The non-trivial random walk of stock prices: sign-size dependence

Gabriele La Spada*[1] and Javier Vicente[2]

[1]*LUISS Guido Carli University, Viale Pola 12, Roma, Italy*

[2]*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501*

We study the inner structure of the random walk of stock prices at transaction-by-transaction level. From our previous works we know the stochastic process underlying the dynamics of stock prices is not a simple random walk: that the signs and the absolute values of the microscopic returns (i.e. the steps fo the random walker) are not independent, but rather there is a subtle relationship between them. We know that this relation cannot be explained in terms of pairwise dependance between a sign and an absolute return; in this paper we show that it depends in some measure on sign predictability: each absolute return depends clearly on the properties of the past sequence of sign. The price stochastic process is a kind of *adaptive random walk*. This means that the response of price to trade depends on the past history of trading activity. When the trading pressure tries to move the price in a given direction, the response of the price is significantly asymmetric. The response of price to trade is significantly smaller in the direction toward which the forcing trading activity tries to move the price. We analyze the highly capitalized stock Astrazeneca (AZN) traded at the London Stock Exchange in the period May 2000–December 2002.

I. INTRODUCTION

The returns of a financial product are unpredictable. This empirical evidence is consistent with the efficient market hypothesis (EMH) in its weak form. *A financial market is said (informationally) efficient if market prices reflect all available information about value.* In the weak form of the EMH "all the available information" is the past prices. Roughly speaking, in the weak form EMH states that there is no way of making a profit on an asset by simply using the recorded history of its price fluctuations. That is, the returns of a given financial asset are (*i*) zero on average and (*ii*) unpredictable from their historical time series. It's well-known that unpredictable time series and stochastic processes are not synonymous, but a widely accepted belief in financial theory is that asset price dynamics can be modeled as a stochastic process. The *random walk* (RW) was the first stochastic process proposed to model price dynamics. Its first formalization was proposed in the doctoral thesis of the French mathematician Louis Bacheliér [1], five years before the famous paper written by Einstein on brownian motion (BM) and the determination of the Avogadro number in 1905. This process assumes that the steps of the random walker are *independent identically distributed* (IID).

From many years it's well known that the stochastic process underlying price changes is not a pure BM. Many empirical studies have shown that financial returns have *fat tails* [6, 10, 11, 15–18, 20] and some non-linear functions of returns are strongly correlated (*clustered volatility*) [3–5, 13, 14, 19]. Then, the steps of the random walker are *not gaussian distributed, neither independent.* Moreover, our recent works suggest that there is also a subtle correlation between the sign of the steps and their magnitude. This result is confirmed by a recent work of Weber [21].

In this paper we study the relation be-

*E-mail: glaspada@luiss.it

tween the signs of the returns and the sizes of the returns at 'transaction-by-transaction' level[1]. We distinguish between transaction-to-transaction returns – price changes driven by one single transaction and the market response – and impacts – price changes driven by one single transaction only. We study both cases. In our previous results we argue that the absolute returns are affected in some way by *sign predictability*. I.e., if signs are in some measure predictable this affects in some way (to reduce volatility) the magnitude of the returns. To test this insight we build a *predictor* of the signs and we study the behavior of the absolute returns as the predictability increases. First, we observe that the next absolute return increases as the predictability increases. Second (and this is the central point), for impacts the next absolute return increases more when the prediction is wrong. This is reasonable, because it means that the liquidity in the book order moves depending on the flow of supply and demand.

We also study the effect of *runs of same-sign returns* on the *first opposite-sign absolute return*. We observe that the longer is the run of equal signs the larger is the price change in the opposite direction. This result confirms that liquidity moves with the flow of supply and demand in order to rebalance the price. Longer is the run of price changes toward a direction, stronger is the response of the market to reduce this imbalance (the opposite price change is larger).

The remainder of the paper is organized as follows. In section II we describe the continuous double auction mechanism, which is at the basis of most of the modern stock markets, and our data set. In section III we deal with the "random walk" of stock prices. First, we introduce the notion of generalized random walk and we present a simple model with signs independent of sizes, and *vice versa*. Then, we show briefly our preliminary results about the overestimation of the expected volatility and the likely relation between step signs and step sizes. In section III we illustrate the core of our project, analyzing directly the relationship between step signs and step sizes. First, we show our result about the key-role of sign predictability and its effect on the next absolute returns. Second, we also study the relationship between long sequences of equal signs and the following absolute returns. Finally, in section IV we conclude with a discussion of our results.

[1] Some of this work is based on research done in collaboration with Doyne Farmer and Fabrizio Lillo, and will be published at a future date.

II. BACKGROUND: THE CONTINUOUS DOUBLE AUCTION MECHANISM

The continuous double auction is the standard mechanism for price formation in most of the modern financial markets. Agents can place different types of orders, which can be grouped into two categories: *market orders*, and*limit orders* or *quotes*.

Market orders, usually submitted by impatient traders, are requests to buy or sell a given number of stock shares immediately at the best available price. More patient traders submit limit orders (quotes), which in addition to the number of shares to sell or to buy also state a limit price π, corresponding to the worst allowable price for the transaction. It's important to note that the word "quote" can be used to refer either 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 with priority called the *limit order book*. Buy limit orders are called *bids* and sell limit orders are called *offers* or *asks*. At any given time t there is a best (lowest) offer to sell with a price $a(t)$ and a best (highest) bid to buy with the price $b(t)$. These are also called the *inside quotes* or the *best prices*. The price gap between them is called the *bid-ask spread* or simply the *spread*: $s(t) = a(t) - b(t)$

It's important to be clear: prices are not continuous, but rather change in discrete quanta called *ticks* whose size is Δp. The number of shares in an order is called either its *size* or its *volume*. In Fig. 1 we illustrate schematically the limit order book and the order placement mechanism.

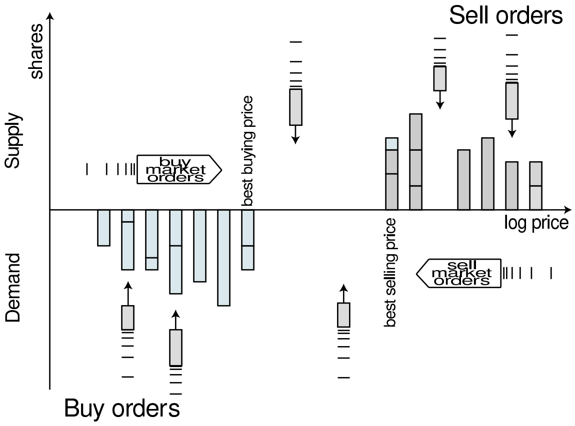

FIG. 1: **Limit order book.** Schematic illustration of limit order placement and market orders matching. Limit orders are stored in the limit order book; as market orders arrive in the market they are matched against limit orders . We adopt the arbitrary convention that buy orders are represented by negative share requests and sell orders by positive share requests.

As market orders arrive they are matched against limit orders of the opposite sign in order of first price and then arrival time. 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 next best is for 300 shares at €60.25; a buy market order for 150 shares will buy all the shares requested at the best price, “leaving” 50 shares at the best ask $a(t) =$ €60, while a buy market order for 250 shares will buy 200 shares at €60 and 50 shares at €60.25, moving the best ask $a(t)$ from €60 to €60.25.

It makes sense to argue that transactions are driven only by matching between market orders and limit orders, but this is not strictly true. In fact, it could happen that a limit order to buy is placed at the same price as the best (lowest) limit order to sell, generating in this way a transaction.

Any given order can always be decomposed into two types: we can call any component of an order that results in immediate execution an *effective market order* and any component that is not executed immediately, but stored in the limit order book, an *effective limit order*. Thus, resuming the previous example, consider the limit order to buy at a price $\pi = a(t)$. Suppose the volume at $a(t)$ is 1,000 shares and the volume of the new limit order is 3,000. Then this limit order is equivalent to an effective market order for 1,000 shares, followed by an effective limit order of 2,000 shares with the same limit price $a(t)$. In either case the same transactions take place and the best prices move to $b(t+1) = a(t)$ and $a(t+1) = a(t) + g(t)$, where $g(t)$ is the price interval to the next highest price level to sell[2]. In the following we will simply call an effective limit order a ‘limit order’ and similarly an effective market order a ‘market order’.

There is no unique notion of price in a real market. Usually it is convenient to use the so called “mid-point price” or “mid-quote price” defined by the best quotes: $m(t) = [a(t) + b(t)]/2$. This price definition is widely used in literature. Another possibility could be the transaction price, but for our work mid-price is a better choice. First, there is very little difference between the two, given that they differ by less than half the spread. Secondly, the mid-price is more convenient because it avoids problems associated with the tendency of transaction prices to bounce back and forth between the best bid and ask.

Given this definition of stock price we can try to explain what causes price changes. Price changes are typically characterized as *returns*: $r_\tau(t) = log\,[m(t)] - log\,[m(t-\tau)]$ At the microstructural level, the arrival of three kinds of events can cause the mid-price to

[2] Note that time counter t advances with events occurring.

change:

- *Market orders* A market order bigger than the opposite best quote widens the spread by increasing the best ask if it is a buy order or decreasing the best bid if it is a sell order.

- *Limit order placements* A limit order placed inside the spread narrows it by increasing the best bid if it is a buy order or decreasing the best ask if it is a sell order.

- *Limit order cancellations* A cancellation of (all) limit order(s) at the best price widens the spread either increasing the best ask or decreasing the best bid.

The price change of each single transaction will be due to (effective) market order execution and this particular price change is called *market impact* or *price impact* or even simply *impact*. Looking at what happens between two consecutive transactions we will see also the effect of quote placements and cancellations. Then, the cumulative effect of these three kinds of events is what happens at the stock price from a transaction to the next one. We refer to this kind of price change as *transaction-to-transaction return*. To summarize, there can be two different *definition of price changes*:

◇ **Impacts**: an impact is defined as the difference between the log mid-price just after the transaction and the log mid-price just before.

◇ **Transaction-to-transaction returns**: a trade to trade price change is defined as the difference between the log mid-price just before the next transaction and the log mid-price just before the current one.

The first definition takes into account *only effective market orders* and the price changes they caused in the market. Thus, *the aggregation of many impacts on a real time interval is not the total price change on that interval*, because for each transaction it doesn't take into account the response of the market, i.e. the placement and the cancellation of limit orders, that can cause other price changes. Otherwise, *the aggregation of many transaction-to-transaction returns on a real time interval is by definition the total price change on that interval.*

III. THE "RANDOM WALK" OF STOCK PRICES

A. The notion of generalized random walk

Let the series of price changes be a time series $r_t = s_t w_t$, where r_t is the t^{th} midpoint price change (or return), s_t is its sign, and w_t is its magnitude. We make the coarsest possible interval of time be the interval between transactions. For convenience we only advance the time counter t for transaction-to-transaction intervals with nonzero price changes, i.e. we ignore any intervals where there are no prices changes. Let us assume: (i) both s_t and w_t *wide sense stationary* (WSS), (ii) s_t and w_t *mutually independent*. We will call such a stochastic process a *generalized random walk* (GRW). This process does not specify the probability density function (PDF) of the steps, neither the autocorrelation function (ACF) of s_t and w_t. Under the assumption that signs are equally distributed, i.e. $\mathbb{P}(s_t = +1) = \mathbb{P}(s_t = -1)$ (*unbiased* GRW), it's easy to get the variance of the process

$$Var(R_n) = n \left[K_{s,w}(n) \sigma_w^2 + K_s(n) \mu_w^2 \right] \quad (1)$$

where n is the number of microscopic nonzero price changes (steps of the walker), μ_w and σ_w^2 are respectively the mean and the variance of the absolute returns, $K_{s,w}(n) = 1 + 2\sum_{l=1}^{n} \left(1 - \frac{l}{n}\right) c_s(l) c_w(l)$ and $K_s(n) = 1 +$

$2\sum_{l=1}^{n}\left(1-\frac{l}{n}\right)c_s(l)$ are two factors that depend on the temporal ACFs of signs and sizes, respectively $c_s(l)$ and $c_w(l)$.

To test this model we divide the series into realtime intervals of length T which we will label with index i. Each real time interval will in general contain a different number of nonzero price changes, and the size of the changes may vary within each interval. We can thus rewrite this equation in the form

$$V_i = n_i\left[K_{s,w}(n_i)\sigma_{w,i}^2 + K_s(n_i)\mu_{w,i}^2\right]. \qquad (2)$$

Now $V_i = Var(R_n)_i$ is the measured sample variance in interval i, which we can also simply call the volatility in that interval. Note that all the variables are labeled with i because they can fluctuate very much from one interval to the other (especially n_i).

To test the reliability of the GRW proposed above we have to compare the expected volatility to an *empirical proxy*. For any real time interval i we choose the square absolute return overall that T-interval, $|R_n|^2$, which is a robust widely used proxy for the volatility. Then, for any real time T-interval we have to compute n_i, $\mu_{w,i}^2$, $\sigma_{w,i}^2$, $K_{s,w}(n_i)$, $K_s(n_i)$, R_{n_i}.

B. Overestimation of volatility

From our previous works we know that the GRW assuming signs and sizes to be independent WSS stochastic processes is not a good approximation of what really happens. It causes a *systematic and significant overestimation* of expected volatility w.r.t. empirical volatility. This result is coherent with a previous work by Gillemot *et al.* (2006) [9], in which authors investigated the inner causes of volatility. We observed the overestimation effect on four highly capitalized stocks traded at LSE in the period May 2000–December 2002: Astrazeneca (AZN), Lloyds Tsb Group (LLOY), Shell Transport & Trading Co. (SHEL) and Vodafone Group (VOD). We report some results in Fig. 2. To investigate the likely causes of such an overestimation we study the single contribution of

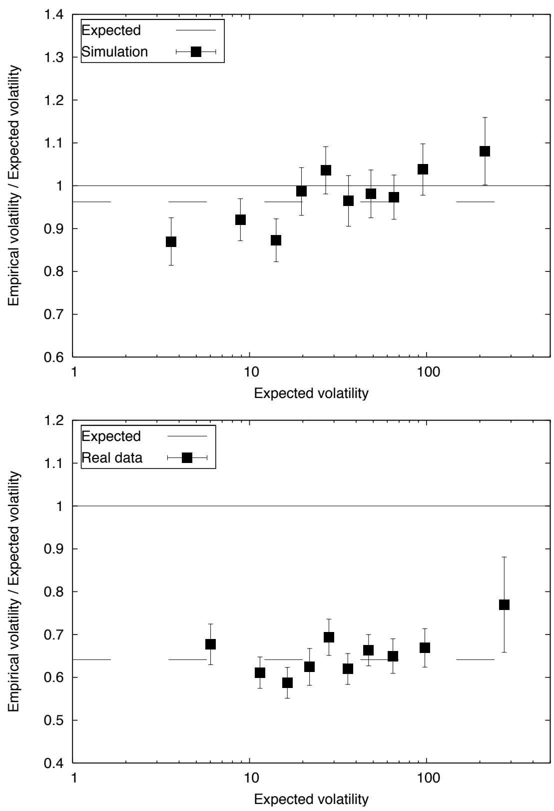

FIG. 2: **Overestimation of volatility** Top: simulation. Bottom: real data. On the x-axis we have the expected volatility. On the y-axis we have the ratio empirical volatility/expected volatility. Data are binned on the x-axis using quantiles. Simulation: virtual log-normal absolute returns long ranged correlated (Hurst exponent $H \sim 0.7$); original series of transaction-to-transaction signs from the AZN data file; 1-hour sub-intervals; original number of non-zero returns for each interval. Real data: AZN, 1 hour sub-intervals. We see that for the simulation the weighted mean value of the data points is close to one, confirming the reliability of our analytical formulas and our experimental procedure. Instead, for real data: we see that the expected volatility is systematically greater than the empirical volatility.

each autocorrelation structure to the volatility. To this aim we perform a set of shuffling experiments on the original time series. These experiments ensured that the cause of such overestimation are not the autocorrelation functions of signs or sizes, neither an equal-time dependence between signs and sizes, but rather a *non-synchronous and non-pairwise relationship between signs and sizes.* We report our results only for AZN and, for the sake of brevity, we do not report all the results, but only the most important.

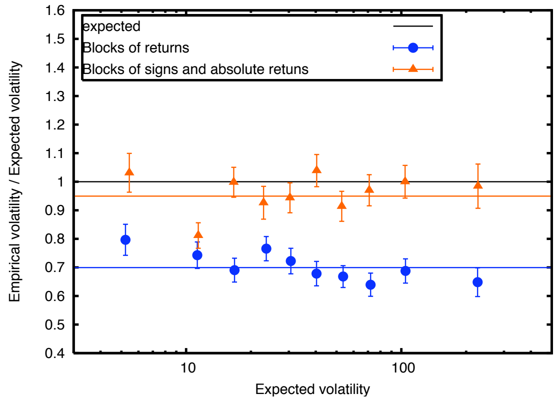

FIG. 3: **Sign-size dependence**. Shuffling blocks of returns vs. shuffling blocks of absolute returns and signs separately. AZN, transaction-to-transaction returns. On the x-axis we have the expected volatility. On the y-axis we have the ratio empirical volatility/expected volatility. Data are binned on the x-axis. Blue circles represent when we shuffle blocks of returns, while orange triangles represent when we shuffle signs and absolute values separately. The blue and orange solid lines are the weighted mean values. In the first case we preserve any correlation between sequences of signs and sequences of absolute returns, while in the second one we destroy it. The size of each block is 100 returns in both the experiments. We see that orange triangles are much closer to the black line at $y = 1$ confirming that a subtle correlation between absolute returns and signs is the cause of the overestimation for real data. We perform the same experiment with impacts instead of transaction-to-transaction returns and we find the same result.

IV. SIGN-SIZE DEPENDENCE: A DIRECT APPROACH

A. Sign predictability

In view of our previous results we argue that the absolute returns are affected in some way by *sign predictability.* This intuition is also supported by other previous works [2, 8, 12] focused on the long-memory of supply and demand (high predictability of impact sign) and its relationship with the linear efficiency of price returns. The point is: if signs are in some measure predictable this affects in some way (to reduce volatility) the magnitude of the returns. To test this insight we build a predictor of the signs and we study the behavior of the absolute returns as the predictability increases.

We use a simple predictor: the *sum* of the signs. For each sequence of signs of length L we compute the sum and we consider the next absolute returns; that is, if the sequence is $s_i, s_{i+1}, \ldots, s_{i+L-1}$, we compute $S_{i,L} = \sum_{k=i}^{i+L-1} s_k$ and take $w_{next} = w_{i+L}$. For each sequence the absolute value of the predictor can range in the discrete subset $\{0, 2, 4, \ldots, L\}$. A large value of the absolute sum means a large number of returns with same sign, that we interpret as high predictability. A small value of the absolute sum means approximately an equal number of positive and negative signs, that we interpret as low predictability. We consider the entire original time series of signs and absolute returns, from May 2000 to December 2002, and we fix a length L for the sequences. We perform our tests on both transaction-to-transaction returns and impacts. For the two definition of return we choose the same length: $L = 20$, even if we know that they have different autocorrelations. For the impacts we also used $L = 50$. We do not consider longer sequences because it does not increase the reliability of the predictor. We have made naive checks of the reliability of the "sum predictor", and we found that is a quite poor predictor in the case of transaction-to-transaction returns, while it works quite well for the impacts. The improvement of the sign predictor (e.g. with autoregressive models and non-linear methods) is certainly an important point for the developments of this line of research, but it is outside the scope of the present work.

We plot the mean value of the next absolute return conditional on the absolute value of the sum predictor $\mathbb{E}\left[w_{i+L}\big||S_{i,L}|\right]$ (y-axis) against the absolute value of the sum predictor $|S_{i,L}|$ (x-axis). We perform two analyses: (i) we compare the result for transaction-to-transaction return to that one for impacts, (ii) we distinguish between right and wrong predictions for both impacts and transaction-to-transaction returns. We report our results in Fig. 4, 5, 6 and 7. From the first analysis we find that for both transaction-to-transaction returns and impacts the next absolute return increases as the predictability increases, but it increases much more significantly for transaction-to-transaction returns than for impacts. This evidence can be summarized as follows:

- if the predictability for the next return sign increases, the return will increase;
- if the predictability for the next impact sign increases, the impact will increase;
- the transaction-to-transaction return increases faster than the impact.

These points give us some information about the reaction of the market to the order flow. For example, let us suppose that the order flow is causing a systematic price change in the upper direction, that is, a long run of buy market orders is causing positive price changes. Let us consider for now only the sell side of the order book. After the execution of a buy market order, who places the limit orders on the sell side could be tempted to place his limit orders at higher prices to get larger gains (translation). If buy market orders still arrive in the market – changing the price even if the sell quotes are getting higher, this strategy will be profitable. In this case such a behavior could explain why the size of transaction-to-transaction returns increases as the probability of future positive transaction-to-transaction returns increases (first item). However, the reaction of the sell side of the limit book will not be rigid, that is, the sell side of the limit will not undergo a rigid translation in the positive direction, but the sizes of the gaps between the quotes will also change in time, causing dilatations and contractions. If the size of the first gap between the quotes on the sell side increases as the predictability increases (dilatation), then the next buy market order (that is expected) will have a larger impact than the previous one (second item). The third item could mean that the translation is faster than the dilatation.

From the second analysis we obtain two different results depending on the definition of returns that we use. For transaction-to-transaction return there is substantially no difference between right and wrong prediction. Instead, by comparing the two curves when analyzing the impacts we find that the next absolute return increases more when the prediction is wrong (Fig. 6 and 7). This is reasonable, because it means that the liquidity in the book order moves depending on the flow of supply and demand. If there is a long run of buy market orders changing the price in the upper direction, we predict to a good agreement that the next impact will be positive; in this case the liquidity (volume and density of limit orders) accumulates on the sell side of the book order, leaving the buy side more sparse and with less volume. Then, if a sell market order instead of the expected buy order occurs, this unexpected outcome prompts a bigger price changes that tends to revert the price of the asset to its original value. We see that this behavior is even more convincing if we use $L = 50$, that is a "longer" predictor[3].

The result obtained for transaction-to-transaction returns is probably affected by the inaccuracy of the predictor. From the efficient market hypothesis we should expect difference between "right prediction" and "wrong prediction" . In fact, the weak form of EMH states that the next return r_t will have zero mean value

$$\mathbb{E}[r_t] = 0\,. \tag{3}$$

This condition holds for the transaction-to-transaction returns, but not necessarily for the impacts. From Eq. (3) we can write

$$p_t^+ \mathbb{E}[r^+] = p_t^- \mathbb{E}[r^-]\,, \tag{4}$$

where p^+ and p^- are the probabilities that r_t will be positive or negative respectively, and $\mathbb{E}[r^+]$ and $\mathbb{E}[r^-]$ are the mean values of the absolute value of r_t if r_t is positive or negative respectively. Then, if a sign is more likely to probable than the other, the average absolute return must be smaller, e.g. if $p_t^+ > p_t^-$, then $\mathbb{E}[r^+] < \mathbb{E}[r^-]$. This relationship is consistent with the result observed for impacts (if we expect the arrive of a market order on a given side of the order book, the return of a market order arriving on the opposite side will be greater), but it is not observed when considering transaction-to-transaction returns. This is a quite surprising result given that relations 3 and 4 hold for the transaction-to-transaction returns, but not (necessarily) for the impacts. To shed more light on this point we perform another test.

[3] We note that such a choice does not improve the reliability, that remains practically the same for all $L \geq 20$.

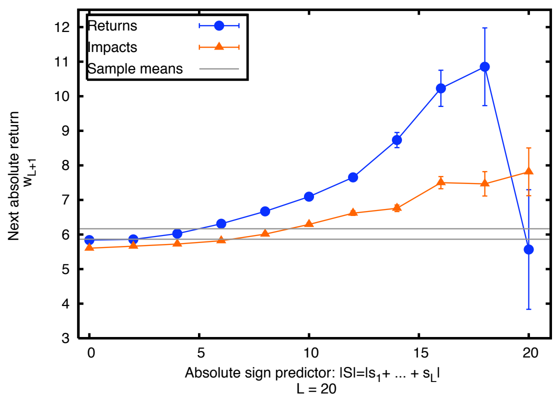

FIG. 4: **Sign predictability and absolute returns.** On the x-axis there is the absolute value of the 'sum' predictor, $|S| = |\sum_{i=1}^{L} s_i|$ with $L = 20$. High predictability corresponds to high values of the absolute sum. On the y-axis there is the mean value of the next absolute returns w_{L+1}. We compare transaction-to-transaction returns to impacts. We compare transaction-to-transaction returns (blue circles) to impacts (orange triangles). In both cases an high predictability of the next sign corresponds to an high value of the next absolute return. This monotonic relationship is much more evident for transaction-to-transaction than for impacts.

B. Another test: runs of same-sign returns

We check the relationship between signs and absolute returns also in another way. We consider runs of equal signs and take the absolute return corresponding to the first sign flip. I.e., let a sequence of signs be $s_i = s_{i+1} = \ldots = s_{i+L-1}$ and $s_{i+L} = -s_i$, we take $w_{next} = w_{i+L}$. We plot the mean value of w_{next} conditional on the length L, $\mathbb{E}[w_{i+L}|L]$ (y-axis), against L (x-axis). For each value of L the standard error is about of the same order of the sample mean; thus, we have preferred to leave it out from our plots in order to get a clearer qualitative trend. Instead, we give some descriptive statistics: min, first quartile, median and third quartile. We perform this test for both transaction-to-

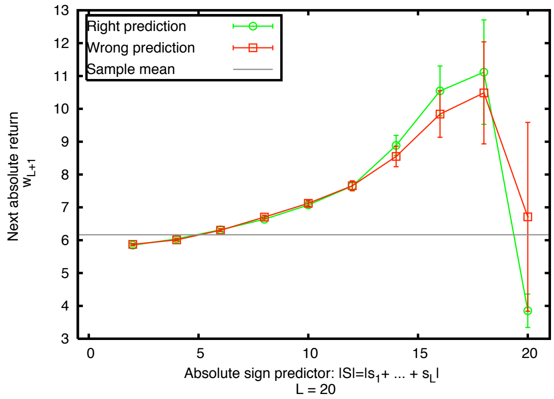

FIG. 5: **Right prediction vs. wrong prediction. Transaction-to-transaction returns.** On the x-axis we put the absolute value of the 'sum' predictor, $|S| = |\sum_{i=1}^{L} s_i|$ with $L = 20$. On the y-axis we put the mean value of the next absolute return w_{L+1}. We distinguish the cases when the predictor works, i.e. the next sign is the same as the prediction (green circles), from the cases when the predictor does not work, i.e. the next sign is the opposite of the prediction (red squares). We see that there is no practically distinction between the two, and both grow monotonically as the predictability increases. But we have to stress that the reliability of the 'sum predictor' is very poor for transaction-to-transaction return.

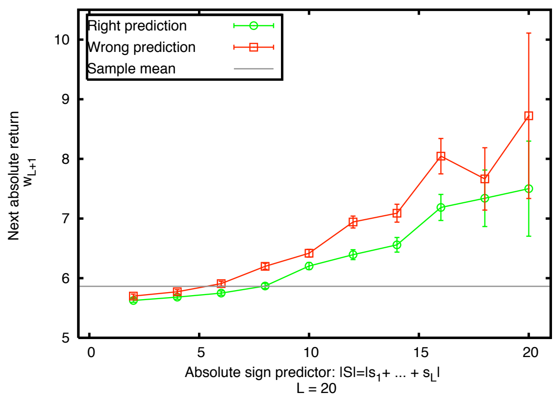

FIG. 6: **Right prediction vs. wrong prediction. Impacts.** On the x-axis we put the absolute value of the 'sum' predictor, $|S| = |\sum_{i=1}^{L} s_i|$ with $L = 20$. On the y-axis we put the mean value of the next absolute return w_{L+1}. We distinguish the cases when the predictor works, i.e. the next sign is the same as the prediction (green circles), from the cases when the predictor does not work, i.e. the next sign is the opposite of the prediction (red squares). We see that both grow monotonically as the predictability increases, but there is *a significant distinction* between the two. When the next return occurs on the opposite side w.r.t. the prediction, its impact is much bigger than when it occurs on the predicted side.

transaction and impacts. We report our results in Fig. 8 and 9.

We see that in both cases the absolute return increases as the length of sequences of identical signs increases. In terms of the sum predictor this corresponds to an increasing of the predictability. From this analysis we have a confirmation of one of our intuitions. The liquidity moves with the run of market orders: the longer the run of penetrating market orders of the same sign, the larger the first absolute impact of the opposite sign. This analysis points out that also transaction-to-transaction returns show a similar behavior: the longer the run of transaction-to-transaction returns of the same sign, the larger is the transaction-to-transaction return of the opposite sign. This result contradicts that one we obtained using the sum predictor (see Fig. 5), which could be affected by the inaccuracy of the predictor.

V. CONCLUSION

Starting from our previous results on the overestimation of volatility we investigated deeper the likely non-synchronous and non-pairwise dependence between signs and sizes for the stock AZN during the period May 2000–December 2002. We observed that sign predictability seems to have a key-role in affecting future absolute returns. We considered a simple predictor of the signs, defined

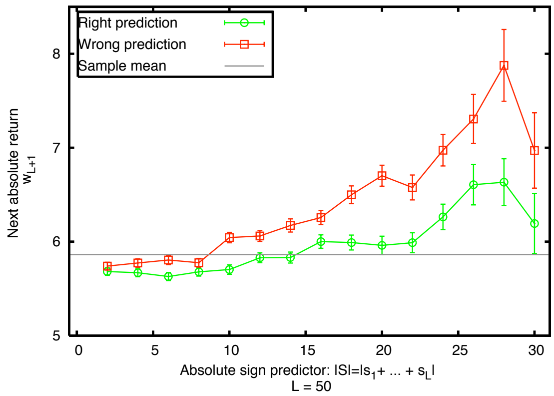

FIG. 7: **Right prediction vs. wrong prediction. Impacts.** On the x-axis we put the absolute value of the 'sum' predictor, $|S| = |\sum_{i=1}^{L} s_i|$ with $L = 50$. On the y-axis we put the mean value of the next absolute return w_{L+1}. We report the same test as in the previous figure, but changing the length of the sum: from $L = 20$ to $L = 50$. This change does not improve significantly the reliability of the sign prediction, but shows more clearly the different effect on the next absolute return between a right prediction (green circles) and wrong prediction (red squares) of the next sign. We see that both grow monotonically as the predictability increases, but when the next return occurs on the opposite side w.r.t. the prediction the impact is much bigger than when it occurs on the predicted side.

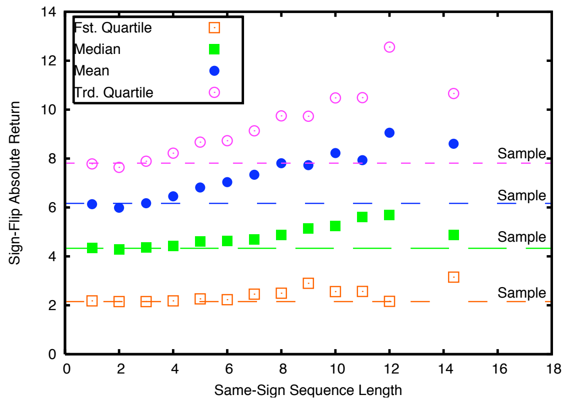

FIG. 8: **Run of identical signs and absolute value of the next opposite return. Transaction-to-transaction.** On the x-axis there is the length of a sequence of identical signs. On the y-axis we report the first quartile, the median, the sample mean and the third quartile of the next absolute return with opposite sign. Data are binned on the x-axis. We observe that median, mean and third quartile increase as the length of the run of identical signs increases. The mean and the third quartile increase significantly. This means that after a long run of returns with the same sign the first opposite return will typically move much deeper the price in the opposite direction. The last data point bins together sequences with $L \geq 13$.

as the sum of the signs on a given length, and we observed that *the next absolute returns is greater when the predictability of the next sign is higher*. This result holds for both transaction-to-transaction returns and impacts. Moreover, we investigated the different effect of right predictions and wrong predictions. For transaction-to-transaction returns we did not observe significant differences, but probably this outcome is due to the inaccuracy of the predictor in this case (because of the short-memory of the signs). Instead, *for impacts we observed a significant difference in the two cases: the absolute return is always greater when the prediction is wrong*. This means that the arrive of "unexpected" orders causes larger price changes, which move the price in the "unexpected" direction. This reverting phenomenon might be explained by liquidity fluctuations: the liquidity in the order book moves depending on the flow of supply and demand, leaving the book more "sparse" on the side where no order arrive is expected. This result is confirmed by the analysis on the runs of equal signs. The longer is the run of returns with the same sign, the larger is the first next absolute return with opposite sign.

In conclusion, from our analyses we find that the signs of the returns and the absolute values of the returns are surely correlated. This correlation is subtle and involves the sign predictability. It is also very likely that this inter-dependance is related to the liquidity imbalance and its time evolution. This conclusion is validated also by other works [7, 8].To gain more information will be necessary to improve the sign predictor, for example by using a simple autoregressive model. This issue will be a central topic to develop in our future works.

Acknowledgments

We are very grateful to Doyne Farmer, Fabrizio Lillo and Austin Gerig for useful discussions and comments. This work was partially supported by the Santa Fe Institute through NSF Grant No. 0200500 entitled "A Broad Research Program in the Sciences of Complexity."

[1] Louis Bacheliér. Thorie de la spculation. *Annales Scientifiques de l'Ecole Normale Suprieure*, III-17:21–86, 1900.

[2] J-P. Bouchaud, Y. Gefen, M. Potters, and M. Wyart. Fluctuations and response in financial markets: The subtle nature of "random" price changes. *Quantitative Finance*, 4(2):176–190, 2004.

[3] P. Cizeau, Y. Liu, M. Meywr, C. K. Peng, and H. E. Stanley. Volatility distribution in the S&P 500 stock index. *Physica A*, 245 (3-4):441–445, 1997.

[4] Z. Ding, C. W. J. Granger, and R. F. Engle. A long memory property of stock returns and a new model. *Journal of Empirical Finance*, 1(1):83–106, 1993.

[5] R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united-kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.

[6] E. F. Fama. The behavior of stock-market prices. *The Journal of Business*, 38(1):34–105, 1965.

[7] J. D. Farmer, A. Gerig, and F. Lillo. Theory for market impact of large trading orders in stock markets. Santa Fe Institute working paper.

[8] J. D. Farmer, A. Gerig, F. Lillo, and S. Mike. Market efficiency and the long-memory of supply and demand: Is price impact variable and permanent or fixed and temporary? *Quantitative Finance*, 6(2): 107–112, 2006.

[9] L. Gillemot, J. D. Farmer, and F. Lillo. There's more to volatility than volume. *Quantitative Finance*, 6:371–384, 2006.

[10] P. Gopikrishnan, M. Meyer, L.A.N. Amaral, and H. E. Stanley. Inverse cubic law for the distribution of stock price variations. *European Physical Journal B*, 3:139–140, 1998.

[11] P. Gopikrishnan, V. Plerou, L.A.N. Amaral, M. Meyer, and H.E. Stanley. Scaling of the distribution of fluctuations of financial market indices. *Physical Review E*, 60: 5305–5316, 1999.

[12] F. Lillo and J. D. Farmer. The long memory of the efficient market. *Studies in Nonlinear Dynamics & Econometrics*, 8(3), 2004.

[13] Y. Liu, P. Cizeau, M. Meyer, C. K. Peng, and H. E. Stanley. Quantification of correlations in economic time series. *Physica A*, 245(3-4):437–440, 1997.

[14] Y. Liu, P. Gopikrishnan, P Cizeau, M. Meyer, C.-K. Peng, and H.E. Stanley. The statistical properties of the volatility of price fluctuations. *Physical Review E.*, 60:1390–1400, 1999.

[15] F. M. Longin. The asymptotic distribution of extreme stock market returns. *The Journal of Business*, 69(3):383–408, 1996.

[16] B. Mandelbrot. The variation of certain speculative prices. *The Journal of Business*, 36(4):394–419, 1963.

[17] R. N. Mantegna and H. E. Stanley. Scaling behavior in the dynamics of an economic

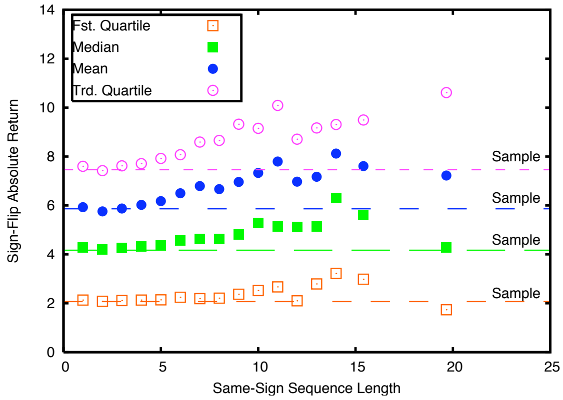

FIG. 9: **Run of identical signs and absolute value of the next opposite return. Impacts.** On the x-axis there is the length of a sequence of identical signs. On the y-axis we report the first quartile, the median, the sample mean and the third quartile of the next absolute return with opposite sign. Data are binned on the x-axis. We observe that median, mean and third quartile increase as the length of the run of identical signs increases. The mean and the third quartile increase significantly. This means that after a long run of impacts (i.e. penetrating market orders) with the same sign the first opposite impact will typically penetrate much deeper in the opposite side of the book order. The last two data points bin together sequences with $15 \leq L < 20$ and $L \geq 20$, respectively.

index. *Nature*, 376(6535):46–49, 1995.

[18] R. R. Officer. Distribution of stock returns. *Journal of the American Statistical Association*, 67(340):807–812, 1972.

[19] M. Pasquini and M. Serva. Multiscaling and clustering of volatility. *Physica A*, 269:140–147, 1999.

[20] V. Plerou, P. Gopikrishnan, L. A. N. Amaral, M. Meyer, and H. E. Stanley. Scaling of the distribution of price fluctuations of individual companies. *Physical Review E*, 60(6):6519–6529, 1999. Part A.

[21] P. Weber. Analysis of aggregated tick returns: Evidence for anomalous diffusion. *Phys. Rev. E*, 75:016105, 2007.