

A fully consistent, minimal model for non-linear market impact

J. Donier,^{1,2} J. Bonart,^{1,3} I. Mastromatteo,⁴ and J.-P. Bouchaud^{1,3}

¹*Capital Fund Management, 23-25 Rue de l'Université, 75007 Paris, France.*

²*Laboratoire de Probabilités et Modèles Aléatoires, Université Pierre et Marie Curie (Paris 6).*

³*CFM-Imperial Institute of Quantitative Finance, Department of Mathematics,
Imperial College, 180 Queen's Gate, London SW7 2RH*

⁴*Centre de Mathématiques Appliquées, CNRS, UMR7641, Ecole Polytechnique, 91128 Palaiseau, France.*

(Dated: March 7, 2015)

We propose a minimal theory of non-linear price impact based on the fact that the (latent) order book is locally linear, as suggested by diffusion-reaction models and general arguments. Our framework allows one to compute the average price trajectory in the presence of a meta-order, that consistently generalizes previously proposed propagator models. We account for the universally observed square-root impact law, and predict non-trivial trajectories when trading is interrupted or reversed. We prove that our framework is free of price manipulation, and that prices can be made diffusive (albeit with a generic short-term mean-reverting contribution). Our model suggests that prices can be decomposed into a transient “mechanical” impact component and a permanent “informational” component.

I. INTRODUCTION

The study of market impact (i.e. the way trading influences prices in financial markets) is arguably among the most exciting current themes in theoretical finance, with many immediate applications ranging from trading cost modelling to important regulatory issues. What is the meaning of the market price if the very fact of buying (or selling) can substantially affect that price? The questions above would on their own justify a strong research activity that dates back to the classic Kyle paper in 1985 [1]. But as often in science, it is the empirical discovery of a genuinely surprising result that explains the recent spree of activity on the subject (see e.g. [2–6] and refs. therein). In strong contrast with the predictions of the Kyle model, market impact appears to be neither *linear* (in the traded quantity Q) nor *permanent*, i.e. time independent [7]. As now firmly established by many independent empirical studies, the average price change induced by the sequential execution of a total volume Q (which we call *meta-order*) appears to follow a sub-linear, approximate \sqrt{Q} law [2, 3, 5, 8–13]. At the end of the meta-order, impact is furthermore observed to decay (partially or completely) towards the unimpacted price [11–14].

Quite strikingly, the square-root law appears to be *universal*, as it is to a large degree independent of details such as the type of contract traded (futures, stocks, options, Bitcoin [14]...), the geographical position of the market venue (US, Europe, Asia), the time period (1995 → 2014), the maturity of the market (e.g. Bitcoin vs. S&P500), etc. While the impact of *single orders* is non universal and highly sensitive to market micro-structure, the impact of *meta-orders* appears to be extremely robust against micro-structural changes. For example the rise of high-frequency trading (HFT) in the last ten years seems to have had no effect on its validity (compare Refs. [2, 8] that uses pre-2004 data with [3, 5, 11] that use post-2007 data). This universality strongly suggests that simple, “coarse-grained” models should be able to reproduce the square-root impact law and other slow market phenomena, while abstracting away from many microscopic details that govern order flow and price formation at high frequencies. This line of reasoning is very similar to many situations in physics, where universal large scale/low frequency laws appear for systems with very different microscopic behaviour. A well known example is the behaviour of weakly interacting molecules which on large length scales can be accurately described by the Navier-Stokes equation, with a single “emergent” parameter (the viscosity) that encodes the microscopic specificities of the system. The Navier-Stokes equation can in fact be derived either from the statistical description of the dynamics of molecules, through an appropriate coarse-graining procedure, or from general considerations based on symmetries, conservation laws and dimensional arguments. Along this path, two pivotal ideas have recently emerged. One is the concept of a *latent* order book [3] that contain the intentions of low-frequency actors at any instant of time, which may or may not materialize in the observable order book. Indeed, since the square-root impact is an aggregate, low-frequency phenomenon, the relevant object to consider cannot be the “revealed” order book, which chiefly reflects the activity of high frequency market-makers. Simple orders of magnitude confirm that the latent liquidity is much higher than the revealed liquidity: whereas the total daily volume exchange on a typical stock is around 1/200th of its market capitalisation, the volume present in the order book at any instant in time is 1000 times smaller than this. Market-makers only act as small intermediaries between much larger volume imbalances present in the latent order book, that can only get resolved on large time scales.

The second idea is that the dynamics of the latent order book can be faithfully modelled by a so-called “reaction-diffusion” model, at least in a region close to the current price where this dynamics becomes universal, i.e. independent

of the detailed setting of the model – and hence, as emphasized above, of the detailed micro-structure of the market and of its high-frequency activity. The reaction-diffusion model in one dimension posits that two types of particles (called B and A), representing in a financial context the intended orders to buy (*bids*) and to sell (*ask*) diffuse on a line and disappear whenever they meet $A + B \rightarrow \emptyset$ – corresponding to a transaction. The boundary between the B -rich region and the A -rich region therefore corresponds to the price p_t . This highly stylized order book model was proposed in the late 90’s by Bak et al. [15, 16] but never made it to the limelight because the resulting price dynamics was found to be strongly mean-reverting on all time scales, at odds with market prices which, after a short transient, behave very much like random walks. However, some of us (together with B. Tóth, [17]) recently realized that the analogue of market impact can be defined and computed within this framework, and was found to obey the square-root law exactly.

This opens the door to a fully consistent theoretical model of non-linear impact in financial markets, which we propose in the present paper. We show how all the previously discussed ingredients can be accommodated in a unifying coarse-grained model for the dynamics of the latent order book that is consistent with price diffusion, with a single emergent parameter – the market *liquidity* \mathcal{L} , defined below. When fluctuations are neglected (in a sense that will be specified below), the impact of a meta-order can be computed exactly, and is found to exhibit two regimes: when the execution rate is sufficiently slow, the model becomes identical to the linear propagator framework proposed in [7, 18], with a bare propagator decaying as the inverse square-root of time. When execution is faster, impact becomes fully non-linear and obeys a non-trivial, closed form integral equation. In the two regimes, the impact of a meta-order grows as the square-root of the volume, but with a pre-factor that depends on the execution rate in the slow regime, but becomes independent of it in the fast regime – as indeed suggested by empirical data. The model predicts interesting price trajectories when trading is interrupted or reversed, leading to effects that are observed empirically but impossible to account for within a linear propagator model. We demonstrate that prices in our model cannot be manipulated, in the sense that any sequence of buy and sell orders that starts and ends with a zero position on markets leads to a non-positive average profit. This is a non trivial property of our modelling strategy, which makes it eligible for practical applications. Finally, we discuss how our framework suggests a clear separation between “mechanical” price moves (i.e. induced by the impact of random trades) and “informational” price moves (i.e. the impact of any public information that changes the latent supply/demand). This is a key point that allows us to treat consistently, within the same model, diffusive prices and memory of the order book – which otherwise leads to strongly mean-reverting prices (see the discussion in [3, 5, 19]). We discuss in the conclusion some of the many interesting problems that our modelling strategy leaves open – perhaps most importantly, how to consistently account for order book fluctuations that are presumably at the heart of liquidity crises and market crashes.

II. DYNAMICS OF THE LATENT ORDER BOOK

Our starting point is the zero-intelligence model of Smith et al. [20], reformulated in the context of the latent order book in [3] and independently in [21]. We assume that each trading (buy/sell) intention of market participant is characterized by a reservation price and a volume.¹ In the course of time, the dynamics of intentions can be essentially of four types: a) reassessment of the reservation price, either up or down; b) partial or complete cancellation of the intention of buying/selling; c) appearance of new intentions, not previously expressed and finally d) matching of an equal volume of buy/sell intentions, resulting in a transaction at a price that delimits the buys/sells regions, and removal of these intentions from the latent order book. It is clear that provided very weak assumptions are met: i) the changes in reservation prices are well behaved (i.e. have a finite first and second moment) and short-ranged correlated in time; ii) the volumes have a finite first moment, one can establish – in the large scale, low frequency, “hydrodynamic” limit – the following set of partial differential equations for the dynamics of the *average* buy (resp. sell) volume density $\rho_B(x, t)$ (resp. $\rho_A(x, t)$) at price level x :

$$\frac{\partial \rho_B(x, t)}{\partial t} = -V_t \frac{\partial \rho_B(x, t)}{\partial x} + D \frac{\partial^2 \rho_B(x, t)}{\partial x^2} - \nu \rho_B(x, t) + \lambda \Theta(p_t - x) - \kappa R_{AB}(x, t); \quad (1-a)$$

$$\frac{\partial \rho_A(x, t)}{\partial t} = \underbrace{-V_t \frac{\partial \rho_A(x, t)}{\partial x} + D \frac{\partial^2 \rho_A(x, t)}{\partial x^2}}_{\text{a- Drift-Diffusion}} - \underbrace{\nu \rho_A(x, t)}_{\text{b- Cancel.}} + \underbrace{\lambda \Theta(x - p_t)}_{\text{c- Deposition}} - \underbrace{\kappa R_{AB}(x, t)}_{\text{d- Reaction}}; \quad (1-b)$$

¹ In fact, each participant may have a full time dependent supply/demand curve with different prices and volumes, with little change in the effective model derived below.

where the different terms in the right hand sides correspond to the four mechanisms a-d, on which we elaborate below, and p_t is the *coarse-grained* position of the price (i.e. averaged over high frequency noise), defined from the condition

$$\rho_A(p_t, t) - \rho_B(p_t, t) = 0. \quad (1)$$

- a- *Drift-Diffusion*: The first two terms model the fact that each agent reassess his/her reservation price x due to many external influences (news, order flow and price changes themselves, other technical signals, etc.). One can therefore expect that price reassessments contain both a (random) agent specific part that contributes to the diffusion coefficient ² D and a *common* component V_t that shifts the entire latent order book. This shift is due to a collective price reassessment due for example to some publicly available information (that could well be the past transactions themselves). The drift component V_t is at this stage very general; one possibility that we will adopt below is to think of V_t as a white noise, such that the price p_t is a diffusive random walk. Since the derivation of these first two terms and the assumptions made are somewhat subtle, we devote Appendix A to a more detailed discussion and alternative models; see in particular Eq. (30).
- b- *Cancellations*: The third term corresponds to partial or complete cancellation of the latent order, with a decay time ν^{-1} independent of the price level x (but see previous footnote²). Consistent with the idea of a common information, cancellation could be correlated between different agents. However, this does not affect the evolution of the average densities $\rho_{B,A}(x, t)$, while it might play a crucial role for the fluctuations of the order book, in particular to explain liquidity crises.
- c- *Deposition*: The fourth term corresponds to the appearance of new buy/sell intentions, modelled by a “rain intensity” λ modulated by an arbitrary increasing function $\Theta(u)$, expressing that buy orders mostly appear below the current price p_t and sell orders mostly appear above p_t . The detailed shape of $\Theta(u)$ actually turns out to be, to a large extent, irrelevant for the purpose of the present paper (see Appendix B for details); for simplicity we will choose below a step function, $\Theta(u > 0) = 1$ and $\Theta(u < 0) = 0$.
- d- *Reaction*: The last term corresponds to transactions when two orders meet with “reaction rate” κ ; the quantity $R_{AB}(x, t)$ is formally the average of the product of the density of A particles and the density of B particles, i.e. $R_{AB}(x, t) \approx \rho_A(x, t)\rho_B(x, t) + \text{fluctuations}$. However, the detailed knowledge of $R_{AB}(x, t)$ will not affect the following discussion. We will consider in the following the limit $\kappa \rightarrow \infty$, which corresponds to the case where latent limit orders close to the transaction price *all become instantaneously visible* limit orders that are duly executed against incoming market orders.

Let us insist that Eqs. (1-a,b) only describe the *average* shape of the latent order book, i.e. fluctuations coming from the discrete nature of orders are neglected at this stage: see Fig. 1 for an illustration. In particular, the instantaneous position of the price $p_t^{\text{inst.}}$ – where the density of buy/sell orders vanishes – has an intrinsic non-zero width even in the limit $\kappa \rightarrow \infty$ [22], corresponding to the average distance $a - b$ between the highest buy order $x = b$ and the lowest sell order $x = a$.³ The instantaneous price can then be conventionally be defined as $p_t^{\text{inst.}} = (a + b)/2$, but will in general not coincide with the coarse-grained price p_t defined by the average shape of the latent order book through Eq. (1). Indeed, as shown in [22], the diffusion width (i.e. the typical distance between $p_t^{\text{inst.}}$ and p_t) is also non zero and actually larger than the intrinsic width, but only by a logarithmic factor.

In the following, we will neglect both the intrinsic width and the diffusion width, which is justified if we focus on price changes much larger than these widths. This is the large scale, low frequency regime where our coarse-grained equations Eqs. (1-a,b) are warranted. Formally, Eqs. (1-a,b) become valid when the market *latent liquidity* \mathcal{L} (defined below) tends to infinity, since both the intrinsic width and the diffusion width vanish as $\mathcal{L}^{-1/2}$ [22].

² In full generality, the diffusion constant D could depend on the distance $|x - p_t|$ to the transaction price. We neglect this possibility in the present version of the model, for reasons that will become clear later: see Appendix B. A similar remark applies to the cancellation rate ν as well.

³ We indeed assume that latent orders become instantaneously visible when close to $p_t^{\text{inst.}}$, in such a way that the latent order book and the observable order book become identical at the best limits. This is of course needed to identify $p_t^{\text{inst.}}$ with the “real” mid-price. It is very interesting to ask what happens if the conversion speed between latent orders and real orders is not infinitely fast, or when market orders become out-sized compared to the prevailing liquidity. As we discuss in the conclusion, this is a potential mechanism for crashes, and the simple coarse-grained framework discussed here has to be adapted to deal with these situations.

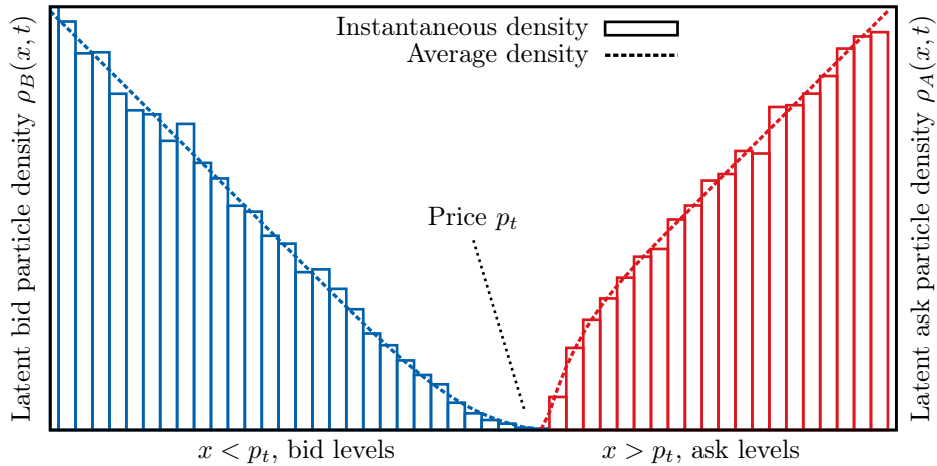


FIG. 1: Snapshot of a latent order book in the presence of a meta-order, with bid orders (blue boxes) and ask orders (red boxes) sitting on opposite sides of the price line and subject to a stochastic evolution. The dashed lines show the mean values of the order densities $\rho_{A,B}(x, t)$, which are controlled by Eqs. (1).

III. STATIONARY SHAPE OF THE LATENT ORDER BOOK

A remarkable feature of Eqs. (1-a,b) is that although the dynamics of ρ_A and ρ_B is non-trivial because of the reaction term (that requires a control of fluctuations, see [22]) the combination $\varphi(x, t) := \rho_B(x, t) - \rho_A(x, t)$ evolves according to a linear equation independent of κ :⁴

$$\frac{\partial \varphi(x, t)}{\partial t} = -V_t \frac{\partial \varphi(x, t)}{\partial x} + D \frac{\partial^2 \varphi(x, t)}{\partial x^2} - \nu \varphi(x, t) + \lambda \text{sign}(p_t - x), \quad (2)$$

where p_t is the solution of $\varphi(p_t, t) = 0$. This solution is expected to be unique for all $t > 0$ if it is unique at $t = 0$ (see also [21]). Note that Eq. (2), without the drift-diffusion terms, has recently been obtained as the hydrodynamic limit of a Poisson order book dynamics in [23].

Introducing $\hat{p}_t = \int_0^t ds V_s$, the above equation can be rewritten in the reference frame of the latent order book $y = x - \hat{p}_t$ as:⁵

$$\frac{\partial \varphi(y, t)}{\partial t} = D \frac{\partial^2 \varphi(y, t)}{\partial y^2} - \nu \varphi(y, t) + \lambda \text{sign}(p_t - \hat{p}_t - y). \quad (3)$$

Starting from a symmetric initial condition $\varphi(y, t = 0) = -\varphi(-y, t = 0)$ such that $p_{t=0} = \hat{p}_{t=0} = 0$, it is clear by symmetry that the equality $p_t = \hat{p}_t$ is a solution at all times, since all terms in the above equation are odd when $y \rightarrow -y$. For more general initial conditions, p_t converges to \hat{p}_t when $t \rightarrow \infty$ and the stationary solution of Eq. (3) reads, in the limit $\mu \rightarrow \infty$:

$$\varphi_{\text{st.}}(y \leq 0) = \frac{\lambda}{\nu} [1 - e^{\gamma y}]; \quad \varphi_{\text{st.}}(y \geq 0) = -\varphi_{\text{st.}}(-y), \quad (4)$$

⁴ The disappearance of κ can be traced to the conservation of $\#A - \#B$ for each reaction $A + B \rightarrow \emptyset$.

⁵ One should be careful with the ‘‘Ito’’ term when V_s is a Wiener noise, which adds a contribution to D , see Appendix A and Eq. (31)

with $\gamma^2 = \nu/D$. This is precisely the solution obtained in [3] which behaves linearly close to the transaction price. But as emphasized in [3] and in Appendix B, this linear behaviour in fact *holds for a very wide range of models* – for example if the appearance of new orders only takes place at some arbitrary boundary $y = \pm L$, as in [17], or else if the coefficients D, ν are non-trivial (but sufficiently regular) functions of the distance to the price $|y|$, etc.

IV. PRICE DYNAMICS WITHIN A LOCALLY LINEAR ORDER BOOK (LLOB)

We will therefore, in the following, “zoom” into the universal linear region by taking the formal limit $\gamma \rightarrow 0$ with a fixed current

$$J = D|\partial_y \varphi_{\text{st.}}|_{y=0} \equiv \lambda/\gamma. \quad (5)$$

This current can be interpreted as the volume transacted per unit time in the stationary regime, i.e. the total quantity of buy (or sell) orders that get executed per unit time. As a side remark, it is important to realize that if the drift V_t contains a Wiener noise component, or jumps, this drift does in fact contribute to J and does not merely shift the latent order book around without any transactions (see Appendix A).

In the limit $\nu, \lambda \rightarrow 0$ with $\lambda/\gamma = J$ fixed, the stationary solution $\varphi_{\text{st.}}(y)$ becomes exactly linear:

$$\varphi_{\text{st.}}(y) = -Jy/D. \quad (6)$$

This is the regime we will explore in the present paper, although we will comment below on the expected modifications induced by non-zero values of ν, λ . Note that $\mathcal{L} = J/D \equiv \lambda\sqrt{D}/\nu$ can be interpreted as the *latent liquidity* of the market, which is large when deposition of latent orders is intense (λ large) and/or when latent orders have a long lifetime (ν small). The quantity \mathcal{L}^{-1} is the analogue, within a LLOB, of Kyle’s “lambda” for a flat order book.

In terms of order of magnitudes, it is reasonable to expect that the latent order book has a memory time ν^{-1} of several hours to several days [3] – remember that we are speaking here of slow actors, not of market makers contributing to the high-frequency dynamics of the revealed order book. Taking D to be of the order of the price volatility, the width of the linear region γ^{-1} is found to be of the order of 1% of the price (see Eq. (4)). Therefore, we expect that restricting the analysis to the *linear* region of the order book will be justified for meta-orders lasting up to several hours, and impacting the price by less than a fraction of a percent. For larger impacts and/or longer execution times, a more elaborate (and probably less universal) description may be needed.

We now introduce a “meta-order” within our framework and work out in detail its impact on the price. Working in the reference frame of the *unimpacted* price \hat{p}_t defined above, we model a meta-order as an extra current of buy (or sell) orders that fall exactly on the transaction price p_t . Introducing $y_t \equiv p_t - \hat{p}_t$, the corresponding equation for the latent order book reads, within a LLOB that precisely holds when $\nu, \lambda \rightarrow 0$:

$$\begin{cases} \frac{\partial \varphi(y, t)}{\partial t} = D \frac{\partial^2 \varphi(y, t)}{\partial y^2} + m_t \delta(y - y_t) \\ \frac{\partial \varphi(y \rightarrow \pm\infty, t)}{\partial y} = -\mathcal{L}, \end{cases} \quad (7)$$

where m_t is the (signed) trading intensity at time t ; $m_t > 0$ corresponding to a buy meta-order. Note that the meta-order will be assumed to be small enough not to change the behaviour of the rest of the market (i.e. the parameters D, ν and λ), so that \mathcal{L} is a fixed parameter in the above equation. Of course, this assumption might break down when the meta-order is out-sized, leading to a sudden increase of the cancellation rate ν and a corresponding drop of the liquidity \mathcal{L} , which might in turn result in a crash (see the discussion in the conclusion).

We will now consider a meta-order that starts at a random time that we choose as $t = 0$, with no information on the state of the latent order book. This means that at $t = 0$, there is no conditioning on the state of the order book that can be described by its stationary shape, $\varphi_{\text{st.}}(y) = -Jy/D$. For $t > 0$, the latent order book is then given by the following exact formula:

$$\varphi(y, t) = -\mathcal{L}y + \int_0^t \frac{ds m_s}{\sqrt{4\pi D(t-s)}} e^{-\frac{(y-y_s)^2}{4D(t-s)}}, \quad (8)$$

where y_s is the transaction price (in the reference frame of the book) at time s , defined as $\varphi(y_s, s) \equiv 0$. This leads to a self-consistent integral equation for the price at time $t > 0$:

$$y_t = \frac{1}{\mathcal{L}} \int_0^t \frac{ds m_s}{\sqrt{4\pi D(t-s)}} e^{-\frac{(y_t-y_s)^2}{4D(t-s)}}. \quad (9)$$

This is the central equation of the present paper, which we investigate in more detail in the next sections.⁶

As a first general remark, let us note that provided impact is small, in the sense that $\forall t, s, |y_s - y_t|^2 \ll D(t-s)$, then the above formula exactly boils down to the *linear propagator model* proposed in [7, 18] (see also [24]), with a square-root decay of impact:

$$y_t = \frac{1}{\mathcal{L}} \int_0^t \frac{ds m_s}{\sqrt{4\pi D(t-s)}}. \quad (10)$$

This linear approximation is therefore valid for very small trading rates m_s , but breaks down for more aggressive executions, for which a more precise analysis is needed. An ad-hoc non-linear generalisation of the propagator model was suggested by Gatheral [24], but is difficult to justify theoretically (and leads to highly singular optimal trading schedules in the continuous time limit [25]). We believe that Eq. (9) above is the correct way to generalize the propagator model, such that all known empirical results can be qualitatively accounted for.

Note that one can in fact define a volume dependent “bid” (or “ask”) price $y_t^\pm(q)$ for a given volume q as the solution of:

$$\int_{y_t^-(q)}^{y_t^+(q)} dy \varphi(y, t) = - \int_{y_t}^{y_t^+(q)} dy \varphi(y, t) = q. \quad (11)$$

Clearly, in the equilibrium state, and for q small enough, $y_t^\pm(q) = y_t \pm \sqrt{2q/\mathcal{L}}$. After a buy meta-order, however, we will find that strong asymmetries can appear.

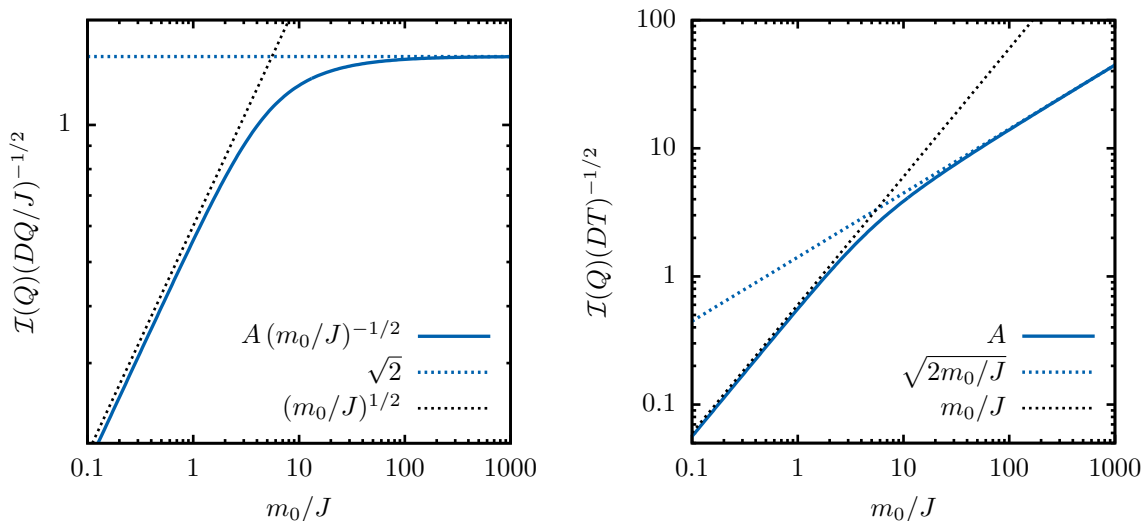


FIG. 2: Left: Dependence of the ratio $A/\sqrt{m_0/J}$ upon the trading rate parameter m_0/J . (This ratio coincides with the empirically used Y ratio if σ^2 is identified with D and V with J). The curve interpolates between a $\sqrt{m_0/J}$ dependence observed at small trading trading rates and an asymptotically constant regime $\approx \sqrt{2}$ for large m_0/J . This is consistent with the weak dependence of Y upon the trading rate observed in CFM empirical data. Right: Dependence of the impact $\mathcal{I}(Q)$ on Q for a *fixed* execution time T – i.e. a variable $m_0 = Q/T$. Note the crossover between a linear behaviour at small Q and a square-root behaviour for large Q .

⁶ When m_s has a non-trivial time dependence, the above equation may not be easy to deal with numerically. It can be more convenient to iterate numerically the Eq. (7) and find the solution of $\varphi(y_t, t) = 0$.

V. THE SQUARE-ROOT IMPACT OF META-ORDERS

The simplest case where a fully non-linear analysis is possible is that of a meta-order of size Q executed at a constant rate $m_0 = Q/T$ for $t \in [0, T]$. In this case, it is straightforward to check that $y_s = A\sqrt{Ds}$ is an *exact* solution of Eq. (9), where the constant A is the solution of the following equation:

$$A = \frac{m_0}{J} \int_0^1 \frac{du}{\sqrt{4\pi(1-u)}} e^{-\frac{A^2(1-\sqrt{u})}{4(1+\sqrt{u})}}. \quad (12)$$

It is easy to work out the asymptotic behaviour of A in the two limits $m_0 \ll J$ and $m_0 \gg J$. In the first case, one finds $A \approx m_0/J\sqrt{\pi}$, while in the second case $A \approx \sqrt{2m_0/J}$. The impact \mathcal{I} of a meta-order of size Q , is defined as:⁷

$$\mathcal{I}(Q) = \langle \varepsilon \cdot (p_{t+T} - p_t) | Q \rangle, \quad (13)$$

where $\langle \dots | Q \rangle$ denotes an average over all meta-orders of sign ε and volume Q , executed over the time interval $[t, t+T]$.

We assuming for now that the meta-order is uninformed, in the following sense:

$$\langle \varepsilon \cdot (\hat{p}_{t+T} - \hat{p}_t) | Q \rangle = 0, \quad (14)$$

such that the only contribution is the “mechanical” impact on the dynamics of y_t . The case of informed meta-orders will be treated in Sect. IX. The mechanical impact at the end of the meta-order is then given by $y_T = A\sqrt{DT}$, i.e.:⁸

$$\mathcal{I}(Q) = \frac{A}{\sqrt{m_0}} \sqrt{DQ} \approx \sqrt{\frac{m_0}{J\pi}} \times \sqrt{\frac{Q}{\mathcal{L}}} \quad (m_0 \ll J); \quad \mathcal{I}(Q) \approx \sqrt{2\frac{Q}{\mathcal{L}}} \quad (m_0 \gg J), \quad (15)$$

i.e. precisely a square-root impact law.

In fact, the empirical result is often written as $\mathcal{I}(Q) = Y\sigma\sqrt{Q/V}$ where σ is the daily volatility and $V \equiv JT_d = D\mathcal{L}T_d$ the daily traded volume ($T_d \equiv 1$ day), and Y a constant of order unity. Assuming that $\sigma^2 \propto DT_d$ (which is the case if $D_0 = 0$, see Appendix A), we see that Eq. (15) exactly reproduces the empirical result, with Y proportional to $\sqrt{m_0/J}$ for small trading intensity m_0 and becoming independent of m_0 for larger trading intensity – see Fig. 2.⁹ CFM’s empirical data indeed suggests that Y only very weakly depends on the trading intensity, which is nicely explained by the present framework.

VI. IMPACT DECAY: BEYOND THE PROPAGATOR MODEL

The next interesting question is impact relaxation: how does the price behave after the meta-order has been executed, i.e. when $t > T$. Mathematically, the impact decay is given by the solution of:

$$y_t = \frac{Dm_0}{J} \int_0^T \frac{ds}{\sqrt{4\pi D(t-s)}} e^{-\frac{(y_t - A\sqrt{Ds})^2}{4D(t-s)}}, \quad (t > T) \quad (16)$$

In the small m_0/J limit, the linear propagation model is appropriate and predicts the following impact relaxation:

$$\frac{\mathcal{I}(Q, t > T)}{\mathcal{I}(Q)} = \frac{\sqrt{t} - \sqrt{t-T}}{\sqrt{T}}, \quad (17)$$

that behaves as $1 - \sqrt{(t-T)/T}$ very shortly after the end of the meta-order and as $\sqrt{T}/t/2$ at long times.

The analysis of Eq. (16) at large m_0/J is more subtle, in particular at short times. The full analysis is given in Appendix C and reveals that the rescaled initial decay of impact is, quite unexpectedly, still exactly given by Eq. (17), independently of m_0/J . For large times, $y_t \rightarrow 0$, which implies that asymptotically $|y_t - A\sqrt{Ds}| \ll \sqrt{Dt}$,

⁷ Note that $\mathcal{I}(Q)$ is a slight abuse of notations since the impact in fact depends in general in the whole trajectory m_s .

⁸ The results in the two limits are (up to prefactors) those obtained in [17] within an explicit reaction-diffusion setting.

⁹ Note that in agreement with our interpretation of the latent order book, the quantity JT must be interpreted as the volume of “slow” orders executed in a time T , removing all fast intra-day activity that averages out and therefore cannot withstand (other than temporarily) the incoming meta-order.

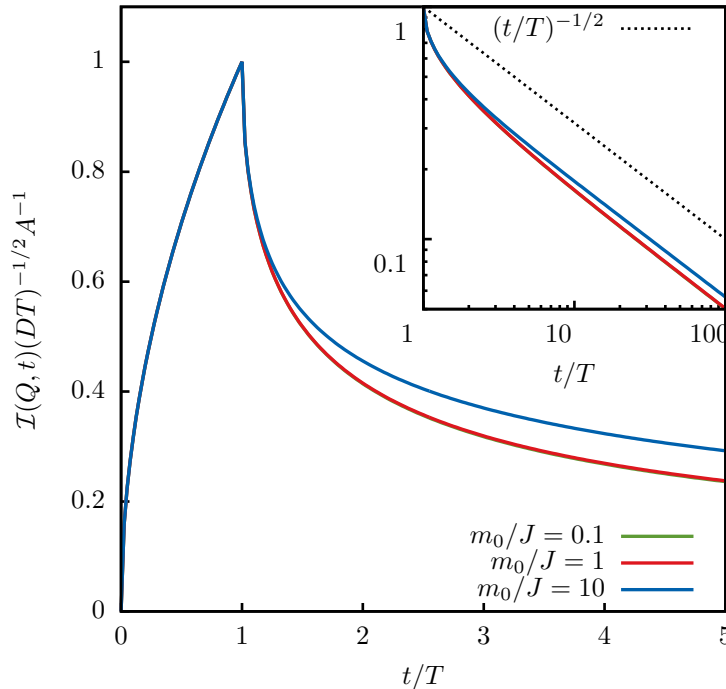


FIG. 3: Impact $\mathcal{I}(Q, t)$ as a function of rescaled time for various trading rate parameters m_0/J . The initial growth of the impact follows exactly a square-root law, and is followed by a regime shift suddenly after end of the meta-order. While for $t = T^+$ the slope of the impact function becomes infinite, at large times one observes an inverse square relaxation $\sim \sqrt{T/t}$ with an m_0/J dependent pre-factor. Note that the curves for $m_0/J = 0.1$ and $m_0/J = 1$ are nearly indistinguishable.

i.e. the exponential term in Eq. (16) is approximately equal to one, leading to an asymptotic rescaled impact decay as $\sqrt{m_0 T}/2\pi J t/4$. We plot in Fig. 3 the normalized free decay of impact for different values of m_0/J for the “mid-price” p_t , and in Fig. 4 the corresponding evolution of the effective “bid-ask” $p_t^\pm(q)$ for a given volume q , illustrating how the latent order book becomes more and more asymmetric as m_0/J increases.

The above analysis can be extended to the case where trading is reverted after time T , i.e. $m_t = m_0$ for $t \in [0, T]$ and $m_t = -m_0$ for $t \in [T, 2T]$. This case is particularly interesting since it puts the emphasis on the lack of liquidity behind the price for large execution rates. Within the linear propagator approximation, it is easy to show that the time needed for the price to come back to its initial value (before continuing to be pushed down by the sell meta-order) is given by $T/4$. In the non linear regime $m_0 \gg J$, the price goes down much faster, and reaches its initial value after a time given by $JT/2m_0 \ll T/4$ – see Figs. 5, 6. Such an asymmetry is indeed seen empirically [26], and means that such a simple round-trip is necessary costly, since the average sell price is below the average buy price. We shall see below that this property (called absence of price manipulation in [27, 28]) holds in full generality within our framework.

VII. PRICE TRAJECTORY AT LARGE TRADING INTENSITIES

Our general price equation Eq. (9) is amenable to an exact treatment in the large trading intensity limit $m_t \gg J$, provided m_t does not change sign and is a sufficiently regular function of time. In such a case, the change of price is large and therefore justifies a saddle-point estimate of the integral appearing in Eq. (9). This leads to the following

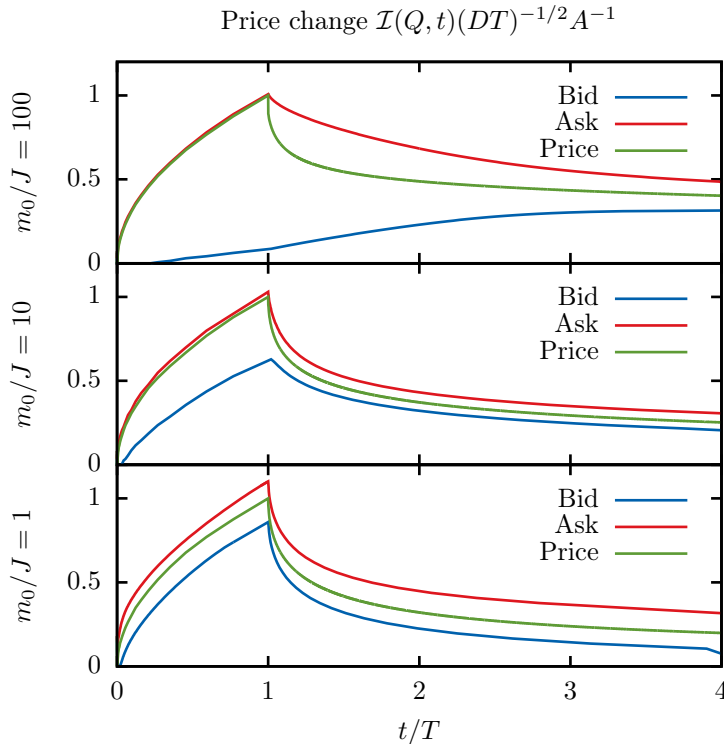


FIG. 4: Evolution in time of the bid $p_t^-(q)$ (blue line) and the ask $p_t^+(q)$ (red line) while executing a meta-order at a rate $m_0/J \in \{1, 10, 100\}$. The price p_t (green line) is also shown for comparison. The three curves correspond to the execution of a constant volume $Q = m_0 T$, while the threshold q has been set by $q = 10^{-3} Q$. The plot illustrates how a large execution rate m_0/J induces a locally asymmetric liquidity profile around the price, see also Fig. 5.

asymptotic equation of motion:

$$\mathcal{L}y_t|\dot{y}_t| \approx m_t \left[1 + D \left(3 \frac{\ddot{y}_t}{\dot{y}_t^3} - 2 \frac{\dot{m}_t}{m_t \dot{y}_t^2} \right) + O\left(\frac{J^2}{m^2}\right) \right]; \quad (18)$$

see Appendix D for details of the derivation and for the next order term, of order J^2/m^2 .

When m_t keeps a constant sign (say positive), the leading term of the above expansion therefore yields the following average impact trajectory:

$$y_t \approx \sqrt{\frac{2}{\mathcal{L}} \int_0^t ds m_s}, \quad (19)$$

i.e. a price impact that only depends on the total traded volume, but not on the execution schedule. This is a stronger result than the one obtained above, where impact was found to be independent of the trading intensity for a uniform execution scheme. This path independence is in qualitative agreement with empirical results obtained at CFM. Using Eq. (18), systematic corrections to the above trajectory can be computed (see Appendix D). Perhaps surprisingly, the execution cost of a given quantity Q is found to be *independent* of the trading schedule even to first-order in J/m – see Appendix D for a proof. Exploring the optimal execution schedule within the full non-linear price equation Eq. (9), and comparing the results with those obtained in Ref. [25], is left for a future study.

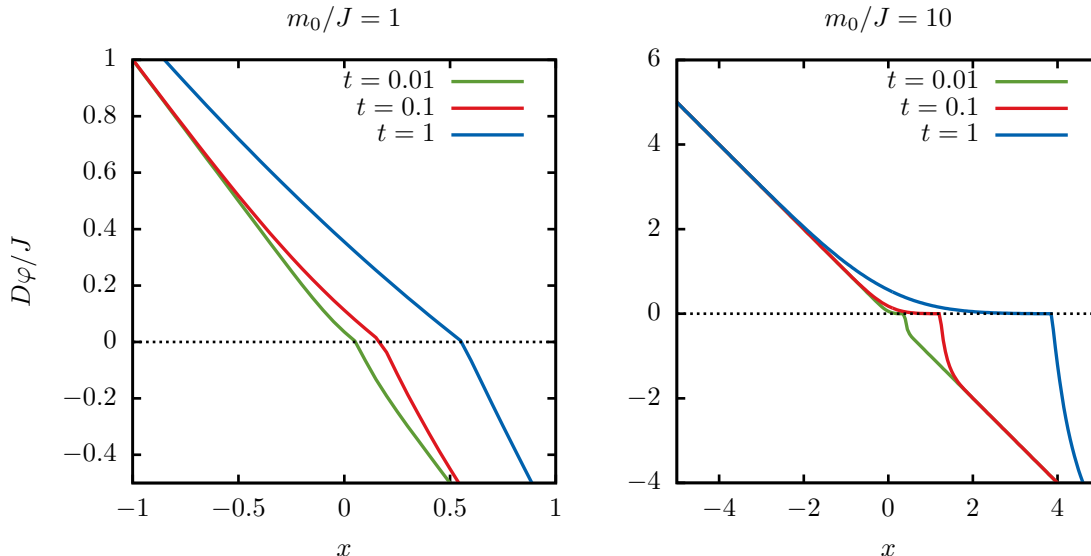


FIG. 5: Evolution of the order book shape $\varphi(x, t)$ during the execution of a meta-order at small trading rate $m_0/J = 1$ (left plot) and large trading rate $m_0/J = 10$ (right plot). The solid lines indicate the profile of the book at $t = 0.01$ (green line), $t = 0.1$ (red line) and $t = 1$ (blue line). While the displacement of the mid-price follows a square root law, the function $D\varphi(x, t)/J + x$ satisfies a scaling relation determined by the parameter m_0/J – see also Appendix C and Fig. 8.

VIII. ABSENCE OF PRICE MANIPULATION

We now turn to a very important issue, that of price manipulation. Although not proven to be impossible in reality, it looks highly implausible that one will ever be able to build a money machine that “mechanically” pumps money out of markets. Any viable model of price impact should therefore be such that mechanical price manipulation, leading to a positive profit after a closed trading loop, is impossible in the absence of information about future prices [27].¹⁰ Here, we show that the non-linear price impact model defined by Eq. (9) is free of price manipulation, generalizing the result of [28] for the linear propagator model, see also [29]. We start by noticing that the average cost of a closed trajectory is given by:

$$\mathcal{C} = \int_0^T ds m_s y_s, \quad \text{with} \quad \int_0^T ds m_s = 0 \quad (20)$$

and y_s given by Eq. (9). The above formula simply means that the executed quantity $m_s ds$ between time s and $s + ds$ is at price y_s .¹¹ Because the initial and final positions are assumed to be zero, there is no additional marked-to-market boundary term. Using Eq. (9), it is not difficult to show that \mathcal{C} can be identically rewritten as a quadratic form:

$$\mathcal{C} = \frac{1}{2} \int_0^T \int_0^T ds ds' m_s M(s, s') m_{s'}, \quad (21)$$

¹⁰ Note that property is highly important for practical purposes as well, since using an impact model with profitable closed trading trajectories in – say – dynamical portfolio algorithms would lead to instabilities.

¹¹ The alert reader might wonder whether m_s is really the *executed* quantity, rather than the *submitted* quantity, as the definition of m_s as a flux of buy/sell orders suggest. However, within the present framework where m_s is deposited precisely at the mid-price p_t , one can check that in the limit $\kappa \rightarrow \infty$, and provided latent and real liquidity are the same close to p_t , the opposite flow of limit orders immediately adapts to absorb exactly the incoming meta-order.

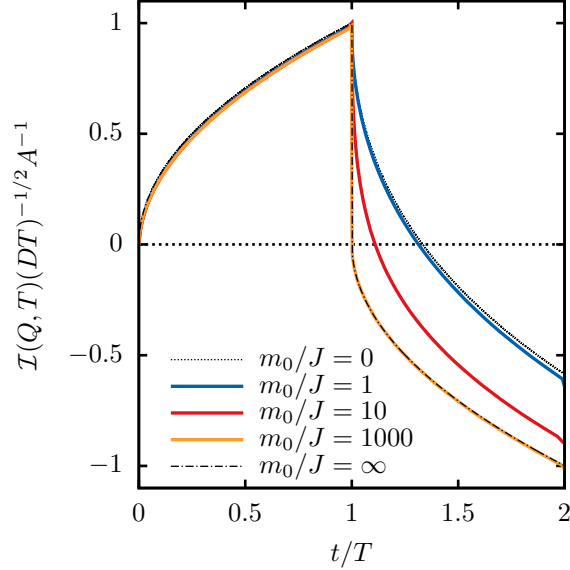


FIG. 6: Trajectory of the average price before and after a sudden switch of the sign of a meta-order. We have considered $m_t = m_0$ for $t < T$ and $m_t = -m_0$ for $t > T$, and plotted the expected price change as a function of time for different values of m_0 . The curves for finite m_0 (solid lines) are also compared with the theoretical benchmark $m_0 = 0$, corresponding to the propagator model (dotted line), and to the $m_0 = \infty$ limit (dot-dashed line). We find that non-linear effects in the large m_0 regime makes to propagator approximation invalid, and increase considerably the impact of the reversal trade.

where $M(s, s')$ is a non-negative operator, since it can be written as a sum of “squares” KK^\dagger , or more precisely:

$$M(s, s') = \frac{D}{\mathcal{L}} \int_{-\infty}^{\infty} dz z^2 \int_{-\infty}^{+\infty} du K_z(s, u) K_z^*(s', u), \quad K_z(s, u) \equiv \Theta(s - u) e^{-Dz^2(s-u) + izy_s}. \quad (22)$$

This therefore proves that $\mathcal{C} \geq 0$ for *any* execution schedule, i.e. price manipulation is impossible within a LLOB (see [6] for loosely related ideas). We note, *en passant*, that this proof extends to a much larger class of Markovian order book dynamics, where the reservation price of latent orders evolves, for example, according to a Lévy process (and not necessarily a diffusion, as assumed heretofore – see Appendix A).

IX. MECHANICAL VS. INFORMATIONAL IMPACT

We now imagine that the agent executing his/her meta-order has some information about the future price, i.e. that the execution flow m_t is correlated with the future motion of the latent order book $V_{t'}$ for $t' > t$. The apparent impact of the meta-order will now contain two contributions that are, within our framework, *additive*. Assuming again, for simplicity, that $m_t = m_0$, one finds that the average price difference can be written as:

$$\langle \varepsilon \cdot (p_t - p_0) | Q \rangle = \langle \varepsilon \cdot (\hat{p}_t - \hat{p}_0) | Q \rangle + \langle \varepsilon \cdot (y_t - y_0) | Q \rangle, \quad (23)$$

where now the first term is non-zero. More explicitly, this leads to:

$$\langle p_t - p_0 \rangle = m_0 \int_0^t ds \int_0^s ds' C(s - s') + A(m_0) \sqrt{Dt}, \quad (t \leq T) \quad (24)$$

where $C(s - s') \propto \langle V_s m_{s'} \rangle$ is a measure of the temporal correlation between meta-orders and future collective latent order moves. Let us insist that we do not assume any causality here: $C(s - s')$ can be interpreted either as the information content of the order that *predicts* future price moves (i.e. the so-called “alpha”), or as the collective reaction of the market to the order flow, i.e. the fact that agents may change their valuation as a result of the trading itself (see [7, 30] for a discussion of this duality).

The second term in the right hand side of Eq. (24) corresponds to the “mechanical” component of the impact discussed above, corresponding to the square-root impact. The first term, on the other hand, may behave very differently as a function of T . For example, if $C(s - s')$ has a range much smaller than T , the first term is expected to grow like Q and not \sqrt{Q} .

When $t > T$, i.e. after the end of the meta-order, the informational contribution adds to the impact decay computed above and can substantially change the apparent evolution of $\langle p_t - p_0 \rangle$. In order to fix ideas, let us assume that $C(s - s') = \Gamma \zeta e^{-\zeta(s-s')}$ (other functional forms would not change the qualitative conclusions below). The behaviour of the “total” impact for $t > T$ is then given by:

$$\mathcal{I}_{\text{tot.}}(Q, t > T) = \mathcal{I}(Q, t > T) + \Gamma Q - \frac{m_0 \Gamma}{\zeta} (1 - e^{-\zeta T}) e^{-\zeta(t-T)} \xrightarrow[t \rightarrow \infty]{} \Gamma Q, \quad (25)$$

which shows that on top of the relaxing mechanical impact (the first term), there is a growing contribution coming from the informational content of the trade (or alternatively from the collective reaction of the market to that trade) that saturates at large time to a finite value proportional to Q – see Fig. 7. This corresponds to a “permanent” component of impact. That the permanent component of impact should be linear in Q conforms well with the assumptions of [1, 2]. However, our calculation shows that the empirical determination of the mechanical component of impact should carefully take into account any possible information content of the analyzed trades, as well as the possible auto-correlation of the trades. This parallels the discussion offered in [11, 13], where attempts are made to measure the decay of mechanical impact $\mathcal{I}(Q, t > T)$ in equity markets, with the conclusion that the mechanical component of impact seems indeed to relax to zero at large times.

The possibility of generating a permanent impact by correlating the collective drift V_t with the flow of meta-orders m_s is in fact important to make our model internally consistent. Absent the permanent impact component, a random flow of meta-orders would give rise to a strongly mean-reverting contribution to the price (on top of the random walk contribution $\hat{p}_t = \int_0^t ds V_s$), and therefore potentially profitable mean-reversion/market making strategies. This profit can however be reduced to zero by increasing the permanent impact component (i.e. the Γ factor above), that acts as an adverse selection bias for market makers. On this point, see the discussion in [31].

X. POSSIBLE EXTENSIONS AND OPEN PROBLEMS

The LLOB framework presented above is surprisingly rich and accounts for many empirical observations, but it can only be a first approximation of a more complex reality. First, we have neglected effects that are in principle contained in Eqs. (1-a,b) but that disappear in the limit of slow latent order books $\nu T \ll 1$ (such that the memory time $1/\nu$ is much longer than the meta-order) and large liquidity \mathcal{L} , such that the meta-order only probes the linear region of the book. Re-integrating these effects perturbatively is not difficult; for example, one finds that the impact $\mathcal{I}(Q)$ of a meta-order of size Q , executed at a constant rate, is lowered by a quantity proportional to νT when $\nu T \ll 1$. In the opposite limit $\nu T \gg 1$, one expects that $\mathcal{I}(Q)$ becomes linear in Q , since impact must become additive in that limit (see [3]). Any other large scale regularisation of the model will lead to the same conclusion. One also expects that the deposition, with rate λ , of new orders behind the moving price should reduce the asymmetry of impact when the trade is reversed. We leave a more detailed calculation of these effects for later investigations.

Another important line of research is to understand the corrections to the LLOB induced by fluctuations, that are of two types: first, as discussed in section II, the theory presented here only deals with the *average* order book $\varphi(x, t)$, from which the price p_t is deduced using the definition $\varphi(p_t, t) = 0$, that allowed us to compute the average impact of a meta-order. However, one should rather compute the impact from the instantaneous definition of the price $p_t^{\text{inst.}}$ (that takes into account the fluctuations of the order book) and then take an average that would lead to $\mathcal{I}(Q)$. The numerical simulations shown in [17] suggest however that the approximation used here is quite accurate for long meta-orders, which is indeed expected as the difference $|p_t^{\text{inst.}} - p_t|$ becomes small compared to $\mathcal{I}(Q)$.

Second, we have assumed that the rest of the market is in its stationary state and does not contribute to the source term modelling the meta-order. One should rather posit that the flow of meta-order m_s has a random component that adds to the particular meta-order that one is particularly interested in. There again, a calculation based on the average order book is not sufficient, since the interaction with other uncorrelated meta-orders then trivially disappears. Following [22], one finds that random fluctuations in m_s do contribute to a strongly mean-reverting term in the variogram of the price trajectory, that should be taken into account in a consistent way. Interestingly, this generic mean-reverting component leads to an excess short-term volatility that is commonly observed in financial markets; more quantitative work on that front would therefore be worthwhile.

Finally, other extensions/modifications may be important in practice: as noted above, the cancellation rate ν is expected to increase with the intensity of meta-orders. Furthermore, the incoming flow of latent orders λ and/or the

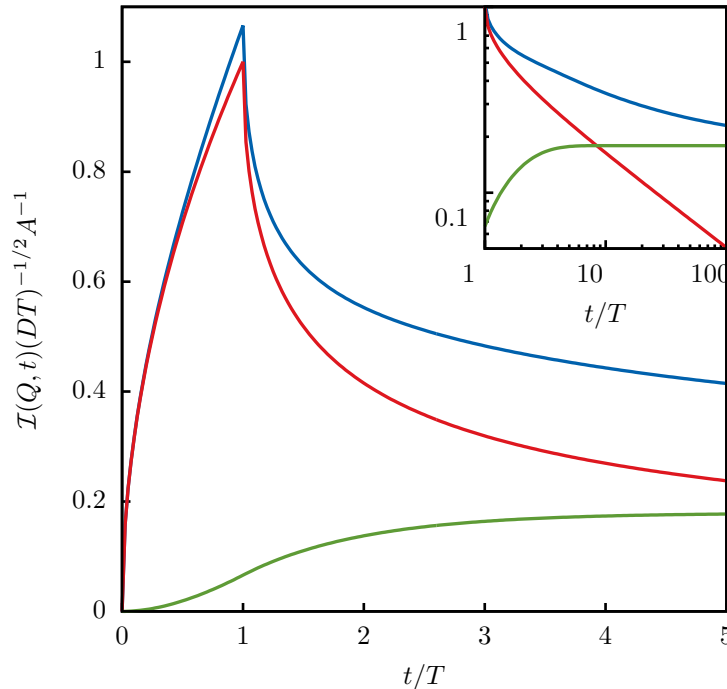


FIG. 7: The figure illustrates the relative rôles of mechanical and informational impact in determining the price trajectory during and after the execution of a meta-order. We have chosen in particular the set of parameters $D = J = \zeta = m = T = 1$ and $\Gamma = 0.1$. The figure indicates that the mechanical part of the impact is the dominating effect at small times. The permanent, informational component of the impact becomes relevant only after the slow decay of the mechanical component, as shown the inset. From a theoretical point of view, the permanent component is important since it counterbalances potential market-making/mean-reversion profits coming from the confining effect of the latent order book on the price.

lifetime of orders $1/\nu$ can be expected to be increasing functions of the distance to the price $|x - p_t|$, i.e. better prices should attract more, and more patient buyers (or sellers), in such a way that the latent order book becomes *convex* at large distances. This would naturally explain why all impact data known to us appear to grow even slower than \sqrt{Q} at large Q ([32, 33], and CFM, unpublished data). Another interesting path would be to allow the “drift” term V_t in Eq. (1-a,b) to become non-Gaussian and thereby study a cumulant expansion of the square-root law.

XI. CONCLUSION

In this paper, we have proposed a minimal theory of non-linear price impact based on a linear (latent) order book approximation, inspired by diffusion-reaction models and general arguments. As emphasized in [3, 5, 17], our modelling strategy does not rely on any equilibrium or fair-pricing conditions, but rather relies on purely statistical considerations. Our approach is strongly bolstered by the universality of the square-root impact law, in particular on the Bitcoin market – as recently documented in [14] – where fair-pricing arguments are clearly unwarranted because impact is much smaller than trading fees.

Our framework allows us to compute the average price trajectory in the presence of a meta-order, that consistently generalizes previously proposed propagator models. Our central result is the dynamical Eq. (9), which not only reproduces the universally observed square-root impact law, but also predicts non-trivial trajectories when trading is interrupted or reversed. Quite surprisingly, we find that the short time behaviour of the free decay of impact is identical to that predicted by a propagator model, whereas the impact of a reversed trade is found to be much stronger. The latter result is in qualitative agreement with empirical observations [26]. We have shown that our model

is free of price manipulation, which makes it the first consistent, non-linear and time dependent theory of impact. Our setting also suggests how prices can be naturally decomposed into a transient “mechanical impact” component and a permanent “informational” component, as initially proposed by Almgren et al. [2], and recently exploited in [11, 13] – see Section IX. Let us insist once again that this decomposition allowed us to construct *diffusive* prices (albeit with a generic short-term mean-reverting contribution).¹²

Although our calculations are based on several approximations (restricting to a locally linear order book and neglecting fluctuations), we believe that it provides a sound starting point for further extensions where the neglected effects can be progressively reinstated. Of particular importance is the potential feedback loop between price moves, order flow and the shape of the latent and of the revealed order books. In particular, we have assumed that latent orders instantaneously materialize in the real order book as the distance to the price gets small: any finite conversion time might however contribute to liquidity droughts, in particular when prices accelerate, leading to an unstable feedback loop. As emphasized in [3, 35, 36] this might be triggered by the anomalous liquidity fluctuations induced by the vanishingly small liquidity in the vicinity of the price. This mechanism could explain the universal power-law distribution of returns that appear to be unrelated to exogenous news but rather to unavoidable, self-induced liquidity crises.

Acknowledgments

We warmly thank M. Abeille, R. Benichou, X. Brokmann, J. de Lataillade, C. Deremble, J. D. Farmer, J. Gatheral, J. Kockelkoren, C. A. Lehalle, Y. Lempérière, F. Lillo, E. Sérié and in particular M. Potters and B. Tóth for many discussions and collaborations on these issues. We also thank P. Blanc, N. Kornman, T. Jaisson, M. Rosenbaum and A. Tilloy for useful remarks on the manuscript. One of us (IM) benefited from the support of the “Chair Markets in Transition”, under the aegis of “Louis Bachelier Finance and Sustainable Growth” laboratory, a joint initiative of École Polytechnique, Université d’Évry Val d’Essonne and Fédération Bancaire Française.

Appendix A: Derivation of the drift/diffusion term

In order to give more flesh to the microscopic assumptions underlying the drift/diffusion equation written in Eqs. (1-a,b), let us assume first that each agent contributes to a negligible fraction of the latent order book, which is probably a good approximation for deep liquid markets. A model for thin markets, where some participants contribute to a substantial fraction of the liquidity, is discussed below, but leads to a very similar final result.

Between t and $t + \delta t$, each agent i revises its reservation price p_i to $p_i + \beta_i \xi_t + \eta_{i,t}$, where ξ_t is common to all i representing some public information (news, but also the price change itself or the order flow, etc.) and $\beta_i > 0$ is the sensitivity of agent i to the news, which we imagine to be a random variable from agent to agent, with a pdf $\Pi(\beta)$ mean normalized to $[\beta_i]_i = 1$. [...] represents a cross-sectional average over agents.] Some agents may over-react, others under-react; β_i might in fact be itself time dependent, but we assume that the *distribution* of β ’s is stationary. The completely idiosyncratic contribution $\eta_{i,t}$ is an independent random variable both across different agents and in time, with distribution $R(\eta)$ of mean zero and rms Σ . We assume that within each price interval $x, x + dx$ lie latent orders from a large number of agents. The density of latent orders $\rho(x, t)$ therefore evolves according to the following Master equation:

$$\rho(x, t + \delta t) = \int_0^\infty d\beta \Pi(\beta) \int_{-\infty}^\infty d\eta R(\eta) \int dy \rho(y, t) \delta(x - y - \beta \xi_t - \eta), \quad (26)$$

or:

$$\rho(x, t + \delta t) = \int_0^\infty d\beta \Pi(\beta) \int_{-\infty}^\infty d\eta R(\eta) \rho(x - \beta \xi_t - \eta, t). \quad (27)$$

Assuming that the price revisions $\beta \xi_t + \eta$ over a small time interval δt are small enough, a second-order expansion Kramers-Moyal of the above equation leads to (see [37] for an in-depth discussion of this procedure):

$$\rho(x, t + \delta t) - \rho(x, t) = -\xi_t \rho'(x, t) + \frac{1}{2} ([\beta^2] \xi_t^2 + \Sigma^2) \rho''(x, t) + \dots \quad (28)$$

¹² A clear theoretical justification of the square-root impact law is also important if one wants to promote the idea of impact discounted mark-to-market accounting rules, as advocated in [34].

At this stage, one can either assume that formally $\xi_t = V_t \delta t$ and $\Sigma^2 = 2D_0 \delta t$ in which case the continuous time limit reads:

$$\frac{\partial \rho(x, t)}{\partial t} = -V_t \frac{\partial \rho(x, t)}{\partial x} + D_0 \frac{\partial^2 \rho(x, t)}{\partial x^2} \quad (29)$$

or that $\xi_t = V_t \sqrt{\delta t}$, where V_t is now a Gaussian white noise of variance σ^2 , and again $\Sigma^2 = 2D \delta t$, in which case the continuous time limit should be written as:

$$d\rho(x, t) = -dW_t \frac{\partial \rho(x, t)}{\partial x} + D_1 dt \frac{\partial^2 \rho(x, t)}{\partial x^2} \quad (30)$$

with dW_t a Wiener noise and $D_1 \equiv D_0 + [\beta^2] \sigma^2 / 2$, to wit, the diffusion constant involves both the idiosyncratic component and the dispersion of reaction to random information. This is the interpretation we will mostly follow in the present paper. A careful derivation of the corresponding equation in the reference frame of the price $\hat{p}_t = \int_0^t dW_s$ finally gives the diffusion part of Eq. (3) in the main text:

$$\frac{\partial \rho(y, t)}{\partial t} = D \frac{\partial^2 \rho(y, t)}{\partial y^2}; \quad D \equiv D_0 + \frac{\sigma^2}{2} \int d\beta \Pi(\beta) (\beta - 1)^2; \quad (31)$$

i.e. only the dispersion of reaction $\beta - 1$ can contribute to the diffusion term, as expected.

Another interpretation of this last equation is to imagine that between t and $t + dt$, a fraction $\phi \in [0, 1]$ (possibly time dependent) of agents collectively change their price estimate by an amount dW_t , with no other idiosyncratic component. This leads to:

$$d\rho(x, t) = \phi [\rho(x - dW_t, t) - \rho(x, t)] = -\phi dW_t \rho'(x, t) + \frac{1}{2} \phi \sigma^2 dt \rho''(x, t), \quad (32)$$

that essentially corresponds to the case above with $\Pi(\beta) = (1 - \phi) \delta(\beta) + \phi \delta(\beta - 1)$. In the price reference frame $\hat{p}_t = \int_0^t \phi dW_s$, one finds Eq. (31) with $D \equiv \phi(1 - \phi) \sigma^2 / 2$. Note that, clearly, these collective price revisions must by themselves induce transactions whenever $0 < \phi < 1$.

If price revisions cannot be considered as small, the resulting evolution of $\rho(x, t)$ should include jumps in the continuous time limit, i.e. one would find an integro-differential equation rather than a partial differential equation for $\rho(x, t)$. However, if the jump process is homogeneous in space, one can diagonalize the evolution operator in Fourier space. This allows one to show that price manipulation is impossible in that case as well.

Appendix B: A generically linear latent order book

Let us consider the case where the deposition flow is not constant. This leads to the following equation for the stationary state of the latent order book:

$$D \frac{\partial^2 \varphi_{\text{st.}}(y)}{\partial y^2} - \nu \varphi_{\text{st.}}(y) + \lambda (\Theta(y) - \Theta(-y)) = 0, \quad (33)$$

with $\varphi_{\text{st.}}(y) = -\varphi_{\text{st.}}(-y)$ (and in particular the market clearing condition $\varphi_{\text{st.}}(y = 0) = 0$).

Let us assume that $\Theta(y) - \Theta(-y)$ behaves, for $y \rightarrow \infty$, as a constant that we can set to unity. The solution $\varphi_{\text{st.}}(y)$ then converges to λ/ν for large y , so we set:

$$\varphi_{\text{st.}}(y) = \frac{\lambda}{\nu} + \Psi(y), \quad (34)$$

where $\Psi(y) = 1 - \Theta(y) + \Theta(-y)$ with $\Psi(y \rightarrow \infty) \rightarrow 0$, and:

$$D \frac{\partial^2 \Psi(y)}{\partial y^2} - \nu \Psi(y) = \lambda \Xi(y), \quad (35)$$

where $\Xi(y \rightarrow \infty) \rightarrow 0$. The boundary condition on $\Psi(y)$ at large y means that we can look at a solution of the form:

$$\Psi(y) = \psi(y) e^{-\sqrt{\nu/D} y}, \quad (36)$$

so that:

$$D\psi''(y) - 2\sqrt{\nu D}\psi'(y) = \lambda\Xi(y)e^{\sqrt{\nu/D}y}. \quad (37)$$

Finally, one finds:

$$\varphi_{\text{st.}}(y) = \frac{\lambda}{\nu} \left[1 - e^{-\sqrt{\nu/D}y} \right] + \frac{\lambda}{D} e^{-\sqrt{\nu/D}y} \int_0^y dy' e^{2\sqrt{\nu/D}y'} \int_{y'}^\infty dy'' e^{-\sqrt{\nu/D}y''} \Xi(y''). \quad (38)$$

The solution given in the main text corresponds to $\Xi \equiv 0$, so that only the first term survives. From the above explicit form, one sees that provided the integral $\int_0^\infty dy'' e^{-\sqrt{\nu/D}y''} \Xi(y'')$ is finite, the behaviour of $\varphi_{\text{st.}}(y)$ is *always* linear in the vicinity of $y = 0$. Only a highly singular the deposition rate, diverging faster than $1/y$ when $y \rightarrow 0$, would jeopardize the local linearity of the latent order book (see also [38], where this property was first discussed).

Appendix C: Shape of the order book during constant rate execution and initial relaxation of impact

When the trading rate is a constant (equal to m_0), one can exhibit an *exact* scaling solution of the time dependent order book of the form $\varphi(x, t) = m_0\sqrt{t/D}F(\frac{x}{\sqrt{Dt}})$, where F is the solution of:

$$2F''(u) + uF'(u) - F(u) = -2\delta(u - A). \quad (39)$$

As we show below, this equation can be solved and gives the exact shape of the book at $t = T$, from which the initial relaxation (after trading has stopped) can be deduced.

Writing $F = uG$ in the above equation, one finds a first order linear equation for $H = G'$:

$$H'(u) + \left(\frac{u}{2} + \frac{2}{u} \right) H(u) = \frac{1}{u} \delta(u - A) \quad (40)$$

which is easily solved as:

$$H(u) = \frac{H_0}{u^2} e^{-u^2/4}, \quad (u < A); \quad H(u) = \frac{H_0 - Ae^{A^2/4}}{u^2} e^{-u^2/4}, \quad (u > A). \quad (41)$$

There are two boundary conditions that are useful. One is the very definition of the price position, $x = A\sqrt{Dt}$ or $A = u$, for which $\phi(x, t) = x$ or $F(A) = AJ/m_0$. The second remark is that when $u = 0$, the integral defining F can be computed, leading to

$$F(0) = \frac{A}{4\sqrt{\pi}} e^{A^2/4} \int_{A^2/4}^\infty \frac{dv}{v^{3/2}} e^{-v}. \quad (42)$$

This allows one to fix H_0 since $G'(u) = F'(0)/u - F(0)/u^2 \approx -F(0)/u^2$ when $u \rightarrow 0$, to be compared with $H(u) \approx H_0/u^2$ in the same limit. Hence $H_0 = -F(0)$.

The final solution for $F(u)$ is easily obtained from integrating H and multiplying by u (see Fig. 8). Using $G(A) = J/m_0$, one finds:

$$G(u) = F(0) \int_u^A \frac{dv}{v^2} e^{-v^2/4} + J/m_0, \quad (u \leq A) \quad (43)$$

and

$$G(u) = -(F(0) + Ae^{A^2/4}) \int_A^u \frac{dv}{v^2} e^{-v^2/4} + J/m_0, \quad (u \geq A) \quad (44)$$

Of special interest is the slope of F for $u = A^\pm$; with $F' = uH(u) + G(u)$ one finds:

$$F'(A^-) = J/m_0 - \frac{F(0)}{A} e^{-A^2/4}; \quad F'(A^+) = J/m_0 - \frac{F(0)}{A} e^{-A^2/4} - 1. \quad (45)$$

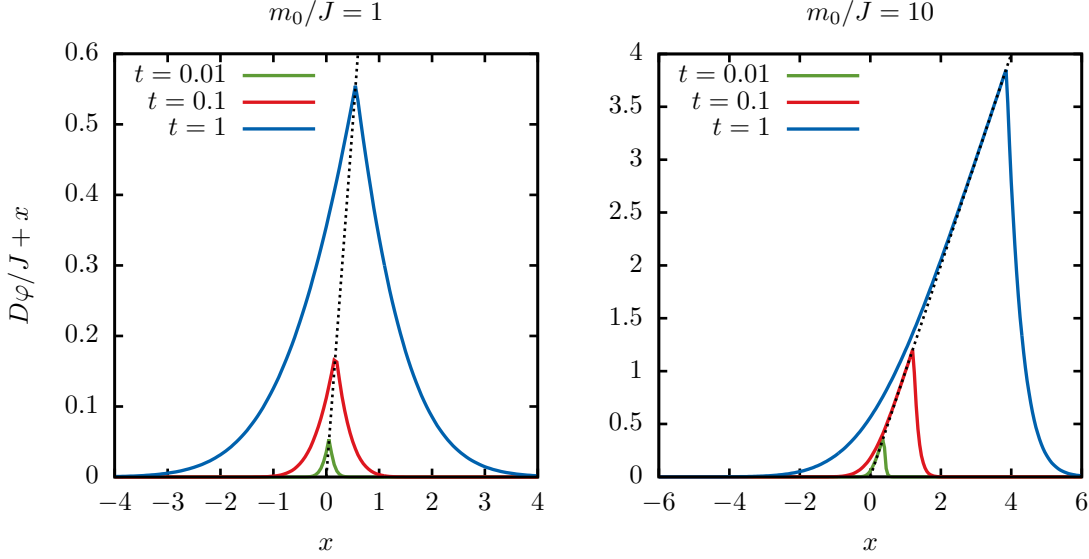


FIG. 8: Evolution of the average order book $\varphi(x, t)$, represented for two different values of the perturbation parameter m_0/J . The curves in different colors are snapshots taken at different times of the difference between the perturbed and the unperturbed average of the book. The scaling form of the book $\varphi(x, t) = m_0 \sqrt{t/D} F(\frac{x}{\sqrt{Dt}})$ is clear from the plots, which illustrate how a flat region of the book is formed at large values of m_0/J .

Now, shortly after the meta-order has stopped, one can look at the solution of the diffusion equation in the vicinity of the final price $p_T = A\sqrt{DT}$, using a piece-wise linear function for the initial condition, with slopes given by $F'(A^\pm)$. The solution then reads, with $t - T = \Delta$ small:

$$\varphi(x, t) = m_0 \sqrt{T\Delta} \left[F'(A^-)z + (F'(A^+) - F'(A^-)) \int_z^\infty du \frac{(u-z)}{\sqrt{4\pi}} e^{-u^2/4} \right], \quad (46)$$

with $z = (p_T - x)/\sqrt{D\Delta}$. The position of the price is given by $p_t = p_T - \sqrt{D\Delta}z^*$, with z^* such that:

$$F'(A^-)z^* + (F'(A^+) - F'(A^-)) \int_{z^*}^\infty du \frac{(u-z^*)}{\sqrt{4\pi}} e^{-u^2/4} = 0. \quad (47)$$

Using the expression for $F(0)$ and the result above for $F'(A^\pm)$, this equation simplifies to:

$$z^* \int_{A^2/4}^\infty \frac{dv}{v^{3/2}} e^{-v} = 2 \int_{z^*}^\infty du (u-z^*) e^{-u^2/4}. \quad (48)$$

Changing variables in the RHS from u to $v = u^2/4$, and integrating by parts, one finds:

$$z^* \int_{A^2/4}^\infty \frac{dv}{v^{3/2}} e^{-v} = z^* \int_{z^{*2}/4}^\infty \frac{dv}{v^{3/2}} e^{-v}, \quad (49)$$

which leads to $z^* = A$ for all m_0/J . [The other solution, $z^* = 0$, is spurious].

Hence, the initial stage of the impact relaxation can be written in a super-universal way:

$$p_t \underset{t \rightarrow T^+}{\approx} p_T \left[1 - \sqrt{\frac{t-T}{T}} \right], \quad (50)$$

i.e. exactly the result from the propagator model, even in the non-linear regime!

Appendix D: A saddle point approximation for large trading rates

In this Appendix we develop a systematic procedure in order to solve perturbatively Eq. (9). The first step in order to find a solution is to introduce an expansion parameter $\epsilon \ll 1$, which we use in order to control the amplitude of the trading rate through $m_t \rightarrow m_t \epsilon^{-1}$. Such substitution implies a scaling of the solution of the form $y_t \rightarrow y_t \epsilon^{-1/2}$, leading to an equation for the price of the form:

$$\mathcal{L}y_t = \int_0^t \frac{ds m_s}{\sqrt{4\pi D\epsilon(t-s)}} e^{-\frac{(y_t - y_s)^2}{4D\epsilon(t-s)}}. \quad (51)$$

which is equivalent to the one which one would have by leaving invariant m_t and by performing the substitutions $D \rightarrow D\epsilon$ and $J \rightarrow J\epsilon$. Hence, the large trading regime is equivalent to the one of slow diffusion.

In this case, it is evident that the integral is dominated by times s close to t , which suggests to Taylor expand both the trading rate m_s and the price y_s around $s = t$, so to insert the resulting series in the integral appearing in Eq. (9). The dominating term results

$$m_t \int_0^\infty \frac{du}{\sqrt{4\pi D u \epsilon}} e^{-\frac{y_t^2 u}{4D\epsilon}} = m_t |\dot{y}_t|^{-1}. \quad (52)$$

The successive corrections to above result can be computed systematically, as they involve developing the square and the exponential function in the Gaussian term in Eq. (9). In particular, by exploiting the identity

$$\int_0^\infty du e^{-z^2 u} u^\alpha = \Gamma(1 + \alpha) |z|^{-2(1+\alpha)} \quad (53)$$

it is possible to derive the expansion

$$\begin{aligned} \mathcal{L}y_t |\dot{y}_t| &= m_t \left[1 + (D\epsilon) \left(\frac{3\ddot{y}_t}{\dot{y}_t^3} - \frac{2\dot{m}_t}{m_t \dot{y}_t^2} \right) \right. \\ &+ (D\epsilon)^2 \left(\frac{6\ddot{m}_t \dot{y}_t^2 - 30\dot{m}_t \ddot{y}_t \dot{y}_t - 10m_t \ddot{y}_t \dot{y}_t + 45m_t \dot{y}_t^2}{m_t \dot{y}_t^6} \right) \\ &+ 5(D\epsilon)^3 \left(\frac{-4\ddot{m}_t \dot{y}_t^3 + 42\ddot{m}_t \ddot{y}_t \dot{y}_t^2 + 28\dot{m}_t \ddot{y}_t \dot{y}_t^2 + 7m_t \ddot{y}_t \dot{y}_t^2}{m_t \dot{y}_t^9} \right) \\ &+ 5(D\epsilon)^3 \left(\frac{-168\dot{m}_t \ddot{y}_t^2 \dot{y}_t - 112m_t \ddot{y}_t \dot{y}_t \dot{y}_t + 252m_t \dot{y}_t^3}{m_t \dot{y}_t^9} \right) \\ &\left. + O(\epsilon^4) \right], \end{aligned} \quad (54)$$

whose first-order terms match the form reported in Eq. (18). Each of the contributions of order ϵ^n can be seen equivalently either as suppressed by the small value of the diffusion constant diffusion (through a D^n factor) or by the large value trading rate (through a factor of the order of m_t^{-n}).

Finally, note that the implicit equation above needs to be inverted in order to obtain a relation yielding y_t as a function of m_t . This is possible by using Eq. (54) as an iterative scheme for y_t , which allows to calculate

$$\begin{aligned} \mathcal{L}y_t |\dot{y}_t| &= m_t \\ &+ (J\epsilon) \left(-3 + \frac{2Q_t \dot{m}_t}{m_t^2} \right) \\ &+ (J\epsilon)^2 \left(-\frac{12Q_t \dot{m}_t}{m_t^3} - \frac{6s \dot{m}_t}{m_t^2} + \frac{4\dot{m}_t}{m_t^2} \int ds \frac{Q_s \dot{m}_s}{m_s^2} + \frac{16Q_t^2 \dot{m}_t^2}{m_t^5} - \frac{4Q_t^2 \dot{m}_t}{m_t^4} \right) \\ &+ O(\epsilon^3), \end{aligned} \quad (55)$$

where $Q_t = \int_0^t ds m_s$.

As a simple application of the above formula, consider the case where $m_t \geq 0, \forall t \in [0, T]$. In this case, \dot{y}_t is also non negative and we can remove the absolute value in the above equation. To order ϵ , the solution of the above equation is (assuming $y_0 = 0$):

$$\frac{1}{2} \mathcal{L}y_t^2 = Q_t + J\epsilon \left(-t - \frac{2Q_t}{m_t} \right), \quad (56)$$

where we have used integration by parts. Now, the cost $\mathcal{C}(Q)$ associated to buying a total quantity Q in time T is given by:

$$\mathcal{C}(Q) = \int_0^T ds m_s y_s, \quad Q = \int_0^T ds m_s. \quad (57)$$

Therefore, to order ϵ :

$$\mathcal{C}(Q) = \sqrt{\frac{2}{\mathcal{L}}} \int_0^T ds m_s \sqrt{Q_s} - \frac{J\epsilon}{\sqrt{2\mathcal{L}}} \int_0^T ds \left[2\sqrt{Q_s} + \frac{sm_s}{\sqrt{Q_s}} \right]. \quad (58)$$

After integrating by parts the last term, one finally finds that the impact cost is, to order ϵ , *independent* of the trading schedule m_s , and given by:

$$\mathcal{C}(Q) = \frac{2}{3} \sqrt{\frac{2}{\mathcal{L}}} Q^{3/2} \left[1 - \frac{3J\epsilon T}{2Q} \right]. \quad (59)$$

The correction term is negative, as expected since a slower trading speed leaves time for the opposing liquidity to diffuse towards the traded price.

-
- [1] A. Kyle, *Continuous auctions and insider trading*, *Econometrica: Journal of the Econometric Society*, 1315-1335 (1985).
[2] R. Almgren, C. Thum, E. Hauptmann, H. Li, *Direct estimation of equity market impact*, *Risk*, 18(7), 5762 (2005).
[3] B. Tóth, Y. Lempérière, C. Deremble, J. De Lataillade, J. Kockelkoren, J.-P. Bouchaud, *Anomalous price impact and the critical nature of liquidity in financial markets*, *Physical Review X*, 1(2), 021006 (2011).
[4] J. D. Farmer, A. Gerig, F. Lillo, and H. Waelbroeck, *How efficiency shapes market impact*, *Quantitative Finance*, 13, 1743 (2013).
[5] I. Mastromatteo, B. Tóth, J.-P. Bouchaud, *Agent-based models for latent liquidity and concave price impact* *Phys. Rev. E* 89, 042805 (2014).
[6] I. Skachkov, *Market Impact Paradoxes*, *Wilmott*, 2014, 7178 (2014).
[7] 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*, Elsevier: Academic Press, 57-156 (2008).
[8] N. Torre, M. Ferrari, *Market impact model handbook*, BARRA Inc., Berkeley (1997), available at <http://www.barra.com/newsletter/nl166/miminl166.asp>, but see also [9].
[9] R. C. Grinold, R. N. Kahn, *Active Portfolio Management* (New York: The McGraw-Hill Companies, Inc., 1999).
[10] 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(6), 066102 (2009).
[11] C. Gomes, H. Waelbroeck, *Is Market Impact a Measure of the Information Value of Trades? Market Response to Liquidity vs Informed Trades*, *Quantitative Finance* (2014), DOI: 10.1080/14697688.2014.963140.
[12] N. Bershova, D. Rakhlin, *The Non-Linear Market Impact of Large Trades: Evidence from Buy-Side Order Flow*, *Quantitative Finance* 13, 1759-1778 (2013).
[13] X. Brokmann, E. Sérié, J. Kockelkoren, J.-P. Bouchaud, *Slow decay of impact in Equity markets*, arXiv:1407.3390, submitted to *Market Microstructure and Liquidity*.
[14] J. Donier, J. Bonart, *A million metaorder analysis of market impact on the bitcoin*, arXiv:1412.4503, submitted to *Market Microstructure and Liquidity*.
[15] P. Bak, M. Paczuski, M. Shubik, *Price variations in a stock market with many agents*, *Physica A: Statistical Mechanics and its Applications*, 246(3), 430-453 (1997).
[16] L. H. Tang, G. S. Tian, *Reaction-diffusion-branching models of stock price fluctuations*. *Physica A: Statistical Mechanics and its Applications*, 264(3), 543-550 (1999).
[17] I. Mastromatteo, B. Tóth, J.-P. Bouchaud, *Anomalous impact in reaction-diffusion models*, *Phys. Rev. Lett.* 113, 268701 (2014).
[18] J.-P. Bouchaud, Y. Gefen, M. Potters, M. Wyart, *Fluctuations and response in financial markets: the subtle nature of "random" price changes*. *Quantitative Finance*, 4(2), 176-190 (2004).
[19] D. E. Taranto, G. Bormetti, F. Lillo, *The adaptive nature of liquidity taking in limit order books*, *J. Stat. Mech.* (2014) P06002.
[20] E. Smith, J. D. Farmer, L. Gillemot, and S. Krishnamurthy, *Statistical theory of the continuous double auction*, *Quantit. Financ.* 3, 481 (2003).
[21] C.-A. Lehalle, O. Guéant, and J. Razafinimanana, *High frequency simulations of an order book: a Two-Scales approach*, in F. Abergel, B.K. Chakrabarti, A. Chakraborti, and M. Mitra (Eds.), *Econophysics of Order-Driven Markets*, New Economic Windows. Springer (2010).

- [22] G.T. Barkema, M.J. Howard, J.L. Cardy, *Reaction-diffusion front for $A + B \rightarrow 0$ in one dimension*, Physical Review E, 53(3), R2017 (1996).
- [23] X. Gao, J. G. Dai, A. B. Dieker, and S. J. Deng, *Hydrodynamic limit of order book dynamics*, arXiv:1411.7502.
- [24] J. Gatheral, *No-dynamic-arbitrage and market impact*. Quantitative Finance, 10, 749-759 (2010).
- [25] G. Curato, J. Gatheral, F. Lillo, *Optimal execution with nonlinear transient market impact*, arXiv:1412.4839.
- [26] R. Benichou, CFM internal report (2013), and R. Benichou et al., in preparation.
- [27] G. Huberman, W. Stanzl, *Price Manipulation and Quasi-Arbitrage*, Econometrica, **72**, 12471275 (2004).
- [28] A. Alfonsi, A. Schied, *Optimal trade execution and absence of price manipulations in limit order book models* SIAM Journal on Financial Mathematics 1 (1), 490-522 (2010).
- [29] A. Alfonsi, P. Blanc, *Dynamic optimal execution in a mixed-market-impact Hawkes price model*, arXiv:1404.0648.
- [30] J.-P. Bouchaud, *Price Impact*, in Encyclopedia of Quantitative Finance, John Wiley & Sons Ltd. (2010).
- [31] M. Wyart, J.-P. Bouchaud, J. Kockelkoren, M. Potters, M. Vettorazzo, *Relation between Bid-Ask Spread, Impact and Volatility in Double Auction Markets*, Quantitative Finance, 8, 41-57 (2008).
- [32] E. Zarinelli, M. Treccani, J. D. Farmer, F. Lillo, *Beyond the square root: Evidence for logarithmic dependence of market impact on size and participation rate*, arXiv:1412.2152.
- [33] C.-A. Lehalle, private communication, and E. Bacry, A. Iugay, M. Lasnierz, C.-A. Lehalle, *Market impacts and the life cycle of investors orders*, arXiv:1412.0217.
- [34] F. Caccioli, J.P. Bouchaud, J. D. Farmer, *Impact-adjusted mark-to- market valuation and the criticality of leverage*, arXiv preprint 1204.0922, published in Risk Magazine, May 2012.
- [35] F. Lillo, J.D. Farmer, *The key role of liquidity fluctuations in determining large price fluctuations*, Fluctuations and Noise Letters, 5, L209 (2005).
- [36] 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).
- [37] C. W. Gardiner, *Stochastic Methods*, 4th Edition, Springer, Berlin (2009).
- [38] J.-P. Bouchaud, M. Mézard, M. Potters, *Statistical properties of stock order books: empirical results and models*, Quantitative Finance 2, 251 (2002)