

Anomalous price impact and the critical nature of liquidity in financial markets

B. Tóth^a, Y. Lempérière^a, C. Deremble^a, J. de Lataillade^a,
J. Kockelkoren^a, J.-P. Bouchaud^a

November 2, 2011

^a *Capital Fund Management, 6, blvd Haussmann 75009 Paris, France*

Abstract

We propose a dynamical theory of market liquidity that predicts that the average supply/demand profile is V-shaped and *vanishes* around the current price. This result is generic, and only relies on mild assumptions about the order flow and on the fact that prices are, to a first approximation, diffusive. This naturally accounts for two striking stylized facts: first, large metaorders have to be fragmented in order to be digested by the liquidity funnel, leading to long-memory in the sign of the order flow. Second, the anomalously small local liquidity induces a breakdown of linear response and a diverging impact of small orders, explaining the “square-root” impact law, for which we provide additional empirical support. Finally, we test our arguments quantitatively using a numerical model of order flow based on the same minimal ingredients.

1 Introduction

Price impact refers to the correlation between an incoming order (to buy or to sell) and the subsequent price change [1, 2, 3]. That a buy (sell) trade should push the price up (down) is intuitively obvious and is easily demonstrated empirically (see [3] for a recent review). Such a mechanism must in fact be present for private information to be incorporated into market prices. But it is also a sore reality for large trading firms for which price impact induces extra costs. Indeed, large volumes must typically be fragmented and executed incrementally. But since each “child order” pushes the price up or down, the total cost of the “metaorder”¹ is quickly dominated, as sizes become large, by the average price impact. Monitoring and controlling impact has therefore become one of the most active domain of research in quantitative finance since the mid-nineties. A huge amount of empirical results has accumulated over the years, in particular concerning the relation between the total size Q of the metaorder and the resulting average price change. These empirical results come either from proprietary trading strategies (and are often not published), or from brokerage firms, who execute on behalf of clients [4, 5, 6, 7, 8], or else from the exchanges, who give exceptional access to identification codes that allow one to reconstruct the metaorders from some market participants [9, 10]. Remarkably, although these data sets are extremely heterogeneous in terms of markets (equities,

¹We call the “metaorder” (or parent order) the bundle of orders corresponding to a single trading decision. A metaorder is typically traded incrementally through several child orders.

futures, FX, etc.), epochs (from the mid-nineties, when liquidity was provided by market makers, to present day’s electronic markets), market participants and underlying trading strategies (fundamental, technical, etc.) and style of execution (using limit or market orders, with high or low participation ratio, etc.), a very similar concave impact law is reported in most studies. More precisely, the average relative price change Δ between the first and the last trade of a metaorder of size Q is well described by the so-called “square-root” law:

$$\Delta(Q) = Y\sigma\sqrt{\frac{Q}{V}}, \quad (1)$$

where σ is the daily volatility of the asset, and V the daily traded volume, both quantities measured contemporaneously to the trade. The numerical constant Y is of order unity. Published and unpublished data suggest slightly different versions of this law; in particular the \sqrt{Q} dependence is more generally described as a power-law relation $\Delta(Q) \propto Q^\delta$, with δ in the range 0.4 to 0.7 [4, 5, 9, 6, 7, 8]. For example, using a large data sample of 700,000 US stock trade orders executed by Citigroup Equity Trading, Almgren et al. [4] extract $\delta \approx 0.6$. Moro et al. [9] report $\delta \approx 0.5$ for trades on the Madrid stock exchange and $\delta \approx 0.7$ for the London stock exchange. We show in Fig. 1 our own proprietary data corresponding to nearly 500,000 trades on a variety of futures contracts, which yields $\delta \approx 0.5$ for small tick contracts and $\delta \approx 0.6$ for large tick contracts, for Q/V ranging from a few 10^{-4} to a few %. Our data on stocks is also compatible with $\delta \approx 0.5$, although more noisy. We note that all these studies differ quite significantly in the details of (a) how the price impact Δ is defined and measured; (b) how different assets and periods are collated together in the analysis; (c) how the fit is performed: over what range of Q/V , adding an intercept or not, etc. But in spite of all these differences and those mentioned above – in particular concerning the strategies motivating the trades –, it is quite remarkable that the square-root impact law appears to hold approximately in all cases.

The aim of the present paper is to provide a theoretical underpinning for such a universal impact law. We first give a general dynamical theory of market liquidity that predicts that the average supply/demand profile is V-shaped around the current price. The anomalously small local liquidity induces a breakdown of linear response and explains the square-root impact law. We then study numerically a stylized model of order flow based on minimal ingredients. The numerical results fully support our analytical arguments, and allow us to get quantitative insights into various aspects of the problem.

2 An intriguing impact law

One should first carefully distinguish the total impact of a given metaorder of size Q from other measures of impact that have been reported in the literature. One is the immediate impact of an *individual market order* of size q , which has been studied by various authors and is also strongly concave as a function of q , i.e. q^α with $\alpha \approx 0.2$, or even $\ln q$ [3]. Another easily accessible measure of impact is to relate the average price change Δ_T in a given time interval T to the *total* market order imbalance \mathcal{Q}_T in the same time period, i.e. the sum of the signed volumes of all market orders. This quantity is estimated using all the trades in the market (i.e. those coming from different market participants) and is clearly different from the impact of a given metaorder (see below). However, there seems to be quite a bit of confusion in the literature and many authors unduly identify the two quantities. If T is very short, such that there are only one or a few trades, one essentially observes the concave impact of individual orders that we just mentioned. But as T increases, such as the number of trades becomes large, the relation between Δ_T and \mathcal{Q}_T becomes more and more *linear* for small imbalances (see e.g. [3], Fig. 2.5), and on time scales comparable to those needed to complete a metaorder, the




Figure 1: The impact of metaorders for CFM proprietary trades on futures markets, in the period June 2007 – December 2010. Impact is measured here as the average execution shortfall of a metaorder of size Q . The data base contains nearly 500,000 trades. We show Δ/σ vs. Q/V on a log-log scale, where σ and V are the daily volatility and daily volume measured the day the metaorder is executed. The blue curve is for large tick sizes, and the red curve is for small tick sizes. For large ticks, the curve can be well fit with $\delta = 0.6$, while for small ticks we find $\delta = 0.5$. For comparison, we also show the lines of slope 0.5 (corresponding to a square-root impact) and 1 (corresponding to linear impact). We have removed a small positive intercept $\Delta/\sigma = 0.0015$ for $Q = 0$, which is probably due to a conditioning effect.

concavity has almost disappeared, except in rare cases when Q_T/V is large, in any case much larger than the region where Eq. (1) holds.

A square-root singularity for small traded volumes is highly non-trivial, and certainly not accounted for in Kyle’s classical model of impact [11], which predicts a linear impact $\Delta \propto Q$. A concave impact function is often thought of as a saturation of impact for large volumes. We believe that the emphasis should rather be placed on the *anomalous high* impact of small trades. Numerically, Eq. (1) means that trading one hundredth of the daily volume moves the price by a tenth of its daily volatility, which is indeed a huge amplification. Mathematically, Eq. (1) implies that marginal impact diverges for small volumes as $Q^{-1/2}$, meaning that the susceptibility of the market to trades of vanishing size is formally infinite. In most systems, the response to a small perturbation is linear, i.e. small disturbances lead to small effects. The breakdown of linear response often implies that the system is at, or close to, a critical point, where very special properties emerge, such as long-range memory or scale invariant avalanches, that accompany this diverging susceptibility. Of course, the mathematical divergence is cut-off in practice – for one thing, the volume Q of a metaorder cannot be smaller than a single lot. Empirical data will never be in the asymptotic limit $Q/V \rightarrow 0$, but this is irrelevant to our discussion. This is in fact also the case for most physical systems for which critical behaviour is observed. The important point here is that the proximity of a critical point can lead to strongly non-linear effects and extreme fragility. As we will argue in detail below, and substantiate within a precise numerical model, we believe that markets operate in a critical regime where liquidity vanishes. This offers a framework to understand many of the anomalies in the behaviour of markets, including the long term memory in order flow and the presence of frequent unexplained jumps in prices, that are – or so we believe – a consequence of the chronic lack of liquidity that lead to micro-crisis. The anomalous high impact of small trades implied by the concave impact law, Eq. (1), is, in our view, another side of the same coin.

Numerous interpretations have been put forth to explain a concave impact law, and can be broadly classified into three types of mechanisms (which are not necessarily exclusive): a) risk-reward optimisation of the liquidity providers/market makers, [12, 13, 14, 15, 16]; b) surprise in order flow and decay of impact [3, 17] and c) locally linear supply profiles [18]. Many of the above stories require that liquidity providers know the fundamental value, carefully monitor the order flow, identify metaorders and adjust their quotes such that they eke out some profit or at least break even. There is no doubt that a fraction of market participants strive to achieve such goals, and develop astute algorithms in this aim. However, liquidity providing is not (anymore) the monopoly of these market participants, who compete with anyone placing limit orders as part of an execution strategy, and the concave impact law in fact holds even when a substantial fraction of the metaorder is executed using limit orders. That individual metaorders can be detected using statistical methods, as advocated in [17], may well be true for large metaorders of unsophisticated traders, but it would be surprising (although not impossible) that our own trades, which are used to obtain Fig. 1, can be systematically detected. The universality of the concave impact law suggests that a robust self-organizing mechanism is at play. Our thesis, that we will substantiate below using both analytical arguments and numerical simulations, is that one can indeed build a theory of impact that relies on minimal assumptions, with no reference whatsoever to notions such as fundamental prices, market maker profit, or adverse selection.

3 A dynamical theory for linear supply/demand profiles

An interesting idea is that the supply/demand curve is a growing function of the difference between the fundamental value p_0 and the price. More precisely, suppose the available volume for sells (resp. buys) at price p or above (below), $\mathcal{V}_{\pm}(p)$, is a linear function, $\pm b_{\pm}(p - p_0)$. The execution of a volume Q of buy orders must therefore take place by moving the price by a quantity Δ such that:

$$Q = \int_{p_0}^{p_0+\Delta} dp \mathcal{V}_+(p) = \frac{b_+}{2} \Delta^2, \quad (2)$$

and similarly for sells. Therefore, if the supply/demand profile is *linear*, the impact is a square-root: $\Delta_{\pm} \approx \sqrt{2Q/b_{\pm}}$. It is indeed tempting to wave hands and argue that the available volume is proportional to the mispricing $|p - p_0|$, but one soon hits a major impediment: what exactly is p_0 ? If the above argument is to make any sense, p_0 is the fundamental value when the metaorder starts trading, and the assumption is that p_0 should not vary too much (compared to Δ) during the execution of the metaorder. But this is absurd: there is no equilibrium price around which the market pauses; prices move all the time, in a diffusive way, in such a way that difference between p_0 and the final price p_T is in fact *much larger* than the impact Δ we try to account for. If the linear profile follows instantaneously the price, the above argument completely falls into pieces since only the locally available volume would play a role. If the linear profile for some reason remains centred for some time around a specific p_0 , why should this price coincide with the price at the beginning of the metaorder?

Still, the basic intuition, that the available volume grows as price excursions get larger, must somehow make sense. The aim of the present section is to propose a dynamic theory of liquidity largely inspired from [19], based on minimal and plausible assumptions, that indeed predicts that the average supply (or demand) is a V-shaped curve that vanishes around the current price p_t . The square-root impact then follows from an argument similar to Eq. (2). These arguments are then tested quantitatively using a numerical model of order flow based on the same minimal ingredients.

Our basic idea is that of a “latent order book” that at any instant of time t aggregates the total *intended* volume for sells at price p or above, $\mathcal{V}_+(p, t)$, and the total *intended* volume for buys at price p or below, $\mathcal{V}_-(p, t)$. We want to emphasize that this is in general *not* the volume revealed in the real (observable) order book, in particular for p remote from the current price p_t . It is rather the volume that *would* reveal itself in the order book, or as market orders, if the price came instantaneously closer to p . But since there is little incentive to reveal one’s intentions too early, most of the volume is latent and not revealed. This is obvious from basic order of magnitude estimates: whereas the total instantaneous volume in the real order book of a typical liquid stock is of the order of 10^{-5} of the market capitalisation, the total transaction volume per day is 10^{-3} , showing that liquidity is a dynamical process. The empirical analysis of Weber and Rosenow [20] shows very clearly how the volume appearing in the order book is indeed stimulated by the trades themselves.

So our latent volumes $\mathcal{V}_\pm(p, t)$ reflect intentions that do not necessarily materialize. How do these volumes evolve with time? Between t and $t+dt$, new buyers/sellers may become interested at levels $p_t \mp u$, at a rate $\lambda(u)$ and with unit volume $q = 1$; while existing buyers/sellers at $p_t \mp u$ might change their price to $p_t \mp u'$ at rate $\nu(u, u')$, or even disappear temporarily (corresponding to $u' = \infty$).² Clearly, $\mathcal{V}_+(p < p_t, t) = 0$ and $\mathcal{V}_-(p > p_t, t) = 0$, meaning that there cannot be unsatisfied seller (buyer) below (above) the current price p_t .

We now assume that the price process p_t is a Brownian walk³, which is only an approximation since in practice (a) at short times, microstructure effects play a role and (b) large jumps are present, and in fact quite frequent (the distribution of price changes is well known to be a power law for large arguments [21]). This approximation however allows us to make precise analytical calculations that illustrate our main point. Since we are interested in phenomena that take place on time scales of a few minutes to a few days, drift effects are completely negligible and we discard them. For the same reason, the difference between an additive and a geometric Brownian motion is irrelevant. A simple equation for the latent volume, averaged over price paths $\overline{\mathcal{V}_\pm(p, t)}$, can be obtained by working in the reference frame moving with the price p_t , provided an extra assumption is made on the rates $\nu(u, u')$. We assume that either $u' = \infty$, with rate $\nu_\infty(u)$, or that the change of price is small, and occurs equally often up or down. We define $D(u) = \int du' (u - u')^2 \nu(u, u')$, where the integral over u' is rapidly convergent (small jumps); it can be interpreted as the (squared) volatility of intentions. With $\mathcal{D}(u) = D(u) + \sigma^2$, where σ is the price volatility, the final equation for $\rho_\pm(u, t) = \overline{\mathcal{V}_\pm(p_t \pm u, t)}$ reads [22]:⁴

$$\frac{\partial \rho_\pm(u, t)}{\partial t} = \frac{1}{2} \frac{\partial^2}{\partial u^2} [\mathcal{D}(u) \rho_\pm(u, t)] - \nu_\infty(u) \rho_\pm(u, t) + \lambda(u); \quad \rho_\pm(u < 0) \equiv 0. \quad (3)$$

Note that if all rates are symmetric for buy orders and sell orders, the long-time, stationary solution $\rho_{st}(u)$ is the same for ρ_+ and ρ_- . It describes the most probable shape of the latent order book, and is such that the right hand side of the above equation is zero. For arbitrary function $D(u)$, $\lambda(u)$ and $\nu_\infty(u)$, the explicit form of $\rho_{st}(u)$ is not known, but provided these functions are regular when $u \rightarrow 0$, one can show that the stationary profile is *linear* close enough to the current price, i.e. $\rho_{st}(u) \approx bu$ when $u \rightarrow 0$, where b is a certain finite constant. In fact,

$$J = \frac{1}{2} \frac{\partial}{\partial u} [\mathcal{D}(u) \rho(u, t)] \Big|_{u=0} \quad (4)$$

²The following equations would not change if we allow q to fluctuate, provided the average of q is finite and set to unity without loss of generality. We furthermore assume complete symmetry between $p > p_t$ and $p < p_t$; i.e. $\lambda_+(u) = \lambda_-(u)$, etc.

³That the price is a diffusive process is a standard assumption in quantitative finance. It is also very well corroborated by data down to very short time scales (see e.g. [25]), at least in liquid markets. This “statistical efficiency” precludes the existence of simple arbitrage strategies.

⁴Here we assume that $F(u) = \int du' (u - u') \nu(u, u') = 0 \forall u$, but adding a non zero drift term in the reconfiguration of orders would not change any of the following conclusions. Only the value of u^* would change.

is the transaction rate per unit time, i.e. the volume of buy/sell market orders per unit time. If we choose the unit of time to be one day, J is precisely what we called V above. Provided $D(u)$ is regular at $u = 0$, the condition $\rho_{st}(u) = 0$ and J positive and finite is enough to impose that the profile is locally linear with $b = 2J/\mathcal{D}(0)$. Therefore, *the hypothesis of a diffusive price with a finite transaction rate immediately leads on average to a locally linear order book.*

As a simple illustration, consider the case where new orders appear uniformly, i.e. $\lambda(u) = \lambda$, and $\mathcal{D}(u) = \mathcal{D}$ independent of u . The exact solution is then:

$$\rho_{st}(u) = \rho_{\infty} [1 - e^{-u/u^*}], \quad (5)$$

with $\rho_{\infty} = \lambda/\nu_{\infty}$ and $u^* = \sqrt{\mathcal{D}/2\nu_{\infty}}$, leading to $b = \rho_{\infty}/u^*$. One sees that even when new orders appear with a finite density around the current price, they have also a large probability to be executed and disappear. This eventually leads to a liquidity trough at $u = 0$ and a linear profile around $u = 0$; u^* can be interpreted as the width of the linear region. It is reasonable to think that D and σ^2 are comparable, and that $1/\nu_{\infty}$, which measures the (volume weighted) average lifetime of an *intended* order, is dominated by slow players and is on the scale of a few days.⁵ Therefore, u^* is of the order of the daily volatility, which shows that the trough in the latent order book extends over a very significant region around the current price. Note that $\tau_{\text{life}} = 1/\nu_{\infty}$ is also the persistence time of the fluctuations of the order book, and the time to reach the stationary state $\rho_{st}(u)$. This time plays a crucial role in the following.

Note that if $\lambda(u)$ is not constant but decays over a price range u_{λ} , the width of the linear region is still given by u^* provided $u^* \ll u_{\lambda}$. In the other limit, on the other hand, one finds $u^* \sim u_{\lambda}$ [23]. We believe that u_{λ} and $\sqrt{\mathcal{D}/2\nu_{\infty}}$ are in fact of the same order of magnitude (a few percent); this means that the players contributing to the ‘true’ liquidity of the market are not sensitive to price changes much smaller than the daily volatility (see also footnote 3).

Eq. (5) contains the central result of the present paper. It predicts that the available volume in the immediate proximity of the current price goes to zero, which is the reason why we say markets are “critical”, i.e. they operate in a regime of vanishing liquidity. This scenario does not arise by *fiat* but is rather a natural consequence of the diffusivity of prices: we believe this is a genuine example Self Organized Criticality [24].

In more concrete terms, the volume at the best quotes (bid or ask), given by $q_{\text{best}} \approx bw^2/2$ (where w is the tick size) is $\sim (w/u^*)$ smaller than the volume $\rho_{\infty}w$ one would expect in the absence of the above sweeping mechanism. This is typically small since $w \sim 0.05\%$ and $u^* \sim 2\%$ (for stocks or futures). One can also compare q_{best} to the the total volume traded in a time T , which is $V = JT \sim \rho_{\infty}u^*$ if one chooses $T = \tau_{\text{life}} \sim$ one day. The ratio is now $q_{\text{best}}/V \sim (w/u^*)^2$, which is very small, as is indeed the case empirically: the immediately accessible volume is typically a 1000 smaller than the daily turnover. This small liquidity compels market participants to fragment their trades, leading to the universally observed long-range correlation in the sign of market orders [25, 26]. It also leads to the square-root impact law, if one trusts the argument after Eq. (2), with here $b = \rho_{\infty}/u^*$. This gives $\Delta \approx \sqrt{\frac{2Q}{b}} \propto \sigma \sqrt{\frac{Q}{J}}$, where we have used $b \sim J/\sigma^2$, with J is the trading rate. This is exactly the square-root impact law, Eq. (1).

However, this last argument is quite naïve since the average impact Δ is much smaller (for small Q) than the typical excursion of the price within the same period. Furthermore, the diffusive behaviour of the price is taken for granted in the above calculation, whereas in fact

⁵Note here that this is where the distinction between latent and revealed orders is crucial: the average lifetime of orders in the order book is much shorter than this, but this is a result of high frequency strategies which are sensitive to minute price changes, but does not relate to changes of intentions from slow players. Correspondingly, the average shape of the true order book is non monotonous [22], and thus very different from the linear prediction, Eq. (5), for the latent liquidity profile.

it results from a subtle compensation [25] between a *confinement effect* created by the linear supply/demand curve (a price movement in one direction hits larger opposing volumes and is more likely to revert) and a *correlation effect*, created by the fragmentation of the trades [27, 28] – itself imposed by the liquidity trough mentioned above.

Since we are unable at this stage to treat these effects consistently within an analytical framework, we now turn to a minimal numerical model that captures all the above effects. We will indeed find a linear demand profile and a concave impact function, and gain considerable insights into the mechanisms leading to such behaviour.

4 A numerical model for an efficient market with long-ranged order flow

The numerical implementation of the above simple Poisson model for intended order flow is quite simple, and we follow the framework proposed in [19] – see also the Appendix below. All orders have unit volume. Limit orders are launched at a uniform rate λ in a finite (but large) interval around the current price. Existing limit orders are individually cancelled at rate ν_∞ . If market orders are themselves completely random, with buys and sells drawn independently with probability $\frac{1}{2}$ and at rate μ , the result price motion is known to be strongly *subdiffusive*, in the sense that the lag dependent diffusion constant $\sigma^2(\ell) \equiv \langle (p_t - p_{t+\ell})^2 \rangle / \ell$ decays when ℓ increases (here time is counted in number of transactions). This was noted in [19], and is a result of the confining effect of the supply/demand curve. The price only becomes diffusive on time scales larger than the memory time τ_{life} . Nothing of that sort is seen in real price dynamics. However, we know from empirical data [25, 26, 3] that the signs of transactions ϵ_t are in fact long-ranged correlated, i.e. that $C(\ell) = \langle \epsilon_t \epsilon_{t+\ell} \rangle$ is decaying as a power-law, $C(\ell) \approx \ell^{-\gamma}$, where $\gamma \approx 0.5$ for single stocks and $\gamma \approx 0.8$ for futures markets [29]. This persistent direction of trading can counterbalance the confinement effect and restore diffusion. In fact, if γ is too small, one expects *superdiffusion*, i.e. $\sigma^2(\ell)$ growing with ℓ .

We therefore want to upgrade the “zero-intelligence” model of [19] to an “ ϵ -intelligence” numerical model which explicitly includes the long ranged correlated nature of the trades, reflecting the fact that large metaorders are fragmented and traded incrementally, see Ref. [27, 28]. We have chosen to work with the LMF specification of the sign process [27], i.e. sequences of $+/-$ signs (buy/sell market orders) are generated, with lengths L drawn from a power-law distribution: $P(L) \sim L^{-(\alpha+1)}$. It can be shown that the sign process has an autocorrelation function that decays asymptotically as $C(\ell) = \ell^{-\gamma}$ with $\gamma = \alpha - 1$. We do not expect that the following results depend much on the precise specification of the model of signs [30].

At this stage, the model contains four parameters: the rate of limit orders λ , the rate of cancellation ν_∞ , the rate of market orders μ and the autocorrelation exponent γ . In fact, only three of them are relevant, up to a change of time unit: the ratios λ/ν_∞ and μ/ν_∞ , and γ . As noted above, the ratio $\rho_\infty = \lambda/\nu_\infty$ determines the supply depth far away from the best quotes, whereas the ratio $r = \mu/\nu_\infty$ tells us whether we are in a “slow” market, where the renewal time τ_{life} of the supply/demand is much longer than the time between individual trades ($r \gg 1$), or if the market is “fast” in the opposite limit $r \ll 1$. It is clear that real markets are in the former limit: if trades take place on a scale of seconds while the renewal time is of the order of hours or days, the ratio r is on the order of 10^4 . As will be discussed later, $r \gg 1$ is a crucial condition for a concave impact law to hold. On the other hand, μ and $\lambda \times w$ are of similar order of magnitude, which means that, as expected, markets are also “deep”: $\rho_\infty \times w \gg 1$.

The problem is that in the limit of deep and slow markets, the above model is always in the subdiffusive phase unless γ is very close to zero (see Fig. 2 below). We therefore need an

extra ingredient to make markets statistically efficient (meaning that prices are diffusive), while keeping γ in the empirical range $\gamma \sim 0.5 - 0.8$. When compared with real markets, the above model is obviously much too simple. For example, the size of both limit and market orders is known to be broadly distributed, whereas we assume, as in [19, 31], that all volumes are unity. The direction and the size of market orders are furthermore strongly conditioned by the volumes at the best quotes q_{best} : the volume of market orders is larger when the offered volume is larger, and the sign of the next market order is anti-correlated with the volume imbalance (i.e. when $q_{\text{bid}} \gg q_{\text{ask}}$, the next trade is most likely to be a buy, and vice-versa).

We have chosen to add one extra parameter both to make a step towards reality and solve the efficiency problem, in the following way: the volume of market orders is chosen to be a random fraction f of the opposite volume at the best, where the distribution of f is given by $P_\zeta(f) = \zeta(1-f)^{\zeta-1}$, where ζ is a parameter ($\zeta > 0$), that determines the typical relative volume of market orders. For $\zeta \rightarrow 0$, the distribution $P_\zeta(f)$ peaks around $f = 1$ and most orders “eat” all the available liquidity; $\zeta = 1$ corresponds to a flat distribution for the fraction of eaten volume; finally the limit $\zeta \rightarrow \infty$ corresponds to very small (unit) volumes and recovers the previous model. The correlation between volume at best and volume of the impinging market order has been reported in many papers (see [32]). This makes perfect sense: since large metaorders must be fragmented because of the small available liquidity, one expects that the executed volume opportunistically follows the offered liquidity.

With the help of the extra parameter ζ we can now tune the model to guarantee diffusive prices for any value of γ , even in the limit of deep and slow markets. We define a measure of pure diffusivity as $\sigma(\ell_1)/\sigma(\ell_2)$ for two time scales ℓ_1 and $\ell_2 > \ell_1$. Subdiffusion corresponds to ratios < 1 , superdiffusion to ratios > 1 and for pure diffusion this ratio must be equal to unity for all ℓ_1, ℓ_2 . In our simulations, we chose $\lambda = 0.5$, $\mu = 0.1$, $\nu_\infty = 10^{-4}$ (corresponding to $\rho_\infty = 5000 \gg 1$ and $r = 1000 \gg 1$). We determine the *diffusion line* in the plane γ, ζ , that separates the subdiffusion regime from the superdiffusion regime, by setting $\ell_1 = 10$, $\ell_2 = 1000$.⁶ The result is shown in Fig. 2. As expected, smaller γ 's favor superdiffusion, whereas larger values of ζ (corresponding to less aggressive market orders) favors subdiffusion. On the boundary between the two regimes we find the purely diffusive, efficient markets we are looking for; in the following we fix the value of $\zeta(\gamma)$ such as to be exactly diffusive.⁷ Since the value of ζ is fixed by the condition of price diffusivity, our model still has three parameters: depth (ρ_∞), inverse speed (r) and trade persistence (γ).

To illustrate the tight connection between the dynamical theory discussed in the previous section and the above described numerical model, in Fig. 3 we compare the stationary density of the book to its expected shape calculated from Eq. (5). The symbols show the stationary density, ρ_{st} , measured in the simulations, while the solid line is Eq. (5) with $u^* = \sqrt{D/2\nu_\infty}$, D being the actual measured price volatility. The latent order book is found to be linear in the immediate vicinity of the price, as predicted. The exponential curve with the analytical expression for u^* leads to a very good fit in the whole range of u .

We now have a model such that (i) the order flow has long range memory but (ii) the price is diffusive. We are thus in a position to test quantitatively our above ideas about the linear supply function and the impact of metaorders. In order to do this, we add to the above order flow an extra agent who becomes active at a certain (arbitrary) time t , chooses a random sign ϵ and a random size Q for his metaorder, which he executes incrementally using market orders until all the volume is transacted. We have considered two execution styles: a) “ ζ -execution”,

⁶The results are not sensitive to the precise choice of ℓ_1 and ℓ_2 , provided $\ell_1, \ell_2 \ll \tau_{\text{life}}$. We have explored a wide range of values of λ, μ, ν_∞ . Provided one remains in the limit of deep and slow markets, the broad picture is unaffected, although the precise location of the diffusion line in Fig. 2 is changed.

⁷It would be interesting to obtain an analytical form for the diffusion line, but we have not attempted to make a theory for this at the present stage.

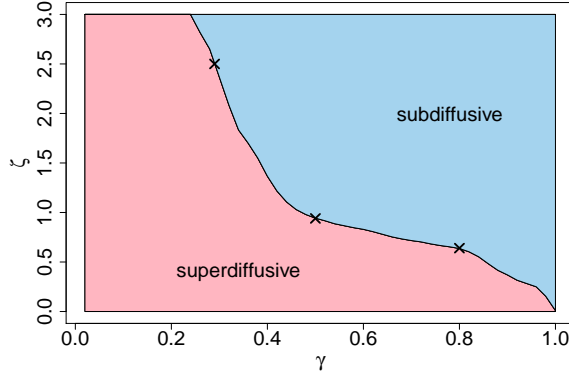


Figure 2: Map of diffusivity in the plane γ, ζ for $\lambda = 0.5, \mu = 0.1, \nu_\infty = 10^{-4}$. The large γ , large ζ region corresponds to subdiffusion, while the small γ , small ζ region corresponds to superdiffusion. The “efficient market” (diffusive) line determined numerically is such that $\zeta(0.8) \approx 0.65$, $\zeta(0.5) \approx 0.95$, and $\zeta(0.3) \approx 2.5$ (crosses). Note that the unit volume limit $\zeta \rightarrow \infty$ corresponds to $\gamma \rightarrow 0$. Without volume fluctuations, deep and slow markets are therefore always found to be in the subdiffusive phase.




Figure 3: Comparison of the stationary density of the book with the expected shape calculated from Eq. (5). The symbols show the empirical shape of the book from the simulations, while the solid line shows the predicted exponential form, with u^* being computed from the actual price volatility. In the simulations we used the parameters, $\gamma = 0.8, \zeta = 0.65, \nu_\infty = 10^{-4}, \lambda = 0.5$. The analytical computations lead to $u^* = \sqrt{D/2\nu_\infty} = 0.49$, while a direct fit to Eq. (5) leads to $u^* \approx 0.48$.




Figure 4: Impact of metaorders on log-log scale. We show Δ/σ vs. Q/V , where σ and V are measured over the timescale τ_{life} . The three curves with symbols show the impact for the ζ -execution with high participation rate, $\Phi = 0.3$, and three values of the sign autocorrelation exponent: $\gamma = 0.3$ (green crosses), $\gamma = 0.5$ (black circles), $\gamma = 0.8$ (red triangles). The two blue lines show the impact for low participation rate $\Phi = 0.05$, for the ζ -execution (full line) and unit-execution (dashed line), both for $\gamma = 0.5$. The purple dots show a case when the time to complete metaorders is longer than the lifetime of the book $T > \tau_{\text{life}}$. In this case we get back a linear impact. The two solid black lines are of slope $\delta = 0.5$ and $\delta = 1$, for comparison, with the choice $Y = 1$ for $\delta = 0.5$.

where the extra agent trades exactly as the rest of the market, by sending a market order of volume $f \times q_{\text{best}}$, where f is chosen according to $P_{\zeta}(f)$ above; b) “unit-execution”, where the market orders are all of unit volume, whenever he trades. (We also studied limit order execution with similar results, see [30].) In both cases, he participates to a fraction Φ of all market orders. We measure, as the real data shown in Fig. 1, the impact Δ as the price paid by the extra agent compared to the price p_t at the beginning of the metaorder, averaged over many different metaorders of size Q . The results are shown in Fig. 4, where we show in a log-log scale Δ/σ on the y-axis versus Q/V on the x-axis, for different values of γ , r and Φ , and for both ζ -execution and unit-execution. We also show two straight lines of slope $\delta = 1/2$, corresponding to a square-root impact, and $\delta = 1$, corresponding to a linear impact. It is clear that for all cases where the execution time T is much smaller than the renewal time τ_{life} , the impact is strongly concave and, *in a first approximation*, independent of γ and of the participation rate Φ , in agreement with empirical observations. The exponent δ is found to be close to $1/2$ for unit-execution, and close to $2/3$ for ζ -execution. The Y factor, defined in Eq. (1), can be measured from the data shown in Fig. 4, and is of order unity, as found empirically. More details about these results will be given in [30].

Some parameters of the model however do influence the value of the effective exponent δ , which might explain why empirical data show some scatter around the value $\delta = \frac{1}{2}$. In particular, when the execution time T increases and becomes comparable to the renewal time τ_{life} , the effective exponent δ increases and the impact becomes linear in the limit $T \gg \tau_{\text{life}}$ – see Fig. 4. This is indeed expected since impact is necessarily additive when all memory of the latent order book has been erased. It is also a direct proof that the persistence of the supply/demand is central to explain the functional form of the impact. The naïve prediction for the price impact, Eq. (2), based on the average slope of the supply curve, is shown in Fig. 5 (left), together with the numerical determination of $\Delta(Q)$ already plotted in Fig. 4. We see that the naïve argument indeed leads to the correct order of magnitude for the impact, but fails to be quantitatively accurate: it underestimates the real impact by a factor ~ 2 . On the same graph, we also show the global measure of impact mentioned above, where the average price change is plotted against the total volume imbalance \mathcal{Q} . We see that the latter is very different from the impact of a given metaorder. The global impact is *linear* in \mathcal{Q} (as observed




Figure 5: (Left) The true impact of ζ -executed metaorders for $\gamma = 0.5$, $\Phi = 0.3$ (black circles), the naïve estimate of the impact from integrating the volume in the average latent order book (blue solid line), and the average price change against the global volume imbalance in a given period (red triangles). We also show a power law fit for the metaorder impact, as $Y(Q/V)^\delta$, with $\delta = 0.7$ and $Y = 1.59$ (solid black line) and a linear fit for the global measure of impact (solid red line). (Right) The decay of impact after the completion of the metaorder for unit-execution for $\gamma = 0.3, 0.5$ and 0.8 and $\Phi = 0.5$. The x-axis is τ/T , the clock-time in units of the time needed to complete the metaorder (with $T \ll \tau_{\text{life}}$), the y-axis is the rescaled price change during the metaorder. After the metaorder, the price appears to relax to a constant level, which is a γ dependent fraction of the temporary impact. The dotted horizontal lines correspond to the actual average execution price of the metaorder, which is found to be $\approx 0.6 \Delta_T$ for all three values of γ .

on empirical data for large enough T) and, for small volume imbalances, much smaller than the impact of an additional metaorder that perturbs the equilibrium flow.

Another interesting aspect of the problem is the temporal structure of the impact of metaorders. In Fig. 5 (right) we show the average difference Δ_τ between the price p_t of the first trade of a metaorder and the price $p_{t+\tau}$ at time τ later, in the direction of the trade. The x-axis is in units of the time needed to complete the metaorder. The exact shape of the build-up in time depends on the execution mechanism, this we will discuss in details in [30]. Here we focus on the price dynamics after the completion of the metaorder. Once the metaorder is completed (i.e. for $\tau > T$), the impact decays down to smaller values, and appears to reach a plateau Δ_∞ , such that Δ_∞/Δ_T depends on the value of γ , and is ~ 0.75 for $\gamma = 0.5$. The cause of the reversion in our model is that during the trading of a metaorder, the order book becomes on average unbalanced: orders on the opposite side become statistically more numerous, as a consequence of the V-shaped liquidity. This results in a partial reversion of the price, once the pressure from the metaorder is over.

Interestingly, a similar behaviour has also been observed in [9], and is predicted by theory of impact recently put forth in [17], which elaborates on the idea of a “fair execution price”. As emphasized in [9, 17], if a metaorder of size Q has an impact that grows as $(q_{\text{exec}}/Q)^\delta$ (where q_{exec} is the volume already executed, i.e. q_{exec}/Q is “volume time”)⁸, then the average execution price of the metaorder is $\Delta_T/(1+\delta)$. In this case, $\Delta_\infty/\Delta_T \equiv 1/(1+\delta)$ ensures that the average execution price of the metaorder is equal to the price long after the execution is completed, in other words that the execution price is fair. Used together with the condition that the price should be a martingale and that the size distribution of metaorders is a power-law, $Q^{-1-\alpha}$, Farmer et al. [17] obtain a concave impact $\Delta \sim Q^\delta$ with $\delta = \alpha - 1 = \gamma$.⁹ Using $\alpha \approx 3/2$ from

⁸Since the impact is $\Delta(Q) = Y\sigma(Q/V)^\delta$, we expect the impact to grow as $\Delta_{q_{\text{exec}}} \approx \Delta(Q)(q_{\text{exec}}/Q)^\delta$ meaning that a metaorder of size Q stopped half-way through execution impacts the price the same way as a metaorder of size $Q/2$, which *a priori* makes sense since there is no way to anticipate when a metaorder is going to stop. Indeed, we find the build-up of the impact in *volume* time to be concave.

⁹Note that in the theory of Farmer et al. [17], the exponents δ and γ are equal, whereas in our model δ is

empirical data on stocks, the square-root impact law is recovered [17], with $\Delta_\infty/\Delta_T = 2/3$. We see however in Fig. 5 (right) that the “fair price” condition does not hold in general within our model, since the plateau value significantly changes with γ , whereas the average price paid is close to $0.6 \Delta_T$ in all cases. Nevertheless, for $\gamma = 0.5$ the plateau is not far from the value $2/3$ predicted by Farmer et al. [17].

5 Conclusion

Although some elements of our model are common with the framework of [17] (broad distribution of metaorders leading to long range correlation of the order flow, and the efficient price condition), others are very different. The theory advocated in [17] requires that metaorders can be identified by market makers, in particular that the very first trades can be detected. As the authors openly admit, the need for market participants to be able to detect the starting and stopping of a given metaorder is potentially a problem, and in fact our model does not rely at all on such a strong assumption. On the contrary, we have shown that the universally observed concave impact law is a consequence of some robust, generic assumptions about market dynamics. In particular, we have provided a dynamical theory of liquidity which leads to a locally linear supply/demand curve, provided high frequency strategies guarantee price diffusivity on all time scales. Our story is purely statistical and does not rely on additional (and sometimes wooly) notions such as fundamental prices, adverse selection or fair pricing. There are no explicit market makers, strategic players or optimising agents in our picture, but rather an ecological equilibrium of indistinguishable traders that interact in a way to make the price statistically efficient.

Although our “ ε -intelligence” numerical model makes an important step towards reality (in particular by including long-range correlations in the order flow and ensuring a diffusive price dynamics), there are still many assumptions that are ad-hoc. One knows for example that the deposition/cancellation rates of limit orders do strongly depend on the distance from the current price [33], that the assumption of a Poisson process is an oversimplification [34] (these observations apply to the real order book, but probably will also hold for the latent order book), that the volume of limit orders is not at all constant, etc. etc. Building a detailed, realistic model of order flow is of course needed to get fully quantitative predictions for the impact, but was not the scope of the present work. We instead wanted to have a simplified model that would allow us to test our central hypothesis: that the latent order book is locally linear and that this is the crucial ingredient to explain the square-root impact law. We believe that this objective has been reached, and we leave the more ambitious project of building a full-scale model for future work. We also note that our central assumption of a latent order book has other, empirically testable assumptions, which we are currently investigating.

As we emphasized in the introduction, understanding impact has immediate practical implications in terms of trading costs and capacity of strategies. In our view, the most important message of the theory presented above concerns the critical, inherently fragile nature of liquidity. By necessity, a diffusive price leads to a vanishing liquidity in the vicinity of the current price. This naturally accounts for two striking stylized facts: first, large metaorders have to be fragmented in order to be digested by the liquidity funnel, which leads to long-memory in the sign of the order flow. Second, the anomalously small local liquidity induces a breakdown of linear response and a diverging impact of small orders. Furthermore, liquidity fluctuations are bound to play a crucial role when the average liquidity is small, and we expect these fluctuations to be at the heart of the turbulent dynamics of financial markets, as postulated in, e.g. [35, 25, 36].

to a large extent unaffected by the precise value of γ .

Acknowledgements

We want to thank Paul Besson, Rémy Chicheportiche, Zoltán Eisler, Péter Horvai, Fabrizio Lillo, Matteo Marsili, Marc Potters, Emmanuel Sérié, Damien Trouvé and especially Doyme Farmer and Jim Gatheral for very helpful discussions and suggestions.

References

- [1] J. Hasbrouck, *Measuring the information content of stock trades*, Journal of Finance, **XLVI**, 179, (1991)
- [2] J.-P. Bouchaud, *Price Impact*, in Encyclopedia of Quantitative Finance, R. Cont Edt.; John Wiley & Sons Ltd. (2010)
- [3] J.-P. Bouchaud, J. D. Farmer, F. Lillo, *How Markets Slowly Digest Changes in Supply and Demand* in Handbook of Financial Markets: Dynamics and Evolution, pp 57-156. Eds. Thorsten Hens and Klaus Schenk-Hoppe. Elsevier: Academic Press, (2008)
- [4] R. Almgren, C. Thum, E. Hauptmann, H. Li, *Direct estimation of equity market impact*, Risk, **18**, 5762, (2005)
- [5] R. Engle, R. Ferstenberg, J. Russell, *Measuring and Modeling Execution Cost and Risk*, working paper (2006).
- [6] X. Abdobal, *Transaction Cost Estimation*, Lehman Brothers commercial document (2006)
- [7] A. Ferraris, *Market Impact Models*, Deutsche Bank internal document (2007)
- [8] R. Kissel, R. Malamut, *Algorithmic Decision Making Framework*, J.P. Morgan internal document (2005)
- [9] E. Moro, J. Vicente, L. G. Moyano, A. Gerig, J. D. Farmer, G. Vaglica, F. Lillo, R. N. Mantegna, *Market impact and trading profile of hidden orders in stock markets*, Physical Review E **80**, 066102 (2009)
- [10] A. J. Menkveld *High Frequency Trading and The New-Market Makers*, Working Paper (2011)
- [11] A. S. Kyle, *Continuous auctions and insider trading*, Econometrica, **53**, 1315 (1985)
- [12] N. Torre, M. Ferrari, *Market impact model handbook*, (Berkeley, California, Barra, 1997), available at http://www.msibarra.com/research/articles/barra/Market_Impact_Model.pdf
- [13] R. C. Grinold, R. N. Kahn, *Active Portfolio Management* (New York: The McGraw-Hill Companies, Inc., 1999).
- [14] X. Gabaix, P. Gopikrishnan, V. Plerou, and H. E. Stanley, *Institutional investors and stock market volatility*; Quarterly Journal of Economics, **121**, 461 (2006)
- [15] G. C. Chacko, J. W. Jurek, and E. Stafford, *The Price of Immediacy*, The Journal of Finance, **LXIII**, 1253 (2008)
- [16] Y. C. Zhang, *Toward a theory of marginally efficient markets*, Physica **A269**, 30 (1999).

- [17] J. D. Farmer, A. Gerig, F. Lillo, H. Waelbroeck, *How Efficiency Shapes Market Impact*, <http://arxiv.org/abs/1102.5457>
- [18] This is a very old idea. L. Harris, for example, presented a back-of-envelope theory of the square-root impact based on this idea at the conference “Microstructure, confronting many viewpoints”, Paris, December 2010. P. Besson also elaborated on this (private communication).
- [19] E. Smith, J.D. Farmer, L. Gillemot, S. Krishnamurthy, *Statistical Theory of the Continuous Double Auction*, *Quantitative Finance* **3** 481-514,(2003)
- [20] P. Weber, B. Rosenow, *Order book approach to price impact*, *Quantitative Finance*, **5**, 4, 357 - 364, (2005)
- [21] V. Plerou, P. Gopikrishnan, L. A. Amaral, M. Meyer, H.E. Stanley, *Scaling of the distribution of price fluctuations of individual companies*, *Phys. Rev. E* **60** 6519 (1999); P. Gopikrishnan, V. Plerou, L. A. Amaral, M. Meyer, H. E. Stanley, *Scaling of the distribution of fluctuations of financial market indices*, *Phys. Rev. E* **60** 5305 (1999)
- [22] see J.-P. Bouchaud, M. Mézard, M. Potters, *Statistical properties of the stock order books: empirical results and models*, *Quantitative Finance*, **2**, 251 (2002) for an alternative presentation.
- [23] J. Gatheral, private communication (July 2011).
- [24] see e.g.: P. Bak, *How Nature Works: The Science of Self-Organized Criticality*. Copernicus, Springer, New York, 1996.
- [25] 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**, 176-190 (2004); J.-P. Bouchaud, J. Kockelkoren, and M. Potters, *Random walks, liquidity molasses and critical response in financial markets*, *Quantitative Finance*, **6**, 115-123 (2006)
- [26] F. Lillo and J. D. Farmer, *The long memory of the efficient market*, *Studies in Nonlinear Dynamics & Econometrics*, **8**, 1 (2004)
- [27] F. Lillo, Sz. Mike, J. D. Farmer, *Theory for Long Memory in Supply and Demand*, *Phys. Rev. E* **71**, 066122, (2005)
- [28] B. Tóth, F. Lillo, J. D. Farmer, *Segmentation algorithm for non-stationary compound Poisson processes*, *Eur. Phys. J. B* **78**, 235-243 (2010)
- [29] Y. Lempérière, M. Vettorazzo, J. Kockelkoren, J.-P. Bouchaud, CFM report (unpublished)
- [30] B. Tóth, Y. Lempérière, C. Deremble, J. de Lataillade, J. Kockelkoren, J.-P. Bouchaud, *long version in preparation*
- [31] see e.g. R. Cont, A. Larrard, *Price dynamics in a Markovian limit order market*, <http://arxiv.org/abs/1104.4596>
- [32] J.D. Farmer, L. Gillemot, F. Lillo, Sz. Mike, A. Sen, *What Really Causes Large Price Changes?* *Quantitative Finance*, **4**, 383-397, (2004)
- [33] M. Potters, J.-P. Bouchaud, *More statistical properties of order books and price impact*, *Physica A: Statistical Mechanics and its Applications*, **324**(1-2), 133-140, (2003)

- [34] Sz. Mike, J. D. Farmer, *An Empirical Behavioral Model of Liquidity and Volatility*, J Economic Dynamics and Control **32**(1), 200-234, (2008)
- [35] F. Lillo and J. D. Farmer, *The key role of liquidity fluctuations in determining large price fluctuations*, Fluctuations and Noise Letters, **5**, L209 (2005).
- [36] A. Joulin, A. Lefèvre, D. Grunberg and J.-P. Bouchaud, *Stock Price Jumps: News and Volume Play a Minor Role*, Wilmott Magazine, September/October, 2008; J.-P. Bouchaud, *The Endogenous Dynamics of Markets: Price Impact, Feedback Loops and Instabilities*, in “Lessons from the 2008 Crisis”, A. Berd Edt., Risk Publications (2011).

A Appendix

This Appendix summarizes the main steps of the numerical model discussed in the paper. In the model we consider that the price axis is very large and is discretized by the tick size (i.e. the minimum increment of price). The minimum and maximum price are K tick away, with $K \gg 1$. Time is measured in discrete steps, and in each time step, three types of events can happen: new limit orders are placed, market orders are placed or extant limit orders are canceled. Many of these events can happen simultaneously during a single time step.

- Limit orders are orders to buy or sell, that do not trigger an immediate trade. According to this, limit orders are placed on a support of size K that is in practice infinite, i.e. much larger than all other price scales in the system. Limit orders placed below the current midpoint price are considered as limit orders to buy, while those placed above the midpoint price are considered orders to sell. All limit orders have unit volume in the present version of the model (but this is in no way an essential ingredient). Limit orders arrive with a uniform rate λ per unit time per unit price. In practice this means that in each time step, there is a probability $\lambda^n e^{-\lambda}/n!$ that n new limit orders fall in each of the K bins of the price axis.
- Market orders are orders that trigger an immediate transaction (with existing limit orders on the opposite best price level). They arrive with a rate μ per unit time. The sign (direction) of market orders is generated using the algorithm proposed in [27] with one active agent at any instant of time. The volume of market orders is chosen to be a random fraction f of the opposite volume at the best where the distribution of f is given by $P_\zeta(f) = \zeta(1-f)^{\zeta-1}$, where ζ is a parameter ($\zeta > 0$), that determines the typical relative volume of market orders. When $\zeta \rightarrow 0$, market orders take all the prevailing volume on the opposite best.
- In each time step, each limit order in the book has the probability ν_∞ of being cancelled. This Poisson process of cancellation defines the typical lifetime of a limit order (as ν_∞^{-1}).

The order of the processes in each time step is first the placement of limit orders, then the possible execution of market orders, finally cancellation of some limit orders. As shown in Fig. 2, with the help of the parameters γ and ζ we can tune the system to guarantee diffusive prices for any value of the parameters μ , λ and ν_∞ . This diffusive market allows us to test the predictions of our analytical model, about the locally linear profile around the current midpoint price. To do this, we add an extra “agent” to the model, who wishes to transact a metaorder that is typically larger than the available volume on the first level of the book and thus has to be split up and traded incrementally. It is the price change from the beginning to the end of such a metaorder, that we would like to study, in order to compare the predictions of the model

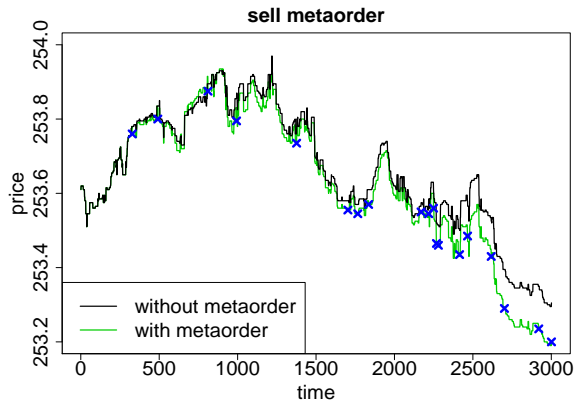


Figure 6: The price process with and without the metaorder trading, with the same random seed used for the two simulations. The black curve shows the price dynamics without the metaorder, while the green curve shows it with the metaorder, trades of the metaorder are denoted by blue crosses. In the simulations we used the parameters, $\gamma = 0.8$, $\zeta = 0.65$, $\nu_\infty = 10^{-4}$, $\lambda = 0.5$. The participation rate of the metaorder is $\Phi = 0.05$.

to empirical results. The only remaining point is to define how the extra agent trades. In the present paper we allowed two trading strategies for the execution of the metaorder: a) “ ζ -execution”, where the extra agent trades exactly as the rest of the market, by sending a market order of volume $f \times q_{\text{best}}$, where f is chosen according to $P_\zeta(f)$ above and b) “unit-execution”, where the market orders are all of unit volume, whenever he trades.

In Fig. 6 we show a segment of the price process for a simulation run, chosen only for illustration purposes. We plot the price as a function of time (simulation steps). The two simulations shown are governed by the same random seed, however in one case an extra agent trading a metaorder *to sell* is added to the market. The black curve shows the price dynamics without the metaorder, while the green curve shows it with the metaorder trades (the metaorder starts at time 0 and finishes at time 3000). Blue crosses indicate the moments when a trade by the extra agent was made. The agent follows the ζ -execution strategy in this particular run. For a while the two curves are identical, then after the first trade of the metaorder the two curves start to deviate from each other. The coarse-grained dynamics of the two curves are similar, however the final price in case of the green curve is pushed down as a result of the extra sell metaorder.