for time intervals longer than 2000 transactions, corresponding to a period of about two trading days. This doesn't appear to be an important restriction, since the average mechanical impact at  $\tau = 2000$  is less than three orders of magnitude less than its initial value, and events where the mechanical impact is non-zero for  $\tau > 2000$  are fairly rare.

## B. Duration and size

We define the *duration* of the mechanical impact of an order as the largest  $\tau$  for which that order has a non-zero impact. Figure 5 shows a histogram of the durations in transaction time. The most common duration is zero – about 33% of the events have no impact at all in transaction time. The duration probability decreases rapidly and roughly monotonically. To get a better view in Figure 5(b) we show this for a longer time period in double logarithmic scale and fit a power law  $K\tau^{-\delta}$  over what we subjectively deem to be the tail. For AZN we estimate  $\delta = 1.5$ , for VOD  $\delta = 1.7$  and for LLOY  $\delta = 1.7$ .

Because mechanical impact is transitory, we define a notion of size in terms of the impact integrated over time. The integrated size  $S_\tau^M(t)$  up to time  $\tau$  associated with the impact event at time  $t$  is defined as

$$S_\tau^M(t) = \sum_{i=1}^{\tau} s_t \Delta p_i^M(t).$$

The distribution of integrated sizes is shown in Figure 6 in double logarithmic scale. If we

<sup>8</sup> In addition to statistical fluctuations, there are problems caused by the slow convergence to asymptotic scaling. Careful testing of the power law hypothesis and proper assignment of error bars for scaling exponents is beyond the scope of this paper.




FIG. 5: Distribution of mechanical impact durations. The top panel is the probability of each duration, and the bottom panel covers a longer time period on double logarithmic scale. The x's correspond to the empirical measurements, and the line to a fit using ordinary least squares.

fit a power law of the form  $K\tau^{-\alpha}$  for AZN we find  $\alpha = 3.3$ , for VOD  $\alpha = 3.4$ , and for LLOY  $\alpha = 3.4$ . Given the small size of the scaling region and the large size of the scaling exponents, it is not at all clear that the integrated impact size asymptotically scales as a power law.




FIG. 6: The distribution of integrated mechanical impact sizes  $S_{\tau}^M(t)$ , with  $\tau = 2000$ . The x's are the real data and the solid line is a fit to a power law. The units of the  $x$  axis are price  $\times$  time, where price differences are measured in units of average spread and time is measured in transactions.

## V. AMPLIFICATION BY LONG-MEMORY

The effect of mechanical impact on prices is amplified by the long-memory of the order signs  $s_t$ . As we explain in more detail below, the long-memory of order signs refers to the strong tendency of buy orders to be followed by more buy orders, and sell orders to be followed by more sell orders. When we compute mechanical impact we subtract the price sequence associated with two series of orders, both of which contain long-memory. The long-memory is thus removed. Nonetheless, since buy orders generated positive mechanical impacts and sell orders generate negative mechanical impacts, the long-memory of orders amplifies the mechanical impact of any given order. In this section we make all this more precise, compute the amplification due to long-memory, and show that it is not




FIG. 7: An illustration of the long-memory of order signs. The ratio  $P_{\tau}^{+}/P_{\tau}^{-} - 1$  is plotted on double logarithmic scale (x's) and compared to a power law (line) with exponent  $\gamma = 0.59$ .

strong enough to create a persistent impact from the transitory nature of individual mechanical impacts.

To make the definition of the long-memory of order signs more precise, define

$$P_{\tau}^{+} = P(s_t = 0 \ \& \ s_{t+\tau} = 0) \\ + P(s_t = 1 \ \& \ s_{t+\tau} = 1)$$

as the probability that transactions at time  $t$  and  $t + \tau$  have the same sign and similarly  $P_{\tau}^{-}$  as the probability that they have the opposite sign. For all the stock markets examined so far (Paris, London, NYSE) the sequence of signs has long-memory, i.e. the autocorrelation function decays as a power law  $\tau^{-\gamma}$  [1, 2, 11, 19]. This also implies that  $P_{\tau}^{+}/P_{\tau}^{-} - 1 \sim \tau^{-\gamma}$ . The exponent  $\gamma$  appears to vary somewhat depending on the market and the stock, but it is consistently less than one. Long-memory is illustrated for AZN in Figure 7. Lillo, Mike, and Farmer [21] hypothesized that long-memory is caused by strategic order splitting and presented results supporting the hypothesis. More recent

results that make use of transaction identity codes strongly support this [26].

The long-memory of order signs creates a puzzle because it naively suggests that prices should be inefficient. Figure 7 makes it clear that order signs are predictable based on their past history. Given that buyer-initiated trades have positive price impact and seller-initiated trades have negative impact, this suggests that price movements should also be predictable. This is not the case: price changes are essentially uncorrelated.

Two explanations have been offered to explain how transactions can have long-memory while prices are efficient. One is due to Bouchaud et al. [1, 2]. They postulate the existence of a bare impact function  $G(\tau)$ , such that the total impact is the sum of the bare impact of each trade. If  $G(\tau)$  decays as a power law  $G(\tau) \sim \tau^{-\beta}$ , providing  $\beta = (1 - \gamma)/2$ , due to the long-memory the bare impacts will accumulate so that the total impact asymptotically approaches a constant, i.e. the decay of the bare impact and the amplification due to long-memory cancel each other. Providing the constant is less than half the spread (which is observed in practice), the market is efficient [28].

The temporary nature of the average mechanical impact response function shown in Figure 3 suggests that it might provide a fundamental explanation for the decay of the bare impact  $G(\tau)$ . However, the average mechanical impact decays too fast for this to be true. According to the formula of the previous paragraph, using  $\gamma = 0.6$  implies a bare propagator exponent of  $\beta = 0.2$ . This is much smaller than the exponent  $\lambda = 1.6$  measured for mechanical impact. Thus it seems that this cannot be the explanation. Instead, the full explanation involves the existence of a liq-

uidity imbalance between buying and selling<sup>9</sup>. Nonetheless, the decaying nature of the mechanical impact may play an important role in making the bare impact temporary. The details of this remain to be worked out.

When we compute the mechanical impact, because the perturbed transaction sequence and the reference sequence in Equation 2 both have long memory, when we subtract them any effects due to long-memory disappear. Nonetheless, it is clear that long-memory will amplify mechanical impact, for the same reasons it amplifies the bare impact. Let the cumulative impact of an order at time  $\tau$  be  $C(\tau)$  and consider, for example, a buy order. This generates a positive average impact at  $\tau = 1$  of  $R^M(1)$ . At  $\tau = 2$  its average impact is reduced to  $R^M(2)$ , but we also need to take into account that the order at  $\tau = 1$  is more likely to be a buy order than a sell order, so that there is an additional contribution  $(P_1^+ - P_1^-)R^M(1)$ . Similarly at  $\tau = 3$  the average cumulative impact  $C(3) = R^M(3) + (P_1^+ - P_1^-)R^M(2) + (P_2^+ - P_2^-)R^M(1)$ . Noting that by definition  $P_0^+ = 1$  and  $P_0^- = 0$ , the general expression is

$$C(\tau) = \sum_{i=1}^{\tau} (P_{\tau-i}^+ - P_{\tau-i}^-) R^M(i) \quad (3)$$

In Figure 8(a) we compare the raw mechanical impact, the cumulative mechanical im-

act  $C(\tau)$  including amplification by long-memory, the informational impact, and the total impact. Plotting the mechanical impact adjusted for long-memory on double logarithmic scale makes it clear that it is still decaying to zero, albeit slower than the raw mechanical impact. Fitting a power law to the long-time behavior of the form  $K\tau^{-\eta}$ , as shown in Figure 8(b) gives an exponent of about  $\eta = 0.6$ .

In conclusion, while it is clear that the correlations associated with the long-memory of order signs amplify mechanical impact, they are not sufficiently strong to overcome the rapid decay of the mechanical impact to make it permanent.

## VI. CORRELATION BETWEEN MECHANICAL AND INFORMATIONAL IMPACTS

We now study the correlation between mechanical and informational impact. Because mechanical impact is temporary, it is not obvious what feature of mechanical impact is likely to be most important. Studying the correlation between mechanical and informational impact is motivated in part by the idea that prices are informative, i.e. as agents observe price changes induced by others, they update their own information. Since mechanical impacts are temporary, it is not clear what others will respond to, particularly if the response is not instantaneous. We somewhat arbitrarily use the integrated size, though one could easily argue that other properties of the mechanical impact might be more reasonable. See the discussion in Section VII A.

In Figure 9 we plot the informational impact at  $\tau = 20$  against the size of the mechanical impact integrated up to  $\tau = 20$ . The relationship for individual transactions is very

---

<sup>9</sup> The hypothesis of a liquidity imbalance between buying and selling was offered by Lillo and Farmer[19] and modified and demonstrated to be effective by Farmer et al. [11]. When buyer-initiated transactions become more likely, the liquidity for buying increases and the liquidity for selling decreases (though with some time lag), making the price responses to buy orders smaller on average than those to sell orders. This damps the effect of the long-memory and keeps the market efficient.




FIG. 8: (a) A comparison of the average mechanical, informational, and total impacts with and without adjustments for long-memory. In ascending order we show the average mechanical impact  $R^M(\tau)$ , the mechanical impact  $C(\tau)$  adjusted for long-memory, the informational impact adjusted for long-memory  $R^T(\tau) - C(\tau)$ , the informational impact  $R^T(\tau) - R^M(\tau)$ , and the total impact  $R^T(\tau)$ , all as a function of time. In (b) we show  $C(\tau)$  over a longer time period in double logarithmic scale.




FIG. 9: The informational impact at  $\tau = 20$  vs. mechanical impact integrated up to  $\tau = 20$ ,  $S_{20}^M(t)$ . In (a) the dots correspond to individual transactions, and the line is a regression using ordinary least squares. In (b) the data is binned based on the integrated mechanical impacts, illustrating that even though the relationship is noisy, the correlation increases consistently with increasing size. Note that the black square corresponds to the case where the mechanical impact is strictly zero.

noisy, but the positive association between mechanical and informational impact is quite clear. The correlation is  $\rho = 0.14$ . A linear regression of the form  $\Delta p_{20}^I(t) = aS_{20}(t) + b$  yields a positive slope  $a = 0.187 \pm 0.003$ . The error bar corresponds to a t-statistic of 55. It is computed under the assumption of normally distributed IID data, which is not valid in this case and is certainly too optimistic. Binning the data based on the integrated size of the mechanical impact makes it clear that the relationship is essentially monotonically increasing – large integrated mechanical impacts are associated with large informational impacts. The positive association between informational and mechanical impacts is clearly highly statistically significant.

We find the surprising result that for cases where the mechanical impact is strictly zero, the informational impact has the opposite sign that one would normally expect, e.g. buy orders with strictly zero mechanical impact tend to have negative total price impacts. (Recall that when the mechanical impact is zero informational and total impact are the same). *Strictly zero* means that the mechanical impact is zero for all times in units of event-time rather than transaction time, i.e. it does not include situations where there is a non-zero impact in event time that dies out before the first transaction. For AZN, for example, 13% of effective market orders have strictly zero mechanical impact, in contrast to 33% in transaction time. To illustrate this in Figure 9(b) we have added the strictly zero case by placing two values over 0 on the x-axis. The first is consistent with the rest of the figure and corresponds to  $S_{20}^M = 0$ , while the other corresponds to the case where the mechanical impact is strictly zero. The average informational impact in the latter case is  $\langle s_t \Delta p_{20}^I(t) \rangle_t = -0.4$ , in contrast to the case where  $S_{20}^M = 0$ , which has roughly  $\langle s_t \Delta p_{20}^I(t) \rangle_t = 0$ . Thus, as soon as there is




FIG. 10: Informational impact as a function of time for transactions with different levels of integrated mechanical impact. We show the average informational impact from  $\tau = 1$  to  $\tau = 20$  for three equal sized quantiles of the transactions with nonzero integrated mechanical impact  $S_{20}(t)$ , as well as for all transactions where the mechanical impact is strictly zero.

any nonzero mechanical impact at all, the average informational impact jumps from a strongly negative value to a positive value.

The relationship between mechanical and informational impact remains consistent through time. In Figure 10 we divide the mechanical impact into four different groups based on their mechanical impact and track the informational impact through time. The first group consists of the cases whose impact is strictly zero and the others to the nonzero mechanical impacts, which are sorted based on  $S_{20}^M$  into three groups with an equal number of events. As before, the behavior of the informational impact when the mechanical impact is strictly zero is quite different than that when it is nonzero. When the mechanical impact is strictly zero, the informational impact has a slightly positive value at

$\tau = 1$ , but then becomes increasingly negative as  $\tau$  increases. We hypothesize that this occurs because orders with no mechanical impact at all tend to be associated with price reversals. For example, consider the case where we remove a buy market order which is too small to remove the best ask, in a situation in which the price immediately drops. New sell orders will fill in below the best ask. In this case it is very unlikely that removal of the buy market order will make any difference to the midprice. In contrast, had the price continued to go up, the rearrangement in the level of sell limit orders triggered by the removal of the buy market order is much more likely to generate a perturbation in the price series.

## VII. SUMMARY

We have introduced a precise definition of the mechanical impact of a transaction and demonstrated how it can be measured. The time behavior of the mechanical impact is highly variable. On average, however, for market orders it is strongest immediately after the order is placed, and then decays to zero over time. It is initially the dominant component of the total price impact, but the relative fraction decreases as the mechanical impact shrinks and the informational impact grows to approach its asymptotic value. Initially the mechanical impact is about half as large as the asymptotic informational impact.

### A. Causality

We have defined the informational impact as what is left over of the total impact once the purely mechanical impact is removed. The justification for this is that the remainder depends on correlations between events. I.e., unless there is something correlating the

information in the reference event  $\omega_t$  to the subsequent sequence of events  $\Omega_{t+1}^{t+\tau}$ , there will be no informational impact. This leaves open the question of causality. One hypothesis is that event  $\omega_t$  causes changes in the subsequent events  $\Omega_{t+1}^{t+\tau}$ . The alternative hypothesis is that there is a common cause for  $\omega_t$  and any aspects of  $\Omega_{t+1}^{t+\tau}$  that are correlated with its presence. Based on the information available here we cannot distinguish these two hypotheses.

The idea of a common cause is easy to understand. Suppose, for example, that an external event causes a group of investors to decide to buy, and they submit a series of buy orders without paying any attention to each other. This will cause a rise in price associated with each buy order, due both to the mechanical impact of each buy order, and the mechanical impact of all the buy orders that are correlated with it. The correlated part of the price rise will be measured as an informational impact (reflecting the external information).

The causal hypothesis is more interesting, and is connected to the role of trading and its associated price changes in transmitting information. In a world where individual agents have private information, trading provides a mechanism for disseminating that information. If an agent receives new private information, this may cause him or her to trade. Trading affects the price, which is visible to everyone. An intelligent agent with different information will observe the change in price, and will infer that his or her valuation must be wrong. As a result, each agent will arrive at a valuation that is based partly on private information and partly on price and other public information. Prices are thus a mechanism for making private information public. This idea is well-accepted in economic theory [12, 13].

Insofar as the causal hypothesis is correct,

the temporary nature of mechanical impact suggests that private information is made public through a highly dynamic process. As each trade happens, it causes a mechanical impact, which is a signal visible to all. Before it decays away, it can cause other trades of the same sign to occur, or it can cause cancellations to occur, or (probably most important) it can cause changes in the limit prices of subsequent orders. Each of these events has its own mechanical impact, creating a cascade of impacts giving rise to a permanent change. The avalanche-like nature of this process suggests possible analogies to self-organized criticality.

At this stage we are unable to say to what extent common cause or causality are at play. It seems likely that both are acting. We do know that these series display long-memory, and as discussed in Section V, evidence suggests that this is essentially an exogenous phenomenon, which from this point of view is a common cause.

### B. Generalizations

The mechanical impact can be defined in a way that is more general than what we have presented here. We have defined mechanical impact by fully removing an event, but there are many other possible modifications of the order book that are worth considering. Equation 2 can be generalized to read

$$\Delta p_\tau^M(t) = \Pi(b_t, \Omega_{t+1}^{t+\tau}) - \Pi(\tilde{b}, \Omega_{t+1}^{t+\tau}) \quad (4)$$

where  $b_t$  is the real order book at time  $t$  and  $\tilde{b}$  is any modification of it that may yield a use-

ful interpretation. For example,  $\tilde{b}$  might be the true order book with an additional order added, or it could be an infinitesimal modification of a given order, e.g. with the size of the most recent order  $\omega_t$  slightly enhanced or diminished. This potentially allows one to make a quasi-continuous analysis, effectively estimating the derivative of the mechanical changes in prices with respect to changes in supply or demand.

### C. Future work

We have left many questions unanswered in this work. In addition to the generalization discussed above, there are many topics that remain to be investigated. For example, what is the dependence of mechanical impact on the size of orders? Previous results have shown that the total impact is remarkably independent of order size [10, 27]; is this also true for mechanical impact? Is the mechanical impact symmetric under increases or decreases in the size of  $\omega_t$ ? What does the price impact of cancellation and limit orders look like? We hope to investigate these and other questions in the future.

### Acknowledgments

We would like to thank Barclays Bank and the Capital Markets CRC for supporting this research, J.-P. Bouchaud, Marc Potters, Austin Gerig, and Adlar Kim for useful comments, and the London Stock Exchange for supplying data.

---

[1] J.-P. Bouchaud, Y. Gefen, M. Potters, and M. Wyart. Fluctuations and response in fi-

nancial markets: The subtle nature of “random” price changes. *Quantitative Finance*,

- 4(2):176–190, 2004.
- [2] J-P. Bouchaud, J. Kockelkoren, and M. Potters. Random walks, liquidity molasses and critical response in financial markets. *Quantitative Finance*, 6(2):115–123, 2006.
- [3] D. Challet and R. Stinchcombe. Nonconstant rates and overdifusive prices in simple models of limit order markets. *Quantitative Finance*, 3:165, 2003.
- [4] L. K.C. Chan and J. Lakonishok. Institutional trades and intraday stock price behavior. *Journal of Financial Economics*, 33:173–199, 1993.
- [5] L. K.C. Chan and J. Lakonishok. The behavior of stock prices around institutional trades. *The Journal of Finance*, 50(4):1147–1174, 1995.
- [6] T. Chordia, R. Roll, and A. Subrahmanyam. Order imbalance, liquidity and market returns. *Journal of Financial Economics*, 65:111–130, 2002.
- [7] D. M. Cutler, J. M. Poterba, and L. H. Summers. What moves stock prices? *The Journal of Portfolio Management*, 15(3):4–12, 1989.
- [8] M. D. D. Evans and R. K. Lyons. Order flow and exchange rate dynamics. *Journal of Political Economy*, 110(1):170–180, 2002.
- [9] J. D. Farmer. Slippage 1996. Technical report, Prediction Company, 1996. URL <http://predict.com/jdf/slippage.pdf>.
- [10] J. D. Farmer, L. Gillemot, F. Lillo, S. Mike, and A. Sen. What really causes large price changes? *Quantitative Finance*, 4(4):383–397, 2004.
- [11] 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):107–112, 2006.
- [12] S. J. Grossman. *The Informational Role of Prices*. MIT Press, Cambridge, 1989.
- [13] S. J. Grossman and J. E. Stiglitz. On the impossibility of informationally efficient markets. *The American Economic Review*, 70(3):393–408, 1980.
- [14] J. Hasbrouck. Measuring the information content of stock trades. *The Journal of Finance*, 46(1):179–207, 1991.
- [15] J. A. Hausman, A. W. Lo, and A. C. Mackinlay. An ordered probit analysis of transaction stock prices. *Journal of Financial Economics*, 31(3):319–379, 1992.
- [16] C. Hopman. Do supply and demand drive stock prices? *Quantitative Finance*, 2006. To appear.
- [17] French I. and R. Roll. Stock return variances: The arrival of information and the reaction of traders. *Journal of Financial Economics*, 17:5–26, 1986.
- [18] A. Kempf and O. Korn. Market depth and order size. *Journal of Financial Markets*, 2(1):29–48, 1999.
- [19] F. Lillo and J. D. Farmer. The long memory of the efficient market. *Studies in Nonlinear Dynamics & Econometrics*, 8(3), 2004.
- [20] F. Lillo and R. N. Mantegna. Power-law relaxation in a complex system: Omori law after a financial market crash. *Physical Review E*, 68(1), 2003. Part 2.
- [21] F. Lillo, S. Mike, and J. D. Farmer. Theory for long memory in supply and demand. *Physical Review E*, 7106(6 pt 2):287–297, 2005.
- [22] S. Mike and J.D. Farmer. An empirical behavioral model for price formation. *Journal of Economic Dynamics and Control*, 2006. to appear.
- [23] V. Plerou, P. Gopikrishnan, X. Gabaix, and H. E. Stanley. Quantifying stock price response to demand fluctuations. *Physical Review E*, 66(2):article no. 027104, 2002.
- [24] M. Potters and J-P. Bouchaud. More sta-

- tistical properties of order books and price impact. *Physica A*, 324:133–140, 2003.
- [25] N. Torre. *BARRA Market Impact Model Handbook*. BARRA Inc., Berkeley, 1997.
- [26] Vaglica, F. Lillo, E. Moro, and R. Mantegna. Unpublished work, 2006.
- [27] P. Weber and B. Rosenow. Large stock price changes: volume or liquidity? *Quantitative Finance*, 6(1):7–14, 2006.
- [28] M. Wyart, J.-P. Bouchaud, J. Kockelkoren, M. Potters, and M. Vettorazzo. Relation between bid-ask spread, impact and volatility in double auction markets. Technical report, 2006. URL [arXiv:physics/0603084](https://arxiv.org/abs/physics/0603084).
- [29] I. Zovko and J. D. Farmer. The power of patience; a behavioral regularity in limit order placement. *Quantitative Finance*, 2(5):387–392, 2002.