

What really causes large price changes?

J. Doyne Farmer,¹ Laszlo Gillemot,^{1,2} Fabrizio Lillo,^{1,3} Szabolcs Mike,^{1,2} and Anindya Sen^{1,4}

¹*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501*

²*Budapest University of Technology and Economics, H-1111 Budapest, Budafoki út 8, Hungary*

³*Istituto Nazionale per la Fisica della Materia, Unità di Palermo, Italy*

⁴*Mathematics Dept., University of Chicago*

(Dated: December 30, 2003)

We study the cause of large fluctuations in prices in the London Stock Exchange. This is done at the microscopic level of individual events, where an event is the placement or cancellation of an order to buy or sell. We show that price fluctuations caused by individual market orders are essentially independent of the volume of orders. Instead, large price fluctuations are driven by liquidity fluctuations, variations in the market's ability to absorb new orders. Even for the most liquid stocks there can be substantial gaps in the order book, corresponding to a block of adjacent price levels containing no quotes. When such a gap exists next to the best price, a new order can remove the best quote, triggering a large midpoint price change. Thus, the distribution of large price changes merely reflects the distribution of gaps in the limit order book. This is a finite size effect, caused by the granularity of order flow: In a market where participants placed many small orders uniformly across prices, such large price fluctuations would not happen. We show that this explains price fluctuations on longer timescales. In addition, we present results suggesting that the risk profile varies from stock to stock, and is not universal: lightly traded stocks tend to have more extreme risks.

Contents

I. Introduction	1
A. Background: Continuous double auction	2
B. What causes returns?	3
II. Data	3
III. Fluctuations in liquidity drive the tails	4
A. Large returns are not caused by large orders	4
B. Volume dependence of mean market impact	6
C. Correlations between order size and liquidity	6
D. Liquidity fluctuations drive price fluctuations	7
IV. Granularity of supply and demand	8
V. Limit orders and cancellations	10
VI. Price changes on longer timescales	10
VII. Correlations in occupied sites	11
VIII. Conclusions	12
Acknowledgments	13
References and Notes	13

I. INTRODUCTION

It has now been known for more than forty years that price changes are fat-tailed [1–13], i.e. there is a higher probability of extreme events than in a normal distribution. This is important for financial risk, since it means

that large price fluctuations are much more common than one might expect. There have been several conjectures about the origin of fat tails in prices [14–25]. Most of these theories are either generic mechanisms for generating power laws, such as multiplicative noise or maximization of alternative entropies or agent-based models that make qualitative predictions that are not very specific.

Two theories that deserve special mention because they make testable hypotheses about the detailed underlying mechanism are the subordinated random process theory of Clark [14] and the recent theory of Gabaix et al. [25]. Clark's proposal is that because order arrival rates are highly intermittent, aggregating in a fixed time interval leads to fat tails in price returns. Gabaix et al.'s proposal is that high volume orders cause large price movements. We show that neither of these theories describe large price changes in the London Stock Exchange¹.

Instead, we show that large price fluctuations are driven by fluctuations in liquidity, i.e. variations in the response of prices to changes in supply and demand. The number of agents that participate in the market at any given time, and thus the number of orders to buy or sell, is rather small. Even for a heavily traded stock, the typical number of orders in the book at any given time is generally around 30. While it is in some cases a good approximation to regard the market as a statistical system, which can be treated using mathematical methods from statistical mechanics [27–32], markets are far from the thermodynamic limit, and display strong finite size effects. Fluctuations in orders are important, but it is

¹ For a critique of the paper of Gabaix et al. see Reference [26].

not the size of orders that drives large price changes, but rather the uniformity of their coverage of price levels. Revealed supply and demand curves at any instant in time are irregular step-like functions with long flat regions and large jumps. The market can be regarded as a granular medium, in which the incremental changes in supply and demand are the grains. The statistical properties of prices depend more on the fluctuations in revealed supply and demand than on their mean behavior.

A common assumption is that large price returns $r(t)$ are asymptotically distributed as a power law. In more technical terms, letting $m(t)$ be the midprice at time t , this means that $r(t) = \log m(t) - \log m(t - \tau)$ satisfies $P(r > x) \sim x^{-\alpha}$. $P(r > x)$ is the probability that $r > x$, τ is an arbitrary time interval, and $f(x) \sim g(x)$ means that $f(x)$ and $g(x)$ scale the same way in the limit $x \rightarrow \infty$ ². α is called the *tail exponent*. It was initially thought that $\alpha < 2$, which is significant because this would imply that the standard deviation of price returns does not exist, and that under aggregation independent price returns should converge to a Levy stable distribution [1, 2]. However, most subsequent studies indicate that $\alpha > 2$ is more common [3–9, 11–13]. Nonetheless, it still remains controversial whether a power law is always the best description of price returns. In this paper we do not attempt to resolve this debate. However, we will use power laws as a useful way to describe the asymptotic behavior of price changes, in particular to compare the distribution of price changes for different stocks. Our results suggest that tail exponents vary from stock to stock, and are not universal.

In the remainder of this section we review the continuous double auction, explaining what we define as an “event”, and introducing notation that we will use throughout the rest of the paper. Section II presents some summary statistics for the data set. Section III demonstrates that there is very little difference in the price response of large and small orders, and that large price fluctuations are driven by fluctuations in liquidity. This is made more explicit in Section IV, where we show how gaps in the occupied price levels in the orderbook lead to large price changes, and show that the gap distribution closely matches the return distribution. In Section V we compare the returns triggered by market orders, limit orders, and cancellations, and show that they are very similar. Section VI demonstrates that the tail behavior of returns on the event scale matches the tail behavior on longer timescales, and that returns in event time and real time are similar. Section VII studies the behavior of the gap distribution in more detail, demonstrating that the (at least approximate) power law behavior we observe is not driven by fluctuations in the

number of orders in the book, but rather depends on correlations in occupied price levels. We conclude and offer a few speculations about the ultimate explanation of large price fluctuations in Section VIII.

A. Background: Continuous double auction

To understand our results it is essential that the reader understand the double continuous auction, which is the standard mechanism for price formation in most modern financial markets. Agents can place different types of orders, which can be grouped into two categories: Impatient traders submit *market orders*, which are requests to buy or sell a given number of shares immediately at the best available price. More patient traders submit *limit orders*, or *quotes* which also state a limit price π , corresponding to the worst allowable price for the transaction. (Note that the word “quote” can be used either to refer to the limit price or to the limit order itself.) Limit orders often fail to result in an immediate transaction, and are stored in a queue called the *limit order book*. Buy limit orders are called *bids*, and sell limit orders are called *offers* or *asks*. At any given time there is a best (lowest) offer to sell with price $a(t)$, and a best (highest) bid to buy with price $b(t)$. These are also called the *inside quotes* or the *best prices*. The price gap between them is called the *spread* $s(t) = a(t) - b(t)$. Prices are not continuous, but rather change in discrete quanta called *ticks*, of size Δp . The number of shares in an order is called either its *size* or its *volume*.

As market orders arrive they are matched against limit orders of the opposite sign in order of first price and then arrival time, as shown in Fig. 1. Because orders are placed for varying numbers of shares, matching is not necessarily one-to-one. For example, suppose the best offer is for 200 shares at \$60 and the the next best is for 300 shares at \$60.25; a buy market order for 250 shares buys 200 shares at \$60 and 50 shares at \$60.25, moving the best offer $a(t)$ from \$60 to \$60.25. A high density of limit orders per price results in high *liquidity* for market orders, i.e., it implies a small movement in the best price when a market order is placed.

There are a variety of different order types defined in real markets, whose details differ from market to market. For our purposes here, any given order can always be decomposed into two types: We will call any component of an order that results in immediate execution an *effective market order*, and any component that is not executed immediately, and is stored in the limit order order book, an *effective limit order*. For example, consider a limit order to buy whose limit price $\pi = a(t)$. Suppose the volume at $a(t)$ is 1000 shares, and the volume of the new limit order is 3000 shares. Then this limit order is equivalent to an effective market order for 1000 shares, followed by an effective limit order of 2000 shares with limit price $a(t)$. In either case the same transactions take place, and the best prices move to $b(t + 1) = a(t)$ and

² $f(x) \sim g(x)$ means that $\lim_{x \rightarrow \infty} L(x)f(x)/g(x) = 1$, where $L(x)$ is a slowly varying function. A slowly varying function $L(x)$ is a positive function that for every t satisfies $\lim_{x \rightarrow \infty} L(tx)/L(x) = 1$.




FIG. 1: A schematic illustration of the continuous double auction mechanism. Limit orders are stored in the limit order book. We adopt the arbitrary convention that buy orders are negative and sell orders are positive. As a market order arrives, it has transactions with limit orders of the opposite sign, in order of price (first) and time of arrival (second). The best quotes at prices $a(t)$ or $b(t)$ move whenever an incoming market order has sufficient size to fully deplete the stored volume at $a(t)$ or $b(t)$, when new limit orders are placed inside the spread (when the limit price satisfies $b(t) < \pi < a(t)$) or when all the orders at the best prices are cancelled. Bids and offers that fall inside the spread become the new best bids and offers.

$a(t+1) = a(t) + g(t)$, where $g(t)$ is the price interval to the next highest occupied price level. Throughout the remainder of the paper we will simply call an effective limit order a “limit order”, and an effective market order a “market order”.

When a market order arrives it can cause changes in the best prices. This is called *market impact* or *price impact*. Note that the price changes are always in the same direction: A buy market order will either leave the best ask the same or make it bigger, and a sell market order will either leave the best bid the same or make it smaller. The result is that buy market orders can increase the midprice $m(t) = (a(t) + b(t))/2$, and sell orders can decrease it.

There is no unique notion of price in a real market. We will let π be the *limit price* of a limit order, and $m(t) = (a(t) + b(t))/2$ be the *midpoint price* or *mid-price* defined by the best quotes. All the results of this paper concern the midprice, rather than transaction prices, but at longer timescales this makes very little difference, since the midpoint and transaction prices rarely differ by more than half the spread. The midpoint price is more convenient to study because it avoids problems associated with the tendency of transaction prices to bounce back and forth between the best bid and ask. Price changes are typically characterized as *returns* $r_\tau(t) = \log m(t) - \log m(t - \tau)$, though we will also characterize them in terms of *ticks* $k = (m(t+1) - m(t))/\Delta p$.

It is also often convenient to characterize the difference of two limit prices, π_1 and π_2 , in terms of ticks, $k = (\pi_1 - \pi_2)/\Delta p$.

B. What causes returns?

In this paper we study changes in the midprice at the level of individual events. The arrival of three kinds of events can cause the midprice to change:

- *Market orders.* A buy market order of sufficient size can increase the best ask, and a sell market order can decrease the best bid, widening the spread.
- *Limit orders.* A limit order that falls inside the spread will increase the best bid if it is a buy order, or decrease the best ask if it is a sell order, narrowing the spread.
- *Cancellations.* A cancellation of a limit order at the best price can either increase the best ask or decrease the best bid, widening the spread.

We prefer to study individual events for several reasons: (1) It removes any ambiguity about time scale, and makes it easier to compare stocks with different activity levels. (2) It minimizes problems associated with clustered volatility (the positive autocorrelation of the absolute value of price changes). Much of clustered volatility is driven by variations in event arrival rates, so the effect is much weaker when the time lag between price changes is measured in terms of number of intervening events. (3) Individual events are the most fundamental level of description – returns on any longer scale can be built out of returns on an event scale. We will generally measure time in terms of *event time*: the time interval between events t_1 and t_2 is measured as the number of intervening events plus one. (One is added so that two adjacent events comprise an event time interval of one).

Analysis at the event scale is particularly useful under the assumption that large price fluctuations are asymptotically power law distributed because of the invariance properties of power laws under aggregation. If two independent variables X and Y both have power law tails with exponents α_X and α_Y , then the variable $Z = X + Y$ also has a power law tail with exponent $\min(\alpha_X, \alpha_Y)$. Thus, if we show that returns at the level of individual events have power law tails with exponent α and if there are not strong and persistent correlations, returns on any longer event time scale should also have power law tails with exponent α . In fact, α appears to be the same whether returns are measured in event time or real time (see Section VI).

II. DATA

In order to have a representative sample of high volume stocks we select 16 companies continuously traded

TABLE I: Summary statistics of the 16 stocks we study for the period 1999-2002. The columns give the number of events of each type, in thousands. All events are “effective” events – see the discussion in Section IA.

tick	market orders	limit orders	cancellations	total
AZN	652	2,067	1,454	4,173
BAA	226	683	487	1,397
BLT	297	825	557	1,679
BOOT	246	711	501	1,458
BSY	404	1,120	726	2,250
DGE	527	1,329	854	2,709
GUS	244	734	518	1,496
HG.	228	676	472	1,377
LLOY	723	1,664	1,020	3,407
PRU	448	1,227	821	2,496
PSON	373	1,063	734	2,170
RIO	381	1,122	771	2,274
RTO	276	620	389	1,285
RTR	479	1,250	820	2,549
SBRY	284	805	561	1,650
SHEL	717	4,137	3,511	8,365
total	6,505	20,033	14,196	40,734

on the London Stock Exchange (LSE) in the 4-year period 1999-2002. The stocks we analyzed are Astrazeneca (AZN), Baa (BAA), BHP Billiton (BLT), Boots Group (BOOT), British Sky Broadcasting Group (BSY), Diageo (DGE), Gus (GUS), Hilton Group (HG.), Lloyds Tsb Group (LLOY), Prudential (PRU), Pearson (PSON), Rio Tinto (RIO), Rentokil Initial (RTO), Reuters Group (RTR), Sainsbury (SBRY), Shell Transport & Trading Co. (SHEL). Table I gives a summary of the number of different events for the 16 stocks.

Trading begins each day with an opening auction. There is a period leading up to the opening auction in which orders are placed but no transactions take place. The market is then cleared and for the remainder of the day (except for occasional exceptional periods) there is a continuous auction. We remove all data associated with the opening auction, and analyze only orders placed during the continuous auction.

An analysis of the limit order placement shows that in our dataset approximately 35% of the effective limit orders are placed *inside the book* ($\pi > a(t)$ or $\pi < b(t)$). 33% are placed at the best prices ($\pi = b(t)$ or $\pi = a(t)$), and 32% are placed inside the spread ($b(t) < \pi < a(t)$). This is roughly true for all the stocks except for SHEL for which the percentages are 71%, 18% and 11%, respectively. Moreover for all the stocks the properties of buy and sell limit orders are approximately the same.

In this dataset cancellation occurs roughly 32% of the time at the best price and 68% of the time inside the book. This is quite consistent across stocks and between the cancellation of buy and sell limit orders. As for the case of the placement of limit orders, the only significant deviation is SHEL for which the percentages are 14% and 86%.

III. FLUCTUATIONS IN LIQUIDITY DRIVE THE TAILS

Given the continuous double auction mechanism for price formation described above, the assumption that large price changes are caused by large market orders is very natural. A very large market order will dig deeply into the limit order book, causing transactions at many price levels, increasing the spread, and changing the mid-price. Surprisingly, this is *not* the cause of most large price changes. Instead, as we will demonstrate in this section, most large price changes are due to discrete fluctuations in liquidity, manifested by gaps in filled price levels in the limit order book. Large price changes caused by large orders are very rare, and play an insignificant role in determining the statistical properties of price changes.

In this section we will focus on price changes caused by market orders, and in Section V we will discuss price changes due to limit orders and cancellations.

A. Large returns are not caused by large orders

We first demonstrate that most large returns are not caused by the arrival of large market orders. The probability density function of price returns that are caused by market orders can trivially be written as $p(r) = \int p(r|\omega)p(\omega)d\omega$, where $p(r)$, $p(\omega)$ and $p(r|\omega)$ are the probability density functions for returns r , market order size ω , and returns given market order size. The conditional probability $p(r|\omega)$ characterizes the response of prices to new orders, and can be viewed as the probability density of market impacts, or alternatively, as characterizing the distribution of liquidity for market orders. When a market order of size ω arrives the midprice will move if ω is larger than or equal to the volume at the matching best price (i.e. the bid for sell market orders and the ask for buy market orders). In the limit of continuous prices we can trivially write

$$p(r|\omega) = (1 - g(\omega))\delta(r) + g(\omega)f(r|\omega) \quad (1)$$

where $\delta(\cdot)$ is the Dirac delta function³ and $g(\omega)$ is the probability that the midprice moves as a function of the order size ω . The function $f(r|\omega)$ is the probability of a price shift r , conditioned on the midprice moving in response to an order of volume ω . $g(\omega)$ and $f(r|\omega)$ behave quite differently. $g(\omega)$ depends strongly on ω . For the LSE it scales roughly as $g(\omega) \sim \omega^{0.3}$, about the same as the average market impact. In contrast, $f(r|\omega)$ is surprisingly independent of the volume ω , in the sense that the unconditional fluctuations in r dominate the dependence on ω .

³ The Dirac delta function $\delta(x)$ is defined so that $\int \delta(x) = 1$ over any domain that includes 0 and $\int \delta(x) = 0$ otherwise.




FIG. 2: Dependence of returns on order size. $F(r > X|\omega)$ is the probability of a return $r > X$ conditioned on the order size ω and on the fact that the price shift is nonzero. Results shown are for buy orders, but similar results are seen for sell orders. The orders are sorted by size into five groups with roughly the same number of orders in each group. Specifically we divide the sets in quintiles according to the order volume. Ranging from low volume to high volume, the curves are black, red, green, blue, and magenta. In panel (a) we show the result for AZN and in panel (b) we show the average over the pool of 16 stocks described in Table I. In this case for each stock we normalize the order volume to the sample mean. Each curve approximately approaches a power law for large r independent of ω , illustrating that the key property determining large price returns is fluctuations in market impact, and that the role of the volume of the order initiating a price change is minor.

To demonstrate this in Figure 2 we show the cumulative probability $F(r > X|\omega) = \int_X^\infty f(r|\omega)dr$ for several different ranges of market order size ω . The distributions for each range of ω are roughly similar, both for individual stocks such as AZN, and for the pool. Although we do not present them here, using the TAQ database we have obtained similar results for a small sample of stocks




FIG. 3: $p(\omega|r > X)$, the distribution of order sizes conditioned on generating a return greater than X , for the stock AZN. The values of X correspond to the 50 percentile (red), 90 percentile (blue), and 99 percentile (green) of the return distribution. The black line is the unconditional distribution $p(\omega)$. The jumps are due to the tendency to place orders in round numbers of shares.

traded in the New York Stock Exchange, suggesting that this behavior is not specific to the LSE.

To reinforce this point, in Figure 3 we show $p(\omega|r > X)$, the probability of an order of size ω conditioned on the return being greater than a certain threshold X . We do this for the stock Astrazeneca (AZN) for several different values of X , getting virtually the same curve independent of X . The distribution of order sizes that generate large returns are essentially the same as those that generate small returns. Similar results are obtained for all the stocks in our sample. Thus, it seems that order size does not play an important role in generating large returns.

Gabaix et al. [25] have recently proposed that large returns can be explained by assuming the market impact density function $p(r|\omega)$ is sharply peaked around a central value⁴ $k\omega^{1/2}$, so that it can be approximated as a Dirac delta function $\delta(r - k\omega^{1/2})$. The result of Figure 2

⁴ Analyses of several markets make it clear that the mean response of prices to orders varies considerably from market to market, and is not in general well characterized by a square root law [26, 29, 31, 33]. It is also worth noting that the volume distribution in Figure 3 does not appear to have a power law tail. This is true of all the stocks in our sample. In contrast, power law tails are observed for stocks in the New York Stock Exchange [34]. One possible explanation is that the data set we use to study the LSE here contains only orders placed in the electronic order book, omitting trades in the upstairs market, whereas the TAQ data set used to study the NYSE is based on transactions and includes trades in the upstairs market, where the largest transactions take place (see also the discussion in [35]).

makes it clear that the sharply peaked assumption is a poor approximation – the distribution is quite broad, in the sense that the conditional distribution $p(r|\omega) \approx p(\omega)$.

B. Volume dependence of mean market impact

From a variety of previous studies it is clear that the mean market impact is an increasing function of ω [29, 33, 36–44]. How do we reconcile this with our claim that the distribution of market impact is almost independent of volume? The key is that the main effect of changing order volume is to change the probability that the price will move, with very little effect on how much it moves. From Eq. (1) one easily obtains

$$E(r|\omega) = g(\omega) \int f(r|\omega)r dr \propto g(\omega) \quad (2)$$

where for the last proportionality relationship we have used the result from Figure 2 that $f(r|\omega)$ is almost independent of ω . Thus, the expected price change scales like the probability of a price change, a relationship that we have verified for both the LSE and the NYSE. However this variation is still small in comparison with the intrinsic variation of returns; the mean market impact of a very large order is less than the average size of the spread, but the largest market impacts are often more than ten times this large.

C. Correlations between order size and liquidity

One possible explanation of the independence of price response and order size is that there is a strong correlation between order size and liquidity. There is an obvious strategic reason for this: Agents who are trying to transact large amounts split their orders and execute them a little at a time, watching the book, and taking whatever liquidity is available as it enters. Thus, when there is a lot of volume in the book they submit large orders, and when there is less volume, they submit small orders. This effect tends to even out the price response of large and small orders. We will see that this effect indeed exists, but it is only part of the story.

In fact, the unconditional correlation between market order size and volume at the best is rather small. For AZN, for example, it is about 1%. However, if we restrict the sample to orders that change the midprice, the correlation soars to 86%. The reason for this is that for orders that do not change the price, there is essentially no correlation between order size and volume at the best. For orders that do change the price, in contrast, most market orders exactly match the volume at the best. As shown in Table II, for the stocks in our sample, 86% of the buy orders that change the price exactly match the volume at the best price. (This is 85% for sell orders.)

The relationship between the volume of market orders and the best price becomes more evident with a nonlin-

tick	% of nonzero return equal to first gap		% of nonzero return with $\omega = V_{best}$	
	sell	buy	sell	buy
AZN	94.3	99.6	83.7	90.1
BAA	95.9	99.1	86.2	87.0
BLT	95.0	99.2	85.8	85.6
BOOT	95.9	99.2	85.7	84.7
BSY	94.3	99.7	85.0	88.0
DGE	96.3	99.7	86.5	87.4
GUS	95.8	99.5	85.0	83.5
HG.	95.9	99.5	85.8	83.0
LLOY	97.3	99.8	88.4	88.6
PRU	95.9	99.5	85.5	78.1
PERSON	93.1	99.6	81.2	86.2
RIO	95.7	99.7	84.6	86.4
RTO	96.1	99.5	84.4	85.3
RTR	93.0	99.7	83.5	85.4
SBRY	95.6	99.4	85.3	83.2
SHEL	98.7	99.9	93.0	92.8
average	95.6	99.5	85.6	86.0

TABLE II: Summary table of the percentage of the time that nonzero changes in the best prices are equal to the first gap (left) and that the market order volume ω exactly matches the volume at the corresponding best price V_{best} (right). Assuming a Bernoulli process the sample errors are the order of 0.1 – 0.2%.

ear analysis. Figure 4 shows $E(\omega|V_{best})$, where ω is the market order size and V_{best} is the volume available to the corresponding best price. This behavior can easily be understood by using the model of Eq. (1), and making the approximation that market orders are never bigger than the volume at the best price. If $g(\omega)$ is the probability that the price moves because of an arriving market order of size ω and $p(\omega)$ is the probability that a market order has size ω , we can write

$$E(\omega|V_{best}) \approx \int_0^{V_{best}} [1 - g(\omega)]\omega p(\omega) d\omega + V_{best} \int_0^\infty g(\omega)p(\omega) d\omega. \quad (3)$$

The first term is the contribution when $\omega < V_{best}$, i.e. when the price does not move. The second term gives the contribution when the price moves because $\omega = V_{best}$. From Table II the case $\omega > V_{best}$ is sufficiently rare that it can be neglected. $p(\omega)$ can be approximated very well by a Gamma distribution, and $g(\omega)$ can be approximated as $g(\omega) = (\omega/\bar{\omega})^\beta / (1 + (\omega/\bar{\omega})^\beta)$. This form is chosen because it matches the empirical result $g(\omega) \approx \omega^{0.3}$ for small ω but saturates for large ω , ensuring that the probability $g(\omega) < 1$ for all ω . This gives a good fit to the result of Figure 4.

However, in the following section we will see that the dependence of order size on liquidity is not strong enough to substantially suppress large price fluctuations.




FIG. 4: The dependence of order size on liquidity. Liquidity is measured as the volume at the best price, which is binned into deciles with roughly equal numbers of events. The vertical axis shows the mean market order volume for each decile, and the horizontal shows the volume at the best. The units of both axes are thousands of shares. The ranges shown around each point indicate one standard deviation (and are not standard errors). The dashed line has slope 1 and serves as a point of comparison.

D. Liquidity fluctuations drive price fluctuations

In this section we demonstrate in concrete terms that price fluctuations are driven almost entirely by liquidity fluctuations. To do this we study the virtual market impact, which is a useful tool for probing the supply and demand curves defined by the limit order book. Whereas the true market impact $p(r|\omega)$ tells us about the distribution of impact of actual market orders, as discussed in Section III A, the virtual market impact is the price change that *would* occur at any given time if a market order of a given size were to be submitted. More formally, at any given time t the limit orders stored in the order book define a revealed supply function $S(\pi, t)$ and revealed demand function $D(\pi, t)$. Let $V(\pi, t)$ be the total volume of orders stored at price π . The revealed supply function is

$$S(\pi, t) = \sum_{\pi} V(\pi, t) \quad (4)$$

The revealed supply function is non-decreasing, and so for any fixed t has a well defined inverse $\pi(S, t)$. The virtual market impact is the price shift caused by a hypothetical order of size S , e.g. for buy orders it is $\pi(S, t) - a(t)$. The virtual market impact for sell market orders can be defined in terms of the revealed demand in a similar manner. By sampling at different values of t , for any fixed hypothetical order size we can create a sample distribution of virtual market impacts. This naturally depends on the sampling times, but these can be




FIG. 5: The cumulative probability distribution $P(r > x)$ of virtual market impacts for the stock AZN, based on samples of the limit order book taken just before the arrival of a market order. From left to right the continuous lines are the virtual market impact for hypothetical buy order sizes $\omega = 100$ shares (red), $\omega = 1,600$ shares (green), $\omega = 8,700$ shares (blue), and $\omega = 25,000$ shares (magenta); these correspond to the 0.1, 0.5, 0.9, and 0.99 quantiles of market order size. These are compared to the average distribution of returns (black dashed), which is very similar to the virtual market impact for the 0.5 quantile. This demonstrates the distribution of returns is essentially determined by fluctuations in the limit order book, and that the typical price return corresponds to the price response to an order of typical size. (Note that we have discarded cases in which the virtual market impact is undefined due to excessively large hypothetical market order volume, but these comprise only about 3% of the events.)

chosen to match any given set of price returns.

In Figure 5 we show the cumulative distribution of virtual market impacts for the stock AZN for several different values of D , corresponding to different quantiles of market order size. The cumulative distributions define a set of approximately parallel curves. They are shifted as one would expect from the fact that larger hypothetical order sizes tend to have larger virtual market impacts. These curves are similar in shape to the distribution of returns. Most striking, for the median market order size, the curves are almost identical. We see similar results for all the stocks in our sample. This demonstrates quite explicitly that the distribution of returns is determined by properties of the limit order book, and not by market order size. The fact that the price distribution is so close to the virtual market impact of a typical order shows that correlations between order size and liquidity are not important in determining price fluctuations.




FIG. 6: A typical configuration of the limit order book before and after a large price fluctuation, for AZN. The two panels plot the volume (in shares) of limit orders at each price level; sell limit orders are shown as positive, and buy limit orders as negative. In panel (a) we see that there is a gap of 43 ticks between the best ask price and the next highest occupied level. The arrival of a crossing limit order to buy removes the volume at the best ask, giving the new limit order book configuration shown in panel (b), which has a much higher best ask price than previously. The only thing atypical about this event is that in most cases the volume at the best price is removed by a market order of exactly the right size.

IV. GRANULARITY OF SUPPLY AND DEMAND

At first sight, the behavior described in the previous section seems baffling: How can market order size be so unimportant to price response? In this section we show how this is due to granularity of revealed supply and demand, which causes large fluctuations in liquidity.

The cause of this puzzling behavior is fluctuations in occupied price levels in the limit order book. In particular, one can define the size of the *first gap* g as the absolute difference between the best log price π_{best} and the log price of the next best quote π_{next} , $g = |\log \pi_{best} - \log \pi_{next}|$. In Figure 6 we show a typical set of events before and after a large price fluctuation. A large first gap appears on the sell side of the limit order book, and an effective market order that removes all the volume at the best ask causes a large price change.

As already shown in Table II, typically about 85% of market orders that result in price changes exactly match the volume at the best price. Furthermore, as shown in the first two columns of this table, when they do exceed the volume at the best price, it is quite rare that they

penetrate the next occupied level. For buy (sell) market orders, on average 95.5% (99.5%) of the nonzero shifts in the best price are exactly equal to the first gap. Interestingly there is a significant asymmetry between buy and sell. Traders act to minimize their transaction costs, so that jumps of more than one gap are rare, particularly for sell orders. (However, as we stressed earlier, the return distribution can be generated by a constant order of median size – this correlation is interesting, but not essential). Note that the trader initiating the change *does not* pay a large spread – that would only happen to the next trader, if she were to immediately place a market order⁵.

This is by far the most common scenario that generates large price changes. In Figure 7(a) we compare the distribution of the first gaps to the distribution of price returns for the stock AZN. We see that the distributions are very similar. For the first gap size the tail index $\alpha = 2.52 \pm 0.07$, and for the return distribution $\alpha = 2.57 \pm 0.08$, showing that the scaling behaviors are similar. However, the similarity is not just evident in the scaling behavior – the match is good throughout the entire range, illustrating that most of the large price changes are caused by events of this type. Panel (b) of Figure 7 shows the same comparison for the pool of 16 stocks. The agreement in this case is even more striking⁶.

To demonstrate that the correspondence in the above figure is not just a coincidence for AZN, we have computed the tail exponents for returns and first gap size for all 16 stocks in our data set. We do this using a Hill estimator [45] by considering the largest \sqrt{n} returns, where n is the size of the sample. The results are shown in Figure 8, where we plot the tail exponents of returns against those of first gaps. There is a strong positive correlation ($R^2 = 0.93$) between the gap tail exponents and the return tail exponents; a least squares linear fit has a slope of 0.98 ± 0.05 , so we are unable to reject the null hypothesis that the tail exponents are the same.

It is worth noting that lightly traded stocks tend to have a smaller tail exponent than heavily traded stocks. To investigate this more quantitatively, in Figure 9 we plot the tail exponent of price changes against the number of market orders for each stock in the sample. While the relationship is noisy, there is a clear positive trend; the slopes of the linear fits to positive and negative returns are both highly statistically significant. The tail exponents vary from about 1.6 to 2.8, whereas the error

⁵ In general after a large shift in the bid or ask price, the next orders tend to be limit orders, but we have not yet been able to study the statistical properties of the sequence of subsequent events in detail.

⁶ We have also studied the distribution of higher order gaps. Moving away from the best price π_0 , the n^{th} order gap for $i = 1, 2, \dots$ can be recursively defined as $g_n = |\log \pi_{n-1} - \log \pi_n|$, where π_n is the n^{th} occupied price level. Interestingly, we find that the tail behavior of higher order gaps is the same as that of g_1 .




FIG. 7: The cumulative distribution $P(g > x)$ of the size of first gaps g (red continuous line), compared to the cumulative distribution of returns generated by market orders $P(r > x)$ (black dashed line). Panel (a) refers to buy market orders for AZN, in double logarithmic scale to highlight the tail behavior. The two distributions are very similar. The result is even more impressive when we consider the average over the 16 stocks described in Table I (Panel (b))

bars are more than a factor of ten smaller than the range of variation⁷. This suggests that the tail exponent is not universal. Not surprisingly, stocks that are more heavily traded tend to display less extreme risk than stocks that are lightly traded.

⁷ Table III gives error bars for tail exponent estimations based on absolute returns (including both buy and sell events), while Figures 7 and 8 are based on buy and sell returns taken separately, but the error bars are comparable.




FIG. 8: A comparison of tail exponents α for returns (horizontal axis) vs. tail exponents for the first gap (vertical axis). The first gaps are sampled immediately preceding the market order that triggers the return. This is done for all the stocks in Table I, and includes both the buy and sell side of the limit order book. We use red circles for low liquidity stocks, blue squares for medium liquidity stocks, and green triangles for high liquidity stocks. Empty symbols refer to sell market orders and filled symbols to buy market orders. The black dashed line is the linear regression. We see that there is a clear positive relationship between the number of orders and the tail exponent.




FIG. 9: Dependence of the tail exponent of price changes on the number of market orders, for the data of Figure 8. Black circles are for positive returns (caused by buy market orders), and red triangles are for negative returns (caused by sell market orders). The steeper red curve is the fit to negative returns and the black curve the fit to positive returns.




FIG. 10: A comparison of the distribution of absolute returns (including both buy and sell orders) caused by limit orders, market orders, and cancellations for AZN. The black circles are for market orders, the blue triangles are for limit orders, and the red squares are for cancellations.

V. LIMIT ORDERS AND CANCELLATIONS

We have so far considered only returns caused by market orders. In this section we will discuss returns caused by limit orders and cancellations, and show that they are statistically indistinguishable from those caused by market orders.

The distribution of returns caused by limit orders, market orders, and cancellations for AZN is shown in Figure 10. The distributions are quite similar. To make the analysis more quantitative, we compute the tail exponents for the 16 stocks in the sample. Table III shows the Hill estimates for returns caused by market orders, limit orders and cancellations. The table shows clearly that the tail exponents for the three distributions are very close: For all 16 stocks the estimated tail exponents are well within the 95% confidence intervals. On the other hand Table III shows that there is a much larger variation across stocks.

Why are the distributions of different events so similar? The correspondence between returns caused by market orders and returns caused by cancellations is not surprising. In terms of the effect on best prices, removal of the volume at the best price by cancellation is equivalent to removal by a market order, in both cases creating a price change equal in size to the first gap. However, for a limit order this is not so obvious: A limit order that falls inside the spread decreases the spread, and causes a price change in the opposite direction from market orders and cancellations. To investigate this we have studied the distribution of the spread, which also appears to be a power law, but with a larger tail exponent. At this point the reason for the close correspondence between the returns generated by limit and market orders remains un-

TABLE III: Hill estimator of the return caused by market orders, limit orders and cancellations. We analyzed the top percentile and give 95% confidence intervals. The confidence intervals assume the samples are uncorrelated, and so may be overoptimistic.

tick	Market orders	Limit orders	Cancellations
AZN	2.57 ± 0.08	2.56 ± 0.07	2.57 ± 0.06
BAA	1.77 ± 0.06	1.77 ± 0.08	1.76 ± 0.08
BLT	1.8 ± 0.1	1.8 ± 0.1	1.8 ± 0.1
BOOT	1.80 ± 0.07	1.80 ± 0.07	1.79 ± 0.08
BSY	2.28 ± 0.04	2.25 ± 0.06	2.26 ± 0.06
DGE	2.31 ± 0.06	2.31 ± 0.07	2.31 ± 0.07
GUS	1.80 ± 0.09	1.80 ± 0.07	1.80 ± 0.07
HG.	1.7 ± 0.1	1.7 ± 0.1	1.66 ± 0.07
LLOY	2.72 ± 0.03	2.72 ± 0.04	2.72 ± 0.05
PRU	2.20 ± 0.06	2.20 ± 0.07	2.20 ± 0.06
PSON	1.9 ± 0.1	1.9 ± 0.1	1.91 ± 0.09
RIO	1.83 ± 0.09	1.82 ± 0.08	1.83 ± 0.07
RTO	1.73 ± 0.09	1.73 ± 0.08	1.73 ± 0.07
RTR	2.45 ± 0.04	2.44 ± 0.05	2.45 ± 0.03
SBRY	1.89 ± 0.07	1.88 ± 0.07	1.88 ± 0.07
SHEL	2.62 ± 0.04	2.61 ± 0.07	2.62 ± 0.04

explained.

VI. PRICE CHANGES ON LONGER TIMESCALES

One can naturally ask whether the microscopic event level analysis we have presented here explains the statistical properties of price changes over longer and/or fixed time horizons. The timescale for the analyses we have presented here is quite fast: For Astrazeneca, for example, within the period of this dataset there are on average about eight market orders every five minutes, but due to highly uneven rates of event arrival, it is not uncommon that many events arrive within the same second. It is natural to question whether events on this timescale reflect the properties of longer timescales, e.g. on the daily timescale of many other studies. Given that order arrival is highly clustered in time, it is also natural to ask whether a description in event time also provides an explanation in real time. We now address both of these questions.

The permanence of these price movements can be seen through the continuity between price returns on the single event scale and the multiple event scale. Figure 11 shows the return density function for AZN for 1, 2, 4, 8, 16, and 32 market order arrivals, as well as for 5, 10, and 20 minute timescales. The regularity of the change, and the similarity of the tail behavior across a range of timescales, indicates that the tail properties on the single event scale are reflected on larger scales (at least as large as 20 minutes). Also, this figure makes it clear that the behavior in real time is essentially the same as the behavior in event time. We see similar behavior for all the stocks in our sample. Since other studies have shown the




FIG. 11: Price returns aggregated on different timescales and for different number of trades. The price return is defined as $\log m(t) - \log m(t - \tau)$, where τ is from left to right 1, 2, 4, 8, 16 and 32 market orders (continuous lines) and 5, 10, and 20 minutes (dashed lines). We see that the tails of all the distributions are quite similar, though the crossover from the central behavior to the tail behavior increases on longer timescales. The data is for AZN, based on absolute returns (both buy and sell).

continuity between timescales from 10 minutes to as long as a month [8, 12, 13], it seems clear that an explanation at the level of single events is sufficient⁸.

The fact that we observe essentially identical distributions when we aggregate over individual transactions or over fixed time rules out any explanation of fat tails for price returns based on a subordinated random process, as was suggested by Clark [14]. While fluctuations in the number of events in a given length of time might be quite important for other phenomena, such as clustered volatility, they are clearly not important in determining the price return distribution.

VII. CORRELATIONS IN OCCUPIED SITES

A natural hypothesis about the origin of the large gap distributions involves fluctuations in the number of occupied sites in the order book. If there are only a few orders in the book, so that it is mostly empty, we naturally expect large gaps to exist. Perhaps the approximate power for the gap distribution can be explained by a similar power law in the number of occupied sites? In this section we show that this is not the case, but rather that the





FIG. 12: The effect of the finite number of orders, for AZN. Panel (a) shows the unconditional distribution of the number of occupied price levels for buy orders, plotted on double logarithmic scale. The dashed line is the best fit of the low N region with a functional form $p(N) \propto N^\beta$. Panel (b) shows the probability density of first gap sizes conditional on the number of occupied price levels, $p(g|N)$. We divide the sample in four subsamples according to the value of N . Specifically, we have $0 < N \leq 15$ (black circles), $15 < N \leq 30$ (red squares), $30 < N \leq 45$ (green diamonds), and $45 < N \leq 60$ (blue triangles). Surprisingly, the first gap size shows approximate power law behavior even when the number of occupied sites is very large.

power law behavior depends on correlations in occupied price levels.

Interestingly, we do find evidence of power law scaling in the number of occupied price levels. Figure 12 (a) shows the probability density function $p(N)$, where N is the number of occupied price levels in the order book at any given time. The low N region is well fit by a power-law $p(N) \sim N^\beta$ over about two decades.

However, the fluctuations in the number of orders stored in the book cannot explain the power law behavior of prices, because the tail exponent of the power law

⁸ We have studied the response to individual events, and see some evidence for mean reversion of large price changes. However, the results of Figure 11 make it clear that these are not sufficient to undo the initial change.

behavior of the gap distribution near zero is much too large. If we assume that the location of order deposition is uncorrelated, then for order arrival on a bounded domain gap size should be proportional to $1/N$. But with $p(N) \sim N^\beta$ and $g \sim 1/N$, $p(g) \sim g^{-(\beta+2)}$. For $\beta \approx 5.5$ this results in $p(g) \sim g^{-7.5}$. In contrast, the empirically observed scaling exponent for the density for AZN is about 3.5.

Surprisingly, the distribution of gap size shows approximate power law scaling, independent of N . This is illustrated for AZN in Figure 12(b), where we plot the probability of gap size conditioned on N , $p(g|N)$, for several different values of N . The approximate power law behavior is evident even with as many as 60 occupied sites on one side of the book. Considering that it is rare for orders to be placed more than 100 price ticks away from the best price, this illustrates that occupied sites display nontrivial correlations, which are essential for explaining large price fluctuations. This conclusion is reinforced by studies we have done comparing the real distribution of occupied levels with models based on IID order placement, which do not reproduce the power law behavior.

VIII. CONCLUSIONS

For the London Stock Exchange, we have shown that large fluctuations in prices are unrelated to large transactions, or to the placement of large orders. Instead, large price fluctuations occur when there are gaps in the occupied price levels in the limit order book. Large changes occur when a market order removes all the volume at the best price, creating a change in the best price equal to the size of the gap.

At a higher level, these results demonstrate that large price changes are driven by fluctuations in liquidity. There are times when the market absorbs changes in supply and demand smoothly, and other times when a small change in supply or demand can result in a very large change in the price. This is due to the fact that supply and demand functions are not smooth, but rather have large, irregular steps and jumps. The market is granular, due to the presence of only a finite number of occupied price levels in the book. This is what in physics is called a finite size effect. Even for an active stock such as AZN, the number of occupied price levels on one side of the book at any given time is typically about 30, and so the system is far from the continuum limit. However, we have shown that this is not a simple matter of fluctuations in the number of occupied price levels; while there is an approximate power law in the limit $N \rightarrow 0$ in the frequency for occupying N price levels, this is not sufficient to explain fluctuations in prices. Instead, the power

law in gaps persists even when the number of occupied levels is quite high, reflecting nontrivial correlations in the positions of orders sitting in the book.

The empirical results that we have presented here raise as many questions as they answer. In particular, what is responsible for the approximate power law distribution of gap sizes, and why do limit orders have the same tail exponents as market orders and cancellations? We have done some modeling to address this question. In particular, we have modified the statistical model of order flow introduced by Daniels et al. [27, 30] to include nonuniform order placement, more closely reflecting the empirical order placement distribution [28, 46], which has an approximate power law tail. For some parameter values our simulations reproduce the power law tails of the gap distribution (and hence of returns), but since we do not yet understand the necessary and sufficient conditions for this, we do not include these results here.

The work presented here suggests that it is important to properly model market institutions. However, at this point it is not clear how these results would change for a different market institution. Fat tails are also observed in markets, such as the London Metals Exchange, that follow very different exchange mechanisms [47]. While the LSE follows a continuous double auction, it is our belief that key elements are likely to persist with other market mechanisms. In particular, we hypothesize that large price fluctuations in any market are driven by liquidity fluctuations, and that the granularity of fluctuations in supply and demand remains the key factor underlying extreme price fluctuations.

These results are important because they reveal the detailed mechanism through which prices display (at least approximate) power law fluctuations. They suggest that many previous models that claimed to explain this phenomenon were misdirected, and provide strong constraints on future models. The tail exponent of large price changes appears to depend on parameters of the market: More lightly traded markets (with lower event rates) tend to display fatter tails, with more extreme risks. This has important practical importance because it gives some understanding of what determines financial risks, and gives some clues about how to reduce them.

Acknowledgments

We would like to thank the James S. McDonnell Foundation for their Studying Complex Systems Research Award, Credit Suisse First Boston, McKinsey Corporation, Bob Maxfield, and Bill Miller for supporting this research. We would also like to thank Marcus Daniels for technical support.

[1] Mandelbrot, B.B., The variation of certain speculative prices, *J. Business*, **36**, 394 (1963).

[2] Fama, E.F., The behavior of stock market prices, *J. Busi-*

- ness, **38**, 34 (1965).
- [3] Officer, R.R., The distribution of stock returns, *J. of the American Statistical Assoc.* **67**, 807-812 (1972).
- [4] Akgiray, V., Booth, G.G. and Loistl, O., Stable laws are inappropriate for describing German stock returns, *Algemeines Statistisches Archiv.* **73**, 115-121 (1989).
- [5] Koedijk, K.G., Schafgans, M.M.A., and De Vries, C.G., The tail index of exchange rate returns, *J. Intl. Economics* **29**, 93 (1990).
- [6] Loretan, M, and Phillips, C.B., Testing the covariance stationarity of heavy-tailed time series: An overview of the theory with applications to several financial datasets, *J. of Empirical Finance* **1**, 211-248 (1994).
- [7] Mantegna, R.N., and Stanley, H.E., Scaling behavior in the dynamics of an economic index, *Nature* **376**, 46 (1995).
- [8] Longin, F., The asymptotic distribution of extreme stock market returns, *J. of Business* **69**, 383-408 (1996).
- [9] Lux, T., 1996, The stable Paretian hypothesis and the frequency of large returns: an examination of major German stocks, *Applied Financial Economics* **6**, 463-475 (1996).
- [10] Ghashghaie, S., et al., Turbulent cascades in foreign exchange markets, *Nature* **381**, 767 (1996).
- [11] Muller, U.A., Dacorogna, M. M., and Pictet, O.V., Heavy tails in high-frequency financial data, in *A Practical Guide to Heavy Tails*, in Adler, R.J., Feldman, R.E., and Taqqu, M.S., Birkhuser, Boston, 398 (1998).
- [12] Plerou, V., Gopikrishnan, P., Amaral, L.A.N., Meyer, M., and Stanley, H.E., Scaling of the distribution of price fluctuations of individual companies, *Physical Review E* **60**, 6519-6529 (1999).
- [13] Mantegna, R.N. and Stanley, H.E., *Introduction to econophysics: correlations and complexity in finance*, Cambridge University Press, Cambridge (1999).
- [14] Clark, P.K., A subordinated stochastic process model with finite variance for speculative prices, *Econometrica* **41**, 135-256 (1973).
- [15] Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R., and Tayler, P., Asset pricing under endogenous expectations in an artificial stock market, in W. B. Arthur, D. Lane, and S.N. Durlauf, *The Economy as an evolving, complex system II*, Addison-Wesley, Redwood City, CA (1997).
- [16] Mandelbrot, B.B, *Fractals and scaling in finance*, Springer-Verlag, New York (1997).
- [17] Sornette, D., Multiplicative processes and power laws, *Phys. Rev. E.* **57**, 4811-4813 (1998).
- [18] Sornette, D. *Critical Phenomena in Natural Sciences* Springer, Berlin (2000).
- [19] Anteneodo C. and Tsallis, C., Multiplicative noise: A mechanism leading to nonextensive statistical mechanics, to appear in *J. Math. Physics* (2004).
- [20] Cont, R. and Bouchaud, J-P., *Macroeconomic Dynamics* **4**, 170 (2000).
- [21] Lux, T. and Marchesi, M., Scaling and criticality in a stochastic multi-agent model of financial market, *Nature* **397**, 498-500 (1999).
- [22] Giardina, I., and Bouchaud, J-P., Bubbles, crashes and intermittency in agent-based market models, <http://xxx.lanl.gov/cond-mat/0206222> (2002)
- [23] Johnson, N.F., Jefferies, P., Hui, P.M., *Financial Market Complexity* Oxford (2003).
- [24] Challet, D., Chessa, A., Marsili, M., and Zhang, Y-C., From minority games to real markets, *Quantitative Finance* **1**, 168 (2001).
- [25] Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H.E. A theory of power-law distributions in financial market fluctuations, *Nature* **423**, 267-270 (2003).
- [26] Farmer, J.D., and Lillo, F., On the origin of power laws in financial markets' to appear in *Quantitative Finance*. Preprint at xxx.lanl.gov/cond-mat/0309416 (2003).
- [27] Daniels, M, Farmer, J.D., Gillemot, L., Iori, G., and Smith, D.E. Quantitative model of price diffusion and market friction based on trading as a mechanistic random process, *Physical Review Letters* **9019**, 10 (2003).
- [28] Bouchaud, J.-P., Mezard, M., and Potters, M., Statistical properties of the stock order books: empirical results and models. *Quantitative Finance* **2**, 251-256 (2002).
- [29] Potters, M. & Bouchaud, J.-P., More statistical properties of order books and price impact, *Physica A* **324**, 133-140 (2003)
- [30] Smith, E., Farmer, J.D., Gillemot, L., and Krishnamurthy, S., Statistical theory of the continuous double auction, *Quantitative Finance* **3**, 481-514 (2003).
- [31] Farmer, J.D., Patelli, P., and Zovko, I.I., The predictive power of zero intelligence in financial markets, xxx.lanl.gov/cond-mat/0309233 (2003).
- [32] Bouchaud, J.-P., Gefen, Y., Potters, M., and Wyart, M., Fluctuations and response in financial markets: the subtle nature of 'random' price changes, xxx.lanl.gov/cond-mat/0307332 (2003).
- [33] Lillo, F. Farmer, J.D. & Mantegna, R.N., Master curve for price-impact Function, *Nature* **421**, 129-130 (2003).
- [34] Gopikrishnan, P., Plerou, V., Gabaix, X., and Stanley, H.E., Statistical properties of share volume traded in financial markets. *Physical Review E* **62**, R4493-R4496 (2000).
- [35] Lillo, F. and Farmer, J.D., The long memory of the efficient market, xxx.lanl.gov/cond-mat/0311053 (2003).
- [36] Hasbrouck, J., Measuring the information content of stock trades, *J. Finance* **46**, 179-207 (1991).
- [37] Hausman, J.A., and Lo, A.W., 'An ordered probit analysis of transaction stock prices, *Journal of Financial Economics* **31**, 319-379 (1992)
- [38] Farmer, J.D., Slippage 1996, *Prediction Company Technical Report*, <http://www.predict.com/jdf/slippage.pdf> (1996).
- [39] Torre, N., *BARRA Market Impact Model Handbook*, BARRA Inc, Berkeley CA, www.barra.com (1997).
- [40] Kempf, A., and Korn, O., Market depth and order size, *J. Financial Markets* **2**, 29-48 (1999).
- [41] Plerou, V., Parameswaran, G., Gabaix, X., and Stanley, H.E., Quantifying stock price response to demand fluctuations, *Phys. Rev. E* **66**, 027104 (2002).
- [42] Rosenow, B., *Int. J. Mod. Phys. C* **13**, 419-425 (2002).
- [43] Evans, M.D. and Lyons, R.K., Order flow and exchange rate dynamics, *J. Political Economy* **110**, 170-180 (2002).
- [44] Takayasu, H., ed., *Empirical Analysis of Financial Fluctuations*, Springer, Berlin, (2002).
- [45] Hill, B.M., A simple general approach to inference about the tail of a distribution, *Annals of Statistics* **3**, 1163-1173 (1975).
- [46] Zovko, I.I., and Farmer, J.D., The power of patience: A behavioral regularity in limit order placement. *Quantitative Finance* **2**, 387-392 (2002).
- [47] One of us (JDF) did an unpublished study in 1995 of daily data for industrial metals such as copper, zinc,

iron, nickel, etc. traded on the London Metals Exchange, whose exchange mechanism resembles a Walrasian mar-

ket. All of these have large kurtosis, indicating fat tails.