QUANTITATIVE FINANCE VOLUME 4 (2004) 176–190
INSTITUTE OF PHYSICS PUBLISHING

Fluctuations and response in financial markets: the subtle nature of 'random' price changes

Jean-Philippe Bouchaud[1,2], Yuval Gefen[3], Marc Potters[2] and Matthieu Wyart[1]

[1] Commissariat à l'Energie Atomique, Orme des Merisiers, 91191 Gif-sur-Yvette Cedex, France
[2] Science and Finance, Capital Fund Management, 109–111 rue Victor Hugo, 92 532 Levallois Cedex, France
[3] Condensed Matter Physics Department, Weizmann Institute of Science, 76 100 Rehovot, Israel

Received 28 August 2003, in final form 21 November 2003
Published 19 December 2003
Online at stacks.iop.org/Quant/4/176 (DOI: 10.1088/1469-7688/4/2/007)

Abstract
Using trades and quotes data from the Paris stock market, we show that the random walk nature of traded prices results from a very delicate interplay between two opposite tendencies: long-range correlated market orders that lead to super-diffusion (or persistence), and mean reverting limit orders that lead to sub-diffusion (or anti-persistence). We define and study a model where the price, at any instant, is the result of the impact of all past trades, mediated by a non-constant 'propagator' in time that describes the response of the market to a single trade. Within this model, the market is shown to be, in a precise sense, at a critical point, where the price is purely diffusive and the average response function almost constant. We find empirically, and discuss theoretically, a fluctuation–response relation. We also discuss the fraction of truly informed market orders, that correctly anticipate short-term moves, and find that it is quite small.

1. Introduction

The efficient market hypothesis (EMH) posits that all available information is included in prices, which emerge at all times from the consensus between fully rational agents, that would otherwise immediately arbitrage away any deviation from the fair price [1, 2]. Price changes can then only be the result of unanticipated news and are by definition totally unpredictable. The price is at any instant of time the best predictor of future prices. One of the central predictions of EMH is thus that prices should be random walks in time, which (to a good approximation) they indeed are. This was interpreted early on as a success of EMH. However, as pointed out by Shiller, the observed volatility of markets is far too high to be compatible with the idea of fully rational pricing [3]. The frantic activity observed in financial markets is another problem: on liquid stocks, there is typically one trade every 5 s, whereas the time lag between relevant news is certainly much larger. More fundamentally, the assumption of rational, perfectly informed agents seems intuitively much too strong, and has been criticized by many [4–6]. Even the very concept of the fair price of a company appears to be somewhat dubious.

There is a model at the other extreme of the spectrum where prices also follow a pure random walk, but for a totally different reason. Assume that agents, instead of being fully rational, have zero intelligence and take random decisions to buy or to sell, but that their action is interpreted by all the others agents as potentially containing some information. Then, the

1469-7688/04/020176+15$30.00 © 2004 IOP Publishing Ltd PII: S1469-7688(04)68104-8

mere fact of buying (or selling) typically leads to a change of the ask $a(t)$ (or bid $b(t)$) price and hence to a change of the midpoint $m(t) = [a(t) + b(t)]/2$. In the absence of reliable information about the 'true' price, the new midpoint is immediately adopted by all other market participants as the new reference price around which new orders are launched. In this case, the midpoint will also follow a random walk (at least for sufficiently large times), even if trades are not motivated by any rational decision and devoid of meaningful information[4]. This alternative, random trading model has recently been the object of intense scrutiny, in particular as a simplified approach to the statistics of order books [8–15]. Since the order flow is a Poisson process, this assumption is quite convenient and leads to tractable analytical models [14, 16]. Perhaps surprisingly, many qualitative (and sometimes quantitative) properties of order books can be predicted using such an extreme postulate [12–14, 17, 18].

Of course, reality should lie somewhere in the middle: clearly, the price cannot wander arbitrarily far from a reasonable value, and trades cannot all be random. The interesting question is to know which of the two pictures is closest to reality and can be taken as a faithful starting point around which improvements can be perturbatively added.

In this paper, we want to argue, based on a series of detailed empirical results obtained on trade by trade data, that the random walk nature of prices is in fact highly non-trivial and results from a fine-tuned competition between two populations of traders, liquidity providers ('market-makers') on the one hand, and liquidity takers (sometimes called 'informed traders', but we will argue below that this might be a misnomer). For reasons that we explain in more detail below, liquidity providers act so as to create anti-persistence (or mean reversion) in price changes that would lead to a sub-diffusive behaviour of the price, whereas liquidity takers' action leads to long-range persistence and super-diffusive behaviour. Both effects very precisely compensate and lead to an overall diffusive behaviour, at least to a first approximation, such that (statistical) arbitrage opportunities are absent, as expected. However, one can spot the vestiges of this subtle compensation from the temporal structure of the market impact function (which measures how a given trade affects future prices on average).

The organization of this paper is as follows. We first present (section 2) our empirical results on the statistics of trades, market impact and fluctuations. We show in particular that the order flow exhibits long-range (power-law) autocorrelations in time, but that this does not lead to any predictability in price changes, as also recently noticed in [7]. Then, we introduce in section 3 a simple model that expresses the price as a linear superposition of the impact of each trade. We show that this model allows us to rationalize our empirical findings, provided a specific relation between the temporal autocorrelation of the sign of the trades (i.e. buyer initiated or seller initiated) and the temporal response to a single trade is satisfied. Finally, in section 4, we give intuitive arguments that allow one to understand the market forces at the origin of this subtle balance between two opposite effects, which dynamically leads to absence of statistical arbitrage opportunities. We argue that, in a very precise sense, the market is sitting on a critical point; the dynamical compensation of two conflicting tendencies is similar to other complex systems such as the heart [19], driven by two antagonistic systems (sympathetic and para-sympathetic), or certain human tasks, such as balancing of a long stick [20]. The latter example illustrates very clearly the idea of dynamical equilibrium, and shows how any small deviation from perfect balance may lead to strong instabilities. This near instability may well be at the origin of the fat tails and volatility clustering observed in financial data (see e.g. [21–26]). Note that these two features are indeed present in the 'balancing stick' time series studied in [20].

[4] That this simplistic model also leads to a random walk behaviour for prices has also very recently been pointed out in [7].

2. Market impact and fluctuations

2.1. Presentation of the data and definitions

In this study, we have analysed trades and quotes data from liquid French stocks in the years 2001 and 2002, although qualitatively similar results were also obtained on British stocks as well. The advantage of the French market, however, is that it is fully electronic, whereas only part of the volume is traded electronically in the London stock exchange. We will illustrate our results mainly using the France-Telecom stock, which is one of the most actively traded stocks, for which statistics are particularly good.

There are two data files for each stock: one gives the list of all successive *quotes*, i.e. the best buy (bid, b) and sell (ask, a) prices, together with the available volume, and the time stamp accurate to the second. A quote can change either as a result of a trade, or because new limit orders appear, or else because some limit orders are cancelled. The other data file is the list of all successive *trades*, with the traded price, traded volume and time stamp, again accurate to the second. Sometimes, several trades are recorded at the very same instant but at different prices: this corresponds to a market order of a size which exceeds the available volume at the bid (or at the ask), and hits limit orders deeper in the order book. In the following, we have grouped all these trades together as a single trade. This allows one to create chronological sequences of trades and quotes, such that between any two trades there is at least one quote.

The last quote before a given trade allows one to define the sign of each trade: if the traded price is above the last midpoint $m = (a + b)/2$, this means that the trade was triggered by a market order (or marketable limit order) to buy, and we will assign to that trade a variable $\varepsilon = +1$. If, on the other hand, the traded price is below the last midpoint $m = (a + b)/2$, then $\varepsilon = -1$. With each trade is also associated a volume V, corresponding to the total number of shares exchanged.

Trades appear at random times, the statistics of which is itself non-trivial (there are intraday seasonalities and also clustering of the trades in time). We will not be interested in this aspect of the problem and always reason in terms of trade time, i.e. time advances by one unit every time a new trade (or a series of simultaneous trades) is recorded. We have also

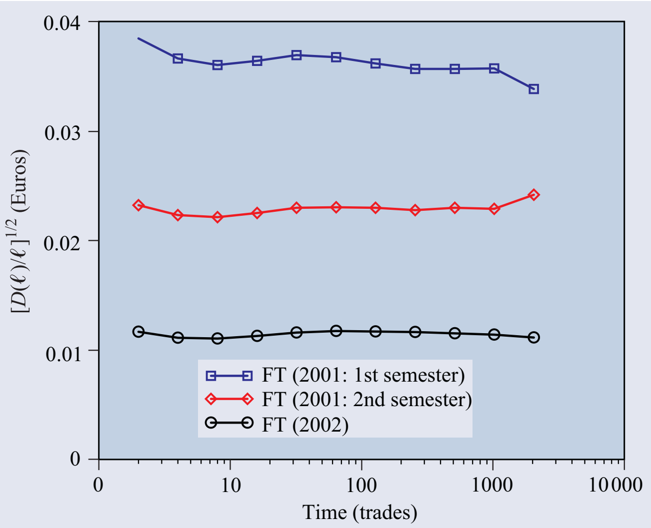

Figure 1. Plot of $\sqrt{\mathcal{D}(\ell)/\ell}$ as a function of ℓ for France-Telecom, during three different periods. The variation of $\mathcal{D}(\ell)/\ell$ with ℓ is very small, in particular in the small-tick (0.01 Euros) period (July 2001–December 2002). For the large-tick-size period (0.05 Euros; January 2001–June 2001), there is a systematic downward trend: see also figure 2.

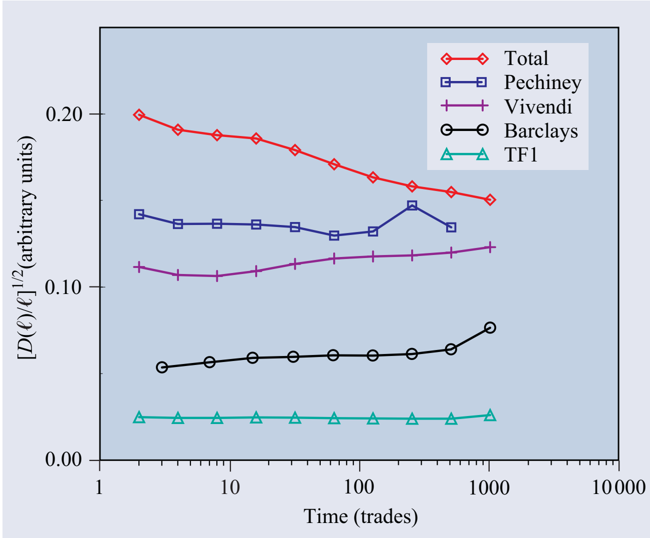

Figure 2. Plot of $\sqrt{\mathcal{D}(\ell)/\ell}$ as a function of ℓ for other stocks during the year 2002, except Barclays (May–June 2002). The y-axis has been rescaled arbitrarily for clarity. We note that stocks with larger tick size tend to reveal a stronger mean-reverting effect.

systematically discarded the first ten and the last ten minutes of trading in a given day, to remove any artifacts due to the opening and closing of the market. Many quantities of interest in the following are two-time observables, that is, compare two observables at (trade) time n and $n + \ell$. In order to avoid overnight effects, we have restricted our analysis mostly to intraday data, i.e. both n and $n + \ell$ belong to the same trading day. We have also assumed that our observables only depend on the time lag ℓ.

In the example of France-Telecom, on which we will focus mostly, there are on the order of 10 000 trades/day. For example, the total number of trades on France-Telecom during 2002 was close to 2×10^6; this allows quite accurate statistical estimates of various quantities. The volume of each trade was found to be roughly log-normally distributed, with $\langle \ln V \rangle \simeq 5.5$ and a root mean square of $\Delta \ln V \simeq 1.8$. The range of observed values of $\ln V$ is between 1 and 11.

2.2. Price fluctuation and diffusion

The simplest quantity to study is the average mean square fluctuation of the price between (trade) time n and $n+\ell$. Here, the price p_n is defined as the mid-point just before the nth trade: $p_n \equiv m_{n^-}$. In this paper, we always consider detrended prices, such that the empirical drift is zero. We thus define $\mathcal{D}(\ell)$ as

$$\mathcal{D}(\ell) = \langle (p_{n+\ell} - p_n)^2 \rangle. \tag{1}$$

As is well known, in the absence of any linear correlations between successive price changes, $\mathcal{D}(\ell)$ has a strictly diffusive behaviour, i.e.

$$\mathcal{D}(\ell) = D\ell, \tag{2}$$

where D is a constant. In the presence of short-range correlations, one expects deviations from this behaviour at short times. However, on liquid stocks with relatively small tick sizes such as France-Telecom (FT), one finds a remarkably linear behaviour for $\mathcal{D}(\ell)$, even for small ℓ. The absence of linear correlations in price changes is compatible with the idea that (statistical) arbitrage opportunities are absent, even for high-frequency trading. In fact, in order to emphasize the differences from a strictly diffusive behaviour, we have studied the quantity $\sqrt{\mathcal{D}(\ell)/\ell}$ (which has the dimension of Euros). We show this quantity in figure 1 for FT, averaged over three different periods: the first semester of 2001 (where the tick size was 0.05 Euros), the second semester of 2001 and the whole of 2002 (where the tick size was 0.01 Euros). One sees that $\mathcal{D}(\ell)/\ell$ is indeed nearly constant, with a small 'oscillation' on which we will comment later. Similar plots can be observed for other stocks (see figure 2). We have noted that for stocks with larger ticks a slow decrease of $\mathcal{D}(\ell)/\ell$ is observed, corresponding to a slight anti-persistence (or sub-diffusion) effect. Note that since we study the mid-point, the anti-correlations reported here are not related to the trivial bid–ask bounce.

The conclusion is that the random walk (diffusive) behaviour of stock prices appears even at the trade by trade level, with a diffusion constant D which is of the order of the typical bid–ask squared. From figure 1, one indeed sees that $\sqrt{\mathcal{D}(1)} \sim 0.01$ Euros, which is precisely the tick size, whereas FT has an average bid–ask spread equal to two ticks. Hence, each transaction typically moves the mid-point by half the bid–ask.

2.3. Response function and market impact

In order to better understand the impact of trading on price changes, one can study the following *response function* $\mathcal{R}(\ell)$, defined as

$$\mathcal{R}(\ell) = \langle (p_{n+\ell} - p_n)\varepsilon_n \rangle, \tag{3}$$

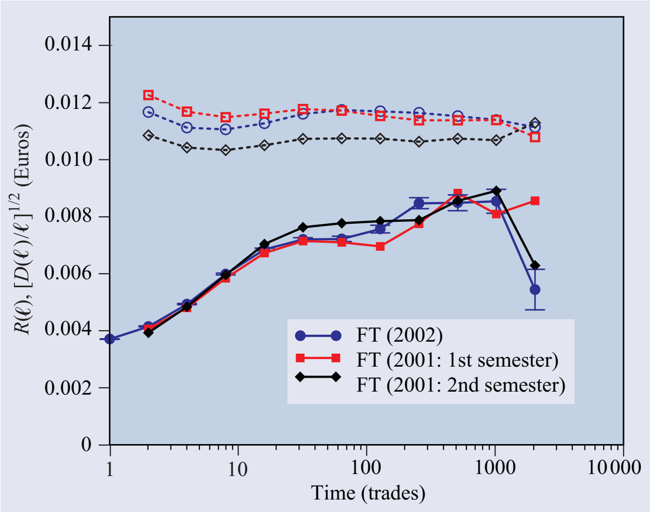

Figure 3. Average response function $\mathcal{R}(\ell)$ for FT, during three different periods (full symbols). We have given error bars for the 2002 data. For the 2001 data, the y-axis has been rescaled to best collapse onto the 2002 data. Using the same rescaling factor, we have also shown the data of figure 1. The fact that the same rescaling works approximately for $\mathcal{D}(\ell)$ as well will be considered further in section 2.5 below.

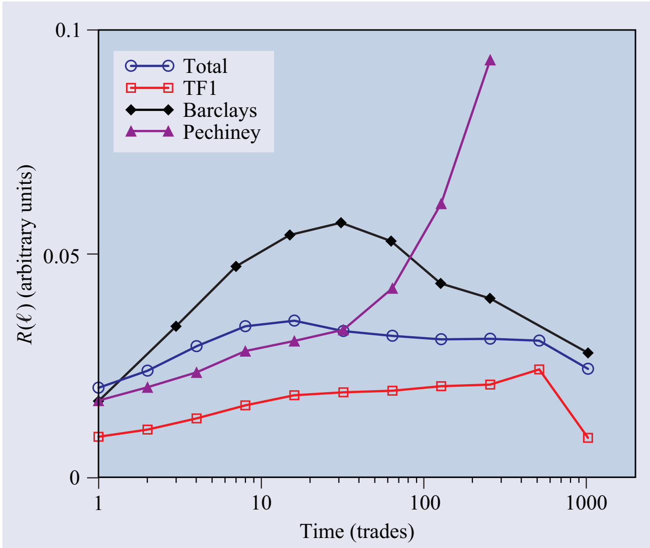

Figure 4. Average response function $\mathcal{R}(\ell)$ for a restricted selection of stocks, during the year 2002.

where ε_n is the sign of the nth trade, introduced in section 2.1. The quantity $\mathcal{R}(\ell)$ measures how much, on average, the price moves up conditioned to a buy order at time zero (or a sell order moves the price down) a time ℓ later. As will be clear below, this quantity is however not the market response to a single trade, a quantity that will later be denoted by G_0. A more detailed object can in fact be defined by conditioning the average to a certain volume V of the nth trade:

$$\mathcal{R}(\ell, V) = \langle (p_{n+\ell} - p_n)\varepsilon_n \rangle|_{V_n = V}. \tag{4}$$

Previous empirical studies have mostly focused on the volume dependence of $\mathcal{R}(\ell, V)$, and established that this function is strongly concave as a function of the volume [7, 27–30]. In [31], a thorough analysis of US stocks was performed in terms of a piecewise power-law dependence for $\mathcal{R}(\ell = 1, V) \propto V^{\alpha}$, with an exponent $\alpha \simeq 0.4$ for small volumes, and a smaller value ($\alpha \simeq 0.2$) for larger volumes (see also [32]). In a previous publication [33], two of us have proposed that this dependence might in fact be logarithmic (see also a footnote in [30]): $\mathcal{R}(\ell = 1, V) = R_1 \ln V$ (where R_1 is a stock dependent constant), a law that seems to satisfactorily account for all the data that we have analysed. The empirical determination of the temporal structure of $\mathcal{R}(\ell, V)$ has been much less investigated (although one can find in [30] somewhat related results on a coarse-grained version of $\mathcal{R}(\ell, V)$). Preliminary empirical results, published in [33], reported that $\mathcal{R}(\ell, V)$ could be written in a factorized form (first suggested on theoretical grounds in [12]):

$$\mathcal{R}(\ell, V) \approx \mathcal{R}(\ell) f(V); \qquad f(V) \propto \ln V, \tag{5}$$

where $\mathcal{R}(\ell)$ is a slowly varying function that initially increases up to $\ell \sim 100$–1000 and then is seen to decrease back, with a rather small overall range of variation. The initial increase of $R(\ell)$ was reported in [27] and has also recently been noticed by Lillo and Farmer [17]. Here, we provide much better data that supports both the above assertions. We show for example in figure 3 the temporal structure of $\mathcal{R}(\ell)$ for France-Telecom, for different periods. Note that $\mathcal{R}(\ell)$ increases by a factor $\sim$2 between $\ell = 1$ and $\ell^* \approx 1000$, before decreasing again. Including overnights allows one to probe larger values of ℓ and confirm that $\mathcal{R}(\ell)$ decreases, and even becomes *negative* beyond $\ell \simeq 5000$ (why this may be so will be explained by our model). Similar results have been obtained for many different stocks as well: figure 4 shows a small selection of other stocks, where the non-monotonic behaviour of $\mathcal{R}(\ell)$ is shown. However, in some cases (such as Pechiney), the maximum is not observed. One possible reason is that the number of daily trades is in this case much smaller ($\sim$1000), and that ℓ^* is beyond the maximum intraday time lag. On the other hand, the model discussed below does also allow for monotonic response functions.

The existence of a timescale ℓ^* beyond which $\mathcal{R}(\ell)$ decreases is thus both statistically significant, and to a large degree independent of the considered stock. On the other hand, the amplitude of the change of $\mathcal{R}(\ell)$ seems to be stock dependent. As will be clear later, the fact that $\mathcal{R}(\ell)$ slowly increases before decreasing back to negative values is a non-trivial result that requires a specific interpretation.

Turning now to the factorization property of $\mathcal{R}(\ell, V)$, equation (5), we illustrate its validity in figure 5, where $\mathcal{R}(\ell, V)/f(V)$ is plotted as a function of ℓ for different values of V. The function $f(V)$ was chosen for best visual rescaling, and is found to be close to $f(V) = \ln V$, as expected. Note that for the smallest volume (open circles) the long-time behaviour of $\mathcal{R}(\ell, V)$ seems to be different, which is probably due to the fact that small volumes are in fact more likely to be large volumes chopped up into small pieces.

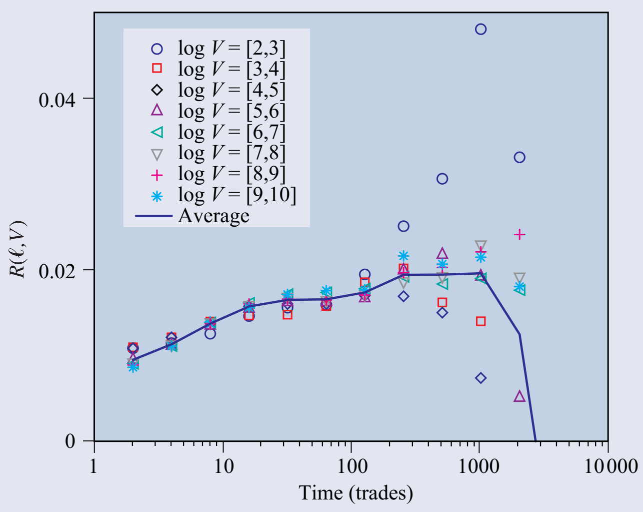

Figure 5. Average response function $\mathcal{R}(\ell, V)$, conditioned to a certain volume V, as a function of ℓ. Data for different V have been divided by $f(V) \propto \ln V$ so as to obtain good data collapse. The thick curve corresponds to $\mathcal{R}(\ell)$ (unscaled).

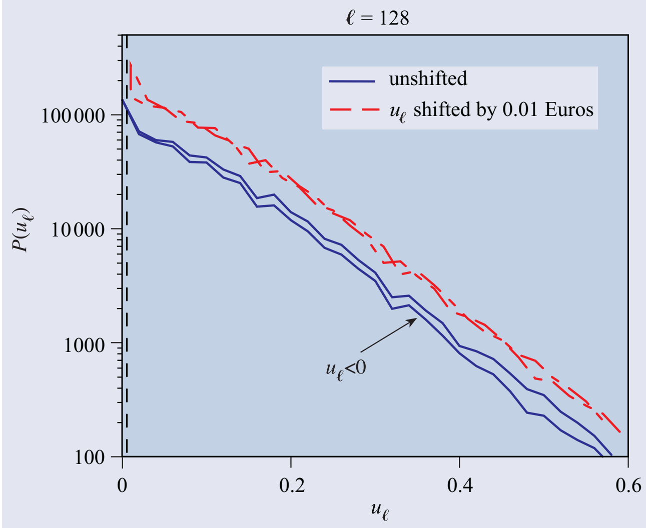

Figure 6. Probability distribution $P(u_\ell)$ of the quantity $u_\ell = (p_{n+\ell} - p_n)\varepsilon_n$ (in Euros), for $\ell = 128$. The data are again FT during 2002. The negative part of the distribution has been folded back to positive u_ℓ in order to highlight the small positive skew of the distribution (which is seen to increase slightly with $|u_\ell|$). The average value $\mathcal{R}(\ell) = \langle u_\ell \rangle$ is shown as the vertical dashed curve. The dashed–dotted curve corresponds to the distribution of $u_\ell - \nu$ with $\nu = 0.01$ Euros. This curve has been shifted upwards for clarity.

2.4. Fraction of informed trades

One has to keep in mind that the above response function $\mathcal{R}(\ell)$ captures a *small*, but systematic effect that relates the average price change to the sign of a trade. The fluctuations around this small signal are large, and increase with ℓ. A way to see this is to introduce the random variable $u_\ell = (p_{n+\ell} - p_n)\varepsilon_n$. By definition, $\mathcal{R}(\ell)$ is the average of u_ℓ, and $\mathcal{D}(\ell)$ is the average of u_ℓ^2. Since $\mathcal{R}(\ell)$ is roughly constant whereas $\mathcal{D}(\ell)$ grows linearly with ℓ, one sees that the impact of a given trade (as measured by $\mathcal{R}(\ell)$) rapidly becomes lost in the fluctuations.

In figure 6, we show the whole empirical distribution $P(u_\ell)$ of u_ℓ for $\ell = 128$ (other values of ℓ will be discussed below). This distribution is found to be only slightly skewed in the direction of positive u_ℓ. In fact, if one considers the shifted variable $u_\ell - \nu$, where $\nu = 0.01$ Euros, the distribution becomes nearly symmetric. Note that 0.01 Euros is equal to half the typical bid–ask spread and can therefore be seen as the minimal cost of a market order.

Empirical price changes are known to be highly kurtic, with sudden jumps separated by less volatile, random-walk-like periods. Market orders, that incur an immediate cost, are usually interpreted as 'informed' trades. A naive interpretation of an informed trade is a trade that correctly anticipates a price jump thanks to some private information. This foresight of future profits would then justify paying a liquidity cost for immediate execution. On the other hand, noise-induced trades, with no information content, should only temporarily impact the price, and be uncorrelated with the long-run value of the stock. Hence, in this picture, the positive tail of the distribution $P(u_\ell)$ (corresponding to informed trades) should be significantly fatter than the negative tail. Note that this conditional asymmetry is not in contradiction with the diffusive behaviour of price changes.

Now, the nearly symmetric shape of $P(u_\ell - \nu)$ in figure 6 shows that one can hardly detect the statistical presence of such informed trades, that correctly anticipate the sign of the price change on a short-term basis, so as to cover at least minimal trading costs. One can study this asymmetry more systematically as a function of the time horizon ℓ. Figure 7 shows a measure of the skewness ς of $P(u_\ell)$ (defined as $\varsigma = f - 1/2$, where f is the fraction of events with $u_\ell > 0$) as a function of ℓ.[5] We find that $\varsigma \approx \ell^{-1/2}$, as expected if the initial impact of the trade for $\ell \sim 1$ is followed by a pure random walk, uncorrelated with the sign of that trade. This behaviour extends to $\ell = 15\,000$ (two trading days), beyond which the skewness becomes indistinguishable from zero. Our finding that $\varsigma \approx \ell^{-1/2}$ is hard to reconcile with the idea that a substantial fraction of market orders represent informed trades, in which case ς should decay more slowly than $\ell^{-1/2}$. Correctly predicting a move L in the future indeed leads to a skewness first decaying as $\ell^{-1/2}$ before saturating, for $\ell \sim L$, to a finite value proportional to the fraction of informed trades. The putative information contained in these trades might of course reveal itself on longer timescales (weeks or months) for which we lack statistical accuracy, but we believe that this possibility is extremely unlikely, since the most reliable information should concern short time horizons. Our result is consistent with the conclusion of other studies, where it is established that investors 'trade too much' [34], and that the uninformed price pressure is large [7]. Note that $P(u_\ell)$ as defined above gives an equal weight to all trades, independently of their volume. We have also considered the volume-weighted $P(u_\ell)$, which leads to the same qualitative conclusion (see figure 7).

[5] We have used other definitions of the skewness, for example by restricting to only 'large' events $|u_\ell| > A\sqrt{\ell}$, with similar conclusions—see figure 7.

Is the above result in contradiction with the idea that the market is informationally efficient? Not necessarily, as explicitly demonstrated by Kyle's insider trading model [35]. In this model, an insider knows exactly the price of the stock at a certain date L in the future, and manages to place optimally winning trades against noise traders without creating any serial correlation in the price time series. The resulting market is informationally efficient at each instant of time, and the insider's information is progressively included in the price, with no jumps. The skewness ς then behaves in this model exactly as for a random walk, i.e. as $\ell^{-1/2}$. Correspondingly, the fraction f of informed trades behaves as $L^{-1/2}$ and is not observable if L is large. Typically, for information one day in the future, $L \approx 10^4$ and therefore $f \approx 1\%$, which could be consistent with our observations. However, there are several conceptual inconsistencies in Kyle's model, which makes it ineligible to describe a real stock market (some improvements have been considered in later work; see [36] for a review). For example, the insider assumed by Kyle is able to sell back all his stocks at time L and take his profit with no impact, which might be justified in the case of a takeover where the transaction price of the stock is agreed upon. But both in Kyle's model and in real markets, selling back will push the price down; within Kyle's model the actual profit of the insider for the round trip is actually exactly zero. In fact, Kyle's model should really be seen as a model of self-fulfilling prophecies, where one trader decides the future value of the stock and trades in such a way that the price indeed reaches this value at the final time. It is clear that within such a framework no profit can ever be made. Another strong assumption of Kyle is that the parameters (such as the number of noise traders or the root mean square of the price variation known to the insider) are known and constant in time. If this is not so, the insiders would be able to profit from anomalously liquid periods and/or anomalously strong information, which should then be revealed in an enhanced skewness ς. In summary, we find much more plausible the idea according to which prices are random because they evolve (at least on short timescales) through the mere effect of uninformed trading, rather than being informationally efficient thanks to a vanishingly small fraction of informed trades that delicately balance noise trades. This last scenario, suggested by Kyle's model, seems to us to be too fragile to apply to real markets. As we will argue below, random price changes emerge naturally from the competition between liquidity takers and liquidity providers, even if the former category has zero information.

2.5. A fluctuation–response relation

In the study of Brownian particles, a very important result that dates back to Einstein relates the diffusion coefficient D to the response of the particle to an external force. That a similar relation might also hold in financial markets was first suggested by Rosenow [37][6], and substantiated there by some empirical results. We have performed an analysis related to, but different from, that of Rosenow. For any given trading

[6] Somewhat related results, although not discussed in these terms, can also be found in [23].

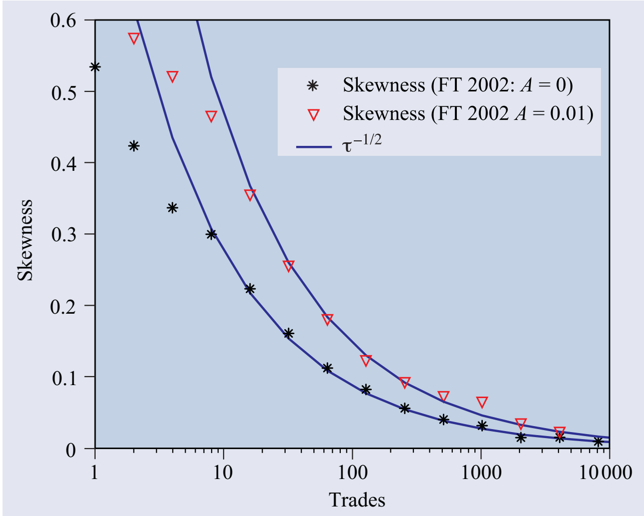

Figure 7. Skewness ς of the distribution $P(u_\ell)$ (now weighted by the volume of each trade) as a function of ℓ. Note that transaction costs are not subtracted here. Data across different days have been used in this analysis. We show for comparison the prediction of a model where the skewness is entirely due to the initial impact of the trade, followed by a pure random walk uncorrelated with the sign of the trade. In this case, $\varsigma \approx \ell^{-1/2}$, in very good agreement with the data. We also show the skewness restricted to price changes larger (in absolute value) than $0.01\sqrt{\ell}$ Euros, corresponding to 1σ moves.

day, one can compute the average local diffusion constant $\mathcal{D}(\ell)$ over a given timescale, say $\ell = 128$, and the average local price response $\mathcal{R}(\ell)$ over the same timescale. Rosenow, on the other hand, computes a 'susceptibility' as the slope of the average price change over a given time interval versus the volume imbalance during the same time interval (see [30]), and relates this susceptibility to the diffusion constant. The analogue of Rosenow's result [37] (which was motivated by a Langevin equation for price variations—see [38]) is a linear relation between $\mathcal{R}^2(\ell)$ and $\mathcal{D}(\ell)$, which we illustrate in figure 8 for FT, for two different periods (first semester of 2001, and 2002). A similar result can also be read from figure 3. As will be clear in the following, such a relation will appear naturally within the simple model that we introduce in section 3.

2.6. Long-term correlation of trade signs

All the above results are compatible with a 'zero-intelligence' picture of financial markets, where each trade is random in sign and shifts the price permanently, because all other participants update their evaluation of the stock price as a function of the last trade. As shown in [9, 10, 12–15], a model of the order book based on a purely random order flow indeed allows one to go quite far in the quantitative understanding of financial markets. In this context, the concave shape of the impact as a function of the volume can be understood as an order book effect, where the average size of the queue increases with depth[7].

This totally random model of the stock market is however qualitatively incorrect for the following reason. Although, as

[7] However, other effects are probably important to understand this concavity, such as the conditioning of large market orders to the size of the order book—see [7, 33].

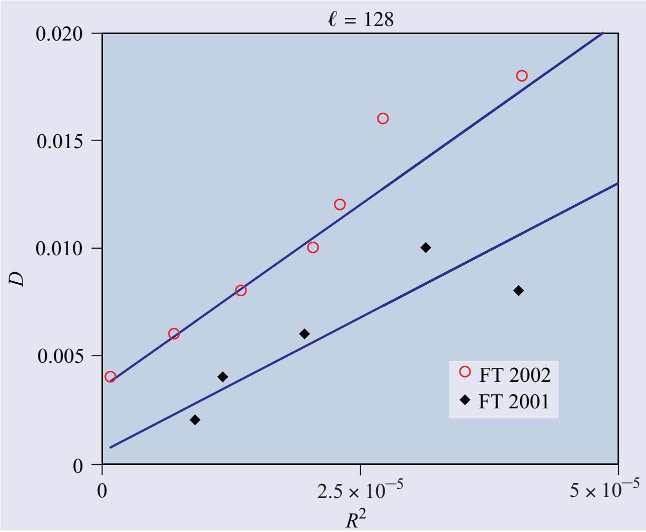

Figure 8. Average diffusion constant $D = \mathcal{D}(\ell)/\ell$, computed for $\ell = 128$, and conditioned to a certain value of $\mathcal{R}^2(\ell)$, also computed for $\ell = 128$ (FT). The open symbols correspond to 2002, whereas the black symbols are computed using the first semester of 2001, where the tick size was five times larger. Correspondingly, the x-axis was rescaled down by a factor of 25 and the y-axis by a factor of five for this data set.

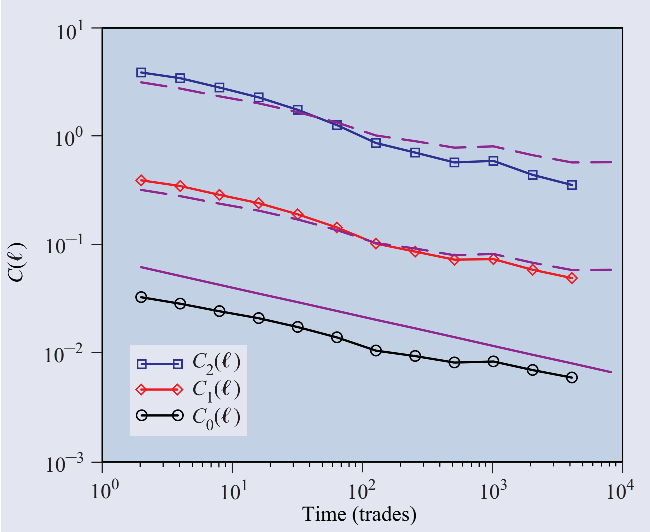

Figure 9. Volume-weighted sign autocorrelation functions as a function of time lag: $\mathcal{C}_0$, $\mathcal{C}_1$ and $\mathcal{C}_2$ (see the text for definitions). The straight line corresponds to $\ell^{-\gamma}$ with $\gamma = 1/5$. The dotted curves correspond to the simple approximation given by equations (11). Using data across different days allows one to extend the power-law decay at least up to 15 000 trades.

mentioned above, the statistics of price changes reveals very little temporal correlation, the correlation function of the sign ε_n of the trades, on the other hand, reveals very slowly decaying correlations as a function of trade time. This correlation has been mentioned in some papers before; see e.g. [7, 27]. Here, we propose that these correlations decay as a power law of the time lag, at least up to $\ell \approx 15\,000$ (two trading days) beyond which we do not have sufficiently accurate data.

More precisely, one can consider the following correlation function:

$$\mathcal{C}_0(\ell) = \langle \varepsilon_{n+\ell}\varepsilon_n \rangle - \langle \varepsilon_n \rangle^2. \tag{6}$$

If trades were random, one should observe that $\mathcal{C}_0(\ell)$ decays to zero beyond a few trades. Surprisingly, this is not what happens: on the contrary, $\mathcal{C}_0(\ell)$ is strong and decays very slowly toward zero, as an inverse power law of ℓ (see figure 9):

$$\mathcal{C}_0(\ell) \simeq \frac{C_0}{\ell^\gamma}, \qquad (\ell \geqslant 1). \tag{7}$$

The value of γ seems to be somewhat stock dependent. For example, for FT, one finds $\gamma \approx 1/5$, whereas for Total $\gamma \approx 2/3$. In their study, Lillo and Farmer found a somewhat larger value of $\gamma \approx 0.39$ for Vodafone [17]. In any case, the value of γ is found to be smaller than one, which is very important because the integral of $\mathcal{C}_0(\ell)$ is then *divergent*. This is in fact the precise definition of 'long-term' correlations. Now, as will be shown more precisely in the next section, the integral of $\mathcal{C}_0(\ell)$ can intuitively be thought of as the effective number N_e of correlated successive trades. Hence, out of—say—1000 trades, one should group together

$$N_e \simeq 1 + \sum_{\ell=1}^{1000} \mathcal{C}_0(\ell) \approx 1 + \frac{C_0}{1-\gamma} 1000^{1-\gamma} \tag{8}$$

'coherent' trades. For FT, $\gamma \approx 1/5$ and $C_0 \approx 0.2$, which means that the effect of one trade should be amplified, through the correlations, by a factor $N_e \approx 50$! In other words, both the response function $\mathcal{R}$ and the diffusion constant should increase by a factor of 50 between $\ell = 1$ and 1000, in stark contrast with the observed empirical data. This is the main puzzle that one should try to elucidate: how can one reconcile the strong, slowly decaying correlations in the sign of the trades with the nearly diffusive nature of the price fluctuations, and the non-monotonic response function?

Before presenting a mathematical transcription of the above question and proposing a possible resolution, let us comment on two related correlation functions that will naturally appear in the following, namely

$$\mathcal{C}_1(\ell) = \langle \varepsilon_{n+\ell}\varepsilon_n \ln V_n \rangle, \tag{9}$$

and

$$\mathcal{C}_2(\ell) = \langle \varepsilon_{n+\ell} \ln V_{n+\ell}\varepsilon_n \ln V_n \rangle. \tag{10}$$

We have found empirically that these two 'mixed' correlation functions are proportional to $\mathcal{C}_0(\ell)$ (see figure 8):

$$\mathcal{C}_1(\ell) \approx \langle \ln V \rangle \mathcal{C}_0(\ell); \qquad \mathcal{C}_2(\ell) \approx \langle \ln V \rangle^2 \mathcal{C}_0(\ell). \tag{11}$$

There are however small systematic deviations, which indicate that (i) small volumes contribute more to the long-range correlations than larger volumes and (ii) $\ln V - \langle \ln V \rangle$ is a quantity exhibiting long-range correlations as well.

3. A micro-model of price fluctuations

3.1. Set up of the model

In order to understand the above results, we will postulate the following trade superposition model, where the price at time

n is written as a sum over all past trades of the impact of one given trade propagated up to time[8] n:

$$p_n = \sum_{n'<n} G_0(n-n')\varepsilon_{n'} \ln V_{n'} + \sum_{n'<n} \eta_{n'} + \epsilon_n, \quad (12)$$

where $G_0(\cdot)$ is the 'bare' impact function (or propagator) of a single trade, that we assume to be a fixed, non-random function that only depends on time differences (which is a rather strong assumption, see [17]). The η_n are also random variables, assumed to be independent from the ε_n, and model all sources of price changes not described by the direct impact of the trades: the bid–ask can change as the result of some news, or of some order flow, in the absence of any trades. We will in the following assume that the η_n are also uncorrelated in time, although this assumption can easily be relaxed. Finally, the ϵ_n are independent (zero-mean) random variables that can be seen as a high-frequency noise on the position of the bid and of the ask, inducing some noise on the determination of the midpoint (see [16] for a more precise discussion of the origin of this term).

The bare impact function $G_0(\ell)$ represents by definition the average impact of a single trade after ℓ trades. It could in principle be measured empirically by launching on the market a sequence of real trades of totally random signs, and averaging the impact over this sample of trades (a potentially costly experiment!)[9]. As will be clear below, the difference between the quantity $\mathcal{R}(\ell)$ introduced in the previous section and $G_0(\ell)$ in fact comes from the strong autocorrelation of the sign of the trades. In order to understand the temporal structure of $G_0(\ell)$, note that a single trade first impacts the midpoint by changing the bid (or the ask). But then the subsequent limit order flow due to that particular trade might either centre on average around the new midpoint (in which case $G_0(\ell)$ would be constant), or, as we will argue below, tend to mean revert toward the previous midpoint (in which case $G_0(\ell)$ decays with ℓ). As discussed below (see equation (26)), the asymptotic behaviour of the bare impact function in fact reveals the average cost of a single market order: if $G_0(\ell \gg 1)/G_0(1)$ is small, the cost is large since the initial impact of the trade is only temporary, and is not followed by a true long-term change of the price.

Using this representation, the price increment between an arbitrarily chosen initial time 0 and time ℓ is

$$\begin{aligned} p_\ell - p_0 = &\sum_{0\leqslant n<\ell} G_0(\ell-n)\varepsilon_n \ln V_n \\ &+ \sum_{n<0}[G_0(\ell-n) - G_0(-n)]\varepsilon_n \ln V_n \\ &+ \sum_{0\leqslant n<\ell} \eta_n + \epsilon_\ell - \epsilon_0. \end{aligned} \quad (13)$$

If the signs ε_n were independent random variables, both the response function and the diffusion would be very easy to compute. For example, one would have[10]

$$\mathcal{R}_t(\ell) = \langle \ln V\rangle G_0(\ell), \quad (14)$$

i.e. the observed impact function and the bare response function would be proportional. Similarly, one would have

$$\begin{aligned} \mathcal{D}_t(\ell) = &\langle \ln^2 V\rangle \left(\sum_{0<n\leqslant\ell} G_0^2(n) + \sum_{n>0}[G_0(\ell+n) - G_0(n)]^2 \right) \\ &+ D_\eta \ell + 2D_0, \end{aligned} \quad (15)$$

where D_η is the variance of the η and D_0 is the variance of the ϵ. In the simplest case of a constant bare impact function, $G_0(\ell) = \Gamma_0$ for all $\ell > 0$, one then finds (for $D_0 = 0$) a pure diffusive behaviour, as expected:

$$\mathcal{D}_t(\ell) = \ell[\langle \ln^2 V\rangle \Gamma_0^2 + D_\eta]. \quad (16)$$

This case (no correlations between the ε and a constant bare impact function) corresponds to the simplest possible zero-intelligence market, where agents are memoryless. However, we have seen that in fact the ε have long range correlations. In this case, the average response function reads

$$\begin{aligned} \mathcal{R}_t(\ell) = &\langle \ln V\rangle G_0(\ell) + \sum_{0<n<\ell} G_0(\ell-n)\mathcal{C}_1(n) \\ &+ \sum_{n>0}[G_0(\ell+n) - G_0(n)]\mathcal{C}_1(n). \end{aligned} \quad (17)$$

Note in passing that our trade superposition model, equation (12), together with equation (11) leads to the factorization property mentioned above (see figure 5):

$$\mathcal{R}_t(\ell, V) = \frac{\ln V}{\langle \ln V\rangle}\mathcal{R}_t(\ell). \quad (18)$$

Now, one sees more formally the paradox discussed in the previous section: assuming that the impact of each trade is permanent, i.e. $G_0(\ell) = \Gamma_0$, leads to

$$\mathcal{R}_t(\ell) = \Gamma_0 \left[\langle \ln V\rangle + \sum_{0<n<\ell} \mathcal{C}_1(n) \right]. \quad (19)$$

If $\mathcal{C}_1(n)$ decays as a power law with an exponent $\gamma < 1$, then the average impact $\mathcal{R}(\ell)$ should *grow* like $\ell^{1-\gamma}$, and therefore be amplified by a very large factor as ℓ increases, at variance with empirical data. The only way out of this conundrum is (within the proposed model) that the bare impact function $G_0(\ell)$ itself should decay with time, in such a way as to offset the amplification effect due to the trade correlations.

3.2. A relation between the bare propagator and the sign correlation function

In order to get some guidance, let us now look at the general formula for the diffusion. After a few lines of calculations, one finds

$$\begin{aligned} \mathcal{D}_t(\ell) = &\langle \ln^2 V\rangle \\ &\times \left[\sum_{0\leqslant n<\ell} G_0^2(\ell-n) + \sum_{n>0}[G_0(\ell+n) - G_0(n)]^2 \right] \\ &+ 2\Delta(\ell) + D_\eta \ell + 2D_0, \end{aligned} \quad (20)$$

[8] A similar model, but linear in V_n and truncated beyond only a few trades in the past, was discussed in [27].

[9] However, such an experimental protocol might induce 'copy-cat' trades and still lead to a difference between the measured response function and G_0.

[10] In the following, we will use the subscript 't' to denote the theoretical expressions for the response function or diffusion.

where $\Delta(\ell)$ is the correlation-induced contribution:

$$\begin{aligned}\Delta(\ell) = &\sum_{0\leqslant n<n'<\ell} G_0(\ell-n)G_0(\ell-n')\mathcal{C}_2(n'-n)\\ &+\sum_{0<n<n'}[G_0(\ell+n)-G_0(n)][G_0(\ell+n')\\ &-G_0(n')]\mathcal{C}_2(n'-n)+\sum_{0\leqslant n<\ell}\sum_{n'>0}G_0(\ell-n)\\ &\times[G_0(\ell+n')-G_0(n')]\mathcal{C}_2(n'+n).\end{aligned} \tag{21}$$

The constraint from empirical data is that this expression must be approximately linear in ℓ. As shown in the appendix, the requirement that $\mathcal{D}_t(\ell)$ is strictly linear in ℓ for all ℓ in fact allows one to express $G_0(\ell)$ as a function of $\mathcal{C}_2(\ell)$. Here, we present a simple asymptotic argument. If we make the ansatz that the bare impact function $G_0(\ell)$ also decays as a power law,

$$G_0(\ell)=\frac{\Gamma_0\ell_0^\beta}{(\ell_0+\ell)^\beta}\qquad(\ell\geqslant 1) \tag{22}$$

then one can estimate $\mathcal{D}_t(\ell)$ in the large-ℓ limit. When $\gamma<1$, one again finds that the correlation-induced term $\Delta(\ell)$ is dominant, and all three terms scale as $\ell^{2-2\beta-\gamma}$, provided $\beta<1$. In other words, the Hurst exponent of price changes is given by $2H=2-2\beta-\gamma$. Therefore, the condition that the fluctuations are diffusive at long times ($H=1/2$) imposes a relation between the decay of the sign autocorrelation γ and the decay of the bare impact function β that reads

$$2\beta+\gamma=1\longrightarrow\beta_c=\frac{1-\gamma}{2}. \tag{23}$$

For $\beta>\beta_c$, the price is *sub-diffusive* ($H<1/2$), which means that price changes show anti-persistence; while for $\beta<\beta_c$, the price is *super-diffusive* ($H>1/2$), i.e. price changes are persistent. For FT, $\gamma\approx 1/5$ and therefore $\beta_c\approx 2/5$.

As shown in the appendix, one can in fact obtain an exact relation between $G_0(\ell)$ and $\mathcal{C}_2(\ell)$ if one assumes that price changes are strictly uncorrelated (i.e. that $\mathcal{D}(\ell)$ is linear in ℓ for all ℓ). The asymptotic analysis of this relation leads, not surprisingly, to the same exponent relation $\beta_c=(1-\gamma)/2$ as above.

At this stage, there seems still to be a contradiction with empirical data, for if one goes back to the response function given by equation (17), one finds that whenever $\beta+\gamma<1$ (which is indeed the case for $\beta=\beta_c$ and $\gamma<1$), the dominant contribution to $\mathcal{R}_t(\ell)$ should behave as $\ell^{1-\beta-\gamma}$ and thus *grow* with ℓ. For example, for $\gamma\approx 1/5$ and $\beta\approx 2/5$, one should find that $\mathcal{R}_t(\ell)\propto\ell^{2/5}$, which is incompatible with the empirical data of figures 3 and 4. But the surprise comes from the numerical prefactor of this power law. One finds, for large ℓ,

$$\begin{aligned}\mathcal{R}_t(\ell)\simeq\langle\ln V\rangle\Gamma_0C_0&\frac{\Gamma(1-\gamma)}{\Gamma(\beta)\Gamma(2-\beta-\gamma)}\\ &\times\left[\frac{\pi}{\sin\pi\beta}-\frac{\pi}{\sin\pi(1-\beta-\gamma)}\right]\ell^{1-\beta-\gamma}.\end{aligned} \tag{24}$$

Therefore, only when $\beta=\beta_c$ is the prefactor exactly zero, and leads to the possibility of a nearly constant impact function! For faster decaying impact functions (larger β), this prefactor is negative, whereas for more slowly decaying impact functions this prefactor is positive[11]. Interestingly, even if the bare response function $G_0(\ell)$ is positive for all ℓ, the average response $\mathcal{R}_t(\ell)$ can become negative for large enough β, as a consequence of the correlations between trades.

3.3. Fitting the average response function

Since the dominant term is zero for the 'critical' case $\beta=\beta_c$, and since we are interested in the whole function $\mathcal{R}_t(\ell)$ (including the small-ℓ regime), we have computed $\mathcal{R}_t(\ell)$ numerically, by performing the discrete sum equation (17) exactly, and fitted it to the empirical response $\mathcal{R}$. The results are shown in figure 10. We have fixed the parameters γ and C_0 to the values extracted from the behaviour of $\mathcal{C}_1(\ell)$ (see figure 8): $\gamma=0.24$ and $C_0=0.20$. The overall scaling parameter Γ_0 is adjusted to match the value of $\mathcal{R}(\ell=1)$.[12] The values of β and ℓ_0 are fitting parameters: we show in figure 10 the response function computed for different values of β in the vicinity of $\beta_c=0.38$, using $\ell_0=20$.

The results are compared with the empirical data for FT, showing that one can indeed satisfactorily reproduce, when $\beta\approx\beta_c$, a weakly increasing impact function that reaches a maximum and then decays. One also sees, from figure 10, that the relation between β and γ must be quite accurately satisfied, otherwise the response function shows a distinct upward trend (for $\beta<\beta_c$) or a downward trend ($\beta>\beta_c$).[13] We have tried other simple forms for $G_0(\ell)$, such as a simple exponential decay toward a possibly non-zero asymptotic value, but this leads to unacceptable shapes for $\mathcal{R}(\ell)$. Of course, a more precise fit of the initial increase of $\mathcal{R}(\ell)$ seen in figure 10 could be achieved by choosing a more complicated function $G_0(\ell)$, that first increases with ℓ before decaying to zero.

It is also interesting to use the propagator G_0^* determined in the appendix from the assumption of a purely diffusive price process for all ℓ. This propagator is plotted in figure 11, and compared to the G_0 determined above from the fit of $\mathcal{R}(\ell)$. As shown in figure 10, the use of G_0^* does *not* lead to a very good fit of $\mathcal{R}(\ell)$. Since the latter quantity is in fact *very sensitive to the chosen shape* for G_0, it does reveal small, but systematic, deviations from a purely diffusive price process. (Note that if one had $\mathcal{C}_2(\ell)=\mathcal{C}_1(\ell)$, the resulting $\mathcal{R}(\ell)$ should be strictly constant.)

3.4. Back to the diffusion constant

As we showed above, the reason for the fine tuning of β is the requirement that price changes are almost diffusive. We can therefore also compute $\mathcal{D}_t(\ell)$ for all values of ℓ using the very same values of γ, β, C_0, ℓ_0 and Γ_0. Now, in order to fit the data one has two extra free parameters: D_η and D_0 (see equation (20)). With these two extra parameters, one can reproduce the empirical determination of $\mathcal{D}(\ell)/\ell$

[11] Note that although this prefactor increases (in absolute value) with β for $\beta>\beta_c$, the power of ℓ decreases, which means that for large ℓ the amplitude of $\mathcal{R}_t(\ell)$ decreases with β, as intuitively expected.

[12] The numerical value of Γ_0 is found to be such that: $\Gamma_0\ell_0^\beta=2.8\times10^{-3}$ Euros.

[13] The former scenario might actually explain the different behaviour of Pechiney seen in figure 4.

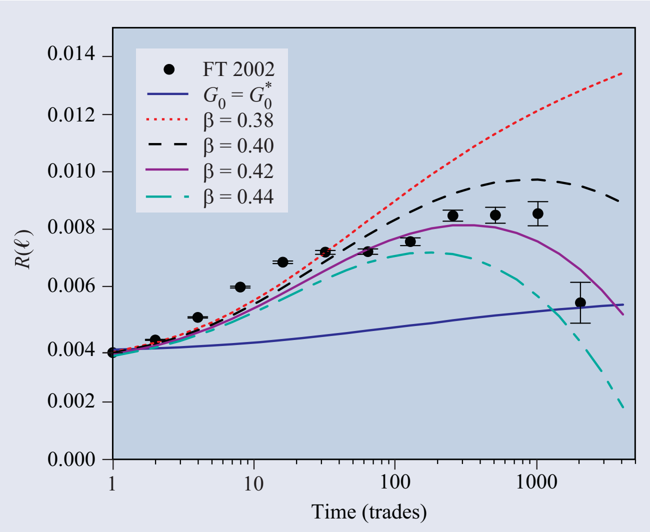

Figure 10. Theoretical impact function $\mathcal{R}_t(\ell)$, from equation (17), and for different values of β close to $\beta_c = 0.38$. The shape of the empirical response function can be quite accurately reproduced using $\beta = 0.42$. The only remaining free parameter is $\ell_0 = 20$. The thick plain curve is $\mathcal{R}_t(\ell)$ computed using the 'pure diffusion' propagator G_0^* determined in the appendix, equation (34). Note that for $\beta > \beta_c$, the response function actually becomes *negative* at long times, as indeed observed empirically for $\ell > 5000$.

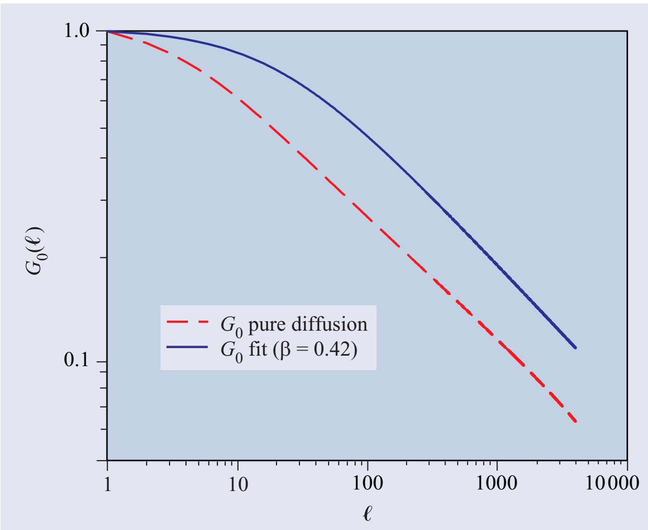

Figure 11. Shape of the bare propagator G_0, determined either by the fit of $\mathcal{R}$, with $\beta = 0.42$ and $\ell_0 = 20$, or using the exact relation, equation (34), derived in the appendix from the assumption of a purely diffusive process.

(see figure 12). The small deviations of this quantity from a horizontal line at finite ℓ are due to the difference between G_0 and G_0^* and/or to the possible autocorrelations between the η_n variables, which we have neglected here. Note that the contribution of the term D_η turns out to be a factor of two larger than that of the impact contribution, equation (20), which means that the small increase of the 'impact contribution' with ℓ (lower graph of figure 11) is hardly detectable in $\mathcal{D}(\ell)/\ell$.

Coming back to the fluctuation–response relation discussed in section 2.5, we see that our model predicts, for

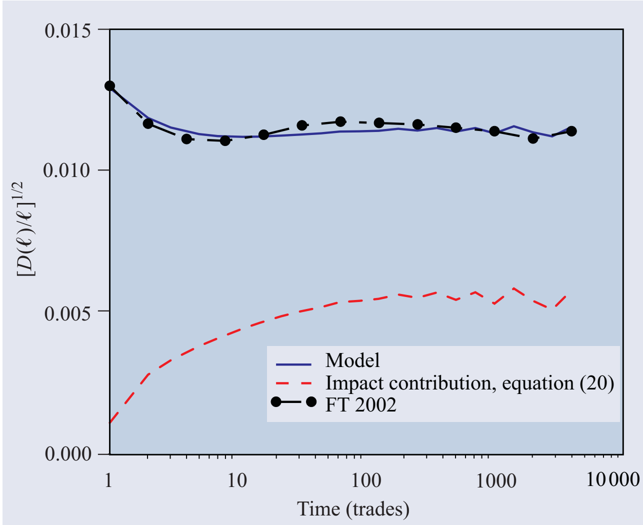

Figure 12. Diffusion constant $\mathcal{D}(\ell)/\ell$, using equation (20), with the values of γ, β, C_0, ℓ_0 and Γ_0 determined from $\mathcal{R}(\ell)$. Two extra parameters were used: $D_\eta = 10^{-4}$ and $D_0 = 3.3 \times 10^{-5}$ (both in Euro squared). The lower graph is the 'impact contribution' to $\mathcal{D}_t(\ell)$, given by equation (20) with $D_\eta = 0$. The 'oscillations' at long times are a numerical artefact.

$\ell \gg 1$ where the effect of D_0 can be neglected,

$$\frac{\mathcal{D}_t(\ell)}{\ell} = Z\langle \ln V\rangle^2 C_0 \Gamma_0^2 + D_\eta, \qquad \mathcal{R}_t(\ell) = Z'\langle \ln V\rangle \Gamma_0 C_0, \tag{25}$$

where Z and Z' are numerical constants. Assuming that from one day to the next both the average (log-)traded volume and the impact Γ_0 of each individual trade might change, while C_0 is fixed, immediately leads to the affine relation between $\mathcal{D}$ and $\mathcal{R}^2$ reported in section 2.5.

3.5. Discussion

The conclusion of this section is that our 'micro-model' of prices, equation (12), can be used as a theoretical canvas to rationalize and interpret the empirical results found in the previous section. Most surprising is the constraint that the empirical results impose on the shape of the 'bare' response function G_0, which is found to be a slowly decaying power law which must precisely cancel the slowly decaying autocorrelation of the trades, but reveals systematic deviations from a pure diffusion process, hardly noticeable in the diffusion constant itself. The fact that the bare impact function decays with time (at least on intraday timescales) in a finely tuned way to compensate the long memory in the trades is the central result of this paper. This effect is lost in the zero-intelligence models of Poissonian order flows, where, after decreasing during a short transient, the impact of each trade becomes permanent: $G_0(\ell) \to G_\infty > 0$. On this point, see the model studied in [14], where it is shown that prices are confined (sub-diffusive) on timescales shorter than the lifetime of limit orders, essentially as a consequence of the shape of the order book. In fact, both the long-time memory of the trades and the slowly relaxing impact function reported here must be the consequence of the strategic behaviour of market participants, that we discuss below in order to get an intuitive understanding of the mechanisms at play.

Although our detailed analysis concerns FT, it is clear that our conclusions are more general, since the strong autocorrelations in the trade signs, the near constancy of the average response function and the diffusive nature of price changes have been observed on all stocks, with only quantitative changes (see figures 2 and 4). It would be interesting to document these quantitative differences, and relate these to liquidity, or to the size of the bid–ask spread.

Finally, it would be very interesting to know whether the bare response function levels off to a finite value for large time lags; this will require more data to go beyond the analysis of the present paper to enlarge the available range of ℓ values. However, it seems reasonable to expect that $G_0(\ell)$ should indeed reach a finite asymptotic value for values of ℓ corresponding to a few days of trading[14].

4. Critical balance of opposite forces: market orders versus limit orders

Although trading occurs for a large variety of reasons, it is useful to recognize that traders organize in two broad categories.

- One is that of 'liquidity takers', that trigger trades by putting in market orders. The motivation for this category of traders might be to take advantage of some 'information', and make a profit from correctly anticipating future price changes. Information can in fact be of very different nature: fundamental (firm based), macro-economical, political, statistical (based on regularities of price patterns), etc. Unfortunately, information is often hard to interpret correctly—except of course for insiders—and it is probable that many of these 'information'-driven trades are misguided (on this point, see [7, 34] and references therein, and the discussion in section 2.4). For example, systematic hedge funds which take decisions based on statistical pattern recognition have a typical success rate of only 52%. There is no compelling reason to believe that the intuition of traders in market rooms fares much better than that. Since market orders allow immediate execution, many impatient investors, who want to liquidate their position or hedge etc, might be tempted to place market orders, even at the expense of the bid–ask spread $s(t) = a(t) - b(t)$.
- The other category is that of 'liquidity providers' (or 'market makers', although on electronic markets all participants can act as liquidity providers by putting in limit orders), who offer to buy or to sell but avoid taking any bare position on the market. Their profit comes from the bid–ask spread s: the sell price is always slightly larger than the buy price, so that each round turn operation leads to a profit equal to the spread s, at least if the midpoint has not changed in the mean time (see below).

This is where the game becomes interesting. Assume that a liquidity taker wants to buy, so that an increased number of buy orders arrive on the market. The liquidity provider is tempted to increase the offer (or ask) price a because the buyer might be informed and really know that the current price is too low and that it will most probably increase in the near future. Should this happen, the liquidity provider, who has to close his position later, might have to buy back at a much higher price and experience a loss. On the other hand, liquidity takers obviously need to place small orders, since large orders would trigger a sudden increase of a and make their trade costly. This is the rationale for dividing one's order into small chunks and dispersing these as much as possible over time so as not to appear on the 'radar screens'. Doing so liquidity takers *necessarily create some temporal correlations* in the sign of the trades. Since these traders probably have a somewhat broad spectrum of volumes to trade [39][15], and therefore of trading horizons (from a few minutes to several weeks), this can easily explain the slow, power-law decay of the sign correlation function $\mathcal{C}_0(\ell)$ reported above.

Now, if the market orders in fact do *not* contain useful information but are the result of hedging, noise trading, misguided interpretations, errors etc, then the price should not move up in the long run, and should eventually mean revert to its previous value. Liquidity providers are obviously the active force behind this mean reversion, again because closing their position will be costly if the price has moved up too far from the initial price. More precisely, a computation of the liquidity provider average gain per share $\mathcal{G}$ can be performed [16], and is found to be, for trades of volume V,

$$\mathcal{G} = s + \mathcal{R}(0, V) - \mathcal{R}(\infty, V) \approx s + \ln V[\mathcal{R}(0) - \mathcal{R}(\infty)], \quad (26)$$

where $\mathcal{R}(0, V)$ is the immediate average impact of a trade, before new limit orders set in. We have in fact checked empirically that $\mathcal{R}(0, V) \approx \mathcal{R}(1, V)$. From the above formula, one sees that *it is in the interest of liquidity providers to mean revert the price*, so as to make $\mathcal{R}(\infty)$ as small as possible. However, this mean reversion cannot take place too quickly, again because a really informed trader would then be able to buy a large volume at a modest price. Hence, this mean reversion must be slow. From the quantitative analysis of section 3, we have found that there is hardly any mean reversion at all on short timescales $\ell < \ell_0$ (perhaps even a small trend following effect, see figure 10), and that this mean reversion must be described as a slow power law for larger ℓ. Actually, the action of liquidity providers and liquidity takers must be such that no (or very little) linear correlation is created in the price changes, otherwise statistical arbitrage opportunities would be created at the detriment of one or the other population.

To summarize: liquidity takers must dilute their orders and create long-range correlations in the trade signs, whereas liquidity providers must correctly handle the fact that liquidity takers might either possess useful information (a rare situation, but that can be very costly since the price can jump as a result

[14] Hopman quotes three days as the time beyond which the autocorrelation of the trades sign falls to zero [7], whereas we find that the power-law decay of this correlation persists up to at least two days of trading.

[15] We associate the power-law distribution of investors' size with the long-range correlations reported here rather than with the power-law tail of return, as advocated in the above cited paper.

of some significant news), or might not be informed at all and trade randomly. By controlling the order flow so as to slowly mean revert the price, liquidity providers minimize the probability that they either sell too low, or have to buy back too high. The delicate balance between these conflicting tendencies conspires to put the market at the border between persistence (if mean reversion is too weak, i.e. $\beta < \beta_c$) and anti-persistence (if mean reversion is too strong, i.e. $\beta > \beta_c$), and eliminate arbitrage opportunities. Therefore, the mere fact of trading so as to minimize impact for liquidity takers, and to optimize gains for liquidity providers, does lead to a random walk dynamics of the price, even in the absence of any real information underlying the trades.

It is actually enlightening to propose a simple model that could explain how liquidity providers enforce mean reversion[16]. Assume that upon placing limit orders, liquidity providers induce a systematic bias toward some moving average of past prices. If this average is for simplicity taken to be an exponential moving average, the continuous time description of this will read

$$\begin{aligned} \frac{\mathrm{d}p_t}{\mathrm{d}t} &= -\Omega(p_t - \bar{p}_t) + \xi_t \\ \frac{\mathrm{d}\bar{p}_t}{\mathrm{d}t} &= \kappa(p_t - \bar{p}_t), \end{aligned} \tag{27}$$

where ξ_t is the random driving force due to trading, Ω the inverse timescale for the strength of the mean reversion and $1/\kappa$ the 'memory' time over which the average price $\bar{p}_t$ is computed. The first equation means that liquidity providers tend to mean revert the price toward $\bar{p}_t$, while the second describes the update of the exponential moving average $\bar{p}_t$ with time. This set of linear equations can be solved, and leads to a solution of the form $p_t = \int^t \mathrm{d}t'\, G_0(t-t')\xi'_t$, with a bare propagator given by

$$G_0(t) = (1 - G_\infty)\exp[-(\Omega+\kappa)t] + G_\infty, \tag{28}$$

i.e. an exponential decay toward a finite asymptotic value $G_\infty = \kappa/(\Omega+\kappa)$. Note that, interestingly, it is the self-referential effect that leads to a non-zero asymptotic impact. If the fundamental price is known to all, $\kappa = 0$ and $G_\infty = 0$. In the opposite limit where $\kappa \gg \Omega$, the last price is taken as the reference price, and $G_\infty \to 1$. A way to obtain $G_0(t)$ resembling a power law is to assume that different liquidity providers use different time horizons to compute a reasonable reference price. This leads to a $G_0(t)$ which is written as the sum of time exponentials with different rates which can easily mimic a pure power law.

The message of the above model is actually quite interesting from the point of view of efficient markets: it suggests that nobody really knows what the correct reference price should be, and that its best proxy is in fact its own past average over some time window (the length of which is itself distributed over several timescales among market participants).

[16] We have in fact directly checked on the data that the evolution of the midpoint between trades (resulting from the order flow) is indeed anticorrelated with the impact of the trades. On this point, see [7], where the limit order flow subsequent to a trade is studied. See however the very recent preprint of Lillo and Farmer [17], where they show that the dominant cause of the mean reversion is not due to a change of bid–ask, but to a bias on the bid and ask volumes in such a way to favour a bounce-back of the midpoint.

5. Summary and conclusion

The aim of this paper was to study in detail the statistics of price changes at the trade by trade level, and to analyse the interplay between the impact of each trade on the price and the volatility. Empirical data show that

(a) the price (midpoint) process is close to being purely diffusive, even at the trade by trade scale,
(b) the temporal structure of the impact function first increases and reaches a maximum after 100–1000 trades, before decreasing again, with a rather limited overall variation (typically a factor of two) and
(c) the sign of the trades shows surprisingly long-range (power-law) correlations, at least up to 15 000 trades (two trading days).

The paradox is that if the impact of each trade were permanent, the price process should be strongly super-diffusive and the average response function should increase by a large factor as a function of the time-lag.

As a possible resolution of this paradox, we have proposed a micro-model of prices, equation (12), where the price at any instant is the causal result of all past trades, mediated by what we called a bare impact function, or propagator G_0. All the empirical results can be reconciled if one assumes that this bare propagator also decays as a power law in time, with an exponent which is precisely tuned to a critical value, ensuring simultaneously that prices are diffusive on long timescales and that the response function is nearly constant. Therefore, the seemingly trivial random walk behaviour of price changes in fact results from a fine-tuned competition between two opposite effects, one leading to super-diffusion (the autocorrelation of trades) and the other leading to sub-diffusion (the decay of the bare impact function). The cancellation is however not exact: the non-trivial behaviour of the average response function allows one to detect small, but systematic deviations from a purely diffusive behaviour, deviations that are hardly detectable on the price fluctuations themselves.

In financial terms, the competition is between liquidity takers, that create long-range correlations by dividing their trading volume into small quantities, and liquidity providers, that adapt the order flow so as to mean revert the price and to optimize their gains (see equation (26)). The resulting absence of linear correlations in price changes, and therefore of arbitrage opportunities, is often postulated *a priori* in the economics literature, but the details of the mechanism that removes these arbitrage opportunities are rather obscure. The main message of this paper is that the random walk nature of price changes on short (intraday) timescales is not due to the unpredictable nature of incoming news, but appears as a dynamical consequence of the competition between the above antagonistic constraints on market participants.

The role of *real* (and correctly interpreted) information appears to be rather thin: the fact that the intraday volatility of a stock is nearly equal to its long-time value suggests that the volatility is mostly due to the trading activity itself, which is dominated by noise trades. This result is most probably one of the mechanisms needed to explain the excess volatility puzzle first raised by Shiller [3]. We have in fact defined a model-independent indicator of the fraction of 'informed' trades as the asymmetry (or skewness) of the probability distribution of the *signed* price variation, where the sign is that of the trade at the initial time. If some market orders correctly anticipate large events, this should generate a detectable positive skew of this distribution. Due to these trades, the skewness of the distribution should decay more slowly than $1/\sqrt{\ell}$, which is the benchmark result for a simple random walk impacted by the initial trade. However, consistently with other studies [7, 34], our empirical results only show very weak asymmetry, barely sufficient to cover trading costs. The skewness is found to decay as $1/\sqrt{\ell}$. This suggests that only a very small fraction of these market orders can *a posteriori* be described as truly informed, whereas most trades can be classified as noise. The remarkable long-range persistence of the sign of these trades is in this respect even more striking, since one could have expected that 'noise' trading would be uncorrelated in time.

The conclusion that price changes are to a large extent induced by the trading activity itself seems to imply that the price will, in the long run, wander arbitrarily far from the fundamental price, which would be absurd. But even if one assumes that the fundamental price is independent of time, a typical 3% noise-induced daily volatility would lead to a significant (say a factor of two) difference between the traded price and the fundamental price only after a few years [40]. Since the fundamental price of a company is probably difficult to determine better than within a factor to two, say (see e.g. [5, 41]), one only expects fundamental effects to limit the volatility on very long timescales (as indeed suggested by the empirical results of de Bondt and Thaler [42]), but that these are probably negligible on the short (intraday) timescales of interest in this paper. We have in fact seen (cf equation (27)) that the reference price that market participants seem to have in mind is in fact a short-time average of the past price itself, rather than any fundamental price.

From a more general standpoint, our finding that the absence of arbitrage opportunities results from a critical balance between antagonistic effects is quite interesting. It might justify several claims made in the (econo-)physics literature that the anomalies in price statistics (fat tails in returns described by power laws [21, 22], long-range self-similar volatility correlations [25, 26] and the long-range correlations in signs reported here and in [17]) are due to the presence of a critical point in the vicinity of which the market operates (see e.g. [43], and in the context of financial markets [44, 45]). If a fine-tuned balance between two competing effects is needed to ensure absence of arbitrage opportunities, one should expect that fluctuations are crucial, since a local imbalance between the competing forces can lead to an instability. In this respect, the analogy with the balancing of a long stick is quite enticing [20]. In more financial terms, the breakdown of the conditions for this dynamical equilibrium is, for example, a *liquidity crisis*: a sudden cooperativity of market orders, that lead to an increase of the trade sign correlation function, can out-weigh the liquidity providers' stabilizing (mean-reverting) role, and lead to crashes. This suggests that one should be able to write a mathematical model, inspired by our results, to describe this 'on–off intermittency' scenario, advocated (although in a different context) in [20, 46, 47].

Acknowledgments

YG thanks Science&Finance/Capital Fund Management for hospitality while this work was completed. We thank Jelle Boersma, Lisa Borland and Bernd Rosenow for inspiring remarks, and J Doyne Farmer for communicating his results on the autocorrelation of trades before publication [17] and for insightful discussions and comments. Xavier Gabaix has made some crucial remarks on the first version of the manuscript, that in particular led to the material contained in the appendix, and has drawn our attention to several helpful references, such as [7, 35]. The final version of this paper owes a lot to him. We also thank S Picozzi for an interesting discussion and for pointing out the possible relevance of [20] to financial markets. Finally, we take the opportunity of this paper to acknowledge the countless efforts of Gene Stanley to investigate the dynamics of complex systems and to bring together different fields and ideas—as testified by the papers cited in the references, which inspired the present work.

Appendix. The case of a strictly diffusive process

This appendix was inspired by a remark of Xavier Gabaix. There is one particular case of our micro-model of prices, equation (12), where prices are purely diffusive *at all times* (rather than only asymptotically). This is the case provided a specific relation between the bare propagator G_0 and the sign correlation function $\mathcal{C}_2(\ell)$ holds. In order to show this, let us assume that the random variable $q_n \equiv \varepsilon_n \ln V_n$ can be written as

$$q_n = \sum_{m \leqslant n} K(n-m)\xi_m, \tag{29}$$

where ξ_n are uncorrelated random variables ($\langle \xi_n \xi_m \rangle = \langle \ln^2 V \rangle \delta_{n,m}$), and $K(\cdot)$ a certain kernel. In order for the q_n to have the required correlations, the kernel $K(\cdot)$ should obey the following equation:

$$\mathcal{C}_2(n) = \langle \ln^2 V \rangle \sum_{m \geqslant 0} K(m+n)K(m). \tag{30}$$

In the case where $\mathcal{C}_2$ decays as $\ell^{-\gamma}$ with $0 < \gamma < 1$, it is easy to show that the asymptotic decay of $K(n)$ should also be a power law $n^{-\delta}$ with $2\delta - 1 = \gamma$. Note that $1/2 < \delta < 1$.

Inverting equation (29) allows one to obtain a set of uncorrelated random variables ξ_n from a set of correlated variables q_n:

$$\xi_n = \sum_{m \leqslant n} Q(n-m)q_m, \tag{31}$$

where Q is the matrix inverse of K, such that $\sum_{m=0}^{n} K(n-m)Q(m) = \delta_{m,n}$. Equations (29) and (31) in fact form the basis of linear filter theories, and ξ_n can be seen as the prediction error on the next variable q_n.

Introducing discrete Laplace transforms,

$$\hat{K}(E) = \sum_{n\geqslant 0} K(n)\mathrm{e}^{-nE} \qquad \hat{Q}(E) = \sum_{n\geqslant 0} Q(n)\mathrm{e}^{-nE}, \quad (32)$$

one finds $\hat{K}(E)\hat{Q}(E) = 1$. For a power-law kernel $K(\cdot)$, one obtains $\hat{Q}(E) \propto E^{1-\delta}$ for $E \to 0$, and therefore $Q(n) \propto n^{\delta-2}$ for large n. It is useful to note that in this case $\hat{Q}(E = 0) = \sum_{n\geqslant 0} Q(n) = 0$.

Now, it is clear that if one defines the price process p_n as

$$p_n = \sum_{m<n} \xi_m, \quad (33)$$

then p_n is a diffusion process with a strictly linear $\mathcal{D}(\ell)$, since the ξ are by construction uncorrelated. The price defined in this way can also be written, using equation (31), as a linear combination of past q_m, as assumed in our micro-model equation (12), with

$$G_0^*(\ell) \equiv \sum_{m=0}^{\ell-1} Q(m). \quad (34)$$

This is an exact relation between $\mathcal{C}_2$ (that allows one to compute in turn K and Q) and the response function G_0^* for all ℓ, where the star indicates that strict diffusion is imposed.

In the case of power-law kernels, one finds from the above relation and from $Q(n) \propto n^{\delta-2}$ for large n

$$G_0^*(\ell) \propto \ell^{\delta-1} \longrightarrow \beta = 1-\delta = \frac{1-\gamma}{2}, \quad (35)$$

which is, not surprisingly, the relation obtained in the main text from the assumption that prices are diffusive on long timescales.

Equation (34) can be used to construct G_0^* from the empirical determination of $\mathcal{C}_2$, shown in figure 11. In order to obtain this curve, we have fitted $\mathcal{C}_2(n)$ as

$$\mathcal{C}_2(0) = 33.5; \qquad \mathcal{C}_2(n) = \frac{7.16}{(1.95+n)^{0.285}}, \quad (36)$$

and used the Levinson–Durbin recursion algorithm for solving a Toeplitz system (see, e.g., [48]).

References

[1] Fama E F 1970 Efficient capital markets: a review of theory and empirical work *J. Finance* **25** 383

[2] Samuelson P A 1965 Proof that properly anticipated prices fluctuate randomly *Industrial Management Rev.* **6** 41

[3] Shiller R J 1981 Do Stock Prices move too much to be justified by subsequent changes in dividends? *Am. Econ. Rev.* **71** 421
Shiller R J 2000 *Irrational Exuberance* (Princeton, NJ: Princeton University Press)

[4] Arthur W B 1995 Complexity in economic and financial markets *Complexity* **1** 1

[5] Shleifer A 2000 *Inefficient Markets, an Introduction to Behavioral Finance* (Oxford: Oxford University Press)

[6] Orléan A 1999 *Le Pouvoir de la Finance* (Paris: Jacob)
Orléan A 2001 A quoi servent les marchés financiers? *Qu'est-ce que la Culture?* (Paris: Jacob)

[7] Hopman C 2002 Are supply and demand driving stock prices? *MIT Working Paper*

[8] Bak P, Paczuski M and Shubik M 1997 Price variations in a stock market with many agents *Physica* A **246** 430

[9] Maslov S 2000 Simple model of a limit order-driven market 2000 *Physica* A **278** 571
Maslov S and Millis M 2001 Price fluctuations from the order book perspective—empirical facts and a simple model *Physica* A **299** 234

[10] Challet D and Stinchcombe R 2001 Analyzing and modelling 1 + 1d markets *Physica* A **300** 285

[11] Slanina F 2001 Mean-field approximation for a limit order driven market model *Phys. Rev.* E **64** 056136

[12] Daniels M, Farmer J D, Iori G and Smith E 2002 Demand storage, market liquidity, and price volatilty *SFI Working Paper* 02-01-001
This paper appeared in final, but truncated form as Daniels M G, Farmer J D, Iori G and Smith E 2003 Quantitative model of price diffusion and market friction based on trading as a mechanistic random process *Phys. Rev. Lett.* **90** 108102

[13] Bouchaud J P, Mézard M and Potters M 2002 Statistical properties of stock order books: empirical results and models *Quant. Finance* **2** 251

[14] Smith E, Farmer J D, Gillemot L and Krishnamurthy S 2003 Statistical theory of the continuous double auction *Quant. Finance* **3** 481–514
Smith E, Farmer J D, Gillemot L and Krishnamurthy S 2002 *Preprint* cond-mat/0210475

[15] Challet D and Stinchcombe R 2003 Non-constant rates and over-diffusive prices in a simple model of limit order markets *Quant. Finance* **3** 155

[16] Wyart M and Bouchaud J P 2003 in preparation

[17] Farmer J D 2003 The long memory of efficient markets *ICBI Global Derivatives Trading Summit (Barcelona, May 2003)*
F Lillo and J D Farmer 2003 The long memory of efficient markets *Preprint* cond-mat/0311053

[18] Farmer J D, Patelli P and Zovko I 2003 The predictive power of zero intelligence models in financial markets *Preprint* cond-mat/0309233

[19] Peng C K, Mietus J, Hausdorff J, Havlin S, Stanley H E and Goldberger A L 1993 Long-range anticorrelations and non-Gaussian behavior of the heartbeat *Phys. Rev. Lett.* **70** 1343–6
Bernaola-Galvan P, Ivanov P Ch, Amaral L A N and Stanley H E 2001 Scale invariance in the nonstationarity of human heart rate *Phys. Rev. Lett.* **87** 168105
The qualitative analogy with financial markets was recently discussed in Struzik Z 2003 Taking the pulse of the economy *Quant. Finance* **3** 78

[20] Cabrera J L and Milton J G 2002 On–off intermittency in a human balancing task *Phys. Rev. Lett.* **89** 158702

[21] Lux T 1996 The stable Paretian hypothesis and the frequency of large returns: an examination of major German stocks *Appl. Financial Economics* **6** 463

[22] Plerou V, Gopikrishnan P, Amaral L A, Meyer M and Stanley H E 1999 Scaling of the distribution of price fluctuations of individual companies *Phys. Rev.* E **60** 6519
Gopikrishnan P, Plerou V, Amaral L A, Meyer M and Stanley H E 1999 Scaling of the distribution of fluctuations of financial market indices *Phys. Rev.* E **60** 5305

[23] Plerou V, Gopikrishnan P, Amaral L A N, Gabaix X and Stanley H E 2000 Diffusion and economic fluctuations *Phys. Rev.* E **62** 3023

[24] Mantegna R and Stanley H E 1999 *An Introduction to Econophysics* (Cambridge: Cambridge University Press)
Bouchaud J-P and Potters M 2003 *Theory of Financial Risks and Derivative Pricing* 2nd edn (Cambridge: Cambridge University Press)

[25] For a recent review, see Cont R 2001 Empirical properties of asset returns: stylized facts and statistical issues *Quant.*

Finance **1** 223 and references therein

[26] Muzy J-F, Delour J and Bacry E 2000 Modelling fluctuations of financial time series: from cascade process to stochastic volatility model *Eur. Phys. J.* B **17** 537–48
For a recent review, see Bouchaud J P and Muzy J F 2004 Financial time series: from Bachelier's random walk to multifractal 'cascades' *Kolmogorov's Legacy* (Berlin: Springer) at press

[27] Hasbrouck J 1991 Measuring the information content of stock trades *J. Finance* **46** 179

[28] Barra 1997 *Market Impact* unpublished report

[29] Kempf A and Korn O 1999 Market depth and order size *J. Financial Markets* **2** 29

[30] Plerou V, Gopikrishnan P, Gabaix X and Stanley H E 2002 Quantifying stock price response to demand fluctuations *Phys. Rev.* E **66** 027104

[31] Lillo F, Mantegna R and Farmer J D 2002 Single curve collapse of the price impact function for the New York stock exchange *Preprint* cond-mat/0207428
Lillo F, Mantegna R and Farmer J D 2003 Master curve for price-impact function *Nature* **421** 129

[32] Lillo F and Farmer J D 2003 On the origin of power-law tails in price fluctuations *Preprint* cond-mat/0309416

[33] Potters M and Bouchaud J P 2003 More statistical properties of order books and price impact *Physica* A **324** 133

[34] Odean T 1999 Do investors trade too much? *Am. Econ. Rev.* **89** 1279

[35] Kyle A S 1985 Continuous auctions and insider trading *Econometrica* **53** 1315

[36] Biais B, Foucault T and Hillion P 1998 *Microstructure des Marchés Financiers* (Paris: PUF)

[37] Rosenow B 2002 Fluctuations and market friction in financial trading *Int. J. Mod. Phys.* C **13** 419

[38] Bouchaud J-P and Cont R 1998 A Langevin approach to stock market fluctuations and crashes *Eur. J. Phys.* B **6** 543

[39] On this point, see Gabaix X, Gopikrishnan P, Plerou V and Stanley H E 2003 A theory of power-law distributions in financial markets fluctuations *Nature* **423** 267

[40] Wyart M and Bouchaud J P 2003 Self-referential behaviour, overreaction and conventions in financial markets *J. Econ. Behaviour Organisation* submitted (*Preprint* cond-mat/0303584)

[41] Black F 1986 Noise *J. Finance* **41** 529

[42] de Bondt W and Thaler R 1985 Does the market overreact? *J. Finance* **40** 793

[43] Bak P 1996 *How Nature Works: the Science of Self-Organized Criticality* (New York: Copernicus, Springer)

[44] Challet D, Chessa A, Marsili M and Zhang Y C 2001 From minority games to real markets *Quant. Finance* **1** 168 and references therein;
Challet D, Marsili M and Zhang Y C 2003 *The Minority Game* in preparation

[45] Giardina I and Bouchaud J P 2003 Bubbles, crashes and intermittency in agent based market models *Eur. Phys. J.* B **31** 421
See also the discussion in Bouchaud J P 2001 Power-laws in economics and finance: some ideas from physics *Quant. Finance* **1** 105

[46] Lux T and Marchesi M 2000 Volatility clustering in financial markets: a microsimulation of interacting agents *Int. J. Theor. Appl. Finance* **3** 675

[47] Krawiecki A, Hoyst J A and Helbing D 2002 Volatility clustering and scaling for financial time series due to attractor bubbling *Phys. Rev. Lett.* **89** 158701

[48] Percival D B 1992 Simulating Gaussian random processes with specified spectra *Comput. Sci. Stat.* **24** 534