



Modeling metaorder impact with a Non-Markovian Zero Intelligence model

Adele Ravagnani ^{a,b} ^{*}, Fabrizio Lillo ^{a,c}

^a Scuola Normale Superiore, Piazza dei Cavalieri 7, Pisa, 56126, Italy

^b Dipartimento di Economia Politica e Statistica, Università di Siena, P.zza S.Francesco 7-8, Siena, 53100, Italy

^c Dipartimento di Matematica, Università di Bologna, Piazza di Porta San Donato 5, Bologna, 40126, Italy

ARTICLE INFO

Dataset link: lobsterdata.com

Keywords:

Limit order book
Market microstructure
Price impact
Market simulators

ABSTRACT

Devising models of the limit order book that realistically reproduce the market response to exogenous trades is extremely challenging and fundamental in order to test trading strategies. We propose a novel explainable model for small tick assets, the Non-Markovian Zero Intelligence, which is a variant of the well-known Zero Intelligence model. The main modification is that the probability of limit orders' signs (buy/sell) is not constant but is a function of the exponentially weighted mid-price return, representing the past price dynamics, and can be interpreted as the reaction of traders with reservation prices to the price trend. With numerical simulations and analytical arguments, we show that the model predicts a concave price path during a metaorder execution and to a price reversion after the execution ends, as empirically observed. We analyze in-depth the mechanism at the root of the arising concavity, the components which constitute the price impact in our model, and the dependence of the results on the two main parameters, namely the time scale and the strength of the reaction of traders to the price trend.

1. Introduction

In electronic markets, trade intentions of market participants are handled by limit order books (LOBs). They are collections of buy and sell orders, which are matched on a price-time priority basis, with a continuous-time double auction mechanism. The resulting dynamics are complex, extremely fast, they lead to a self-organized price formation process and summarize the interactions between a diversity of agents with different objectives and trading strategies. In order to reproduce these dynamics and their statistical properties (so-called *stylized facts*), researchers have been engaged in this challenging modeling task since the early 2000s, resulting in a vast amount of papers [1]. Obtaining a realistic simulator of the LOB dynamics is fundamental both for academic and practical interests. It allows to gain theoretical insight about financial markets and their functioning. On the other hand, it allows to precisely test trading strategies before they are used in real markets. These tests are usually performed by relying on the so-called *backtesting*, which consists in simulating the performance of the trading strategies by relying on historical data, under several scenarios. However, even though the backtesting approach is commonly adopted, it is inaccurate: contrary to employing LOB simulators, it does not allow to consider the market response to each action of the strategy we aim to test. In the literature, several refined and realistic models that allow to perform these tests have been introduced. In our work, we do not want to propose a new simulator instead, we aim to show how starting from a simple LOB model based on Poisson processes, the introduction of a non-Markovian mechanism allows to capture realistically the shape of the market response when exogenous orders are submitted. Our model is explainable and this is fundamental to us since we aim to shed light on the mechanisms at the root of market impact dynamics.

^{*} Corresponding author.

E-mail addresses: adele.ravagnani@sns.it, adele.ravagnani@unisi.it (A. Ravagnani), fabrizio.lillo@sns.it (F. Lillo).




Fig. 1. Expected average shape of the market impact during and after the execution of a buy metaorder. Q is the metaorder size (in shares), T is its duration.

In this paper we are mostly interested in the execution of large orders, the so-called *metaorders*. As it is known since the seminal Kyle paper [2], it is optimal to split the metaorder into small *child market orders* that are executed progressively. This allows to limit the *market impact*, which is the positive correlation between the sign of an incoming order and the subsequent price change, and which represents the main source of trading costs for medium and large size investors. In a LOB simulator, the desired behavior of a metaorder execution's responsiveness is its concavity, as it is empirically observed and studied in several papers [3]. Concerning the behavior of the price at the end of the metaorder execution, several works empirically and theoretically studied this issue. It emerged that, following the peak price reached at the end of the execution, a reversion occurs, with a decay and converging to a plateau [4–7]. A graphical representation of the expected average shape of the market impact during and after the execution of a buy metaorder is provided in Fig. 1. The Non-Markovian Zero Intelligence LOB model we propose in this paper, aims to reproduce the empirically observed market impact shape. Our model is an extension of a well-established simulator, the so-called Zero Intelligence model [8–10], which leads to a linear market impact when executing a metaorder.

1.1. Related literature

Since the first studies of the early 2000s, a vast amount of research articles [1] has been devoted to the challenging task of modeling the dynamics of the Limit Order Book (LOB). Following [1], four broad categories of LOB simulators can be identified: Point Processes, Agent-Based, Deep Learning Based and Stochastic Differential Equations Based models. Among the Point Processes models, the Zero Intelligence simulators [9,11–13] consider orders arrival as independent Poisson processes. The Zero Intelligence model [8–10] (aka Santa Fe model) which is extended in this paper, belongs to this class of models. Despite the simplicity of the Zero Intelligence simulators, they allow to reproduce several stylized facts, such as the spread distribution. However, especially given the unrealistic assumption of orders' independence, when metaorders are executed by employing these simulators, the market impact results to be linear in time and in addition, other important properties as the inter-arrival times and cross-excitation behavior, cannot be replicated. In order to overcome this major limitation, Poisson models with variable order intensity, as [14–16], were proposed. For instance, the Queue-Reactive model [15] considers intensities which depend on the current state of the order book. An impressive achievement of this model, which makes it stand out among Point Process simulators, is its ability to reproduce the concavity of the market impact against both time and volume. However, this model still lacks of endogenously excited order flow. This issue motivated the development of another set of simulators: the Hawkes Process based, as [17–22]. They are fully explainable and able to correctly reproduce relevant microstructural properties as the volatility clustering and the Epps effect. However, their calibration is not straightforward and the choice of the kernel can be crucial for the model's performance.

A different approach is employed by Agent Based simulators [23–26]. The LOB dynamics result from the interaction of several categories of investors (e.g. informed, uninformed, high frequency, trend followers, mean reverters, noise and algorithmic traders) and each of them is modeled differently. Some of these models, as [24–26], return a concave market impact. On the other hand, two main drawbacks of this class of simulators can be outlined: the modeling of the different types of investors' behavior is mostly heuristic and the simulation of this plethora of market agents leads to high computational costs.

Given the high complexity and nonlinearity of the LOB's dynamics, in the last years, machine learning has been increasingly employed in order to devise LOB simulators. The architectures proposed in the literature are mostly based on Long Short-Term Memory Networks, Recurrent Neural Networks, and Generative Adversarial Networks, e.g., [27–33]. These models produce satisfying performances in the reproduction of the stylized facts. Nonetheless, if we consider the refined and recent contribution [31], the extended study about a metaorder execution which is performed by the authors, highlights the inability of the model to capture the well-known concavity. Moreover, the models based on machine learning lack of explainability, they sensitively depend on several hyperparameters that need to be calibrated carefully and on an enormous amount of parameters.

Parsimony and explainability are instead the features of the Stochastic Differential Equations Based models as [34–38]. They model the LOB evolution as a set of differential equations and allow to focus on the steady state of the dynamics. Despite their

mathematical tractability, an explicit solution cannot be derived by most of these models. Moreover, they highly depend on the prior assumptions they use and the resulting LOB simulators require high computational costs.

The majority of the models we revised above, manage to reproduce the dynamics of *large-tick* assets. These stocks have a spread which is most of the times equal to the tick size and dense LOBs, i.e., all levels of the LOB are occupied. This is in contrast with *small-tick* assets, which are such that the mean spread is much larger than the tick size and so, their LOB is sparse. Consequently, small-tick assets dynamics are more complex and challenging to reproduce.

1.2. Contributions of the paper and outline

The aim of this article is to propose an extension of the Zero Intelligence (ZI) model [8], that is able to reproduce the concavity of the market impact during the execution of a metaorder and the subsequent price reversion after its end. As we have seen in Section 1.1, in the literature, models of the LOB which are more sophisticated than the ZI, have been proposed. Nonetheless, we focus on this Point Process model since, even if the analytical treatment of the full model is extremely difficult, several relations can still be derived by adopting some assumptions. Moreover, this model can simulate the dynamics of both small-tick and large-tick assets. Our focus is on the former class, since their modeling is little explored in the literature. Additionally, the computational costs of the ZI model are contained and, most importantly, its explainability allows to provide insights into the mechanism underlying market impact, which is a major objective of this work.

Our model places itself as the small-tick assets counterpart of the Queue-Reactive (QR) model [15], which was devised for large-tick assets. The QR model well reproduces the concavity of the market impact, it is Markovian and the probability of the orders' arrival at a given queue depends on the queue's size. If, as in our case, the focus is on small-tick assets, considering orders' arrivals which depend on the queues' sizes is unreasonable: the LOBs for this class of assets usually have queues with small sizes, often equal to the minimum size of a single order. Therefore, in order to devise a LOB model that can reproduce the concavity of the market impact, we propose a non-Markovian variant of the ZI model. The idea is that given a limit order, its sign is sampled with a probability that is a function of an indicator of the past price trend. Financially, this choice links our model to the Locally Linear Order Book model [39] since it can be interpreted by considering that in real markets, long term investors have *reservation prices* and their decisions to submit or not submit limit orders are driven by these values and the current price trend. In our model, the reaction of these traders occurs with a time scale that corresponds to the memory we employ in the evaluation of the price dynamics, and its strength is tuned by a parameter that enters the definition of the probability that a limit order is a sell/buy. Throughout the paper, we study how our results and market impact dynamics depend on these two parameters.

The paper is organized as follows. Section 2 provides a brief overview of limit order books. In Section 3, the Zero Intelligence model is revised. Section 4 presents our non-Markovian version of the model with theoretical insights and numerical results. Finally, conclusions are drawn in Section 5. In the Supplementary material, details about the data employed for parameter estimation and the approach adopted are provided; also, additional investigations about our model and some figures, which are explained in the main text, are reported.

2. Limit order books

In LOBs several types of events can occur. A *buy limit order* of Y shares at price y (multiple of a given monetary value, called the *tick size*, that is set by the regulators for each exchange) represents the intention of an agent to buy Y shares at a price that is lower than or equal to y . On the other hand, a *sell limit order* of Z shares at price z represents the intention of an agent to sell Z shares at a price that is greater than or equal to z . Buy and sell limit orders (LOs) constitute the bid and ask side of the LOB respectively and they are separated by the *bid-ask spread*, that is the distance between the highest and lowest prices of the bid and ask sides (*best bid* and *best ask*). Instead, the mean of these two prices is the so-called *mid-price*, which is usually employed as a proxy of the asset price. Agents can also cancel their limit orders and submit *market orders* (MOs), which are instantaneously executed at the best price. Therefore, a trade occurs when there are a buyer and a seller that agree on a price. We note that in the majority of LOBs, the order book matching is done on a price/time priority basis (*First In First Out*): for a given price level, the first order to be executed is the first in the queue, i.e. the first which has been posted. A graphical representation of the LOB for a given stock at a given time is provided in Fig. A1 in the Supplementary material.

3. The Zero Intelligence model and its limitations

The Zero Intelligence (ZI) model [8–10] was designed with the aim of being as analytically tractable as possible. It allows to simulate the LOB by relying on an extremely simple description: LOs, cancellations and MOs are considered as three independent Poisson processes, with rates that are independent of prices and sizes. Each order has unitary size and bid or ask sides are equally probable, as well as the price levels. Therefore, only four parameters (the LO rate λ , the MO rate μ , the cancellation rate δ and the unitary volume q_0) need to be estimated from empirical data and this can be performed straightforwardly via elementary computations. We adopt the approach illustrated in [3], that we present in Subsection C.1 of the Supplementary material.

Despite its simplicity, the ZI model reproduces fairly well the mean and the distribution of the bid-ask spread (it slightly underestimates the mean spread), and the increasing mean volume profile near to the best quotes. On the other hand, it does not capture the decrease in mean volume in limit order book's deeper levels; it overestimates volatility for small-tick assets and underestimates it for large-tick assets. Most importantly, this LOB simulator leads to a flat response function (as we see in the next subsection) and profitable market-making strategies. Contrary to real markets where the market orders' signs are strongly and positively autocorrelated, the order flow in the ZI model is uncorrelated. This makes the impact of a market order to be instantaneous and permanent.




Fig. 2. Comparison between the response function $R(\tau) = \langle (m_{t+\tau} - m_t)e_t \rangle_{t,MO}$ (in ticks) computed on empirical data and on data simulated with the ZI model.

3.1. A linear model of market impact

One important limitation of the ZI model is that it provides a linear model of market impact, both as a function of time and of the metaorder size. Indeed, the response function, or lag- τ impact, is constant. This function is defined as

$$R(\tau) = \langle (m_{t+\tau} - m_t)e_t \rangle_{t,MO},$$

where m_t is the mid-price before the market order at time t , e_t is the sign of the market order at time t , and $\langle \cdot \rangle_{t,MO}$ represents the mean over all market orders. In real markets, the response function is concave, positive, increasing and saturates for large τ while in the ZI model, as reported in [3], it is well approximated by:

$$R(\tau) \simeq \frac{1}{2} \mathbb{P}(q_{best} = 1) \langle x_1 \rangle \equiv k \quad (1)$$

where q_{best} is the volume (in units of q_0) at the best quote, x_1 is the first gap size, i.e., the difference between the best bid/ask and the second price level, and $\langle \cdot \rangle$ denotes its average. For small-tick assets the probability that q_{best} is unitary is close to 1 and so, the response function is about half the mean first gap size.

Consequently, since in the ZI model the impact of a single trade is time-independent and equal to k , if we consider a metaorder that is executed in $[0, T]$ with trading speed v_t and total volume Q , the mid-price impact is

$$m_T - m_0 = \int_0^T k v_t dt = kQ \quad (2)$$

where m_t is the mid-price at time t .

We confirm these predictions by performing several simulations of the ZI calibrated on real data. We use the LOBSTER dataset¹ for the small-tick asset Tesla on January 5, 2015, which has a tick-size equal to 0.01\$. After the pre-processing,² we end up with 170,416 events. The average mid-price and its standard deviation are: 209.62\$ and 1.44\$. The mean spread is 16.71 ticks. Limit orders constitute the 49.90% of the market events, cancellations the 50.09% and market orders the 0.01%. The parameters of the ZI are estimated³ as illustrated in Subsection C.1 of the Supplementary material obtaining: $\lambda = 0.0131$, $\mu = 0.0441$, $\delta = 0.1174$, $q_0 = 101$.⁴

After the estimation of the ZI model, we run several simulations (without any metaorder execution) to estimate the response function $R(\tau)$ and the parameter k characterizing the impact. Fig. 2 shows the response function obtained from the simulations of the ZI model together with its empirical counterpart. As expected, the two curves are very different since the ZI model predicts a constant response function, while real data shows a slow increase.

The estimated average value of the response function of the ZI model is equal to $\hat{k} = (4.917 \pm 0.002)$ ticks. The theoretical value predicted by Eq. (1) is $\hat{k}^{th} = (5.107 \pm 0.02)$ ticks which is slightly larger than $\hat{k}$ but in overall agreement with it.

We now consider the price trajectory during a metaorder execution. We simulate a calibrated ZI model and superimpose the execution of a buy metaorder with constant speed and one unit, i.e., q_0 , executed for each child MO. This last choice is motivated

¹ lobsterdata.com

² For details about the LOBSTER data and the pre-processing steps we performed, see Section B of the Supplementary material.

³ Code related to this calibration step and the models employed in the following is available at: github.com/adeleravagnani/non-markovian-zero-intelligence-lob-model.

⁴ We consider only one trading day for parameter estimation since our interest is to obtain parameters that are representative of small-tick stocks. When we estimate the model parameters for all days in January 2015, their variability is relatively small, hence the day we choose is somewhat representative of the whole month.




Fig. 3. ZI model. On the left, we plot the average path of the mid-price increment versus the number of the child MOs of the metaorder. More precisely, the plotted quantity is: $P(\tau) = \langle (m_{(\tau-1)(\Delta+1)} - m_0)\epsilon \rangle_{sim}$ with $\tau = 1, \dots, Q$. The metaorder has total volume $Q = 100$, buy direction, i.e. $\epsilon = +1$, and different values of trading interval Δ (in event time). The inset plot represents the participation rate as a function of the trading interval. On the right, we plot the estimated slope of $P(\tau)$. The black line is the coefficient of the single trade response function obtained on simulated data without metaorder execution. The blue line is the coefficient obtained on simulated data without any metaorder execution and by employing Eq. (1).




Fig. 4. ZI model. Mid-price path for an execution with total volume $Q = 100$, buy direction, and different values of trading interval Δ (in event time). Each black line corresponds to a simulation. The red line is the mean path. The vertical dashed blue lines are the beginning and the end of the metaorder execution.

by the fact that in the ZI model, each order has unitary size. In order to investigate how the results depend on the trading speed and the total volume, we carry out two experiments. (1) We set the total volume equal to $Q = 100$ (in units of q_0) and we vary the number of events between two child orders that we denote as *trading interval* Δ (in event time), which is the inverse of the trading speed v ; (2) we set the trading interval $\Delta = 50$ and vary the total volume executed Q . All the results are averaged over 200 simulations.

Experiment 1. The left panel of Fig. 3 shows the estimated path of the mid-price increment $m_{(\tau-1)(\Delta+1)} - m_0$ as a function of $\tau = 1, \dots, Q$ that is the number of child orders of the metaorder executed in $[0, T]$ with $T = (Q-1)(\Delta+1)$. Intuitively, if we define the *participation rate* as $100 * Q / N_{MO}^{ex}$ where N_{MO}^{ex} is the number of MOs executed by all market in the period of the metaorder execution, this variable increases as the trading interval decreases (see inset plot in the left panel of Fig. 3). The right panel shows the fitted slopes as a function of the trading interval. This illustrates the equality between the permanent impact parameter and the constant response function of the ZI model $\hat{k}$, which is estimated by running the ZI model without any execution of a metaorder (black line in the right plot of Fig. 3). Although the estimated coefficients are not perfectly constant, they fluctuate quite close to the theoretical value $\hat{k}$.

Fig. 4 shows the mid-price paths before, during, and after the executions. In this case the time on the abscissa is the market event time. By focusing on the average path (red line), we observe that during the execution there is a linear increase in the mid-price while before and after the execution, the mid-price remains approximately constant at the level reached by the last trade of the metaorder.

Experiment 2. Results are reported in Fig. C2 in the Supplementary material. As expected, the larger the total volume, the stronger the increase of the mid-price during the execution of the metaorder. By fitting the difference between the mid-price at the end and at the beginning of the execution, i.e., $m_T - m_0$, as a function of Q and referring to Eq. (2), we obtain a coefficient equal

to (5.063 ± 0.011) ticks. This is consistent with $\hat{k}$ estimated without any execution and the coefficients estimated with a metaorder execution and different trading intervals.

In conclusion, the numerical investigation of the ZI model shows the main limitations in the behavior of the price dynamics during and after the execution of a metaorder. As seen, during execution, the price increases linearly with time, and after the end of the metaorder, the price remains constant. These properties are at odds with what observed empirically (displayed schematically in Fig. 1). The next Section introduces a modified ZI model which provides more realistic price patterns.

4. The Non-Markovian Zero Intelligence model

In the ZI model, each event is simulated without considering any information about past events and, as we previously recalled, this makes the order flow uncorrelated and the impact constant. Inspired by [40], we propose a non-Markovian model which is analogous to the ZI, except for one aspect. Contrary to the ZI model where the LOs' signs are picked randomly with equal probability, we set

$$\mathbb{P}(\text{sell LO at time } t | \text{LO at time } t) = \frac{1}{1 + \exp(-\alpha \bar{R}_t)} \quad (3)$$

where $\bar{R}_t$ is the exponentially weighted mid-price return and α is an intensity parameter ($\alpha > 0$). The former represents the past price trend and it is defined as

$$\bar{R}_t = \sum_{s=1}^t e^{-\beta(t-s)} r_s, \quad (4)$$

where $\beta > 0$ is an inverse characteristic time, $r_s = m_s - m_{s-1}$, m_s is the mid-price before event s and the initial time is $s = 0$ with a given m_0 .⁵ It is important to stress that the specific functional forms of the probability in Eq. (3) and of the indicator of the past price trend in Eq. (4) could be generalized to alternative forms.

Our choice can be justified by considering that in real markets, long term investors have *reservation prices*: they are not willing to buy (sell) at prices that are higher (lower) than a given value. These values are usually motivated by considerations that rely on fundamental analysis. Consequently, the decisions to submit or not submit LOs by these traders are driven by their observation of the price path. For instance, let us consider a trader who wants to sell Z shares with a reservation price equal to 12.5\$, the best bid is 10\$ and the best ask is 12\$. Then, if the price shows an upward trend, the trader may decide to submit a LO at 12.5\$ in order to gain time priority and ensure she sells as soon as the price reaches her target. In the framework of our model, in order to understand whether the price has an upward or downward trend, traders have a memory that is dictated by the parameter β . On the other hand, α corresponds to the intensity of trading reaction of these investors to the price trends. This interpretation connects our model to the Latent Liquidity Theory introduced in [41] and in particular, to the Locally Linear Order Book (LLOB) model [39]. This reaction-diffusion model is grounded on the idea that most of the liquidity in the book remains latent, i.e., it is not displayed in the book, until the market price becomes close to the reservation price. In other words, this latent liquidity represents trading intentions of long term investors who have reservation prices (to buy or to sell) that are updated in time because of price changes, incoming news, etc. However, contrary to our model, the LLOB is Markovian and the adjustment of the reservation price is independent from the market dynamics, except for the mechanism that triggers the liquidity to become visible.

An alternative interpretation of our model is related with *contrarian* investment strategies [42,43]: investors buy past losers and sell past winners. This trading behavior is opposed to *momentum* trading and it is widely explored in the literature. Several works, as [44–51], provide empirical evidences that the contrarian behavior is commonly adopted, especially among individual investors. It is driven by the mean-reversion paradigm and partially rooted in the *disposition effect*, i.e., the investors' tendency to hold on to losers stocks and sell the winners.

We notice that if $\alpha = 0$, we recover the standard ZI model, i.e., independently of the exponentially weighted mid-price return, sell and buy LOs are equally probable.

Finally, the dynamics of the LOB, following the Non-Markovian Zero Intelligence (NMZI) model,⁶ can be simulated according to Algorithm 1, displayed in Subsection C.1 of the Supplementary material. We stress that the non-Markovianity of our model is due to the introduction of the time-varying probability that a LO is a sell: given a state of the LOB, each time we sample a LO, the evolution of the LOB and so, its next state depends on previous states and in particular, on past mid-prices.

In the following, the Section is organized as follows. In Section 4.1, we provide an overview of the dynamics of the mid-price that we obtain and the intuition of the mechanism which is at the root of the market impact concavity in our simulator. During the execution, the mid-price path is first concave and then, it reaches a stationary regime where it converges to a straight line. After the execution, it exponentially decays. We analyze in depth each component of this behavior. In order to do so, in Section 4.2 we start to focus on the behavior of $\bar{R}_t$, we state two different ways to compute it and derive a condition that links these two descriptions. In Section 4.3, we focus on the stationary regime of the mid-price path. We derive a relation between the stationary values of $\bar{R}_t$ and the slope of the mid-price trajectory. Then, in Section 4.4 we derive two master equations describing the evolution of the

⁵ We note that Eq. (4) could be interpreted as an operative definition in our model. On the other hand, we can state a more formal definition that ensures stationarity: $\bar{R}_t = \sum_{s=-\infty}^t e^{-\beta(t-s)} r_s$.

⁶ The presence of memory and adaptation to past history would suggest calling our model a *minimal intelligence model*, but we choose NMZI since we aim to emphasize its derivation from the ZI model.




Fig. 5. NMZI model. Mid-price path for an execution with total volume $Q = 2,000$ (left) and $Q = 100$ (right), buy direction, and $\Delta = 20$. Each black line corresponds to a simulation. The red line is the mean path. The vertical dashed blue lines are the beginning and the end of the metaorder execution.




Fig. 6. NMZI model. Evolution of $\bar{R}_t$ (t denote the market events since the beginning of the metaorder execution) when a metaorder with total volume $Q = 2,000$, buy direction, and $\Delta = 20$ is executed. In the left panel, all the evolution (during the execution of the metaorder and after ending) is plotted while on the right, a zoom on the evolution until the 9th child MO is provided.

mid-price and the spread in the time windows between two child MOs. In Section 4.5, we analyze the interplay between the impact components that leads to the concavity in our model. Finally, in Section 4.6, we focus on the mid-price path after the execution of the metaorder ends and we show that the decay of the mid-price is exponential.

4.1. A glimpse of market impact dynamics in our model

In order to grasp the outcome of our model, let us start to consider the parameters $\beta = 10^{-3}/(\Delta + 1)$, $\alpha = 10^{-3}$ and $\Delta = 20$. We simulate the LOB dynamics⁷ with the execution of a buy metaorder of $Q = 2,000$ (in units of q_0) and in the left panel of Fig. 5, the mid-price paths before, during and after the execution of the metaorder are plotted. Before the execution starts, the mid-price is approximately constant. During the execution, two regimes can be distinguished: first, the price path is concave (denoted in the following *concave regime* and it is more evident in the right panel of Fig. 5 where we represent the mid-price evolution when a metaorder of size $Q = 100$ is executed) and then, it increases linearly in the so-called *stationary regime*. Finally, after the execution ends, the reversion to a lower price value occurs. This behavior is mirrored by the exponentially weighted mid-price return whose dynamics during and after the execution of a metaorder with volume $Q = 2,000$ is represented in the left panel of Fig. 6. By construction, $\bar{R}_t$ is a measure of the past price trend.

We start to provide the intuition that explains this behavior in our model. In the next subsections, we will deepen it and analyze in detail each component of the mid-price trajectory. Let us consider a metaorder which is composed of Q child MOs, each of them

⁷ The parameters $\lambda, \mu, \delta, q_0$ of the NMZI model are the same of the ZI model we estimate in Section 3.1. Similarly to what we do for the ZI model, in all simulations of the NMZI model we perform, we execute a buy metaorder with constant speed and 1 unit executed for each child MO. All results are averaged over 200 simulations. We note that without any execution of metaorders, the response function we obtain is flat and equal to (4.867 ± 0.05) ticks while, the corresponding value according to Eq. (1) turns out to be equal to (5.114 ± 0.002) ticks.

of unitary size, submitted every Δ steps and with the same sign, e.g. buy. When the first child order of the metaorder is executed, its *immediate price impact* is about equal to half the first gap size in the ask side. This is equivalent to Eq. (1) with $\mathbb{P}(q_{best} = 1) \simeq 1$ that holds for small-tick assets (this is the class of assets under our interest). For the standard ZI model, in the time interval until the next child MO, the mid-price remains stable. This is not the case for our model which leads to a price reversion. In Fig. D3 in the Supplementary material, we provide a stylized representation of the mid-price path between two child MOs in the ZI and in the NMZI models. In our model, the immediate impact component is such that $\bar{R}_t$ increases and so does the probability of sell LOs over buy LOs. This triggers more sell LOs and consequently, more sell LOs in the spread, which make the mid-price decrease on average: this is the *reversion component of the impact* in the interval between two child MOs. As we see in the following and it is represented in the right panel of Fig. 6, $\bar{R}_t$ increases after each child MO and it is approximately constant within each interval between two child MOs. In this way, the entity of the reversion increases as more child MOs are executed and the imbalance between sell and buy LOs increases, making the first gap size in the ask side likely to decrease. This would cause a reduction of the immediate price impact component. However, in the following we see that the variation of the immediate impact component is marginal. To sum up, in the concave regime, we have two components in the price impact of every child MO: an immediate contribution which mildly varies in time and a reversion which increases in time. The difference between these two components is the so-called *overall impact of the single child MO* and, as we see in the following, it decreases as new child MOs are executed, making the concavity in the mid-price trajectory appear.

At a certain point a transition to the stationary regime occurs since $\bar{R}_t$ becomes constant and so does the probability that a LO is a sell. Finally, the decay after the execution ends is a consequence of the reversion mechanism we have in the mid-price trajectory after a child MO.

4.2. The behavior of $\bar{R}_t$ and two equivalent descriptions for its computation

In the preceding subsection we qualitatively explained the mechanisms of our model. Now, we go deeper and, as a first step, we focus on the behavior of $\bar{R}_t$ in the concave regime: we aim to justify the stepwise characteristic shown in the right panel of Fig. 6. In order to do so, we introduce a different way to compute $\bar{R}_t$, namely by using only child MOs' times and in this setting, by construction, $\bar{R}_t$ increases after each child MO of a buy metaorder and it is constant within each interval between two child MOs. Therefore, by showing the conditions which lead to the equality between the two different ways to compute $\bar{R}_t$, it follows that even if we compute $\bar{R}_t$ by using all market events' times, it is approximately a step function as represented in the right panel of Fig. 6. This highlights that $\bar{R}_t$ reflects the price dynamics due to the metaorder execution and this is included in the sampling of the signs of the LOs.

Let us denote as (i) the description for which $\bar{R}_t$ is computed by using all market events' times and as (ii) the description for which $\bar{R}_t$ is computed by using only child MOs' times. They are approximately equivalent for a suited rescaling and a condition on the exponent β , which enters the definition of $\bar{R}_t$ in Eq. (4). If t denotes the market events' times and the child MOs are executed at $t_j = \Delta j + j - 1$ with $j = 1, \dots, Q$, the starting time is 0 and $t_Q = T$, we have the following relations:

$$\begin{aligned} \text{description (i): } \bar{R}_t^{(1)} &= \sum_{s=1}^t e^{-\beta_1(t-s)}(m_s - m_{s-1}) \\ \text{description (ii): } \bar{R}_t^{(2)} &= \sum_{s=1}^t \mathbb{I} \left[\frac{s - \Delta - 1}{\Delta + 1} \in \mathbb{N}_0 \right] e^{-\beta_2(\frac{t^+-s}{\Delta+1})}(m_s - m_{s-\Delta-1}) \\ \text{where } t^+ &= \max_{s \leq t} \left\{ s : \frac{s - \Delta - 1}{\Delta + 1} \in \mathbb{N}_0 \right\}. \end{aligned} \quad (5)$$

The outputs of these relations are approximately equivalent if $\beta_2 = (\Delta + 1)\beta_1$, $\beta_1 \ll 1/\Delta$ and the mid-price differences within the trading interval Δ are negligible with respect to the ones due to the child MOs. We observe that β_2 represents the characteristic time of the description (ii) therefore, requiring $\beta_2 \ll (\Delta + 1)/\Delta$ amounts at including in the computation of $\bar{R}_t$ the price dynamics which cover a large number of past child MOs, exploiting the memory of the model. We provide further details and evidence of this finding in Subsection D.2 of the Supplementary material.

We notice that the economically relevant description is (i): the price trend indicator $\bar{R}_t$ should be computed by relying on all market events' times since *a priori* the child MOs' times are not known.

4.3. The stationary regime

We showed the behavior of $\bar{R}_t$ in the concave regime however, at some point this quantity stabilizes to a stationary value, as represented in the left panel of Fig. 6. In the framework of description (ii), Eq. (5) can be rewritten in the following recursive way:

$$\bar{R}_{t_j+1}^{(2)} = r_{t_j+1} + e^{-\beta_2} \bar{R}_{t_j-\Delta}^{(2)} \quad (6)$$

where $r_{t_j+1} = m_{t_j+1} - m_{t_j-\Delta}$ is the mid-price change due to child MO j . Given what we showed in the preceding subsection, the price trend indicator is non-decreasing during the execution of a buy metaorder and we can identify the *stationary regime* as the set of




Fig. 7. NMZI model. Comparison between the stationary values of $\bar{R}_t$ (left) and r_t (right) obtained with the two different descriptions. The superscript (1) refers to the description for which the exponentially weighted mid-price return $\bar{R}_t$ is computed by using all market events' times; the superscript (2) refers to the description for which $\bar{R}_t$ is computed by using only child MOs' times and $\beta_2 = \beta_1(\Delta + 1)$. The dark line is the bisector of the first quadrant.

times $\{t \in [t^*, T] : \bar{R}_t \simeq \bar{R}_{t-1} \simeq \bar{R}^* > 0\}$. From Eq. (6), it follows that in this regime also $r_t \rightarrow r^*$ and the exponentially weighted mid-price return and the mid-price return are approximately proportional:

$$\bar{R}^* \simeq \frac{r^*}{1 - e^{-\beta_2}}. \quad (7)$$

Let us note that the mid-price return can also be seen as the slope of the mid-price trajectory. Therefore, after a regime in which the mid-price trajectory is concave, it stabilizes to a straight line with slope r^* .

We point out that Eq. (6) expresses an implicit link between the price trend indicator and the intensity parameter α . $\bar{R}_t^{(2)}$ is a step function: after each child MO, it is dampened by $e^{-\beta_2}$ and increases of $m_{t_{j+1}} - m_{t_j - \Delta}$. This quantity is the mid-price difference between two child MOs and it is linked to α through the time-varying probability that a LO is a sell. Indeed, as we see in Section 4.5, α affects the magnitude of the impact between two child MOs.

4.3.1. Numerical results

We verify numerically the findings we illustrated so far. We compare descriptions (i) and (ii). Several combinations of the parameters β_2, α, Δ are considered: $\beta_2 = [10^{-3}, 10^{-2}, 10^{-1}]$, $\alpha = [10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}]$, $\Delta = [20, 50, 100]$. Let us recall that β_2 is the parameter which enters the definition of the exponentially weighted mid-price return $\bar{R}_t$ under description (ii). Therefore, in order to compare the two descriptions, given β_2 , the corresponding parameter of the description (i), i.e. β_1 , is set equal to $\beta_2/(\Delta + 1)$.

We determine the values $\bar{R}^*$, r^* and t^* in the following way. We consider $\bar{R}_t$ during the metaorder's execution and fit it with an exponential function. More precisely, in order to obtain more trustworthy estimations, we fit the set of points $\{\bar{R}_{t_j}\}$ where t_j , $j = 1, \dots, Q$ are the child MOs' times. In this way, we only consider the values of $\bar{R}_t$ at the times of each child MO, we eliminate the step behavior and make the fit easier. The functional form of the fit is of the type: $f(j) = a - be^{-cj}$ and the stationary value $\bar{R}^*$ corresponds to the parameter a . Similarly, the beginning of the stationary regime in terms of the number of the child MOs is identified as $\tau^* = -\log(\sigma_a/b)/c$ where σ_a is the error of the estimated parameter a , i.e., τ^* is the time such that $be^{-c\tau^*} = \sigma_a$. Given that $t_j = (\Delta + 1)j - 1$, $j = 1, \dots, Q$, the value t^* is equal to $(\Delta + 1)\tau^* - 1$. Finally, a linear fit is run on the mid-price path which corresponds to child MOs that occur at times after τ^* . The slope of this linear fit is the stationary value of the mid-price difference r^* .

In Fig. 7, we compare the stationary values $\bar{R}^*$ and r^* , obtained with the two descriptions. They lie around the bisector, highlighting the equivalence between the two descriptions given $\beta_1 = \beta_2/(\Delta + 1)$.

This analogy does not only concern the stationary values but it also holds for the whole evolution, as shown in Fig. D4 in the Supplementary material which displays the evolution in time of the exponentially weighted mid-price return $\bar{R}_t$ for $\alpha = 10^{-3}$, two different choices of trading interval Δ , i.e., $\Delta = 20, 100$, and of the parameter β_2 , i.e., $\beta_2 = 10^{-2}, 10^{-3}$. We observe that, as expected, $\bar{R}_t$ stabilizes to an equilibrium value after a concave regime. Similarly, Fig. D5 in the Supplementary material represents the mid-price paths as a function of the number of the child MOs. They exhibit a concave increase followed by a linear behavior, which mirror the increasing and constant behavior of $\bar{R}_t$ respectively.

Now, we aim to investigate the role that the parameters β_2 and α play in the findings related to the stationary regime. We consider the results obtained by relying on the economically relevant description (i), i.e., $\bar{R}_t$ is computed by using all market events. In Figures D6–D7–D8 in the Supplementary material, we show heat maps representing the stationary values $\bar{R}^*$, r^* and τ^* for several parameters' combinations ($\beta_2 = [10^{-3}, 10^{-2}, 10^{-1}]$, $\alpha = [10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}]$, $\Delta = [20, 50, 100]$).

Role of β_2 . For fixed values of Δ and α , as $1/\beta_2$ increases so do $\bar{R}^*$ and τ^* . This means that considering a longer memory in the evaluation of the price trend leads to a longer concave regime and so, to a better reproduction of the behavior of the market impact we observe in real markets. This strengthens the intuition we build our model on: including the price dynamics in the orders' sampling process of a LOB model is crucial to reproduce the market impact concavity.




Fig. 8. NMZI model. Check of the proportionality relation between the stationary values of r_t and $\bar{R}_t$. The dark line is the bisector of the first quadrant.

Concerning r^* , if $1/\beta_2$ increases, the stationary value of the mid-price differences decreases. Since $\bar{R}^*$ is larger for larger $1/\beta_2$, also the value of the probability that a LO is a sell in the stationary regime is larger and equivalently, the strength of posting sell LOs in response to an upward price trend. As we see in Section 4.5, this implies that the reversion in the interval between two child MOs is stronger and the overall impact of each child MO in the stationary regime, that is r^* , is smaller.

Role of α . For fixed values of Δ and β_2 , as α decreases, $\bar{R}^*$ and τ^* increase. This can be explained by considering that the greater the reaction intensity of traders to price trends, the faster is the convergence of $\mathbb{P}(\text{sell LO at } t | \text{LO at } t)$ to the constant value 1. A measure of the transition speed of this probability, is given by its first derivative with respect to $\bar{R}_t$, that is equal to $\alpha/4$ for $\bar{R}_t = 0$. The limited concave regime we observe for larger values of α , is coherent with our idea that time varying sampling probabilities or similarly, decisions to submit LOs that vary in time and are driven by the price trend, lead to the emergence of the concavity in the mid-price path.

On the other hand, as α increases, r^* decreases. As we highlight before, larger values of α are accompanied with faster transition speeds to the stationary regime and to the value 1 for the probability that a LO is a sell. Therefore, as we see in Section 4.5, if α is larger, the reversion in the interval between two child MOs is stronger and the overall impact of each child MOs in the stationary regime is smaller.

Finally, we verify the proportionality relation between the exponentially weighted mid-price return and the response function, that we derived in Eq. (7). We consider the results obtained by relying on description (i). The stationary values r^* obtained for this description are compared to $\bar{R}^*(1 - e^{-\beta_2})$ and the results are reported in Fig. 8 for the same parameters' combinations considered in Fig. 7. A good agreement between $\bar{R}^*(1 - e^{-\beta_2})$ and r^* is obtained.

From now on, the description under interest is (i), being the economically relevant one, as previously pointed out. Therefore, we intend β_1 for β and $\beta_1(\Delta + 1)$ for β_2 .

4.4. The mid-price path between two child MOs

To understand the origin of the concavity of the price in the NMZI model, we present a model of the mid-price evolution in the intervals between two child MOs. Previously, we showed that $\bar{R}_t$ can be approximated by a constant function in the intervals between two child MOs. Therefore, given the recursive relation which links $\bar{R}_t$ and r_t , if $\bar{R}_t$ is constant so does r_t . Since r_t expresses the mid-price difference, we expect that a linear function is likely to describe the price reversion in the intervals between two child MOs. In the following, we focus on the stationary regime where the probability that a LO is a sell is constant and we set it to $p_1 = 1/(1 + e^{-\alpha\bar{R}^*})$. By means of master equations, we study the evolution of the mid-price in the time windows between two child MOs.

We need to consider that, given a state of the order book at time t with mid-price $m = m(t)$,⁸ it can switch from the value m to one of the values $m' = m + \frac{1}{2}, m + 1, \dots, m + \frac{s-1}{2}$, where $s = s(t)$ is the spread value at time t , because of a buy LO inside the spread at time $t + 1$. On the other hand, if at time $t + 1$ we have a sell LO inside the spread, the mid-price decreases to $m' = m - \frac{1}{2}, m - 1, \dots, m - \frac{s-1}{2}$. Sell (buy) MOs and buy (sell) cancellations make the mid-price decrease (increase) by half the first bid (ask) gap size. In the setting of Algorithm 1 (displayed in Subsection C.1 of the Supplementary material), the first gap size can assume values between 1 and $K - 2$ where K is the number of tick levels in the grid which represents the LOB. Therefore, since when LOs inside the spread occur, the value that the mid-price can assume depends on the spread, first we need to derive the master equation which describes the evolution of this quantity. Then, we will be able to derive the master equation related to the mid-price.

⁸ In this subsection we adopt this notation for the mid-price at time t . However, $m(t)$ is analogous to m_t .

4.4.1. Spread evolution

We derive the master equation which describes the evolution of the spread in the intervals between two child MOs, and we show that it does not depend on the probability that a LO is a sell; therefore, it is also valid for the ZI model. This is due to the fact that the imbalance between sell and buy LOs does not impact the evolution of the spread indeed, the contribution of these two types of orders cancel out in the master equation. This is in contrast to the mid-price: in this case, the imbalance between sell and buy LOs is crucial in determining the reversion in the intervals between two child MOs.

In order to write the master equation describing the evolution of the spread between two child MOs, we need to consider that the spread can change in two ways. Given the spread $s(t)$ at time t , it can decrease to the values $1, 2, \dots, s(t) - 1$ if at time $t + 1$ there is the placement of a LO inside the spread. On the other hand, the spread can increase to the values $s(t) + 1, s(t) + 2, \dots, s_{MAX}$ if at time $t + 1$ a MO or a cancellation occurs and the first gap size is $x_1(t) = s(t + 1) - s(t)$. In the setting of Algorithm 1, the maximum value that the spread can assume is $s_{MAX} = K - 1$. Starting from these observations, we derive the transition rates entering the master equation for the spread evolution which is:

$$\begin{aligned} \mathbb{P}(s, t + 1) - \mathbb{P}(s, t) &= \sum_{s' \neq s} \left[\omega_{s' \rightarrow s} \mathbb{P}(s', t) - \omega_{s \rightarrow s'} \mathbb{P}(s, t) \right] = \\ &= \frac{\mu + \delta}{\Gamma(t)} \mathbb{P}(s - x_1^a(t), t) + \frac{\mu + \delta}{\Gamma(t)} \mathbb{P}(s - x_1^b(t), t) + \\ &+ \frac{2\lambda}{\Gamma(t)} \sum_{s'=s+1}^{K-1} \mathbb{P}(s', t) - \mathbb{P}(s, t) \left[\frac{2\lambda(s-1) + 2\mu + 2\delta}{\Gamma(t)} \right] \end{aligned} \quad (8)$$

where $\mathbb{P}(s, t)$ is the probability of having spread s at time t , $\Gamma(t) = \lambda K + 2\mu + \delta n_{orders}(t)$, $n_{orders}(t)$ is the total number of orders in the book at time t , $x_1^a(t)$ and $x_1^b(t)$ are the first gaps in the ask and bid side at time t . Therefore, we have that

$$\langle s(t) \rangle = \sum_{s'=1}^{K-1} s' \mathbb{P}(s', t).$$

For further details concerning the derivation, see Subsection D.3 of the Supplementary material. We observe that the master Eq. (8) depends on three quantities which are functions of time: $\Gamma(t)$, $x_1^a(t)$ and $x_1^b(t)$. However, as simulations will show, the number of orders in the book $n_{orders}(t)$ as well as the first gap sizes $x_1^{b/a}(t)$ have a modest range of variability in the stationary regime, and if they are considered as constant, it negligibly impacts our results. Moreover, it is evident that the spread evolution does not depend on the probability of sell limit orders making this equation also valid for the standard ZI model.

For $t \rightarrow +\infty$, we expect that the spread reverts to the equilibrium value, which can be roughly determined as $\hat{s}$ such that:

$$\frac{\lambda K}{\Gamma(t)} \frac{(\hat{s} - 1)}{K/2} = \frac{2(\mu + \delta)}{\Gamma(t)} \iff \hat{s} = \frac{\mu + \delta}{\lambda} + 1. \quad (9)$$

This corresponds to the value for which the contributions of LOs inside the spread and cancellations and MOs which are aggressive, i.e., that lead to mid-price changes, are balanced out. It is the mean spread we expect to obtain by simulating both the ZI and the NMZI models, without any execution of metaorders.

4.4.2. Mid-price evolution

Similarly to what we did for the spread and by considering the observations made at the beginning of the subsection, we obtain the master equation describing the evolution of the mid-price in the intervals between two child MOs:

$$\begin{aligned} \mathbb{P}(m, t + 1) - \mathbb{P}(m, t) &= \frac{2\lambda}{\Gamma} \sum_{j=1}^{s(t)-1} \left[p_1 \mathbb{P}\left(m + \frac{j}{2}, t\right) + (1 - p_1) \mathbb{P}\left(m - \frac{j}{2}, t\right) \right] + \\ &+ \frac{\mu + \delta}{\Gamma} \left[\mathbb{P}\left(m + \frac{x_1^b(t)}{2}, t\right) + \mathbb{P}\left(m - \frac{x_1^a(t)}{2}, t\right) \right] + \\ &- \mathbb{P}(m, t) \left[\frac{2\lambda(s(t) - 1) + 2\mu + 2\delta}{\Gamma(t)} \right] \end{aligned} \quad (10)$$

and

$$\langle m(t) \rangle = \sum_{m'=1}^{m_{MAX}} m' \mathbb{P}(m', t)$$

where $\mathbb{P}(m, t)$ is the probability of having a mid-price equal to m at time t and m_{MAX} is the maximum mid-price we can have in the grid which represents the LOB, i.e., $((K - 2) + (K - 1))/2$.

Contrary to the master equation for the spread, the evolution of the mid-price depends on the probability p_1 that a LO is a sell. Therefore, Eq. (10) holds for the standard ZI model only if $p_1 = 0.5$.




Fig. 9. NMZI model, stationary regime. Spread (in ticks) as a function of the time step after child MOs. Mid-price difference $m_{\tau+h+1} - m_{\tau+1}$ (in ticks) as a function of h where h is the time step after child MOs at time τ .

4.4.3. Numerical results

We verify whether the outcomes predicted by the master equations we derived — are consistent with simulations' results. In order to avoid edge effects due to the vicinity to the concave regime, we focus on the *deep stationary regime* that we define as the set of times greater than $t^* + 1/\beta$ and in such a way that 500 child MOs are considered. We compute: the spread, the probability of a sell LO (given a LO) and the number of orders in the book as functions of the time step after child MOs; the mid-price difference $m_{\tau+h+1} - m_{\tau+1}$ as a function of h where h is the time step after child MOs at times τ ; the frequency of the different types of events in each time window between child MOs. They are shown in Figs. 9 and D9–D10 in the Supplementary material for $\beta = 10^{-3}/(\Delta + 1)$, $\alpha = 10^{-3}$ and two choices of Δ . In particular, in Fig. 9, we compare the evolution of the spread and the mid-price differences that we obtain from simulations, to the corresponding evolution resulting from the master Eqs. (8) and (10). When we determine the evolution of the spread, we set the initial condition $\mathbb{P}(s, t = 0) = \delta(s = s_0)$ where s_0 is the integer part of the mean spread after the child MOs in the stationary regime according to simulations. Concerning the mid-price, Eq. (10) depends on the spread and we employ the values obtained with the corresponding master equation. We observe that the probability of a sell LO (given a LO), the number of orders and the first gap in the bid/ask side (displayed in Figures D9–D10) have a modest range of variability and we verified that considering them constant in the master equations negligibly impacts our results.

Focusing on the spread as a function of the time step after child MOs, we see that it reaches a maximum at step 1 and then, it decreases. The maximum represents the spread after a child MO occurs and when $\Delta = 20$, it turns out to be 28.42 ticks while when $\Delta = 100$, it is equal to 21.82 ticks. As expected, they are both larger than the mean spread in the stationary regime which is 21.29 and 14.08 ticks respectively. We observe that the master equation well describes the evolution of the spread between two child MOs. Moreover, if we focus on the larger trading interval $\Delta = 100$, the spread reverts to the equilibrium value according to Eq. (9), i.e., $\hat{s} \approx 13.33$ (dashed black line in Fig. 9). This is also confirmed by computing the spread evolution with the master equation for times much larger than Δ .

Similarly to the spread, the master equation of the mid-price is able to reproduce its behavior. It also emerges that employing a linear function in order to describe the mid-price evolution between two child MOs is a well-grounded choice. This is actually the behavior we expect, as introduced in the beginning of Section 4.4.

4.5. The impact components

In the intervals between two child MOs, we observe a reversion of the mid-price. This is in contrast to the ZI, where there is not any reversion: the only component of the impact is the immediate price impact following each child MO and equal to half the first




Fig. 10. NMZI model. Evolution of the exponentially weighted moving averages (EWMA) of the immediate component I^I of the price impact. The inset plots focus on the first 500 child MOs.




Fig. 11. NMZI model. Evolution of the exponentially weighted moving averages (EWMA) of the reversion component I^R of the price impact.

gap size (for small-tick assets). In the NMZI model, the interplay between the immediate impact and the reversion is at the root of the concavity in the mid-price path and for this reason we now analyze in depth this issue. We see that in the concave regime the main contribution which makes the overall impact decrease, is due to the reversion component, that increases in time. In the stationary regime, the overall child MOs' impact stabilizes to a stationary value.

Let us define the times at which the child MOs are executed as $t_j = \Delta j + j - 1$ with $j = 1, \dots, Q$ and the starting time of the simulation equal to 0. We can describe the evolution of the mid-price during the execution as

$$m_t = G_{j_t}(t - t_{j_t} - 1), \quad j_t = \left\lfloor \frac{t}{\Delta + 1} \right\rfloor, \quad t \in (t_1, t_Q] \quad (11)$$

where each $G_j(x)$ with $x \in [0, \Delta]$ and $j = 1, \dots, Q - 1$ is a non-increasing function describing the mid-price path in the time interval between the child MOs j and $j + 1$. Therefore, we have that:

- the immediate price impact of the child MO j is: $I_j^I = G_j(0) - G_{j-1}(\Delta)$;
- the reversion of the mid-price following the child MO j is: $I_j^R = G_j(0) - G_j(\Delta)$;
- the overall price impact of the child MO j is: $I_j^I - I_j^R$.

Concerning the functional form of the functions $G_j(\cdot)$, we employ linear kernels given the findings of the preceding subsection. However, by studying the evolution after the execution ends with the same approach, in Section 4.6 we see that the linear kernel is actually an approximation of the exponential case for trading intervals which are not too large.

By referring to Eq. (11), we set:

$$G_j(t - t_j - 1) = \eta_j (1 - \rho_j(t - t_j - 1)),$$




Fig. 12. NMZI model. Evolution of the exponentially weighted moving averages (EWMA) of the overall impact $I^I - I^R$.

for

$$t \in [t_j + 1, t_{j+1}], \quad j = 1, \dots, Q - 1.$$

The parameter η_j represents the value of the mid-price after the child MO j while ρ_j is the slope of the straight line which describes the reversion in the interval between the child MOs j and $j + 1$. By performing linear fits on the mid-price paths, we estimate the parameters of these kernels and the corresponding impact components,⁹ which are:

$$I_j^I = \eta_j - \eta_{j-1} + \eta_{j-1} \rho_{j-1} \Delta, \quad I_j^R = \eta_j \rho_j \Delta, \quad I_j^I - I_j^R, \quad j = 2, \dots, Q - 1.$$

In Figs. 10–11–12, we plot their exponentially weighted moving averages (EWMA) for $\Delta = 20$, and several choices of β_2 and α (the results related to additional cases are plotted in Fig. D11–D12–D13 in the Supplementary material). We consider the EWMA since the impacts' evolutions are noisy. As expected, two regimes are identified. In the stationary regime, the impact stabilizes to a constant value.¹⁰ When $\bar{R}_t$ is constant, the probability of sell LOs is constant and so is the reversion of the mid-price in the intervals between two child MOs. In parallel, in the stationary regime, the first gap size is approximately constant and so is the immediate impact component due to the execution of child MOs. Indeed, in Fig. D9 in the Supplementary material we plot the evolution of the first gap size in the intervals between two child MOs in the stationary regime and we observe that the standard errors are negligible with respect to the values of the first gap size (averaged over 500 child MOs). It follows that in the stationary regime the overall child MO impact is approximately constant. On the other hand, in the concave regime of the mid-price trajectory, we observe several facts which draw attention to the mechanism at the root of the concavity's emergence in our model.

4.5.1. Immediate impact component in the concave regime

First, concerning the immediate component (see Fig. 10), it initially increases for the first child MOs. Then, we can have two different behaviors. In the first case, e.g., $\beta_2 = 10^{-3}$ and $\alpha = 10^{-3}$ (green curve in the left panel of Fig. 10), following the initial increase, the immediate impact mildly decreases in time and then, stabilizes to its stationary value. The other possible behavior is that after the initial increase, the immediate impact stabilizes to its stationary value, e.g., $\beta_2 = 10^{-1}$ and $\alpha = 10^{-3}$ (blue curve in the left panel of Fig. 10). Both behaviors are explained by the evolution of the imbalance between the orders in the bid and ask side in the LOB, which is shown in Fig. D14 in the Supplementary material for $\Delta = 20$, $\alpha = 10^{-3}$ and $\beta_2 = 10^{-3}, 10^{-2}, 10^{-1}$. In all three cases, for the first approximately 10 market events after the beginning of the execution, the imbalance is negative, i.e., the number of orders in the bid side is greater than the number of orders in the ask side. For $\beta_2 = 10^{-3}, 10^{-2}$, this is highlighted in the inset plots. Therefore, at the beginning of the execution, the immediate impact increases since the executions of the child MOs deplete the ask side and the time-varying probability that a LO is a sell is still close to $1/2$.

If $\beta_2 = 10^{-1}$, after this increase the imbalance stabilizes to a negative stationary value. In this case, as we see in Section 4.3, the stationary value of $\bar{R}^*$ is small compared to the other cases investigated (e.g. $\beta_2 = 10^{-3}$) and it is such that the stationary value of the probability that a LO is a sell is close to $1/2$. This implies that, as we observe from the blue curve in the left panel of Fig. 11,

⁹ Instead of performing linear fits in the intervals between child MOs, we also estimated the impact components directly from the mid-price path: the immediate impact of child MO j as $m_{t_{j+1}} - m_{t_j}$ and the reversion impact of child MO j as $m_{t_{j+1}} - m_{t_{j+1}}$. The results we obtain are analogous to the outputs of the linear fits, thus further corroborating the choice of linear kernels to describe the mid-price path between two child MOs. Because of that, plots are not displayed in the paper but they are available upon request.

¹⁰ In some cases, e.g., $\beta_2 = 10^{-3}$ and $\alpha = 10^{-4}$ (red curve in the right panel of Fig. 12), the stationary regime is not reached in the plots since the transition to the stationary regime is slow and an higher number of child MOs would be needed. However, in order to better compare the results for different choices of parameters, in the plots we consider a metaorder of volume $Q = 2,000$.

the reversion is small compared to the other cases and the depletion of the LOB due to the metaorder execution is not compensated by the submission of sell LOs which replenish the ask side.

On the other hand, if $\beta_2 = 10^{-3}$ or 10^{-2} , before the stabilization of the immediate impact to a constant value, there is a decrease which corresponds to an increase of the imbalance. If $\beta_2 = 10^{-3}$, the transition to the stationary regime is slower compared to $\beta_2 = 10^{-2}$ and the stationary value of the imbalance is higher. This is mirrored by the path of the immediate impact, which has a more accentuated decrease before the stationary regime. Indeed, a longer memory (smaller β_2) corresponds to a longer concave regime. Anyway, for both $\beta_2 = 10^{-3}$ and $\beta_2 = 10^{-2}$, the decrease of the immediate impact before the stationary regime is mild.

4.5.2. Reversion impact component in the concave regime

The reversion (see Fig. 11) increases and this is attributed to the evolution of our indicator of the past price dynamics $\bar{R}_t$ and its role in the time varying probability that a LO is a sell.

4.5.3. Overall impact in the concave regime

The overall impact (see Fig. 12) decreases, making the price path concave, and the major contribution in the decrease is ascribable to the behavior of the reversion component. Therefore, introducing a model that mimics the behavior of traders who have reservation prices and post LOs in response to the price trend, proves to be pivotal in reproducing the well-known concavity of the price impact during the execution of metaorders.

4.5.4. Effect of β_2 and α on the impact components

So far, we focused on the behavior of the impact components which determine the concavity in our model. However, understanding the role of the parameters β_2 and α in our results, is crucial since it allows to shed light on the financial interpretation of our model. We observe that lower values of β_2 i.e. $\bar{R}_t$ with longer memory, lead to a more pronounced reversion (see left panel of Fig. 11) and indeed, as previously shown in Section 4.3, they correspond to longer concave regimes in the mid-price paths. In the framework of the interpretation of our model in terms of traders with reservation prices, this means that, if traders use longer values of the memory in the evaluation of the price trend, then the reversion mechanism and their response to the upward price trend is stronger.

Concerning α , lower values of this parameter make the transition to the stationary regime slower (see right panel of Fig. 11). As we explain in Section 4.3, the derivative of the probability that a LO is a sell with respect to $\bar{R}_t$ can be seen as a measure of its transition speed to 1. We also observe that low values of α as $\alpha = 10^{-4}$ lead to lower variability of the overall impact (see right panel of Fig. 12 and Fig. D13 in the Supplementary material). This is due to the fact that if α is small such that even though $\bar{R}_t$ varies, $\mathbb{P}(\text{sell LO}|\text{LO})$ does not deviate consistently from 1/2, then this probability, that is at the core of our model, is not enough informative about the price dynamics we aim to consider in the sampling. The reaction of traders to the price trend is extremely mild and so, our model does not consistently differ from the ZI.

Additionally, for $\alpha = 10^{-1}$ and $\beta_2 = 10^{-3}$, the overall impact is about 0 in the stationary regime (see the blue curve in the right panel of Fig. 12), leading to a flat trajectory of the mid-price. This is due to the fact that, as we say in Section 4.3, with high values of α , the probability that a LO is a sell tends fast to 1. Therefore, the reaction of traders to the upward price trend totally compensates the immediate impact due to the child MOs. We also observe that if $\alpha = 10^{-1}$ and $\beta_2 = 10^{-2}$ or $\beta_2 = 10^{-1}$, the overall impact is not null in the stationary regime (see Fig. D13 in the Supplementary material). In these cases, the values of $\bar{R}^*$ are even smaller than what we have for $\beta_2 = 10^{-3}$ and the effect of the high value of the reaction intensity in the time-varying probability of the LOs' sign is weakened.

To sum up these observations related to α , if traders react strongly, the upward trend of the price becomes rapidly compensated by the submission of sell LOs, making the mid-price become approximately constant. If instead, traders mildly react to arising price trends, the reversion entity is moderate and slowly converges to a constant value.

4.6. Price decay after the execution

Empirically, it is well known that a buy metaorder pushes the price up during its execution, with a concave path. When the execution ends, so does the buying pressure and this leads to a price reversion (see [4–7]). This behavior is represented in Fig. 1.

As we have seen in Section 3, the ZI model is a linear model of price impact. During the execution of a buy metaorder, the price increases linearly, with a slope equal to the response function obtained for the model without any execution. During the intervals between two child MOs, the price remains approximately constant and so does after the execution ends. The NMZI model we propose is able to overcome these limitations and as shown in previous subsections, during the execution of a metaorder, the concavity of the price path emerges. In the following, we focus on the behavior of the mid-price after the end of the execution and we show that our model is able to mimic the well-known price reversion.

Similarly to Eq. (11), we can describe the price evolution after the execution of the metaorder ends as

$$m_t = G_{after}(t - t_Q - 1), \quad t > t_Q \quad (12)$$

where $G_{after}(x)$ with $x \in [0, +\infty)$ is the corresponding function. In order to infer the functional form of this kernel, we argue the following. The functional form which describes the reversion of $\bar{R}_t$ after the execution ends, is linked to the functional form that




Fig. 13. NMZI model. Evolution in time of $\bar{R}_t$ and the mid-price after the execution of the metaorder ends. $a = (675.95 \pm 0.27)$ ticks, $b = 0.000199 \pm 0.00001$, $\bar{a} = (-0.1141 \pm 0.001)$ ticks, $\tilde{b} = 0.000219 \pm 0.00001$, $c = (22512.12 \pm 0.04)$ ticks, $a(1 - \gamma e^b) = -0.1004$ ticks, $m_{t_Q+1} + \bar{a}/b_2 = 22523.30$ ticks.

describes the mid-price path after the execution. Given the simulations' results, our hypothesis is that $\bar{R}_t$ decays exponentially i.e. $\bar{R}_{t'} = ae^{-bt'}$ where $t' = (t - t_Q - 1), t > t_Q$ and, since by definition $\bar{R}_t = \bar{R}_{t-1} + r_t$, $t \geq 1$, we have that:

$$ae^{-bt'} - \gamma ae^{-b(t'-1)} = r_{t'} \iff r_{t'} = \bar{a}e^{-bt'},$$

where

$$\bar{a} = a(1 - \gamma e^b), \gamma = e^{-\beta} = e^{-\beta_2/(\Delta+1)}.$$

Since $r_{t'}$ represent mid-price differences, we end up with

$$\begin{aligned} m_t &= m_{t_Q+1} + \int_{t_Q+1}^t ds \bar{a} e^{-b(s-t_Q-1)} = c - \frac{\bar{a}}{b} e^{-b(t-t_Q-1)} \\ c &= m_{t_Q+1} + \frac{\bar{a}}{b} \\ t &\in [t_Q + 2, +\infty) \end{aligned} \quad (13)$$

with $\bar{a} < 0$, $b > 0$, i.e., we expect that the mid-price path decays exponentially with parameters that are linked to the ones which enter the fit of $\bar{R}_t$. We notice that at the end of the decay of $\bar{R}_t$, the model reverts to the standard ZI since the probability that a LO is a sell becomes approximately equal to 1/2. However, a permanent component of the impact can still persist and it corresponds to the difference between the parameter $c = \lim_{t \rightarrow +\infty} m_t$ and the mid-price just before the beginning of the execution.

In Fig. 13, we focus on the case $\beta_2 = \alpha = 0.001$ and two choices of trading intervals, i.e., $\Delta = 20, 100$. The evolution in time of $\bar{R}_t$ and the mid-price after the execution of the metaorder ends, are shown together with the exponential fits

$$\bar{R}_{t-t_Q-1} = ae^{-b(t-t_Q-1)} \text{ and } m_t = c - \frac{\bar{a}}{b} e^{-\tilde{b}(t-t_Q-1)}, \quad t \geq t_Q + 1.$$

We observe that the exponential fits are suitable to describe the reversion of $\bar{R}_t$ and of the mid-price. Moreover, by considering the estimated parameters that are reported in the figure's caption, we note that the agreement with the theoretical findings related to the equality between the parameters of the fits of the two quantities, qualitatively holds, i.e., $b \simeq \tilde{b}$, $\bar{a} \simeq a(1 - \gamma e^b)$ and $c \simeq m_{t_Q+1} + \bar{a}/\tilde{b}$.

The decay after the execution of the metaorder is analogous to the reversion we observe between two child MOs, with the difference that the next child MO is posted with an infinite trading interval and the exponential mid-price reversion cannot be anymore approximated by a straight line. Therefore, the role played by the parameters β_2 and α in determining the post-execution decay is equivalent to what happens for the reversion between two child MOs. In the following, we investigate how the permanent impact, the peak impact and the magnitude and the speed of the reversion, vary as functions of β_2 and α .

As we observe in Section 4.5, when traders use a longer memory in the evaluation of the price trend, i.e., smaller β_2 , the entity of the reversion increases and the concavity in the mid-price trajectory is more pronounced. Similarly, after the end of the execution of the metaorder, the mid-price slowly reverts to its initial value if the price trend indicator $\bar{R}_t$ is computed with a memory that is extremely large. This can be seen by comparing Fig. 5 and Fig. D16 in the Supplementary material. Both represent the mid-price trajectory for the execution of a metaorder with trading interval $\Delta = 20$, volume $Q = 2,000$ (left) and $Q = 100$ (right) however, in the former figure it is $\beta_2 = 10^{-3}$, while in the latter it is $\beta_2 = 10^{-6}$. As expected, when the memory of $\bar{R}_t$ increases, the limiting value of the price c approaches the mid-price before the execution starts and therefore, the permanent impact tends to 0. Analogously, given that the reversion is stronger and the overall impact of each child MO is weaker when β_2 is smaller, the peak impact is smaller if the memory of the price trend indicator $\bar{R}_t$ is larger. This is further highlighted in the left panel of Fig. D17 in the Supplementary

Table 1

NMZI model. Summary of the role played by β_2 and α . The meaning of the symbols $\uparrow$, $\downarrow$ and $\simeq$ is *increases*, *decreases* and *is approximately constant*. The double vertical lines separate three blocks: the first refers to the stationary regime, the second to the impact components and the third to the reversion after the execution ends.

	$\bar{R}^*$	r^*	τ^*	I^I	I^R	$I^I - I^R$	Peak impact	Speed of reversion	Permanent impact
$\beta_2 \downarrow$	$\uparrow$	$\downarrow$	$\uparrow$	$\simeq$	$\uparrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$
$\alpha \downarrow$	$\uparrow$	$\uparrow$	$\uparrow$	$\simeq$	$\downarrow$	$\uparrow$	$\uparrow$	$\downarrow$	$\uparrow$

material, where we plot the evolution of the mid-price before, during and after the execution of a metaorder of volume $Q = 100$, with $\Delta = 20$, $\alpha = 10^{-3}$ and several values of β_2 .

Also the behavior of the permanent impact and the peak impact as a function of α can be explained by the behavior of the reversion. When the intensity of reaction of the traders to the price trends, i.e. α , is weaker, so does the reversion mechanism. This implies that both the permanent impact and the maximum impact increase, as represented in the right panel of Fig. D17 in the Supplementary material where we plot the evolution of the mid-price before, after and during the execution of a meta-order of volume $Q = 100$, with $\Delta = 20$, $\beta_2 = 10^{-3}$, and several values of α .

Moreover Fig. D18 in the Supplementary material sheds light on the role played by both β_2 and α in the decay after the ending of the execution. This figure presents heat-maps representing $-\bar{a}/\bar{b}$ and $\log(2)/\bar{b}$ as functions of β_2 and α for $\Delta = 20, 100$. The ratio $\bar{a}/\bar{b}$ is the amount of the price reversion, whereas $\log(2)/\bar{b}$ is a measure of the decay speed of the mid-price. As expected, when β_2 decreases with α fixed, $-\bar{a}/\bar{b}$ increases, i.e. the magnitude of the reversion after the execution ends, is larger when the memory of $\bar{R}_t$ is longer. Moreover, when β_2 (α) decreases with α (β_2) fixed, $\log(2)/\bar{b}$ increases: the time to revert to a constant value is larger for α and β_2 smaller. This is consistent with what we have previously noted concerning the role of the two parameters in the reversion between two child MOs: as α or β_2 decreases, the transition to the stationary regime is slower.

Additionally, we consider the role of the metaorder size Q in determining the speed of impact decay. By fitting the mid-price trajectories after the execution in Figs. 5 and D16, we observe that the speed of decay $\bar{b}(\simeq b)$ is approximately independent of Q and of whether the stationary regime is reached during the metaorder execution:

- $\beta_2 = 10^{-3}$ (Fig. 5): if $Q = 100$, $\bar{b} = 0.000228 \pm 0.000001$ and if $Q = 2,000$, $\bar{b} = 0.000219 \pm 0.00001$;
- $\beta_2 = 10^{-6}$ (Fig. D16): if $Q = 100$, $\bar{b} = 0.000157 \pm 0.000001$ and if $Q = 2,000$, $\bar{b} = 0.000168 \pm 0.000001$.

We also note that the entity of the reversion after the execution, with respect to the peak impact, decreases as Q increases for given β_2 and α . Indeed, if we consider $\beta_2 = \alpha = 10^{-3}$, $\Delta = 20$ (as in Fig. 5), and several values of the metaorder size i.e. $Q = [10, 100, 10^3, 10^4]$, the reversion amounts at 79.13%, 73.56%, 40.24%, 7.28% of the peak impact, where both the peak impact and the price level at the end of the reversion (c in the exponential fit of the mid-price decay) are normalized with respect to the price before the beginning of the execution. Therefore, if the memory that is considered in the evaluation of the past price trend, manages to cover only part of the dynamics during the execution i.e. $Q \gg 1/\beta_2$, the reaction mechanism to the price trend is weaker. This is coherent with our previous finding that, for a given number of total shares, considering a longer memory leads to a stronger reversion.

To summarize our findings, Table 1 shows the considerations we made throughout the paper about the role played by the parameters β_2 and α in our model. The analogy between the behavior of the reversion between two child MOs and of the decay after the end of the execution is highlighted.

5. Conclusion

In this work we aim to model and reproduce the behavior of the price impact when a metaorder is executed, by introducing a novel statistical model of the limit order book, the Non-Markovian Zero Intelligence. As its name suggests, it is the non-Markovian counterpart of the pre-existing Zero Intelligence model [8–10]. In the original model, limit orders, market orders and cancellations in the LOB are modeled as three independent Poisson processes. Despite its simplicity, this model reproduces well several stylized facts, but it misses to capture the well-known concavity of the market impact, which is a desirable feature for a simulator, since it allows to test trading strategies in a realistic framework. In order to overcome this limitation, we propose our non-Markovian variant: contrary to the ZI model which considers equally probable sell and buy LOs, we define the probability that a LO is a sell, as price path dependent. Financially, this can be interpreted as traders who have reservation prices and submit LOs in response to the price trend. This apparently small modification leads to substantial accomplishments in terms of the price impact. During the execution of a metaorder, the concavity of the market impact emerges. Indeed, two regimes can be distinguished: initially, the price path is concave then, after a large number of child MOs, it stabilizes to a stationary regime which is characterized by a linear impact. In the first regime, the concavity originates because of the interplay between two mechanisms. For a given child MO of the metaorder e.g. buy, there is an immediate impact which makes the mid-price increase and more sell LOs are triggered. Consequently, more sell LOs in the spread occur and so, the mid-price decreases. This causes the price reversion in the intervals between two child MOs and this reversion results to be increasing with time. In parallel, also the immediate impact component varies in time since so does the imbalance between sell and buy LOs, making the first gap size in the ask side change. However, we observe that this

variation plays a marginal role: the pivotal contribution is due to the reversion mechanism and overall, the impact proves to be decreasing in time, causing the concavity. Finally, after the execution of the metaorder, the price path exponentially decreases.

In addition to analyze in depth the mechanisms at the root of the price impact in our model, we investigate how the results depend on the two main parameters. The former represents the inverse of the memory to evaluate the price trend which enters the probability that a LO is a sell. If this parameter is infinite, our model collapses to the ZI. We observe that if the memory is longer, the concave regime is longer, the convergence to the stationary regime is slower, the reversion mechanism is stronger and the permanent impact is smaller. The second parameter also enters the definition of the probability that a LO is a sell but it represents the intensity of the reaction to the price trend. The case with this parameter equal to 0 corresponds to the ZI model. Our results point out that if the reaction intensity is stronger, then the concave regime is shorter, the convergence to the stationary regime is faster, the reversion mechanism is stronger, and the permanent impact is smaller. To sum up, the shape of the market impact during and after the execution of a metaorder is well-captured when the decisions to submit LOs vary in time as a response to the price trend that is evaluated with a long memory.

The ability of our model to capture the market impact's concavity and the price reversion after the execution of a metaorder ends, is accompanied by three other bright sides, which make it an innovative and relevant contribution. This model works for small-tick assets: the modeling of their dynamics is under explored in the literature, given its challenging nature. Also, this model is completely explainable and even though it is not fully analytical, several relations can still be derived. We derive master equations which describe the evolution of the spread and the mid-price in the intervals between two child MOs, and they provide a deeper explanation of the mechanism underlying our model.

The NMZI model inherits several limitations from the standard ZI, e.g., the unrealistic modeling based on Poisson processes and the autocorrelation of MO signs. Also, as the ZI model, the NMZI captures the linear scaling of the aggregate-volume impact that is observed empirically for small imbalances, but it fails to reproduce the saturation for large imbalances and the sigmoidal shape. On the other hand, compared to the ZI, the innovation of our model introduces mean-reversion and this mechanism prolongs the convergence towards diffusive dynamics of the mid-price when small memory parameter values are considered. Nevertheless, diffusive behavior is eventually reached for the long-term dynamics. Concerning the well-known square-root law of market impact, our model fails to capture the independence of the impact from the trading speed, nor the square root scaling as a function of volume. However, similarly to [6], we observe a crossover between an initial steep market impact regime followed by a concave regime.

Several extensions could be interesting to study in future works. In this paper, we focused on the market impact associated with the executions of metaorders with constant speed but undoubtedly, more complex trading strategies could be investigated. Moreover, another version of our model with two types of agents that have different submission and cancellation rates (fast and slow traders) could be investigated, referring to [52]. Also, we consider a specific choice of our indicator of the past price trend $\bar{R}_t$, that is the exponentially weighted mid-price return. However, instead of the exponential damping factor, different kernels, as a power-law function, could be employed. Similarly, we define the time-varying probability that a LO is a sell as a sigmoid function of $\bar{R}_t$ but, of course, other functional forms could be considered. Finally, similarly to what we do for LOs, also the sampling of cancellation signs could be performed via a time-varying probability that depends on the past price trend.

CRedit authorship contribution statement

Adele Ravagnani: Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Fabrizio Lillo:** Writing – review & editing, Supervision, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

FL acknowledges support from the grant PRIN2022 DD N. 104 of February 2, 2022 “Liquidity and systemic risks in centralized and decentralized markets”, codice proposta 20227TCX5W - CUP J53D23004130006 funded by the European Union NextGenerationEU through the Piano Nazionale di Ripresa e Resilienza (PNRR).

Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.physa.2025.131056>.

Data availability

The data used to estimate the model's parameters can be purchased on the website lobsterdata.com.

References

- [1] K. Jain, N. Firoozye, J. Kochems, P. Treleven, Limit order book simulations: A review, 2024, <http://dx.doi.org/10.48550/arXiv.2402.17359>, arXiv preprint.
- [2] A.S. Kyle, Continuous auctions and insider trading, *Econ.: J. Econ. Soc.* (1985) 1315–1335, <http://dx.doi.org/10.2307/1913210>.
- [3] J.P. Bouchaud, J. Bonart, J. Donier, M. Gould, Trades, Quotes and Prices: Financial Markets under the Microscope, Cambridge University Press, 2018.
- [4] 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, *Phys. Rev. E* 80 (6) (2009) 066102, <http://dx.doi.org/10.1103/PhysRevE.80.066102>.
- [5] X. Brokmann, E. Sérié, J. Kockelkoren, J.P. Bouchaud, Slow decay of impact in equity markets, *Mark. Microstruct. Liq.* 01 (02) (2015) 1550007, <http://dx.doi.org/10.1142/S2382626615500070>.
- [6] 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, *Mark. Microstruct. Liquidity* 1 (02) (2015) 1550004, <http://dx.doi.org/10.1142/S2382626615500045>.
- [7] F. Bucci, M. Benzaquen, F. Lillo, J.P. Bouchaud, Slow decay of impact in equity markets: Insights from the ANcerno database, *Mark. Microstruct. Liq.* 04 (03n04) (2018) 1950006, <http://dx.doi.org/10.1142/S2382626619500060>.
- [8] M. Daniels, J. Farmer, L. Gillemot, G. Iori, E. Smith, Quantitative model of price diffusion and market friction based on trading as a mechanistic random process, *Phys. Rev. Lett.* 90 (2003) 108102, <http://dx.doi.org/10.1103/PhysRevLett.90.108102>, 04.
- [9] E. Smith, J.D. Farmer, L. Gillemot, S. Krishnamurthy, Statistical theory of the continuous double auction, *Quant. Finance* 3 (6) (2003) 481–514, <http://dx.doi.org/10.1088/1469-7688/3/6/307>.
- [10] J.D. Farmer, P. Patelli, I.I. Zovko, The predictive power of zero intelligence in financial markets, *PNAS* 102 (6) (2005) 2254–2259, <http://dx.doi.org/10.1073/pnas.0409157102>.
- [11] J.P. Bouchaud, M. Mézard, M. Potters, Statistical properties of stock order books: empirical results and models, *Quant. Finance* 2 (4) (2002) 251–256, <http://dx.doi.org/10.1088/1469-7688/2/4/301>.
- [12] R. Cont, S. Stoikov, R. Talreja, A stochastic model for order book dynamics, *Oper. Res.* 58 (3) (2010) 549–563, <http://dx.doi.org/10.1287/opre.1090.0780>.
- [13] F. Abergel, A. Jedidi, A mathematical approach to order book modelling, in: F. Abergel, B.K. Chakrabarti, A. Chakraborti, M. Mitra (Eds.), *Econophysics of Order-Driven Markets. New Economic Windows*, Milano, Springer, 2011, http://dx.doi.org/10.1007/978-88-470-1766-5_7.
- [14] H. Hult, J. Kiessling, Algorithmic trading with Markov chains, 2010, Preprint: <https://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-25075>.
- [15] W. Huang, C.-A. Lehalle, M. Rosenbaum, Simulating and analyzing order book data: The queue-reactive model, *J. Am. Stat. Assoc.* 110 (509) (2015) 107–122, <http://dx.doi.org/10.1080/01621459.2014.982278>.
- [16] X. Lu, F. Abergel, Order-book modeling and market making strategies, *Mark. Microstruct. Liq.* 4 (01n02) (2018) 1950003, <http://dx.doi.org/10.1142/S2382626619500035>.
- [17] I.M. Toke, Market making in an order book model and its impact on the spread, in: F. Abergel, B.K. Chakrabarti, A. Chakraborti, M. Mitra (Eds.), *Econophysics of Order-Driven Markets. New Economic Windows*, Milano, Springer, 2011.
- [18] B. Zheng, F. Roueff, F. Abergel, Ergodicity and scaling limit of a constrained multivariate hawkes process, 2014, <http://dx.doi.org/10.48550/arXiv.1301.5007>, arXiv preprint.
- [19] E. Bacry, T. Jaisson, J.-F. Muzy, Estimation of slowly decreasing hawkes kernels: application to high-frequency order book dynamics, *Quant. Finance* 16 (8) (2016) 1179–1201, <http://dx.doi.org/10.1080/14697688.2015.1123287>.
- [20] X. Lu, F. Abergel, High-dimensional hawkes processes for limit order books: modelling, empirical analysis and numerical calibration, *Quant. Finance* 18 (2) (2018) 249–264, <http://dx.doi.org/10.1080/14697688.2017.1403142>.
- [21] O. Mounjid, M. Rosenbaum, P. Saliba, From asymptotic properties of general point processes to the ranking of financial agents, 2019, <http://dx.doi.org/10.48550/arXiv.1906.05420>, arXiv preprint.
- [22] K. Lee, B.K. Seo, Modeling bid and ask price dynamics with an extended hawkes process and its empirical applications for high-frequency stock market data, *J. Financ. Econ.* 21 (4) (2023) 1099–1142, <http://dx.doi.org/10.1093/jfinec/nbab029>.
- [23] M.E. Paddrik, R.L. Hayes, A. Todd, S.Y. Yang, W. Scherer, P. Beling, An agent based model of the E-mini S & P 500 and the flash crash, 2012, <http://dx.doi.org/10.2139/ssrn.1932152>, SSRN preprint.
- [24] D. Byrd, M. Hybinette, T. Hybinette Balch, ABIDES: Towards high-fidelity market simulation for AI research, 2019, <http://dx.doi.org/10.48550/arXiv.1904.12066>, arXiv preprint.
- [25] P. Belcak, J.P. Callies, S. Zohren, Fast agent-based simulation framework of limit order books with applications to pro-rata markets and the study of latency effects, 2020, <http://dx.doi.org/10.48550/arXiv.2008.07871>, arXiv preprint.
- [26] M. Dicks, A. Paskaramoorthy, T. Gebbie, A simple learning agent interacting with an agent-based market model, *Phys. A: Stat. Mech. Appl.* 633 (2024) 129363, <http://dx.doi.org/10.1016/j.physa.2023.129363>.
- [27] J. Li, X. Wang, Y. Lin, A. Sinha, M.P. Wellman, Generating realistic stock market order streams, 2020, <http://dx.doi.org/10.48550/arXiv.2006.04212>, arXiv preprint.
- [28] Z. Shi, Y. Chen, Cartledge J., The LOB recreation model: Predicting the limit order book from TAQ history using an ordinary differential equation recurrent neural network, 2021, <http://dx.doi.org/10.48550/arXiv.2103.01670>, arXiv preprint.
- [29] Z. Shi, J. Cartledge, The limit order book recreation model (LOBRM): An extended analysis, in: *Machine Learning and Knowledge Discovery in Databases. Applied Data Science Track: European Conference ECML PKDD 2021*, Bilbao, Spain, September (2021) 13–17, Proc. Part IV 21, Springer, pp. 204–220, http://dx.doi.org/10.1007/978-3-030-86514-6_13.
- [30] A. Coletta, A. Moulin, S. Vyetenko, T. Balch, Learning to simulate realistic limit order book markets from data as a world agent, in: *Proc. Third ACM Int. Conf. AI in Finance*, 2022, pp. 428–436, <http://dx.doi.org/10.1145/3533271.3561753>.
- [31] R. Cont, M. Cucuringu, J. Kochems, F. Prenzel, Limit order book simulation with generative adversarial networks, 2023, <http://dx.doi.org/10.2139/ssrn.4512356>, SSRN preprint.
- [32] P. Nagy, S. Frey, S. Sapor, K. Li, A. Calinescu, S. Zohren, J. Foerster, Generative AI for end-to-end limit order book modelling: A token-level autoregressive generative model of message flow using a deep state space network, in: *Proc. Fourth ACM Int. Conf. AI in Finance*, pp. 91–99, <http://dx.doi.org/10.1145/3604237.3626898>.
- [33] P. Kumar, Deep hawkes process for high-frequency market making, *J. Bank. Financ. Technol.* 8 (2024) 11–28, <http://dx.doi.org/10.1007/s42786-024-00049-8>.
- [34] P. Lakner, J. Reed, S. Stoikov, High frequency asymptotics for the limit order book, *Mark. Microstruct. Liq.* 2 (01) (2016) 1650004, <http://dx.doi.org/10.1142/S2382626616500040>.
- [35] W. Huang, M. Rosenbaum, Ergodicity and diffusivity of Markovian order book models: a general framework, *SIAM J. Financ. Math.* 8 (1) (2017) 874–900, <http://dx.doi.org/10.1137/16M1064337>.
- [36] B. Hambly, J. Kalsi, N. Newbury, Limit order books, diffusion approximations and reflected SPDEs: from microscopic to macroscopic models, *Appl. Math. Finance* 27 (1–2) (2020) 132–170, <http://dx.doi.org/10.1080/1350486X.2020.1758176>.
- [37] R. Cont, P. Degond, L. Xuan, A mathematical framework for modelling order book dynamics, *SIAM J. Financ. Math.* 16 (1) (2025) 123–166, <http://dx.doi.org/10.1137/22M1541538>.

- [38] D. Diana, T. Gebbie, Non-uniformly sampled simulated price impact of an order-book, *J. Comput. Appl. Math.* 456 (2025) 116202, <http://dx.doi.org/10.1016/j.cam.2024.116202>.
- [39] J. Donier, J. Bonart, I. Mastromatteo, J.P. Bouchaud, A fully consistent, minimal model for non-linear market impact, *Quant. Finance* 15 (7) (2015) 1109–1121, <http://dx.doi.org/10.1080/14697688.2015.1040056>.
- [40] A. Fosset, J.P. Bouchaud, M. Benzaquen, Endogenous liquidity crises, *J. Stat. Mech. Theory Exp.* (2020) <http://dx.doi.org/10.1088/1742-5468/ab7c64>.
- [41] B. Tóth, Y. Lempriere, C. Deremble, J. De Lataillade, J. Kockelkoren, J.P. Bouchaud, Anomalous price impact and the critical nature of liquidity in financial markets, *Phys. Rev. X* 1 (2) (2011) 021006, <http://dx.doi.org/10.1103/PhysRevX.1.021006>.
- [42] K.C. Chan, On the contrarian investment strategy, *J. Bus.* 61 (2) (1988) 147–163, <http://dx.doi.org/10.1086/296425>.
- [43] J. Lakonishok, A. Shleifer, R.W. Vishny, Contrarian investment, extrapolation, and risk, *J. Financ.* 49 (1994) 1541–1578, <http://dx.doi.org/10.1111/j.1540-6261.1994.tb04772.x>.
- [44] T. Odean, Are investors reluctant to realize their losses? *J. Financ.* 53 (5) (1998) 1775–1798, <http://dx.doi.org/10.1111/0022-1082.00072>.
- [45] M. Grinblatt, M. Keloharju, The investment behavior and performance of various investor types: a study of Finland's unique data set, *J. Financ. Econ.* 55 (1) (2000) 43–67, [http://dx.doi.org/10.1016/S0304-405X\(99\)00044-6](http://dx.doi.org/10.1016/S0304-405X(99)00044-6).
- [46] S.G. Badrinath, S. Wahal, Momentum trading by institutions, *J. Financ.* 57 (6) (2002) 2449–2478, <http://dx.doi.org/10.1111/1540-6261.00502>.
- [47] L. De Haan, J. Kakes, Momentum or contrarian investment strategies: Evidence from dutch institutional investors, *J. Bank. Financ.* 35 (9) (2011) 2245–2251, <http://dx.doi.org/10.1016/j.jbankfin.2011.01.027>.
- [48] L. Che, Investor types and stock return volatility, *J. Empir. Financ.* 47 (1) (2018) 139–161, <http://dx.doi.org/10.1016/j.jempfin.2018.03.005>.
- [49] M. Baltzer, S. Jank, E. Smajlbegovic, Who trades on momentum? *J. Financ. Mark.* 42 (1) (2019) 56–74, <http://dx.doi.org/10.1016/j.finmar.2018.08.003>.
- [50] M. Grinblatt, G. Jostova, L. Petrasek, A. Philipov, Style and skill: hedge funds, mutual funds, and momentum, *Manag. Sci.* 66 (12) (2020) 5505–5531, <http://dx.doi.org/10.1287/mnsc.2019.3433>.
- [51] R. Bradrania, W. Wu, Foreign institutions, local investors and momentum trading, *J. Empir. Financ.* 73 (2023) 40–64, <http://dx.doi.org/10.1016/j.jempfin.2023.05.005>.
- [52] M. Benzaquen, J.P. Bouchaud, Market impact with multi-timescale liquidity, *Quant. Finance* 18 (11) (2018) 1781–1790, <http://dx.doi.org/10.1080/14697688.2018.1444279>.