Liquidity crises on different time scales

Francesco Corradi,[1] Andrea Zaccaria,[2,1,*] and Luciano Pietronero[1,2]

[1]*"Sapienza" University of Rome, Rome, Italy*
[2]*ISC-CNR, Rome, Italy*

We present an empirical analysis of the microstructure of financial markets and, in particular, of the static and dynamic properties of liquidity. We find that on relatively large time scales (15 min) large price fluctuations are connected to the failure of the subtle mechanism of compensation between the flows of market and limit orders: in other words, the missed revelation of the latent order book breaks the dynamical equilibrium between the flows, triggering the large price jumps. On smaller time scales (30 s), instead, the static depletion of the limit order book is an indicator of an intrinsic fragility of the system, which is related to a strongly nonlinear enhancement of the response. In order to quantify this phenomenon we introduce a measure of the liquidity imbalance present in the book and we show that it is correlated to both the sign and the magnitude of the next price movement. These findings provide a quantitative definition of the effective liquidity, which proves to be strongly dependent on the considered time scales.

PACS number(s): 89.65.Gh, 89.75.−k

I. INTRODUCTION

The application of tools and methods borrowed from the fields of statistical mechanics and the physics of complex systems to the study of financial markets has led to a number of empirical and theoretical results [1,2]. In particular, the analysis of markets' microstructure [3,4] benefits from the huge amount of available data, which permits to ground models and findings on a strong empirical basis [5]. One of the most important issues is related to the subtle relationships among order flow, liquidity, and price movements. The empirical evidence, in fact, is in contrast with the natural guess that price movements should be trivially related to the arrival of orders and, in particular, to their volume. First, orders are strongly correlated in sign [6–8], in sharp contrast with the well-known absence of autocorrelation in price returns, which is at the basis of the efficient markets hypothesis [9]. Moreover, the magnitude of price movements has been shown to be little dependent on the volume of the incoming orders [10–12] (in passing we point out that, as it has been shown in Ref. [13], neither the arrival of new information and news in the market is able to explain the occurrence of large price movements. In the following, we will focus on the endogenous mechanisms that can trigger price jumps, regardless of possible external causes). The possibility to reconstruct the microstructure of financial markets, namely, the so called limit order book (LOB) [14,15], permits us to stress the importance of agents' strategies on the state of the market, for example, on spread dynamics and volatility clustering [16], and the reaction in terms of balancing order flow to explain the magnitude of price movements [17–20]. These findings directly lead us to the concept of *liquidity*. Even if it is widely recognized as "the most important characteristic of well-functioning markets" [3], liquidity does not have a single shared definition. Generally speaking, a market is *liquid* if agents can trade large volumes quickly, with little impact on the price and facing low transaction costs. In other words, it is a measure

of the ease of converting assets into legal tender (which is, by definition, fully liquid) and vice versa. From a practical point of view, and keeping in mind possible applications to the LOB, this situation can be described by a number of features that make a market liquid. Using the suggestions of Sarr and Lybek [21], we can say that a market is liquid if it has the following characteristics: (i) tightness, both implicit and explicit transaction costs are low; (ii) immediacy, the trading speed is high; (iii) depth, the orders already placed in the book cover a wide range of prices; (iv) breadth, the best quoted orders are large enough to guarantee a low impact even for large incoming orders; (v) resilience: the market reacts quickly to possible order imbalances by means of an opposite order flow [22]. These features are linked, even if they capture different aspects of liquidity. In this work we will argue that the relative importance of these different aspects depends on the time scales in play and, in particular, we will focus on the *depth* and the *breadth* of the LOB for small time scales and on its *resilience* for large time scales, introducing suitable measures of liquidity. The importance of the latter is based on a number of empirical studies on the LOB, which singled out liquidity as the main driver of large price fluctuations [23]. In particular, the distribution of returns due to single transactions has been proved to be substantially equal to the distribution of the first gaps of the LOB [18]. Toth *et al.* [24] carried out an extensive study on the post-large-events dynamics that characterizes large price changes, finding evidence for power-law relaxations typical of complex systems [25].

In our work we will study which kind of LOB dynamics may eventually lead to large price changes. The aim is to identify those situations in which an intrinsic state of instability enhances the probability to have a price jump on the short run. The first step is to build a methodology to identify large events in an unbiased way. Then, we study the order flows around these events, finding peculiar patterns, especially at the best prices. The analysis of the statics and the dynamics of the LOB permits us to define in a quantitative way liquidity, whose imbalance is directly related to future price changes. Our focus will change with the time scales in play. In particular, order flow will turn out to be the main driver on large time scales

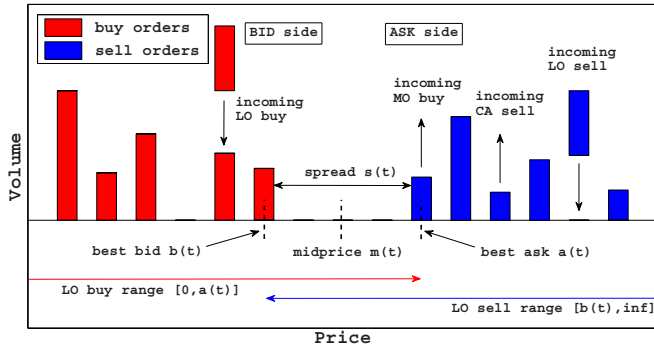*Corresponding author: and.zaccaria@gmail.com

FIG. 1. (Color online) Synthetic view of the different operations available on the electronic market. Incoming limit orders add liquidity to the book, while market orders and cancellations remove such availability of orders. The midprice $m(t)$ can change if, as a consequence of these operations, the best sell limit orders (ask) or the best buy limit orders (bid) change.

while the depletion of one side of the book will turn out to be more important on small time scales. In general, this empirical work confirms the idea that price changes, and in particular large fluctuations, are closely related to liquidity dynamics. To give the reader an idea of the potential role of liquidity we cite the fact that, for a large cap stock, the typically available liquidity is about 1% of the daily traded volume, which is, in turn, only 0.1% of the total capitalization [26,27].

Before proceeding let us briefly recall how a LOB works. The different elements we are going to introduce are depicted in Fig. 1. Traders have at their disposal three different *operations*: (i) they can place a *limit order* (LO), which is the proposal to buy or sell a certain amount (or *volume*) of shares at a fixed price. Since limit orders are usually not immediately traded, they are stored in the book until they can be matched with another order; (ii) they can decide to buy or sell a volume of shares at the best available price (these prices are called best *ask* and best *bid*, respectively), that is, place a so-called *market order* (MO), which will be matched with the best available limit order; (iii) *cancel* (C) a placed LO. Potentially, all three operations may have an impact on the price, which is usually defined as the value at which the last transaction occurred or the average between best bid and best ask (in this last case the price is called *midprice*). For example, an incoming MO to buy a given volume of shares could remove totally the best ask, triggering a price increment. Traders can place sell (buy) LOs also at a price that is lower (higher) than the best bid (ask): in this case one can see all or part of the LO as an *effective* MO, because it will trigger an immediate transaction. In the following, as it is usually done in the literature, we will call MOs the effective MOs. The price of a LO has to be chosen among a discrete set of prices. For example, a LO could be placed at 9.5 or 9.75 dollars, but not at 9.66. The minimum distance in price between two orders is called *tick*. In the previous example, the tick was equal to 0.25. This is the natural measure of distance among the placed limit orders, which constitute the limit order book.

Our empirical analysis is based on a database of various stocks traded at the London Stock Exchange. In particular, we analyzed four liquid stocks (AZN, BP, RBS, and VOD), but

we show, for reasons of space, only the results for AZN. We point out that we performed our analysis also on the other three stocks, finding similar results. Our database covers the whole year 2002. As it is usually done in the literature, we discarded the first 30 min of each trading day to avoid anomalous effects due to the opening of traders' positions.

II. UNBALANCED ORDER FLOWS DURING LARGE PRICE FLUCTUATIONS

A. Selection of large events

In this section we will be mainly interested in large price fluctuations, so we need in the first place a criterion to define a *large event*. A naive approach would select those time windows within which a return above a certain fixed threshold, let us say x%, occurred. However, such an *absolute* filter would be affected by the intraday volatility pattern and will select mostly events from the begin and from the end of the trading day. On the other hand, a *relative* filter, in which one selects those events that are larger than the average fluctuations at the same time of the day, will suffer from the effect of an opposite uneven selection. As a consequence, in this work we use a combination of the two filters, following Refs. [24,28,29]. If one chooses suitable thresholds for both filters, a uniform distribution of events over the day can be recovered.

The authors of the above-mentioned papers do not set a time interval for the large event, because a large return that passes both filters may be defined on different time scales. So they simply set the temporal reference frame in order to have $t = 0$ when both filters are passed. In this way, they can study only the post-event dynamics, since events have different temporal dimensions. Moreover, also a classification as a function of the return is troublesome: since the algorithm that searches for the large events stops when the thresholds are reached, all events will tend to have similar returns, since the only possible variation will be due to the U-shaped intraday volatility pattern. In order to avoid these drawbacks we use time windows of fixed length $\Delta t = 15$ min. In practice, for every operation in the LOB we check if a large price fluctuation is present in the next Δt. The specific value of Δt is chosen to have a reasonable compromise between two needs: to have enough statistics the time interval cannot be too large, but to be able to study the dynamics inside the event Δt cannot be too small. Obviously, by keeping the time windows' size constant we are treating in the same way events that could be very different, the only thing in common being the presence of a large net return between the start and the end of the considered time interval. However, this is the only way to make a pre-event analysis possible and to have a broad distribution of returns to analyze. We would like to stress that the heterogeneity of events is an intrinsic feature of this kind of analysis: for example, using the same criteria of Ref. [24], one could find events with a temporal extension up to 2 h. An important issue is the presence of strong autocorrelations in almost every time series of interest and, in particular, the volumes of incoming orders and the absolute values of returns. To avoid our analysis to be undermined by such effects, we discard all those events with another close event before, that is, we discard those events that are preceded by another large event at a distance Δt.

In practice, for the (relatively) large time scales analysis ($\Delta t = 15$ min) we select those time windows with a price fluctuation larger than $x = 0.5\%$ (in absolute value) and larger than $n = 3$ times the average volatility σ during that time of the day. After the application of the various filters we find 332 positive and 362 negative events for AZN, 358 positive and 381 negative events for BP, 306 positive events and 325 negative events for RBS, 402 positive and 436 negative events for VOD.

B. Order flows before and during large fluctuations

It is a well-known feature of LOBs that the arrival of a significant flow of MOs triggers a liquidity response in terms of LOs, which are usually placed close to the best price of the pushed side [19,30]. For example, if a persistent flow of MOs, such as the ones that result from the splitting of a large hidden order [31], is somehow detected, the other traders will likely provide more liquidity, effectively contrasting the incoming flow [this behavior is sketched in Fig. 2(a)]. As a consequence, the ratio of MOs to LOs remains constant even if the two volumes can fluctuate a lot, keeping the market resilient and contrasting large price fluctuations. In this section we will see that the break of this dynamical equilibrium is a key element of large price fluctuations. In fact, we will show that during large price movements the flow of LOs is no more able to stand in the way of the incoming MOs, as schematically depicted in Fig. 2(b). We note that, in our reasonings, cancellations have not been mentioned, even if their volume is not negligible. This is due to the fact that in the analyses we are going to present their relative flow does not change from normal times to time windows characterized by large price movements.

In order to check this line of reasoning, we empirically study the orders' flow before and during the large events

we defined in the previous section. First of all, we point out that during large events all flows exhibit large fluctuations. In particular, the side under pressure experiences an increase of LO and cancellations up to about three times, and of MO up to five times, with respect to the respective annual averages. In the following we will consider the relative flows of the three possible operations (LO, MO, and cancellations) with respect to their sum during large events. Since we will study one side of the book at a time and positive and negative large events separately, we will have four different relative flows for each operation. For example, the relative flow of LOs on the ask side, denoted by A, during large positive events, denoted by $+$, is calculated in the following way:

$$r_{\text{LO}}^{A+} = \left\langle \frac{Q_{\text{LOS}}}{Q_{\text{LOS}} + Q_{\text{MOB}} + Q_{\text{CS}}} \right\rangle_{A+} \tag{1}$$

where Q_{LOS} is the volume of sell LOs, Q_{MOB} is the volume of buy MOs, and Q_{CS} is the volume of the cancellations of sell LOs. So we consider the relative flows of all the operations that are made on the same side of the book during large events. The averages are performed by considering a total time range of 1 h from the end of the large event, but backwards, in such a way that 15 min belong to the large event and 45 min belong to the pre-event. We divide such a time range in subintervals of 30 s each, and we calculate the averages $\langle \ldots \rangle$ for each subinterval separately. The other relative flows are calculated in a similar way, always considering operation made on the same side of the book. In Fig. 3 we plot the relative flows of MOs (red triangles), LOs (blue squares), and cancellations (black dots) at the best bid and ask for both positive and negative events. The beginning of the large event is indicated by the vertical dashed line, while the horizontal line represents the relative flows calculated by averaging over the whole year. While cancellations remain roughly constant in all four plots, one can notice that the side of the book that is under pressure [that is, the ask side during positive events (a), and the bid side during negative events (b)] shows an excess flow of MOs that is not balanced by a corresponding increase of LO flow. The opposite behavior is evident for the pressing side of the book [(c) and (d)]. This situation is clearly related to the presence of large price fluctuations: since we have a large flow of MOs not opposed by an adequate flow of LOs, the MOs face a low resistance and can penetrate into the side under pressure of the order book, leading to a large price movement.

This situation is, at a first glance, in sharp contrast with the results presented in Ref. [24], where the relative flows remain constant before and after the large event. Moreover, the relative number of MOs and cancellations are practically inverted; while in Ref. [24] one can see a relative number of cancellations roughly equal to 0.35 and a relative flow of MOs around 0.05, in Fig. 3 we observe approximately 0.32 for MOs and 0.15 for cancellations. While the discrepancy between the average values is due to our choice to show only the flows at the best bid and ask, we believe that the sharp increase of MOs and the decline of LOs are visible only if one averages over a coherent set of events. In fact, while we keep our time windows fixed, Toth *et al.* consider events of very different duration, and time windows which could span, in principle,
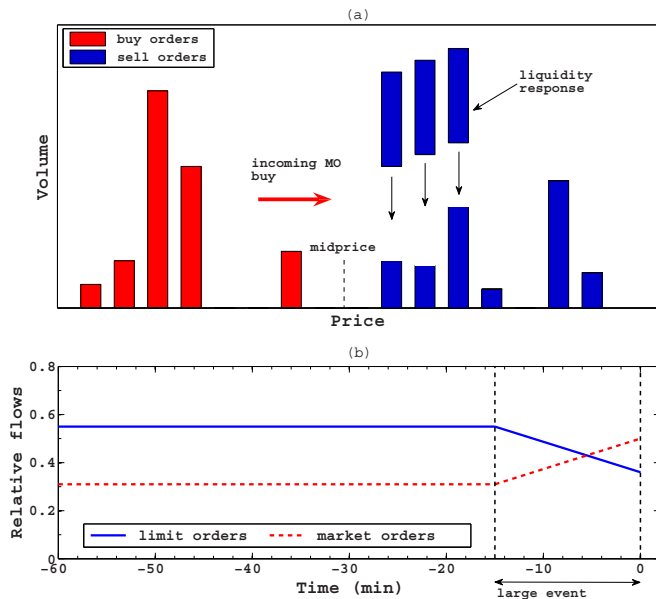

FIG. 2. (Color online) (a) Sketched representation of the reaction, in terms of placement of limit orders, to an incoming flow of market orders. (b) We have a large event when the relative flows change, breaking the dynamical equilibrium.
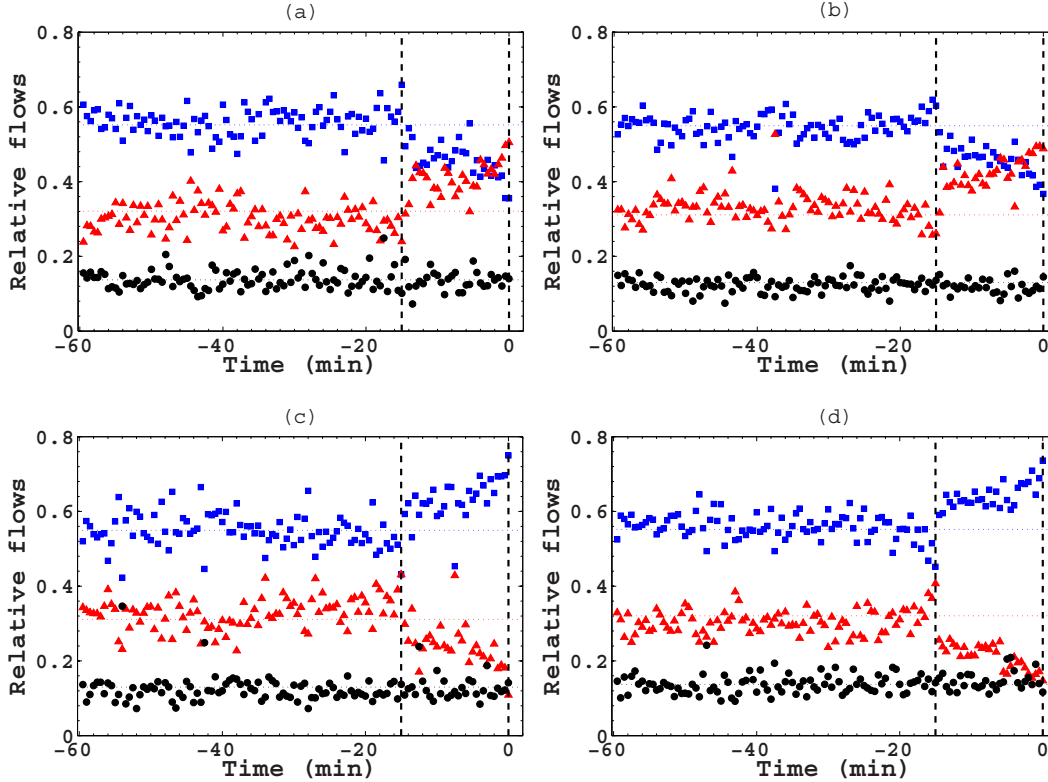
FIG. 3. (Color online) Relative volumes of operations at the best price before and during large price fluctuations, defined in time windows of 15 min. On top, the two sides under pressure: the ask side for positive events (a) and the bid side for negative events (b). The increase of market orders (represented by red triangles) is not balanced by an opposite response of limit orders (blue squares). At the bottom, we show the bid side during positive events (c) and the ask side during negative events (d). The pressing side experiences the opposite situation. Cancellations (black circles) remain constant in all four cases.

from a few minutes up to 2 h. Moreover, we point out that the increase of MOs and the decrease of LOs is much less evident, even if still clearly visible, when one selects the large events by means of our filter but considers the whole book instead of the bests.

C. Lack of display of latent liquidity

The previous analysis suggests that large price fluctuations are due to an unexpected lack of response in terms of liquidity to an incoming MO flow, that is, a loss of *resilience*. In order to confirm this view we studied how the number of LOs placed at the best price varies as a function of the number of the incoming MOs. Let us focus on the ask side of the book (the same results are valid for the bid side). Since we find a clear correlation between the volumes of MOs to buy Q_{MOB} and LOs to sell Q_{LOS} placed at the best ask in the same time window of size $\Delta t = 15$ min, we can use, as a first-order approximation, a linear relationship of the type $Q_{LOS} \approx a Q_{MOB} + b$, where a and b can be seen as proxies for market's resilience. We point out that at this stage of our analysis, both flows being calculated on the same time window, there is no clear causal relationship between the two, even if one usually thinks of the placement of LOs close to the best as a reaction to an incoming MOs flow. To make cause-effect relationships come to light we need a finer temporal resolution, which will be adopted in the following sections.

Now we study how the linear relationship between the volumes of MOs and LOs changes during the large price fluctuations. To do so we consider the order flows Q_{LOS} and Q_{MOB} in each time window, but selecting the latter in three different ways: during normal times, large positive and large negative events. We come up with three sets of points in the Q_{LOS}-Q_{MOB} plane. In Fig. 4 we plot both the parametric and the nonparametric (binning) fits of the three sets of points, finding that the response in terms of sell LOs on the ask side of the book is much less than the average during large positive events. The black dots represent the average response, which is measured considering all time windows. The linear relationship is clear. The blue squares represent the same calculation, but considering only that subsample of windows that have been classified as large positive events, that is, the ones in which the ask side is under pressure and does not resist to the incoming flow of MOs. One can see that both the coefficient a of the linear approximation and the absolute mean values of LOs, given a similar flow of MOs, are lower. On the contrary, during the large negative events (red triangles), in which the considered side of the book is pressing, one can see an excess placement of LOs. We find very similar results for the bid side of the book and the relationship between Q_{LOB} and Q_{MOS}.

This is a direct, empirical evidence of the relationship between large price fluctuations and traders' operations. It is important to recall herein that most of the liquidity of the
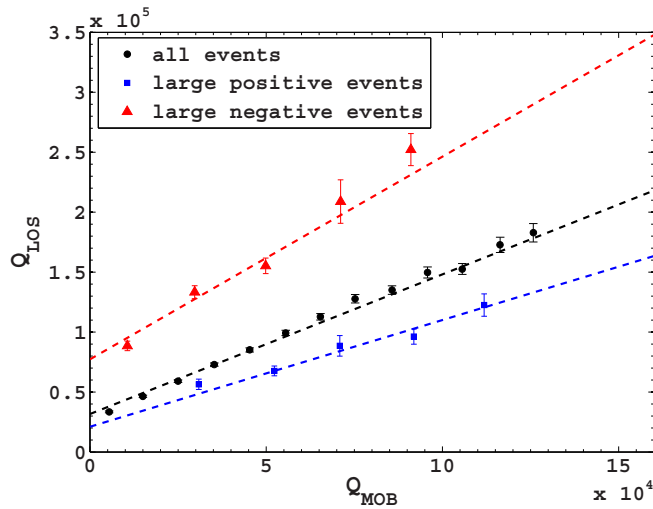
FIG. 4. (Color online) Comparison between market and limit order flows, on the ask side of the book, on normal times and during large price movements. With respect to the average (black dots), during large positive events (blue squares) there is a lack of display of hidden liquidity, namely, too few limit orders to counterbalance the market orders flow. The opposite behavior can be seen during large negative events (red triangles). A similar analysis can be performed on the bid side of the book, finding analogous results.
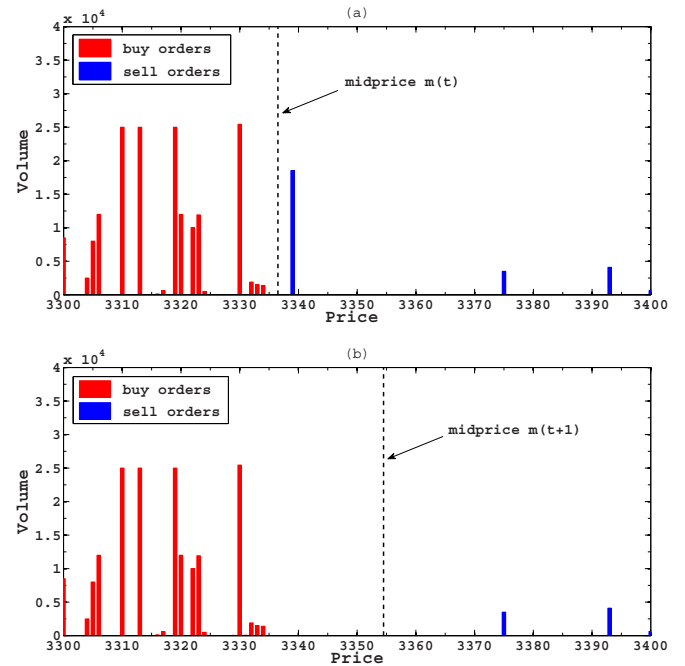

FIG. 5. (Color online) Large price fluctuations are driven mostly by the presence of large gaps in the LOB. (a) A situation of intrinsic fragility of the ask side book. (b) The consequence of an incoming MO, which triggers a large midprice change.

book is *hidden*, in the sense that most traders prefer not to place their limit orders until a clear flow of market orders is detected. Only in this case, and proportionally, additional liquidity is provided. We have seen that during large events this mechanism of automatic compensation stops: a big price jump is therefore related to a lack of revealed liquidity (which is instead present during normal times) and which makes the market less resilient. Similar concepts have been used by Toth *et al.* [27], who proposed a model to investigate the relationship between revealed liquidity and the empirical price impact function, introducing the concept of *latent* order book.

III. DEPLETION OF THE LIMIT ORDER BOOK ON SHORT TIME SCALES

The main empirical properties of financial markets, the so called stylized facts [1,2], are present on very different time scales. In particular, the log returns $r_{\Delta t} = \log p(t + \Delta t) - \log p(t)$ are power-law distributed for values of Δt ranging from seconds to months. Only if one considers single operations this scale invariance breaks down, as expected, because of discreteness effects. In this section we investigate in which terms liquidity crises may trigger large price fluctuations at smaller time scales. We will consider time windows and returns taking *events* of $\Delta t = 30$ s and the state of the LOB just before them. Practically, we study the condition of the LOB at the moment of the last operation before the beginning of the time window that defines the event. We point out that in this way we neglect the orders' flow inside the time windows, adopting a different approach with respect to the analysis we performed in the previous sections. In fact, we believe that on smaller time scales the static structure of the book is the

key element to determine the magnitude and the sign of price movements.

The smallness of the considered time intervals permits us to study the condition of the LOB before the large price fluctuations with a significant number of events at our disposal. First of all, we check if some average property of the LOB is changed before the large events. A natural observable is the LOB profile, that is, the volume present in the form of available LOs as a function of the distance, expressed in price units or ticks, from the respective bests, usually averaged over a certain time interval [32]. Bouchaud *et al.* [33] have found that the LOB profile has an universal shape, even if large temporal fluctuations are present. The study of these fluctuations concerns, using the terminology of Sarr and Lybek [21], the breadth and the depth of the order book, instead of its resilience, which is a property closer to the ones we analyzed in the previous section. In order to illustrate how the fluctuations of the order book profile can influence the magnitude of price movements we consider a real situation, showed in Fig. 5, and in particular we focus on the structure of the LOB before [Fig. 5(a)] and after [Fig. 5(b)] the arrival of a MO. The reconstruction of the order book refers to the stock AZN on 26/04/2002, at time 13:31:26 and 13:31:27, respectively. Before the transaction, we have a different breadth of the two sides of the book. In fact, while the bid side is dense and able to absorb a discrete amount of incoming MOs without affecting the midprice in a relevant way, the ask side is far more fragile, because the presence of large gaps between the order indicate that the arrival of sufficiently large MOs could trigger large movements of the midprice, that is what happens just a second later [Fig. 5(b)].
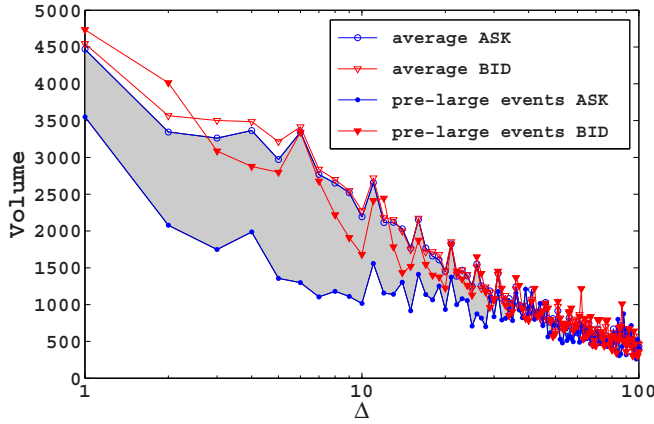
FIG. 6. (Color online) Comparison between the average order book profile with the average profile before a large positive price fluctuation. While the bid profile before the event is roughly similar to the average ask and bid profiles, the ask profile, that is, the side which is going to break down, is much less liquid.

In order to explore this idea in a systematic way, we select the large events with the same methodology we used before, but considering $\Delta t = 30$ s and $x = 0.3\%$ and $n = 6\sigma$ as thresholds for the absolute and relative filters, respectively. Furthermore, we discard those events that are preceded by another large event within 90 s. Again, we consider only large positive events, but our results are qualitatively similar for the negative ones. In Fig. 6 we plot the average profile (empty blue circles for the ask, empty red triangles for the bid) of the two sides of the LOB in comparison with the same sides just before the large positive fluctuations, that is, their state at the beginning of the considered time window (full blue circles for the ask, full red triangles for the bid). One can see that, while the pressing side (which is, for positive returns, the bid side, red triangles) remains substantially unchanged, the side that is about to break down (in this case the ask, blue circles) is characterized by substantially lower volumes, especially near the best. We can state that there is an intrinsic instability of that side of the book that can cause large price jumps even if the incoming flow of MO is not significantly higher than the average. In the following section we introduce a quantitative framework to measure the breadth and the depth of the book and the correlation between the lack of liquidity and the ensuing returns.

A. A microstructural definition of liquidity

As already noted in Refs. [18,23], the presence of gaps in the LOB plays a key role in enhancing the impact of MOs on the price. In particular, the authors have shown that the distribution of returns is to a large extent coincident with the one of the gaps, if one considers placement of single orders. However, when different time scales are considered, also the amount of shares and their distribution at different prices must be taken into account. In order to cope with these needs we study the breadth and depth of the LOB, which we quantify calculating the *exponential liquidity*

$$L(\delta) = \frac{1}{\langle V_N \rangle} \sum_{\Delta=1}^{N} V(\Delta) \exp(-\Delta/\delta), \quad (2)$$
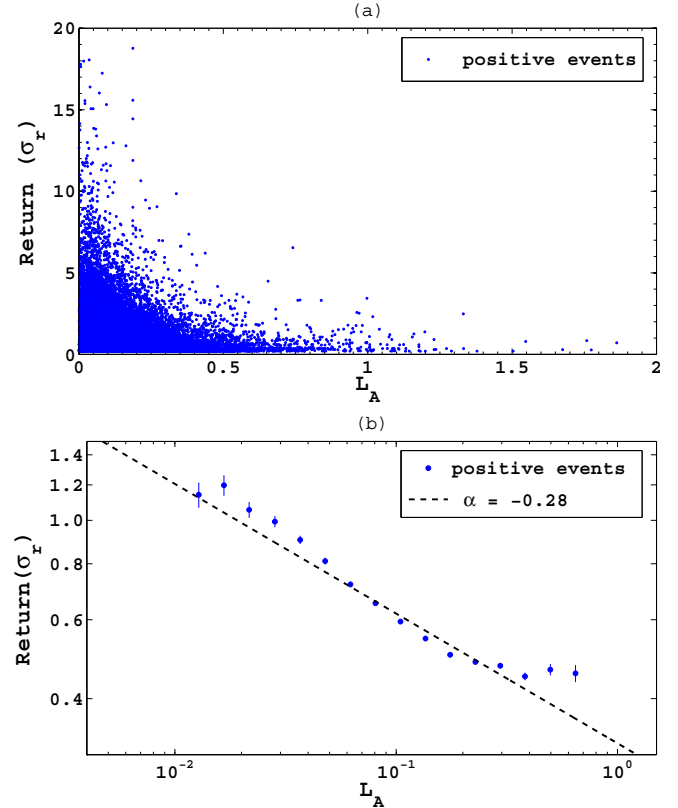

FIG. 7. (Color online) (a) Returns, measured as multiples of the average standard deviation, and liquidity of the ask side of the book for all time windows with positive returns. High returns are present most probably during liquidity crises. (b) Power-law fit and binned version of the cloud of points. The clear deviation at high values of liquidity is due to the presence of a minimum value for the possible returns.

where $\langle V_N \rangle$ is the average volume in the book present within a maximum distance of $N = 100$ ticks from the best, $V(\Delta)$ is the volume placed at tick Δ, where $\Delta = 1$ corresponds to the best price, and δ is a parameter that regulates the exponential weights, defining the relative importance of the orders. The liquidity is calculated on one side of the book and at the beginning of the time window. In the following we will indicate the liquidity that is computed on the ask side with L_A and the one computed on the bid side with L_B. To study the correlation between the state of the LOB and the following return we plot the latter and the ask-side liquidity for all the time windows containing positive returns. The results are presented in Fig. 7(a), where each point corresponds to a different time window of 30 s and the returns are expressed in units of the average standard deviation of the year. The same analysis can be performed using negative events and the bid side of the book, obtaining analogous results. One can see that the number of events prevents a visual inspection of the core and of the distribution of the resulting cloud of points. However, it is evident that high returns occur only if liquidity is low and high values of liquidity prevent large fluctuations. In order to investigate this result in a quantitative way we fit the whole set of points with a power law. We compare this fit with a logarithmic binning in Fig. 7(b). The deviation for
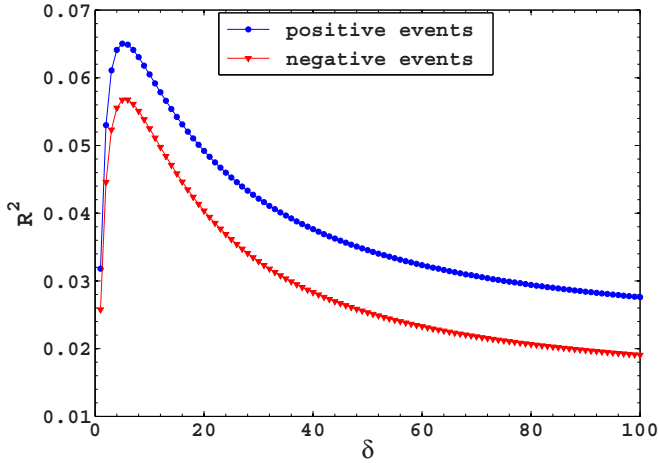
FIG. 8. (Color online) Coefficient of determination relative to the power-law fit of returns versus liquidity for positive (blue dots) and negative (red triangles) events, as a function of the parameter δ, the characteristic distance of the exponential liquidity. The correlation is higher for values of δ around 5–6.

small return values is due to the presence of a minimum value of the possible returns, which is a consequence of the discrete nature of the order book. The minimum return is equal to one tick; this effect is less evident, but present, also in Fig. 7(a).

We have verified that these results are stable for different values of the parameters in Eq. (2), as long as N is large enough. The characteristic distance δ can be chosen, for instance, to maximize the goodness of fit, as measured by R^2 [34]. In Fig. 8 we plot the values of R^2 we obtain when we use different values of δ to calculate the exponential liquidity. One can notice a maximum around $\delta = \delta^\star \approx 5$–6 and then a sharp decline. This is a clear evidence that, for the considered time scale, the volumes near the best influence more the subsequent returns, but all the book must be considered. In a further analysis, which we do not show for reasons of space, we have found that the value of $\delta^\star$ increases with the length of the time window Δt. Obviously, if Δt is too large, the correlation disappears, because at longer time scales price fluctuations are mostly driven by order-flow effects, studied in the previous sections, rather than a static depletion of the book (to give an idea, in a time window of 15 min hundreds of orders are placed or canceled).

At a first look the values of the correlation coefficient could seem too little to be meaningful. However, one has to consider that the number of points is large enough to make the results of our fit statistically significant, as can be seen if one performs a t test and checks the resulting p value. A direct consequence is the low errors associated to the estimated parameters. Assuming a relationship of the kind $r(t) = K L_A^{-\alpha}$ one finds $K = 0.2944 \pm 0.0028$ and $\alpha = 0.2800 \pm 0.0041$ for positive events. The analysis of the negative events gives values of A and α compatible within error bars.

Recently, Yura *et al.* [35] introduced a new mathematical framework to investigate the microstructural dynamics of stock prices, seen as colloids that interact with the various layers of the LOB. In this work we have found qualitatively compatible results, namely, that orders placed near the best should weigh

more. In particular, they use larger time windows (on average) and find the largest correlation at a higher distance from the best with respect to our $\delta^\star \approx 5$–6. Even if a systematic comparison is made difficult by the different datasets we use and by the different definitions of both time windows and liquidity, since we observe an increase of $\delta^\star$ with Δt we can reasonably expect the two analyses to be compatible.

We end this section noticing that, in general, many possible empirical measures of liquidity are possible. According Aitken and Comerton-Forde, the liquidity measures used in literature are actually 68 (see Ref. [36] and references therein), but many of them are widely discussed, leading Goyenko *et al.* to the choice of a meaningful title for their review of the subject: "Do liquidity measures measure liquidity?" [37]. A systematic comparison of some measures of liquidity with the ones we introduce in this article will be presented in a forthcoming work.

B. Liquidity imbalance

In the previous section we measured the liquidity considering only one side of the book and conditioning our analysis to the presence of a future return of a given sign. As a consequence, we left aside returns exactly equal to zero (which are, especially for small time windows, a considerable fraction of the total) and we did not study any possible predictive power on the sign that this kind of approach could have. In order to define a microstructural property of the book that takes into account the static asymmetry between the volumes placed on both sides we define the *liquidity imbalance* as

$$L_{\text{imb}}(\delta) = \frac{L_B(\delta) - L_A(\delta)}{L_B(\delta) + L_A(\delta)}, \tag{3}$$

where $L_A(\delta)$ is the exponential liquidity defined in Eq. (2) and calculated for the ask side, and $L_B(\delta)$ refers instead to the bid side. As before, we measure these quantities just before the beginning of the time window. Given the definition of Eq. (3),
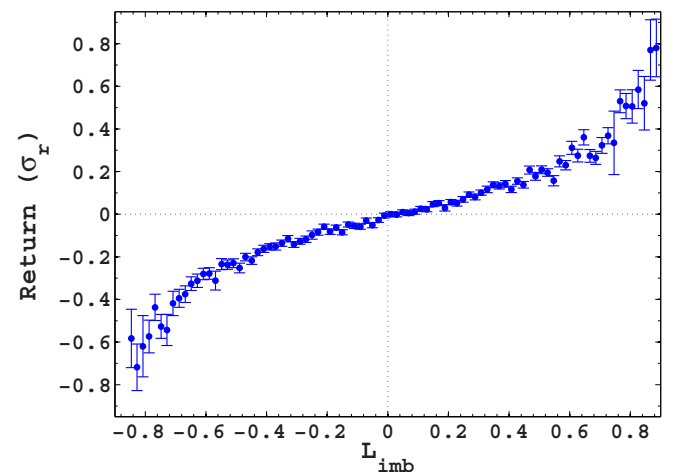

FIG. 9. (Color online) Average return given the liquidity imbalance. The return is measured over all the time windows (which are 30 s long) and the imbalance at the beginning of them. We find a monotonic, nonlinear dependence.
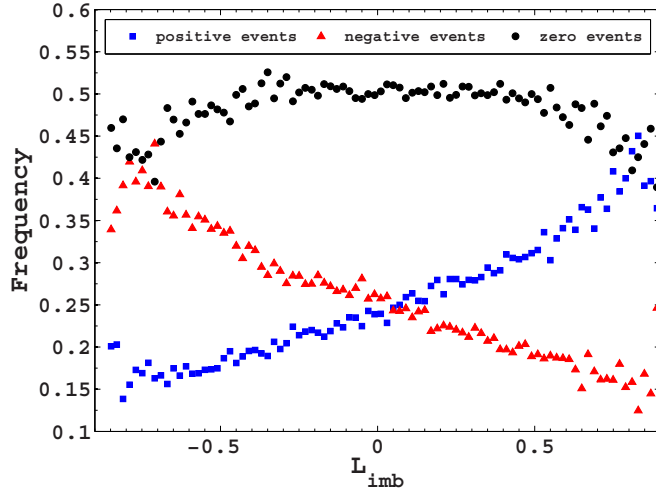
FIG. 10. (Color online) Relative frequencies of return signs as a function of the liquidity imbalance defined by Eq. (3). The imbalance is calculated at the beginning of the time window, which defines the price movement. Positive liquidity imbalances are often associated with positive returns, and vice versa. For instance, when $L_{imb} \approx 0.8$, one finds a number of positive returns, that is roughly double than the number of the negative ones.

we expect that the state of the book facilitates positive returns when $L_{imb}(\delta)$ is positive, and vice versa. If one takes $\Delta t = 30$ s and $\delta = 5$ ticks one obtains a cloud of points in the L_{imb}-returns plane. A nonparametric fit of this set of points is depicted in Fig. 9. Each point represents the average over all the returns that occur on all time windows that follow a given liquidity imbalance. The bars are given by two times the standard error of the average value. We remind that each time window is 30 s long and the liquidity is computed at the beginning of the time interval. One finds a clear, nonlinear correlation between the liquidity imbalance and the return in the following time window.

In Fig. 10 we repeat the same analysis but considering the frequency of three possible outcomes for the return that follows the LOB state as measured by the liquidity imbalance: positive (blue squares), zero (black dots), and negative returns (red triangles). Even if about half of the events have zero return, it is evident that different states of the book lead to different frequencies of subsequent return signs. In particular, high values of L_{imb} are clearly associated to positive returns. For example, if one takes a state of the book in which $L_{imb} \approx 0.8$, the frequency of the positive returns is more than double the frequency of the negative ones. To assess the predictive power of this kind of analysis one should measure this effect on a training set of events and then check its ability to forecast

future price movements on a test set. We aim to perform this kind of test in a future work.

IV. CONCLUSIONS AND DISCUSSION

In this work we present an empirical study of the static and dynamic properties of the limit order book and their relationship with price movements. In particular, we study two effects, both connected to the concept of liquidity, which are present on different time scales. We divide our time series in time windows of length $\Delta t = 15$ min and 30 s. We find that on (relatively) large time scales large price fluctuations are connected to a low resilience of the market, namely, a lack of response, in terms of limit orders, to an incoming flow of market orders. Obviously, at this stage of analysis we are not able to claim a direct cause-effect relationship between order flow imbalance and price jumps, because we do not have a time delay between the two and we do not perform the opposite analysis; that is, we do not consider all the events with a clear order flow imbalance. On smaller time scales, instead, we adopt a static approach, in which the state of the book at the beginning of the time window is analyzed. A suitable measure of liquidity, and in particular of the breadth and the depth of the market, is presented. We find that the liquidity present on one side of the book is correlated with the magnitude of the immediately following return. Moreover, the liquidity imbalance between the two sides of the book is clearly related with the sign of this return.

Such an analysis can be inserted in the wider field that studies the lack of linearity and simple cause-effect mechanisms in financial markets. In fact, we find that liquidity crises can enhance in a strongly nonlinear way the predisposition toward instability of markets. This behavior defines a stylized fact that has to be taken into account in both theoretical modeling and practical applications. In order to cope with these effects we suggest considering measures of liquidity like the ones we proposed in practical situations, such as the evaluation of the expected price impact of an order. This can be accomplished, for example, by considering a two-dimensional price impact function, like the one introduced in Ref. [38], which explicitly takes into account the impact dependence on both volume and liquidity.

[1] J. P. Bouchaud and M. Potters, *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management* (Cambridge University Press, New York, 2003).

[2] R. N. Mantegna and H. Stanley, *An Introduction to Econophysics: Correlation and Complexity in Finance* (Cambridge University Press, New York, 2000).

[3] L. Harris, *Trading and Exchanges* (Oxford University Press, Oxford, 2003).

[4] J. Hasbrouck, *Empirical Market Microstructure* (Oxford University Press, Oxford, 2007).

[5] J.-P. Bouchaud, J. D. Farmer, and F. Lillo, *How Markets Slowly Digest Changes in Supply and Demand.* (Elsevier/Academic Press, New York, 2008).

[6] I. N. Lobato and C. Velasco, J. Bus. Econ. Stat. **18**, 410 (2000).

[7] F. Lillo and J. D. Farmer, Studies Nonlin. Dynam. Economet. **8**(3), Article 1 (2004).

[8] J. P. Bouchaud, Y. Gefen, M. Potters, and M. Wyart, Quant. Fin. **4**, 176 (2004).

[9] E. F. Fama, J. Fin. **25**, 383 (1970).

[10] P. Weber and B. Rosenow, Quant. Fin. **5**, 357 (2005).

[11] P. Weber and B. Rosenow, Quant. Fin. **6**, 7 (2006).

[12] L. Gillemot, J. D. Farmer, and F. Lillo, Quant. Fin. **6**, 371 (2006).

[13] A. Joulin, A. Lefevre, D. Grunberg, and J.-P. Bouchaud, arXiv:0803.1769.

[14] B. Biais, P. Hillion, and C. Spatt, J. Fin. **50**, 1655 (1995).

[15] M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison, Quant. Fin. **13**, 1709 (2013).

[16] A. Zaccaria, M. Cristelli, V. Alfi, F. Ciulla, and L. Pietronero, Phys. Rev. E **81**, 066101 (2010).

[17] T. Chordia, R. Roll, and A. Subrahmanyam, J. Fin. Econ. **65**, 111 (2002).

[18] J. D. Farmer, L. Gillemot, F. Lillo, S. Mike, and A. Sen, Quant. Fin. **4**, 383 (2004).

[19] A. Mendoza, H. Waelbroeck, and C. R. Stephens, available at SSRN 1511205 (2009).

[20] R. Cont, A. Kukanov, and S. Stoikov, J. Fin. Econ. **12**, 47 (2013).

[21] A. Sarr and T. Lybek, *Measuring Liquidity in Financial Markets*, IMF working paper 02/232 (International Monetary Fund, 2002).

[22] These features could be referred to the whole market, but in the following we will focus not on market liquidity in general but on the liquidity of a single stock.

[23] F. Lillo and J. D. Farmer, Fluct. Noise Lett. **05**, L209 (2005).

[24] B. Tóth, J. Kertész, and J. D. Farmer, Eur. Phys. J. B: Condens. Matter Complex Syst. **71**, 499 (2009).

[25] A. Ponzi, F. Lillo, and R. N. Mantegna, Phys. Rev. E **80**, 016112 (2009).

[26] J.-P. Bouchaud, The Endogenous Dynamics of Markets: Price Impact, Feedback Loops and Instabilities, in *Lessons from the 2008 Crisis*, edited by A. Berd (Risk Books, Incisive Media, London, 2011).

[27] B. Toth, Y. Lemperiere, C. Deremble, J. de Lataillade, J. Kockelkoren, and J.-P. Bouchaud, Phys. Rev. X **1**, 021006 (2011).

[28] A. Zawadowski, J. Kertész, and G. Andor, Physica A **344**, 221 (2004).

[29] A. G. Zawadowski, G. Andor, and J. Kertesz, Quant. Fin. **6**, 283 (2006).

[30] B. Toth, Z. Eisler, F. Lillo, J. Kockelkoren, J.-P. Bouchaud, and J. D. Farmer, Quant. Fin. **12**, 1015 (2012).

[31] E. Moro, J. Vicente, L. G. Moyano, A. Gerig, J. D. Farmer, G. Vaglica, F. Lillo, and R. N. Mantegna, Phys. Rev. E **80**, 066102 (2009).

[32] D. Challet and R. Stinchcombe, Physica A **300**, 285 (2001).

[33] J.-P. Bouchaud, Physica A **313**, 238 (2002).

[34] In Fig. 7 we have used those values of exponential liquidity which corresponds to a choice of δ which maximizes R^2.

[35] Y. Yura, H. Takayasu, D. Sornette, and M. Takayasu, Phys. Rev. Lett. **112**, 098703 (2014).

[36] M. Aitken and C. Comerton-Forde, Pacific-Basin Fin. J. **11**, 45 (2003).

[37] R. Y. Goyenko, C. W. Holden, and C. A. Trzcinka, J. Fin. Econ. **92**, 153 (2009).

[38] M. Cristelli, V. Alfi, L. Pietronero, and A. Zaccaria, Eur. Phys. J. B **73**, 41 (2010).