Optimal Market-Making Strategies under Synchronised Order Arrivals with Deep Neural Networks

So Eun Choi[a], Hyun Jin Jang[b,*], Kyungsub Lee[c], Harry Zheng[d]

[a]*Department of Mathematical Sciences, Korea Advanced Institute of Science and Technology (KAIST), Daejeon 34141, Republic of Korea*
[b]*School of Business Administration, Ulsan National Institute of Science and Technology (UNIST), Ulsan 44919, Republic of Korea*
[c]*Department of Statistics, Yeungnam University, Gyeongsan, Republic of Korea*
[d]*Department of Mathematics, Imperial College, London SW7 2AZ, United Kingdom*

Abstract

This study investigates the optimal execution strategy of market-making for market and limit order arrival dynamics under a novel framework that includes a synchronised factor between buy and sell order arrivals. Using statistical tests, we empirically confirm that a synchrony propensity appears in the market, where a buy order arrival tends to follow the sell order's long-term mean level and vice versa. This is presumably closely related to the drastic increase in the influence of high-frequency trading activities in markets. To solve the high-dimensional Hamilton–Jacobi–Bellman equation, we propose a deep neural network approximation and theoretically verify the existence of a network structure that guarantees a sufficiently small loss function. Finally, we implement the terminal profit and loss profile of market-making using the estimated optimal strategy and compare its performance distribution with that of other feasible strategies. We find that our estimation of the optimal market-making placement allows significantly stable and steady profit accumulation over time through the implementation of strict inventory management.

Keywords: Optimal strategy, Order arrival models, Synchrony, High-dimensional Hamilton–Jacobi–Bellman, Deep neural network
JEL: G13, C63, C22

1. Introduction

The emergence of innovative technologies has accelerated the paradigm shift in trading activities in financial markets. In particular, automated trading based on ultra-low latency

*Corresponding author
Email addresses: cksoeun@kaist.ac.kr (So Eun Choi), janghj@unist.ac.kr (Hyun Jin Jang), ksublee@yu.ac.kr (Kyungsub Lee), h.zheng@imperial.ac.uk (Harry Zheng)

electronic systems facilitates order generation, routing, and execution, all within a fraction of a second. It uses computerised algorithmic trading called high-frequency trading (HFT) that has grown dramatically over the past decade and now accounts for 55% and 40% of all trades in the US and European equity markets, respectively; it is also rapidly growing in the Asian market for a variety of asset classes (Miller and Shorter, 2016).

In general, little is publicly known about HFT strategies. According to the concept releases of the US Securities and Exchange Commission (SEC, 2010, 2014), these strategies can be categorised into two groups by risk appetite: passive and aggressive. Passive strategies include market-making and arbitrage trading, which rarely depend on the direction of price movements. Market-making trades mainly provide liquidity to the marketplace, exploiting both bid and ask orders to generate a profit from the bid/ask spread. Arbitrage trading is generally pursued to generate a profit from the price disparities among related securities such as exchange-traded funds and baskets of underlying stocks. Meanwhile, aggressive strategies involve momentum ignition and order anticipation strategies, which use the price direction along with either a long or a short position. A momentum ignition strategy aims to trigger sharp price movements – either up or down – by initiating a series of orders. An order anticipation strategy seeks to identify large institutional orders and then trade ahead of those orders in anticipation that they will move market prices. Although these strategies are not new, the advanced technology now available may enable traders to better identify profit opportunities and execute their strategies more effectively than in the past.

These HFT strategies can lead financial markets to become more synchronised and substantially increase correlations in the price structure because HFT is more likely to occur by tracking price movement patterns than changes in market fundamentals. In addition, orders tend to be submitted as a pair of long and short sides (i.e. round-trip trades) and they are also executed subsequently and repeatedly on one side – either buy or sell. Gerig (2015) proposes a single-period model of synchrony in financial markets caused by HFT by drawing similarities with the behaviour of animal groups such as schooling fish and herding birds. Given the current market circumstances, it is worthwhile to discuss how a model can capture the highly dependent structure of order flow arrivals attributed to HFT activities. In addition, for market-makers that may be designated by a firm or pursue a market-making HFT strategy, the optimal placement of the bid/ask spread is a crucial issue.

Despite the market structure changing towards a hyper-correlation regime, most studies modelling high-frequency dynamics still employ exciting factor-based Hawkes models. However, synchronisation in order flow arrivals might not be temporary. In other words, it can disappear if no subsequent orders are placed, owing to the continued likelihood of

orders being synchronised because of HFT activities. In this context, we propose a model for market order arrivals and limit order book dynamics based on the Hawkes process with the addition of a novel variable, referred to as a *synchronising factor*. This factor makes contemporaneous interactions between buy and sell orders feasible and also enables the interacting effect to remain permanent in the long-term mean of order flows. We then consider a decision-making problem for a market-maker when the synchronised market and limit orders enter a trading book. We finally derive the market-maker's optimal trading strategy by maximising his or her profit and liquidating the inventory over a finite period.

Our order arrival model is based on Hawkes processes (Hawkes, 1971; Hawkes and Oakes, 1974) that have been employed as a tool for modelling price movements in high-frequency dynamics because of their great flexibility and versatility. As a pioneering work, Bowsher (2007) introduces a bivariate Hawkes process to model the joint dynamics of trades and mid-price changes on New York Stock Exchange (NYSE) stocks. Large (2007) formulises the resiliency of limit order books based on a 10-variate Hawkes process by testing stocks on the London Stock Exchange. In addition, many studies have employed Hawkes processes in high-frequency finance (Bacry et al., 2012; Da Fonseca and Zaatour, 2014; Aït-Sahalia et al., 2015; Lee and Seo, 2017). The key feature of Hawkes-based models is the inclusion of exciting factors in their intensity processes. This means that a counting process is more likely to increase when the counting event arrives because its intensity can instantly jump depending on the movement of the original process. Such a design enables us to capture the clustering phenomenon of arrivals in the counting process by using a feedback kernel of its intensity that communicates with the counting arrivals.

In the proposed Hawkes intensity process, the exciting and synchronising components appear to play similar roles, which leads to a strong correlation over time. Nevertheless, from the perspective of the mechanism causing the abnormal impacts, synchronisation is different from excitation. Exciting events are activated by certain external stimulations such as a status change in the original and other relevant processes, and hence they accelerate the arrivals of subsequent orders instantaneously. By contrast, the synchronising factor enhances the likelihood of integrating two processes irrespective of exogenous events. In other words, for the exciting factor, intensity processes are assumed to be independent unless the underlying process changes because order arrivals lead to a temporary increase in intensity. On the contrary, the synchronising factor forces one process to lead the other, and vice versa, implying that the two processes tend to be associated endogenously in the long-term beyond market fundamentals.

This study makes three primary contributions. First, we examine an evidence of the existence of the synchronising factor using market order data. For the six representative stocks in the pool of US large caps (IBM, Chevron, Apple, Amazon, JP Morgan, and

Microsoft), we find a remarkable increase in the synchronising tendency between market buy and sell order arrivals in 2018 compared with in 2008. The increase over the past decade differs across stocks, with the highest change observed for the JP Morgan stock. This finding implies that the frequencies of market buy and sell orders in a unit time tend to interact more dynamically with each other than they used to do in the past.

Second, we adopt the deep neural network (DNN) technique to solve the optimisation problem derived as a high-dimensional partial differential equation (PDE) with the discontinuity terms driven by the exciting factors and discretely varying inventory amount in a market-maker's trading. This is generally known to be difficult to solve analytically. To solve the PDE numerically, we propose the DNN-based approximation inspired by the deep Galerkin method (Sirignano and Spiliopoulos, 2018), a mesh-free simulation suitable for applying to high-dimensional PDE problems. We verify that the existence of a DNN structure that guarantees that the loss function – defined using the boundary and terminal conditions of a given PDE – is sufficiently small in a compact domain. The DNN algorithm enables us to construct an approximate solution with a low numerical error, which is trained to satisfy the fact that the loss function is minimised for the generated random samples over the PDE domain.

Third, we conduct a variety of computational simulations to derive implications from the perspective of wealth management in high-frequency market-making by assuming models with and without the synchrony effect as well as those with different stability levels. In terms of synchrony, we compare the profit performance of posting optimal strategies with presence of the synchronising factor under scenarios in which the order dynamics are fully synchronised. The results show that the strategy considering the synchronising effect produces more gains and less risks than the one that partly considers or does not. In terms of stability, the more unstable the high-frequency market, the more likely market-makers are to obtain higher expected returns. These findings are meaningful to market practitioners given the increase in instability in the high-frequency market for some particular stocks in the past ten years.

The remainder of this paper is organised as follows. Section 2 reviews the relevant literature. Section 3 develops the model of market and limit order dynamics, including the synchronising factor, and builds an optimal execution problem for market-makers with the derivation of the associated PDE. Section 4 presents empirical evidence of the existence of synchrony in the proposed model. Section 5 discusses the DNN to estimate the PDE and its convergence, and Section 6 presents the simulation results under the DNN estimation to see the difference in a market-maker's wealth using the optimal strategies when facing various market situations. Finally, Section 7 concludes. The technical proofs are presented in Appendix A. Additional simulations are displayed in Appendix B and Appendix C.

2. Literature Review

This study is related to prior research that applies DNNs to solve PDEs and optimisation problems for market-makers under highly correlated price dynamics. This section reviews the literature on the application of machine learning to finance, optimal placement strategies for market-making, and models of market price structure with hyper dependency.

The application of machine learning to quantitative finance fields has rapidly expanded in recent years. Studies have examined the problems of pricing and hedging derivatives (Hutchinson et al., 1994; Gramacy and Ludkovski, 2015; De Spiegeleer et al., 2018), predicting financial markets (Sirignano, 2019; Dixon et al., 2019), and managing credit risk (Baesens et al., 2003; Fitzpatrick and Mues, 2016; Loterman et al., 2012; Khandani et al., 2010). Among the variety of machine learning approaches, DNNs employ a multilayer structure in the neural network framework to learn more about complex nonlinear relationships. In particular, DNNs are suitable for modelling high-dimensional nonlinear problems, which allow the estimation of arbitrarily continuous functions on compact sets (Hornik et al., 1989; Hornik, 1991). For example, DNNs are used as a numerical scheme for high-dimensional PDEs to overcome the curse of the dimensionality problem. They also enable the consideration of stochastic optimal control and exotic option pricing beyond the Black-Scholes setting, which can be transformed to a problem of solving nonlinear PDEs. To solving high-dimensional PDEs, one method employs DNNs with a network structure similar to the long short-term memory (LSTM) which is trained to satisfy the given PDE conditions (Sirignano and Spiliopoulos, 2018; Berg and Nystrom, 2018). Another stream is based on the backward stochastic differential equation approach that resembles the spirit of deep reinforcement learning with gradient acting as the policy function (E et al., 2017; Han et al., 2018; Fujii et al., 2019).

The optimal decision of a market-maker under order dynamics has been widely studied in the market microstructure field. Ho and Stoll (1981) discuss the optimal market-making policy by specifying a true price for assets on the supply and demand curves of a public market. They derive the optimal bid/ask quotes around the true price by accounting for the inventory effect. In this spirit, Avellaneda and Stoikov (2008) propose a market-making model in an order book by employing a diffusion process for a mid-price and a Poisson process for executed limit orders. For the exponential utility function, this provides an asymptotic solution for quoting spreads and reservation prices. Rosu (2009) derives an equilibrium transaction price between market and limit orders in terms of utility by considering the trade-off between execution prices and waiting costs at the bounded discrete price levels. Gueant et al. (2013) study the same problem as Avellaneda and Stoikov (2008) by adding inventory volume constraints and then approximate the optimal control with asymptotic limits over an infinite time horizon. Cartea et al. (2014) model the dynamics of

the arrival of market orders and resulting changes in the limit order book's shape with self-exciting and mutually exciting Hawkes processes. Guilbaud and Pham (2013) investigate the optimal market-making policy when an agent uses both limit, and market orders by employing a numerical method that estimates the optimal problem. Cartea and Jaimungal (2013) study modelling price revision and duration for HFT activities using the hidden Markov model with regime switching. Veraart (2010) models two types of market-making actions in foreign exchange markets, removing, and adding liquidity with two-dimensional Brownian motions. Guo et al. (2017) employ a correlated random walk for the best bid/ask prices and solve the optimal placement problem with a reflection principle.

Hyper-dependence in price dynamics has been modelled in the context of cointegration in econometrics and flocking behaviour. Cointegration occurs when two or more non-stationary time series are driven by one or more common nonstationary time series, as proposed in the seminal works by Granger (1981) and Engle and Granger (1987). Many financial data series exhibit cointegration, such as, international stock markets (Cerchi and Havenner, 1988; Duan and Pliska, 2004), foreign exchange rates (Baillie and Bollerslev, 1989; Kellard et al., 2010), futures and spot prices (Ng and Pirrong, 1994), and commodity futures prices (Chiu et al., 2015; Jang et al., 2020). Similarly, flocking is referred to as a collective motion of a large number of self-propelled entities. Reynolds (1987) proposes the break-through algorithm that makes generating realistic computer simulations of flocking agents feasible. Flocking behaviour is observed in many areas in physics, biology, engineering, and human systems, including financial markets (e.g. Rauch et al., 1995; Huepe and Aldana, 2008; Ha et al., 2015). In a similar manner, Gerig (2015) employs the term 'synchrony' in financial markets, which implies that market prices are monitored by several traders who quickly disseminate information to each other, thus possibly resulting in price comovement. Such a phenomenon resembles animal group behaviours, wherein individuals communicate with each other to share information by contemporaneously scanning the environment using many eyes.

3. Market Models and Optimisation Problems

In Section 3.1, we describe in detail the model of arrivals for buy and sell orders that interact with each other in market and limit orders. We then pose an optimal execution problem for a market-maker who seeks to maximize his/her wealth from his/her round-trip trades until the end of the day by penalising inventories. We derive the Hamilton–Jacobi–Bellman (HJB) equation in Section 3.2.

3.1. Model setup

We first define a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$ satisfying the usual conditions. It is assumed that all the stochastic processes introduced in this paper are defined on $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$. Let S_t be the mid-price of the asset at time t with the dynamics

$$\mathrm{d}S_t = \sigma S_t \mathrm{d}W_t,$$

where σ is a positive constant and W is a standard Brownian motion.

Consider a market-maker who continuously posts a limit buy order and sell order of the asset with depth $\delta_t^-, \delta_t^+ \geq 0$, respectively. In other words, the market-maker posts a buy limit order at a price of $S_t - \delta_t^-$, and a sell limit order at a price of $S_t + \delta_t^+$. The market-maker provides liquidity to the market and earns profits from the bid-ask spread.

We assume that transactions only occur when market orders arrive and match with the pending limit orders posted by the market-maker. Let the counting processes M_t^+ and M_t^- with intensities λ_t^+ and λ_t^- denote the arrival of other participants' buy and sell market orders, respectively. We denote the market-maker's filled buy and sell limit orders by the counting processes N_t^- and N_t^+, respectively. As a measure of the probability with which the market-maker's limit buy and sell orders are executed, we consider the fill probabilities $h(\delta_t^\pm, c_t^\pm)$ at time t for limit orders placed $\delta_t^\pm$ away from S_t. Since the market buy orders lift the market-maker's sell limit orders and the market sell orders hit the market-maker's buy limit orders, the probability that limit orders are filled increases as the distance from the mid-price δ decreases, and vice versa. The process c_t can be interpreted as the parameters directly determining the shape of the limit order book.

From this setup, the processes $N_t^\pm$ can be regarded as the pathwise stochastic integral with respect to $M_t^\pm$:

$$N_t^- = \int_0^t I_s^- \mathrm{d}M_s^- \text{ and } N_t^+ = \int_0^t I_s^+ \mathrm{d}M_s^+. \tag{1}$$

Here, $I_t^\pm$ are defined by

$$I_t^- = \begin{cases} 0 & \text{if } M_t^- - M_{t-}^- = 0 \\ \epsilon_t^- & \text{otherwise} \end{cases} \text{ and } I_t^+ = \begin{cases} 0 & \text{if } M_t^+ - M_{t-}^+ = 0 \\ \epsilon_t^+ & \text{otherwise} \end{cases}, \tag{2}$$

where $\epsilon_t^\pm$ are Bernoulli random variables with probability $h(\delta_{t-}^\pm, c_{t-}^\pm)$, respectively. If the expectation of $\lambda_t^\pm$ is bounded for $0 \leq t \leq T$, $M_t^\pm$ have only a finite number of jumps along the time interval $[0, T]$ almost surely. It follows that $I_t^\pm$ are progressively measurable, and therefore $N_t^\pm$ are well defined as the Lebesgue-Stieltjes integrals and are progressively measurable. We address the condition for bounded intensities later in Lemma 1.

In this study, we assume that market order volumes are independent and identically distributed and exponentially distributed and that the shape of the limit order book is flat, similar to the setup used by Avellaneda and Stoikov (2008) and Cartea et al. (2014). Given that market orders arrive, the probability that a limit order at price level $S_t \pm \delta_t^{\pm}$ is executed is assumed to be $h(\delta_t^{\pm}, c_t^{\pm}) = \mathrm{e}^{-\delta_t^{\pm} c_t^{\pm}}$.

From the flatness assumption of the limit order book, c_t can be regarded as the market depth of the order book. Precisely, c_t^- and c_t^+ represent the depth of the limit buy (-) and sell (+) order books at the price levels, respectively. For a small c, this suggests that the volume of accumulating limit orders before where the agent's posted is almost zero, which implies that the posted limit orders are highly likely to be executed by market orders even when they are a long distance from the mid-price δ. Conversely, in the case of a large c, the posted limit orders should wait until the stacking limit orders are executed even for a short distance δ.

The market-maker's execution processes have two driving factors: the frequency at which market orders walk into the limit order book (i.e. λ_t) and the volume depth for the limit order book (i.e. c_t). We model λ_t and c_t by including comovement features and feedback effects in the limit and market order arrival dynamics.

First, we design the intensity processes for market sell/buy orders $\boldsymbol{\lambda}_t = (\lambda_t^-, \lambda_t^+)_{t \geq 0}$ as a process that has interacting, self-exciting, and mutually exciting features:

$$\begin{aligned} \mathrm{d}\lambda_t^- &= \beta(\theta^- - \lambda_t^- + \kappa\lambda_t^+)\mathrm{d}t + \eta \mathrm{d}M_t^- + \nu \mathrm{d}M_t^+, \\ \mathrm{d}\lambda_t^+ &= \beta(\theta^+ - \lambda_t^+ + \kappa\lambda_t^-)\mathrm{d}t + \eta \mathrm{d}M_t^+ + \nu \mathrm{d}M_t^- \end{aligned} \tag{3}$$

where $\beta, \theta^{\pm}$ are strictly positive and η, ν, κ are non-negative coefficients. The market order arrival intensity jumps immediately after any market order arrival, where the parameters η and ν govern the responsiveness of the self-exciting and mutually-exciting components of the intensity due to additional market orders, respectively. However, their states revert to the original mean level with speed β since the exciting impacts from market order arrivals are temporary.

Unlike existing models, we take the stochastic mean-reversion level of $\theta^- + \kappa\lambda_t^+$ for the sell intensity and $\theta^+ + \kappa\lambda_t^-$ for the buy intensity, where $\theta^{\pm}$ are constant mean levels. This indicates that two processes may have a higher endogenous correlation besides the exogenous exciting responses due to the increase in strategy executions related to round-trip trading by HFT traders. Hence, even when no buy or sell market order arrives, the two processes can retain their interaction tendency in dynamics. From a traditional economics perspective, the effect could be insignificant if transactions only exist for typical supply and demand purposes. However, it is necessary to consider this effect to model the current market microstructure in electronic exchange-traded assets with high liquidity. Thus, the

parameter κ represents the degree to which the strength of such an endogenous interactive tendency is retained in buy and sell orders; hereafter, we call this the *synchronising factor*.

Second, we assume that the depth dynamics of the limit buy/sell order book $\boldsymbol{c}_t = (c_t^-, c_t^+)$ are the processes that are excited by market order arrivals and are coupled throughout the synchronising factor:

$$\begin{aligned} \mathrm{d}c_t^- &= \xi(\alpha - c_t^- + \kappa_c c_t^+)\mathrm{d}t + \eta_c \mathrm{d}M_t^- + \nu_c \mathrm{d}M_t^+, \\ \mathrm{d}c_t^+ &= \xi(\alpha - c_t^+ + \kappa_c c_t^-)\mathrm{d}t + \eta_c \mathrm{d}M_t^+ + \nu_c \mathrm{d}M_t^- \end{aligned} \tag{4}$$

where α, ξ are strictly positive and η_c, ν_c, κ_c are nonnegative constants. This indicates that the depth of the limit order book jumps when both market buy/sell orders arrive. This is a one-way effect because market orders cause jumps in subsequent limit orders, whereas jumps in limit orders do not induce jumps in market order arrivals (Large, 2007). In addition, the depth processes of both limit buy/sell orders interact with each other similar to λ_t. Precisely, α is the long-run mean level, κ_c is the synchronising factor, ξ is the mean-reverting speed, and η_c, ν_c are the exciting terms of market orders for the depth processes in the buy/sell sides of the limit order book.

To examine the conditions that guarantee that the intensity processes for market and limit orders are stable, we define the mean future rate, $m_t^{\pm}(u) = \mathbb{E}[\lambda_u^{\pm}|\mathcal{F}_t]$ and $n_t^{\pm}(u) = \mathbb{E}[c_u^{\pm}|\mathcal{F}_t]$, respectively, for $u \geq t$. For the intensity processes $\lambda_t^{\pm}$ ($c_t^{\pm}$) to be stable, $m_t^{\pm}$ ($n_t^{\pm}$) must remain bounded as a function of u for each t. As a stability condition of $\lambda_t^{\pm}$, we obtain $(1-\kappa)\beta > \eta + \nu$. Appendix A.1 provides more detailed results and proofs.

3.2. Optimisation problems for market-makers

This section presents the market-maker's optimisation problem based on the proposed order book dynamic model. We consider a market-maker's cash process X_t that satisfies

$$\mathrm{d}X_t = (S_t + \delta_{t-}^+)\mathrm{d}N_t^+ - (S_t - \delta_{t-}^-)\mathrm{d}N_t^-,$$

which accounts for the cash increase when a sell limit order is lifted by a buy market order, and the cash decrease when a buy limit order is hit by a sell market order. Accordingly, a market-maker's inventory process q_t is given as

$$\mathrm{d}q_t = \mathrm{d}N_t^- - \mathrm{d}N_t^+.$$

A market-maker seeks the strategy $(\delta_t^-, \delta_t^+)_{0 \leq t \leq T}$ that maximises the cash value at the terminal date T. At time T, the market-maker liquidates the terminal inventory q_T using market orders at a price lower than the mid-price to account for liquidity costs as well as

the market orders walking into the limit order book. The performance of the market-maker achieved during $[t, T]$ is given by

$$\Phi_T = X_T + q_T(S_T - \phi q_T) - \psi \int_t^T q_u^2 du, \tag{5}$$

where $\phi \geq 0$ is the cost attributed to liquidity as well as the impact of the market order walking into the limit order book, and $\psi \geq 0$ is the running inventory penalty parameter.

The value function of the market-maker is given by

$$V(t, x, s, q, \boldsymbol{\lambda}, \boldsymbol{c}) = \max_{(\delta_u^+, \delta_u^-)_{t \leq u \leq T}} \mathbb{E}\left(\Phi_T \middle| X_t = x, S_t = s, q_t = q, \boldsymbol{\lambda}_t = \boldsymbol{\lambda}, \boldsymbol{c}_t = \boldsymbol{c} \right), \tag{6}$$

where ϕ, $\psi \geq 0$, and the initial states of the cash amount x, stock price s, inventory amount q, and intensity levels $\boldsymbol{\lambda} = (\lambda^-, \lambda^+)$, $\boldsymbol{c} = (c^-, c^+)$ are given.

The control problem in Eq.(6) can be employed to use the dynamic programming principle to show that the function V solves the following HJB equation

$$\begin{aligned}
&\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 s^2 \frac{\partial^2 V}{\partial s^2} + \beta(\theta^- - \lambda^- + \kappa\lambda^+)\frac{\partial V}{\partial \lambda^-} + \beta(\theta^+ - \lambda^+ + \kappa\lambda^-)\frac{\partial V}{\partial \lambda^+} \\
&+ \xi(\alpha - c^- + \kappa_c c^+)\frac{\partial V}{\partial c^-} + \xi(\alpha - c^+ + \kappa_c c^-)\frac{\partial V}{\partial c^+} \\
&+ \lambda^- \max_{\delta^- \geq 0}\{\mathrm{e}^{-\delta^- c^-}(\Delta^-_{x,q,\lambda,c} V - V) + (1 - \mathrm{e}^{-\delta^- c^-})(\Delta^-_{\lambda,c} V - V)\} \\
&+ \lambda^+ \max_{\delta^+ \geq 0}\{\mathrm{e}^{-\delta^+ c^+}(\Delta^+_{x,q,\lambda,c} V - V) + (1 - \mathrm{e}^{-\delta^+ c^+})(\Delta^+_{\lambda,c} V - V)\} - \psi q^2 = 0
\end{aligned} \tag{7}$$

with the terminal condition

$$V(T, x, s, q, \boldsymbol{\lambda}, \boldsymbol{c}) = x + q(s - \phi q), \tag{8}$$

where the shift operators $\Delta^{\pm}_{x,q,\lambda,c}$ and $\Delta^{\pm}_{\lambda,c}$ are defined as

$$\begin{aligned}
\Delta^-_{x,q,\lambda,c} V(t, x, s, q, \boldsymbol{\lambda}, \boldsymbol{c}) &= V(t, x - s + \delta^-, s, q + 1, \boldsymbol{\lambda} + (\eta, \nu), \boldsymbol{c} + (\eta_c, \nu_c)), \\
\Delta^+_{x,q,\lambda,c} V(t, x, s, q, \boldsymbol{\lambda}, \boldsymbol{c}) &= V(t, x + s + \delta^+, s, q - 1, \boldsymbol{\lambda} + (\nu, \eta), \boldsymbol{c} + (\nu_c, \eta_c)), \\
\Delta^-_{\lambda,c} V(t, x, s, q, \boldsymbol{\lambda}, \boldsymbol{c}) &= V(t, x, s, q, \boldsymbol{\lambda} + (\eta, \nu), \boldsymbol{c} + (\eta_c, \nu_c)), \\
\Delta^+_{\lambda,c} V(t, x, s, q, \boldsymbol{\lambda}, \boldsymbol{c}) &= V(t, x, s, q, \boldsymbol{\lambda} + (\nu, \eta), \boldsymbol{c} + (\nu_c, \eta_c)),
\end{aligned}$$

with the given controls δ^- and δ^+.

To obtain the simpler form of the HJB, we consider the following ansatz solution. Since the initial cash and initial stock price do not affect the market-maker's strategy, we propose

$$V(t, x, s, q, \boldsymbol{\lambda}, \boldsymbol{c}) = x + qs + g(t, q, \boldsymbol{\lambda}, \boldsymbol{c}), \tag{9}$$

which is inspired by observing the terminal condition Eq.(8) for V.

From the substitutions, Eq.(7) is simplified to

$$\begin{aligned}
&\frac{\partial g}{\partial t} + \beta(\theta^- - \lambda^- + \kappa\lambda^+)\frac{\partial g}{\partial \lambda^-} + \beta(\theta^+ - \lambda^+ + \kappa\lambda^-)\frac{\partial g}{\partial \lambda^+} \\
&\quad + \xi(\alpha - c^- + \kappa_c c^+)\frac{\partial g}{\partial c^-} + \xi(\alpha - c^+ + \kappa_c c^-)\frac{\partial g}{\partial c^+} \\
&\quad + \lambda^- \max_{\delta^- \geq 0}\{\mathrm{e}^{-\delta^- c^-}(\delta^- + \Delta^-_{q,\lambda,c} g - g) + (1 - \mathrm{e}^{-\delta^- c^-})(\Delta^-_{\lambda,c} g - g)\} \\
&\quad + \lambda^+ \max_{\delta^+ \geq 0}\{\mathrm{e}^{-\delta^+ c^+}(\delta^+ + \Delta^+_{q,\lambda,c} g - g) + (1 - \mathrm{e}^{-\delta^+ c^+})(\Delta^+_{\lambda,c} g - g)\} - \psi q^2 = 0,
\end{aligned} \tag{10}$$

where the shift operators $\Delta^{\pm}_{q,\lambda,c}$ and $\Delta^{\pm}_{\lambda,c}$ are defined as follows:

$$\begin{aligned}
\Delta^-_{q,\lambda,c} g(t,q,\boldsymbol{\lambda},\boldsymbol{c}) &= g(t,q+1,\boldsymbol{\lambda} + (\eta,\nu), \boldsymbol{c} + (\eta_c,\nu_c)), \\
\Delta^+_{q,\lambda,c} g(t,q,\boldsymbol{\lambda},\boldsymbol{c}) &= g(t,q-1,\boldsymbol{\lambda} + (\nu,\eta), \boldsymbol{c} + (\nu_c,\eta_c)), \\
\Delta^-_{\lambda,c} g(t,q,\boldsymbol{\lambda},\boldsymbol{c}) &= g(t,q,\boldsymbol{\lambda} + (\eta,\nu), \boldsymbol{c} + (\eta_c,\nu_c)), \\
\Delta^+_{\lambda,c} g(t,q,\boldsymbol{\lambda},\boldsymbol{c}) &= g(t,q,\boldsymbol{\lambda} + (\nu,\eta), \boldsymbol{c} + (\nu_c,\eta_c)).
\end{aligned}$$

Thus, we explicitly obtain the optimal controls

$$\begin{aligned}
(\delta^-_t)^* &= \frac{1 - c^-(\Delta^-_{q,\lambda,c} g - \Delta^-_{\lambda,c} g)}{c^-} \mathbb{1}_{\{c^-(\Delta^-_{q,\lambda,c} g - \Delta^-_{\lambda,c} g) < 1\}}, \\
(\delta^+_t)^* &= \frac{1 - c^+(\Delta^+_{q,\lambda,c} g - \Delta^+_{\lambda,c} g)}{c^+} \mathbb{1}_{\{c^+(\Delta^+_{q,\lambda,c} g - \Delta^+_{\lambda,c} g) < 1\}}.
\end{aligned} \tag{11}$$

By substituting the derived optimal controls into Eq.(10), we finally obtain

$$\begin{aligned}
&\frac{\partial g}{\partial t} + \beta(\theta^- - \lambda^- + \kappa\lambda^+)\frac{\partial g}{\partial \lambda^-} + \beta(\theta^+ - \lambda^+ + \kappa\lambda^-)\frac{\partial g}{\partial \lambda^+} \\
&\quad + \xi(\alpha - c^- + \kappa_c c^+)\frac{\partial g}{\partial c^-} + \xi(\alpha - c^+ + \kappa_c c^-)\frac{\partial g}{\partial c^+} \\
&\quad + \lambda^-(\Delta^-_{q,\lambda,c} g - g)\mathbb{1}_{\{c^-(\Delta^-_{q,\lambda,c} g - \Delta^-_{\lambda,c} g) \geq 1\}} + \lambda^+(\Delta^+_{q,\lambda,c} g - g)\mathbb{1}_{\{c^+(\Delta^+_{q,\lambda,c} g - \Delta^+_{\lambda,c} g) \geq 1\}} \\
&\quad + \lambda^- \left(\frac{\mathrm{e}^{c^-(\Delta^-_{q,\lambda,c} g - \Delta^-_{\lambda,c} g)}}{\mathrm{e} c^-} + \Delta^-_{\lambda,c} g - g \right) \mathbb{1}_{\{c^-(\Delta^-_{q,\lambda,c} g - \Delta^-_{\lambda,c} g) < 1\}} \\
&\quad + \lambda^+ \left(\frac{\mathrm{e}^{c^+(\Delta^+_{q,\lambda,c} g - \Delta^+_{\lambda,c} g)}}{\mathrm{e} c^+} + \Delta^+_{\lambda,c} g - g \right) \mathbb{1}_{\{c^+(\Delta^+_{q,\lambda,c} g - \Delta^+_{\lambda,c} g) < 1\}} - \psi q^2 = 0
\end{aligned} \tag{12}$$

with the terminal condition

$$g(T,q,\boldsymbol{\lambda},\boldsymbol{c}) = -\phi q^2.$$

4. Evidence of Synchrony in the Market

Empirical observations can detect the properties of the stochastic process of the order arrival intensities. Figure 1 presents the empirical intensity processes for the IBM stock's market buy and sell orders, which are inferred from the estimated model defined in Eq. (3) on a particular date (3 January, 2018) for the two five-minute time-frames (13:37 to 13:52 and 14:21 to 14:26). The unit time for the intensities is set to one second. Even at first glance, we see that the arrival patterns of buy and sell orders are highly interactive and synchronised. Similar patterns can be observed when the stock, date, and time are chosen randomly. An important feature of the paths is that each market order arrival is highly clustered with the spikes of arrivals tending to disappear quickly. In addition, it appears that neither the buy nor the sell-side order arrivals are independent with no one-sided order arrivals, and the bursts of activities tend to arrive almost contemporaneously.

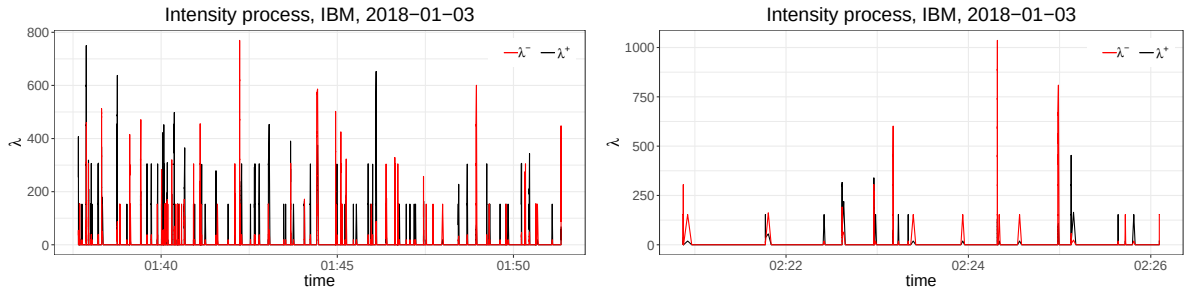

Figure 1: Empirical intensity processes λ^+, λ^- of IBM (NYSE ticker: IBM) for the buy and sell market orders traded from 13:37 to 13:52 (left) and from 14:21 to 14:26 (right) on 3 January 2018

To justify the proposed model, this section aims to capture the synchrony tendency using real market data. Owing to the difficulties in accessing limit order book data, the test is conducted based on transacted prices, which are assumed to be the market order executed by incoming limit orders. To investigate order arrival patterns, we calibrate all the model parameters using the maximum likelihood estimation (MLE) method and then examine the evidence, especially whether synchronising and exciting factors have changed considerably over time.

To perform MLE for the six parameters in the proposed model $(\beta, \theta^-, \theta^+, \kappa, \nu, \eta)$, we employ the Broyden-Fletcher-Goldfarb-Shanno algorithm (Fletcher, 1987), a kind of quasi-Newton method for numerical optimisation that facilitates the adequate estimation of Hawkes-type processes. The numerical procedure for MLE is implemented using an R optimisation function and its wrapper package maxLik[1]. The methodology involves max-

[1] `https://cran.r-project.org/web/packages/maxLik/index.html`

imising the log-likelihood function:

$$L(s) = \int_0^T \log \lambda_t^+ \mathrm{d}M_t^+ + \int_0^T \log \lambda_t^- \mathrm{d}M_t^- - \int_0^T (\lambda_t^+ + \lambda_t^-)\mathrm{d}t, \tag{13}$$

where s is the set of parameters to maximise (Ogata, 1978; Henningsen and Toomet, 2011). The market order arrivals $\lambda_t^{\pm}$ are given as closed forms in Eqs.(A.7) and (A.9) in Appendix A.2, which are used to implement Eq.(13). The MLE algorithm is processed until a tolerance of 10^{-6} is achieved. More details of the calibration procedure of Hawkes-type models appear in Da Fonseca and Zaatour (2014); Bacry and Muzy (2014); Bacry et al. (2016); Lee and Seo (2017); Lu and Abergel (2018).

We select six representative stocks among US large cap equities, Microsoft, Amazon, Apple, JP Morgan, Chevron, and IBM. Our data consist of trades and quotes for NYSE-listed (JP Morgan, Chevron, and Apple) and NASDAQ-listed (Microsoft, Amazon, and IBM) equities[2]. The trade and quote data are recorded at micro- or nanosecond decimal precision. The estimations are conducted using the trade prices and their arrival times reported on the exchanges on which they are listed. To investigate how the dynamic patterns in market order arrivals have changed over the past decade, we collect data from 2008 to 2018. On each day, trade data between 10:30 and 15:30 are analysed, excluding trades in the beginning and ending times to avoid seasonal effects.

For data pre-processing, to classify the transacted prices for buy- and sell-initiated orders, we employ the Lee-Ready algorithm (Lee and Ready, 1991). Using the best bid and ask price flows from the quote data, the mid-price processes are derived and the trade type – whether the transaction is seller-initiated or buyer-initiated – is determined using the relative position of the trade price compared with the mid-price. Precisely, when the trade price is less (greater) than the mid-price, which means that the transaction occurs on the bid (ask) side, the market order is considered to be a sell (buy) market order. When the mid and trade prices are equal, the tick test is applied: if the traded price falls below (rises above) the most recent traded price, which is different from the current trade price, it is considered to be a sell (buy) market order. If it increases, the transaction is regarded as a buy market order.

For the two end-years of 2008 and 2018, Table 1 presents the descriptive statistics for the number of arrivals of market buy/sell orders in a day following the preprocessing procedure. For all the stocks, the buy/sell order arrival frequency was 200,326 in 2008 and 139,039 in 2018; that is, the absolute size of orders in 2008 was larger than that in 2018

[2]The trade and quote information of NYSE and NASDAQ stocks is obtained from the Consolidated Tape Association (`https://www.ctaplan.com/index`) and Unlisted Trading Privileges (`http://www.utpplan.com/`), respectively.

Table 1: Summary of data statistics for the number of market buy and sell order arrivals in a day.

		2008				2018			
		Mean	Std	Max	Min	Mean	Std	Max	Min
Buy	Microsoft	33,445	19,173	136,025	4,536	20,035	9,767	61,549	6,531
	Amazon	9,563	4,192	31,329	1,500	13,788	7,173	38,780	4,071
	Apple	36,111	15,078	103,042	7,576	22,587	9,685	69,497	8,210
	JP Morgan	10,501	4,658	35,598	2,199	5,859	2,734	16,520	2,205
	Chevron	5,809	2,105	15,062	1,894	2,542	1,128	11,207	1,027
	IBM	4,466	1,618	12,367	1,289	2,295	1,294	8,729	762
Sell	Microsoft	33,702	18,763	137,227	5,346	21,159	10,690	68,887	6,558
	Amazon	9,440	4,005	28,434	1,317	14,777	8,082	50,443	4,478
	Apple	35,635	14,744	100,928	8,082	24,382	10,200	72,341	9,778
	JP Morgan	10,773	4,581	35,329	2,202	6,523	2,994	17,922	2,599
	Chevron	6,173	2,179	11,975	1,849	2,659	1,073	9,074	1,114
	IBM	4,708	1,622	10,976	1,137	2,429	1,424	11,654	957

(increase by 140%). One feature of order flows is that the difference between buy and sell orders[3] widened in 2018 compared with 2008: for all the stocks, the difference was 536 in 2008, whereas it was 4,823 in 2018.

We perform MLE based on daily data from 2 January to 31 December for each year (e.g. the estimation is conducted 252 times for the data in 2008 and 250 times in 2018). First, we investigate the result of the synchronising factor κ. Table 2 presents the average of the estimated κ, the number of significant estimates at the 5% significance level in the whole sample in each year, and the average of the stability condition $(1-\kappa)\beta/(\eta+\nu)$ derived in Lemma 1 with the estimated samples. The column 'Overall' contains the average of all the stocks for κ as well as the stability and sum of all the significantly estimated samples. We see that the synchronising level in 2018 jumped by about eight times than that in 2008 on average and that the number of significant estimations increased 5.1 times during the same period; the total number of significant samples was 152 in 2008, but 775 in 2018 for all the stocks, and the number of order arrivals every single day, on average, is specified in Table 1. Further, overall market stability increased marginally due to increase in stability of JP Morgan and Microsoft despite decrease in the other stocks. Note that the stability of the process can be highly guaranteed because all are greater than one[4].

[3]The difference = the number of sell orders - the number of buy orders

[4]Using this condition without the synchronising parameter, the literature has discussed the measure-

Table 2: Results of the mean of the estimated κ, the number of significant estimation for κ, and computing the stability condition $(1-\kappa)\beta/(\eta+\nu)$

	2008			2018		
	$\kappa(\%)$	# significance	stability	$\kappa(\%)$	# significance	stability
IBM	0.1343	8 (3.2%)	3.2342	7.0297	136 (54.2%)	2.9650
Chevron	0.3359	10 (3.9%)	3.3635	5.9676	163 (64.9%)	3.0852
Apple	0.8712	58 (23.1%)	1.8783	1.7943	146 (58.1%)	1.7770
Amazon	1.6632	76 (30.4%)	1.9250	0.2066	58 (23.1%)	1.6500
JP Morgan	0.0000	0 (0%)	2.6078	8.4821	209 (83.3%)	3.4369
Microsoft	0.0000	0 (0%)	1.1912	0.5315	63 (25.2%)	1.6503
Overall	0.5008	152 (10.1%)	2.3667	4.0020	775 (51.5%)	2.4274

Next, we discuss the calibration results of $\theta^+, \theta^-, \eta, \nu$, and β in 2008 and 2018 for the basic statistics – mean, standard deviation, skewness, and kurtosis (see Table B.9). The column 'Average' represents the average value of all the stock's estimations. All the estimations of the five parameters are statistically significant. The mean-reversion levels θ^+ and θ^- for the six stocks dropped in 2018 compared with in 2008 by 57% on average. The exciting levels η and ν declined by 26% and 33%, respectively, and the reverting speed β by 10% on average over the same period. For these parameters, however, there were differences in the increase and decrease for each stock. While Apple and Amazon showed clear increases in η, ν, and β, the others showed decreases in 2018 compared with in 2008. In addition, we can see that the self-exciting factor η has a remarkably higher value than the mutually-exciting factor ν across all the stock samples. We provide more detailed calibration results of different aspects in Appendix B.

Such a consistent tendency in the market in terms of hyper-dependency is presumably closely related to the dramatic increase in HFT activities over the past decade. Although the total number of arrivals in 2018 decreased by 31% compared with in 2008, the synchrony tendency in buy and sell orders increased eight times over the same period. These empirical results may support the argument that the arrival intensities for buy and sell orders dynamically incorporate the synchrony impact as well as the exciting tendency. Thus, we need a model with greater explanatory power to develop more profitable trading strategies.

ment of the degree of market instability and resiliency (e.g.,Danielsson et al., 2012; Filimonov and Sornette, 2012; Hardiman et al., 2013).

5. Neural Network Approximation for the HJB Equation

In this section, we propose an approximation method to find a solution to the high-dimensional HJB equation using DNNs. To do so, we adopt a mesh-free scheme called the deep Galerkin method proposed by Sirignano and Spiliopoulos (2018). Although mesh-based methods, such as the finite difference method are the most common way of solving PDEs, they are computationally infeasible for high-dimensional cases. However, the deep Galerkin method trains batches of randomly sampled time and space points through neural networks instead of forming a mesh, making it flexible and versatile, as required for high-dimensional PDE problems. The deep Galerkin method takes input values from the domain of the PDE and produces a candidate of the PDE solution by composing a smooth activation function repeatedly. The extent of composing represents the depth of the network and the composed function can solve the PDE numerically by tuning the DNN parameters. The training procedure aims to find the best DNN parameters that minimise a loss function, which indicates how close the DNN architecture is to satisfying the PDE's given conditions.

5.1. Neural network approximation

This section explains the procedure of the proposed DNN-based approximating method to solving the HJB equation derived in Section 3.2. We first consider a domain of the PDE such that $\mathcal{D} = (-\infty, \infty) \times [0, \infty)^4 \subset \mathbb{R}^5$ and $\mathcal{D}_T = [0, T) \times \mathcal{D}$. We assume that there exists a unique solution $u(t, q, \boldsymbol{\lambda}, \boldsymbol{c}) \in C^1(\mathcal{D}_T)$ to the following PDE:

$$\begin{aligned} &\frac{\partial u}{\partial t}(t, q, \boldsymbol{\lambda}, \boldsymbol{c}) + \mathcal{L}u(t, q, \boldsymbol{\lambda}, \boldsymbol{c}) = 0, \text{ for } (t, q, \boldsymbol{\lambda}, \boldsymbol{c}) \in \mathcal{D}_T \\ &u(T, q, \boldsymbol{\lambda}, \boldsymbol{c}) = -\phi q^2, \text{ for } (q, \boldsymbol{\lambda}, \boldsymbol{c}) \in \mathcal{D}, \end{aligned} \tag{14}$$

where $\mathcal{L}$ is a nonlinear PDE operator with a one time variable and five state variables defined in Eq.(A.10). If the PDE (12) has a unique solution $g(t, q, \boldsymbol{\lambda}, \boldsymbol{c})$, then

$$u(t, q, \boldsymbol{\lambda}, \boldsymbol{c}) = g(t, q, \boldsymbol{\lambda}, \boldsymbol{c}), \quad \text{for } (t, q, \boldsymbol{\lambda}, \boldsymbol{c}) \in \mathcal{D}_T.$$

According to the deep Galerkin method, u can be approximated with a DNN structure defined by $f(\cdot, \Theta)$, where Θ is a set of the DNN parameters. Among the various architectures of f, we employ a fully-connected feedforward network, which is the most basic type of neural network but has outstanding performance in finding our PDE solution.

Now, we consider a reduced compact domain so that differential operators can be defined with respect to all the variables to use the universal approximation theorem proposed by Hornik (1991). Assume a compact set $\tilde{\mathcal{D}} = [-N_1, N_1] \times [n_2, N_2]^2 \times [n_3, N_3]^2 \subset \mathcal{D}$

for any large positive numbers N_1, N_2, N_3 and small positive numbers n_2, n_3, and let $\tilde{\mathcal{D}}_T = [0,T) \times \tilde{\mathcal{D}}$. In the reduced domain, we construct the following loss function

$$L(f) = \left\| \frac{\partial f}{\partial t}(t,q,\boldsymbol{\lambda},\boldsymbol{c};\Theta) + \mathcal{L}f(t,q,\boldsymbol{\lambda},\boldsymbol{c};\Theta) \right\|^2_{\tilde{\mathcal{D}}_T,\mu_1} + \left\| f(T,q,\boldsymbol{\lambda},\boldsymbol{c};\Theta) + \phi q^2 \right\|^2_{\tilde{\mathcal{D}},\mu_2}, \quad (15)$$

where μ_1 and μ_2 are probability measures of $\tilde{\mathcal{D}}_T$ and $\tilde{\mathcal{D}}$, respectively, which are absolutely continuous with respect to the Lebesgue measure[5]. The loss function represents how far f deviates from the original PDE's operator and terminal conditions. We verify that a feedforward neural network f exists that makes the loss function in Eq.(15) sufficiently small on the compact set $\tilde{\mathcal{D}}_T$ (see Theorem 3, Appendix C).

The goal is to find the DNN parameters Θ such that the error in Eq.(15) is minimised. To do so, we develop Algorithm 1 which states the DNN algorithm using the stochastic gradient descent method on a sequence of time and state points randomly sampled on $\tilde{\mathcal{D}}_T$ and $\tilde{\mathcal{D}}$ with respect to μ_1 and μ_2, respectively. In Step 2 the mean squared error $\tilde{L}(f;\Theta)$ is an unbiased estimator of the loss function $L(f)$. The algorithm proceeds in a descending direction, which means that the loss function decreases after an iteration and the next Θ can be a better parameter estimate than the previous one.

Algorithm 1 The DNN algorithm for solving the PDE (14).

1: Generate random samples $\{(t_i, q_i, \boldsymbol{\lambda}_i, \boldsymbol{c}_i)\}_{i=1}^m$ from the probability measure μ_1 in $\tilde{\mathcal{D}}_T$ and $\{(\tilde{q}_i, \tilde{\boldsymbol{\lambda}}_i, \tilde{\boldsymbol{c}}_i)\}_{i=1}^m$ from the probability measure μ_2 in $\tilde{\mathcal{D}}$ with the batch size m.

2: Calculate the mean squared error $\tilde{L}(f;\Theta)$ as

$$\begin{aligned} \tilde{L}(f;\Theta) = \frac{1}{m}\sum_{i=1}^{m} \Bigg\{ &\left(\frac{\partial f}{\partial t}(t_i,q_i,\boldsymbol{\lambda}_i,\boldsymbol{c}_i;\Theta) + \mathcal{L}f(t_i,q_i,\boldsymbol{\lambda}_i,\boldsymbol{c}_i;\Theta) \right)^2 \\ &+ \left(f(T,\tilde{q}_i,\tilde{\boldsymbol{\lambda}}_i,\tilde{\boldsymbol{c}}_i;\Theta) + \phi\tilde{q}_i^2 \right)^2 \Bigg\} \end{aligned} \quad (16)$$

3: Update the parameters Θ in the opposite direction of the gradient of $\tilde{L}(f;\Theta)$ with regard to the parameters:

$$\Theta \leftarrow \Theta - \ell \, \nabla_\Theta \tilde{L}(f;\Theta)$$

with the learning rate ℓ.

4: Repeat until the gradient approaches zero.

[5]The loss function could be defined using the Sobolev training (Czarnecki et al., 2017; Matthias and Diepold, 2019), which additional work may leave as a future study.

To implement Algorithm 1, we employ the hyperbolic tangent function as the activation function to enable f to become smooth and apply the differential operator to the neural network. We find the following DNN architecture to be effective[6]:

$$\begin{aligned} h_1 &= \tanh(W_0 \boldsymbol{x} + b_0) \\ h_{l+1} &= \tanh(W_l h_l + b_l), \qquad l = 1, \cdots, L \\ f(\boldsymbol{x}; \Theta) &= W_{L+1} h_{L+1} + b_{L+1}, \end{aligned} \tag{17}$$

with $h_l \in \mathbb{R}^n$ and L hidden layers having n hidden units in each hidden layer. The function $\tanh(\cdot)$ is taken to each component of h_l. This network takes the input $\boldsymbol{x} = [t \; q \; \boldsymbol{\lambda} \; \boldsymbol{c}]^T$ and maps this input to the value of the function $f(\boldsymbol{x}; \Theta)$ as a candidate for the solution of the PDE. The parameter Θ consists of weight matrices $W_0 \in \mathbb{R}^{n \times 6}$, $W_l \in \mathbb{R}^{n \times n}$, $W_{L+1} \in \mathbb{R}^{1 \times n}$ and bias vectors $b_0 \in \mathbb{R}^n$, $b_l \in \mathbb{R}^n$, $b_{L+1} \in \mathbb{R}^1$.

For the key hyperparameters L and n in the DNN, we find that three-hidden-layers ($L = 3$) and 900 hidden nodes ($n = 900$) perform well for the given market model parameters. We use a batch size of 25,000 ($m = 25,000$) drawn from the uniform probability measures μ_1 and μ_2 on $\tilde{\mathcal{D}}$[7]. The initial learning rate ℓ is chosen as 10^{-4} and this is reduced by a factor of 10 when the loss function stops decreasing. DNN testing is conducted using TensorFlow in Python, and the DNN parameters are updated using the Adam optimisation algorithm, which is a momentum-based stochastic optimisation method (Kingma and Ba, 2014). The DNN parameters are initialised using the Xavier initialisation. Our computations are performed on a Windows 10 PC with an AMD Ryzen Threadripper 1950X CPU, a 3.40 GHz 16-core processor, 64 GB RAM, and an NVIDIA Titan V GPU. Under the network structure and this computing environment, this DNN-training takes 35.83 seconds per 100 iterations. In total, iterations are run 170,000 times to obtain the target loss level, which is the level that no longer reduces even with more iterations.

The estimation domain $\tilde{\mathcal{D}}_T$ is adopted to $N_1 = 100$, $N_2 = 250$, $N_3 = 150$, $n_2 = 0.5$, and $n_3 = 0.3$, and $T = 10800$ to generate random samples as inputs[8]. To ensure that the five state variables stay in the reduced domain $\tilde{\mathcal{D}}_T$, we confirm that millions of simulated paths $q_t, \lambda_t^{\pm}$, and $c_t^{\pm}$ are contained in $\tilde{\mathcal{D}}_T$. Furthermore, we choose the parameters of the

[6] From the tests of various types of architectures including the classic/modified LSTM, monotonic non-linear transformation, drop-out layers, skip connections models and different activation functions, we chose the best practice.

[7] Noting that a lower sample size produces a big oscillation in the error convergence with large confidence intervals, this batch size is adopted to ensure more highly convincing estimation for the function values over the whole domain, less relying on random sampling of the variables.

[8] The time T represents 10800 seconds because all the parameters in our model are measured at the unit level of seconds, which means that the trading lasts three hours.

market model as follows: $\beta = 100$, $\theta^- = 0.5$, $\theta^+ = 0.4$, $\kappa = 0.2$, $\eta = 20$, $\nu = 15$ (for the intensity processes $\boldsymbol{\lambda}$ in Eq.(3)), $\xi = 40$, $\alpha = 0.3$, $\kappa_c = 0.1$, $\eta_c = 8$, $\nu_c = 5$ (for the depth processes $\boldsymbol{c}$ in Eq.(4)), $\phi = 0.1$ and $\psi = 0.01$ (penalty on the inventory in Eq.(5)).

Figure 2 presents a contour plot of the optimal quote distances δ^- and δ^+ with regard to the inventory q and time t. The intensities for the market order and limit order depth processes are taken as $\lambda^- = 3.5$, $\lambda^+ = 3$ and $c^- = 1.5$, $c^+ = 2$, respectively. The level of optimal distances from the mid-price is outstandingly prominent when the market-maker retains a large amount of inventory that needs to be liquidated or acquired, and little time is left to trade the stocks. As the inventory q becomes positively larger, the market-maker's strategy is more reluctant to buy stocks over time, and such a strategy is likely to place the limit buy order δ^-, which is far from the mid-price, to reduce the chance of executions.

As q becomes negatively larger (i.e. the short position increases), the strategy tends to increase the stocks by placing the limit buy order almost at the mid-price to enhance the chance of executions. In a similar manner, as q becomes negatively larger, the strategy becomes more hesitant to sell stocks by placing the limit sell order δ^+, which is far from the mid-price. As q becomes positively larger, the strategy tends to reduce the stock position by placing the limit sell order almost at the mid-price.

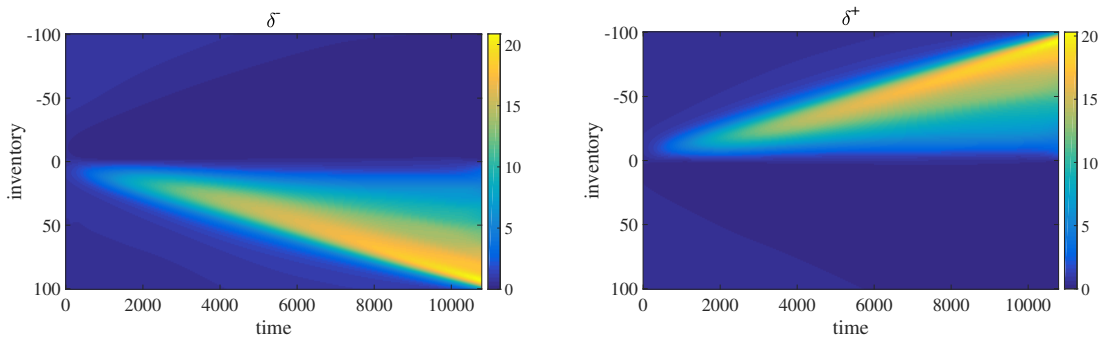

Figure 2: Contour plots of the optimal quotes δ^- for limit buy (left) and δ^+ for limit sell (right) orders depending on the change in inventory q when $\lambda^- = 3.5$, $\lambda^+ = 3$, $c^- = 1.5$, $c^+ = 2$.

5.2. Neural network training loss

This section investigates the convergence of the proposed DNN algorithm in Section 5.1 by reporting the training loss computed as the mean-squared loss $\tilde{L}(f)$ in Eq.(16). The training is run until the loss no longer decreases. In this case, following the training is conducted 170,000 iterations, the loss eventually reduces to 189.7, after initially starting at 180,244. Figure 3 illustrates the training loss against the number of iterations on the estimation domain. The shaded area represents a 95% confidence interval for computing

the expectation of the mean-squared loss under a batch size of 25,000 samples, which represent that the loss is likely to stay in this interval with 95% probability.

When applying the estimated optimal strategies to the generated scenarios, the absolute inventory size varies between ± 7 in the selected parameter set, as discussed in Section 6. To check the accuracy of the ultimately employed values of the estimation f to obtain the optimal strategy, it is worthy to note the training loss in the local domains of $|q|$. To examine the training loss over the more reduced domain for q, additional trainings are conducted by choosing $N_1 = 10$ and 50 while the other conditions remain the same. Figure 4 illustrates the training loss against the number of iteration steps in theses domains of inventory q with $|q| \leq 50$ (left) and $|q| \leq 10$ (right). The case of $|q| \leq 50$ has the ultimate training loss of 47.70 when the number of iterations approaches 80,000, initially starting from 9105.9. Figure 4 (left) displays the training loss after running 20,000 iterations. The case of $|q| \leq 10$ has a loss of 10.62 with 30,000 iterations, starting from 123.7. Hence, the local loss for which the estimations are mainly employed to compute the ultimate values is relatively small.

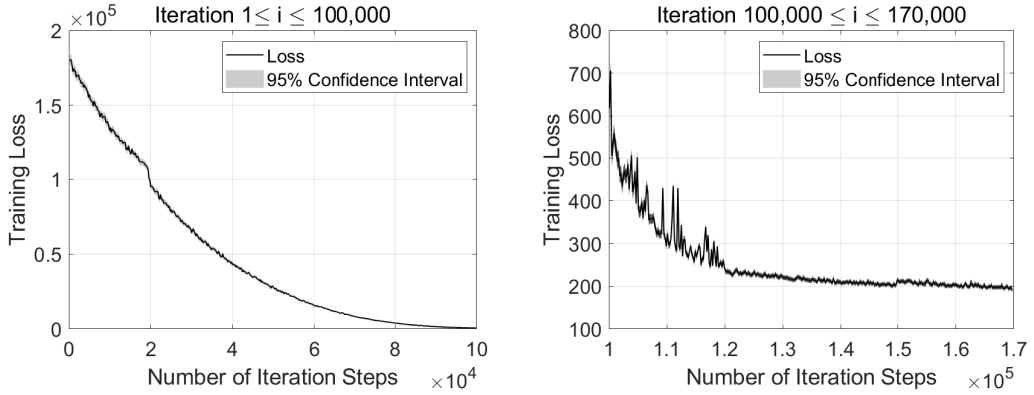

Figure 3: DNN training loss against the number of iteration steps until 170,000 – the loss with iteration from 1 to 100,000 (left) and the loss with iterations from 100,000 to 170,000 (right). The shaded area represents the 95% confidence interval in the computing loss with the batch size of samples.

Next, we present the DNN-approximated solution f for $0 \leq t \leq T$ with respect to the inventory size q with the given ϕ. The other variables $\lambda^{\pm}$ and $c^{\pm}$ are randomly sampled with 100,000 points for each and taken on average for this visualisation. In Figure 5, the left panel illustrates the contour lines of the estimated solution of g over the whole period, and the right panel shows only the part of the estimated value of g under the differential operator $\frac{\partial}{\partial t} + \mathcal{L}$, which has a value of zero in the true solution. In Figure 6, the left panel indicates the true solution g (red line) and the DNN-estimated solution (blue line) at the end time T over the domain q, whereas the right panel shows its error computed as the

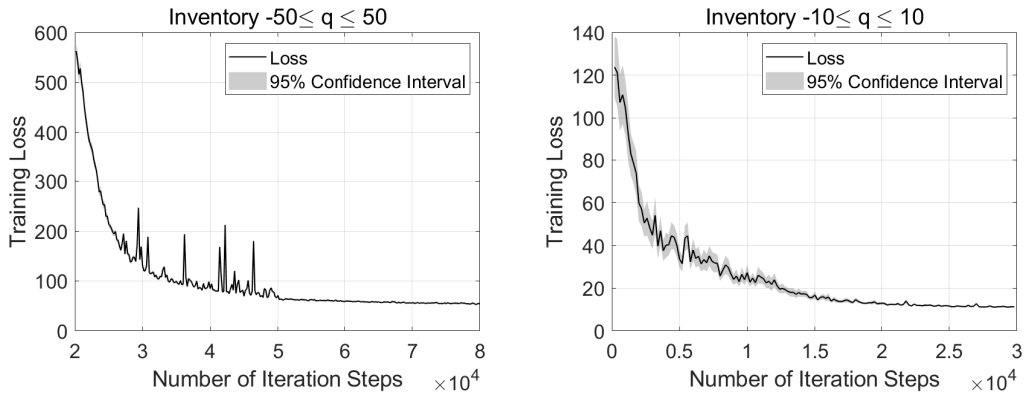

Figure 4: DNN training loss against the number of iteration steps in the estimation domains of $|q| \leq 50$ (left) and $|q| \leq 10$ (right). The shaded area represents the 95% confidence interval in the computing loss with the batch size of samples.

absolute difference between the two curves that appear in the right panel. We can see that the estimation error measured at the terminal boundary rises, as the inventory size either increases or decreases. However, given that the error in the case of $|q| = 100$ is about 14, whereas the true values are -1000, we may judge that the absolute error level may be acceptable. Moreover, to observe more distinct difference in the estimated values, we illustrate f by zooming in on the partial domains of $|q| \leq 50$ and $|q| \leq 10$ from the original graphs. Figure 7 presents a closer look of Figure 5 over the particular domains of q with the corresponding heat maps (i.e. $|q| \leq 50$ (top) and $|q| \leq 10$ (bottom)).

6. Simulation

In this section, we conduct numerical simulations for the market-maker's wealth under the optimal and other various deterministic strategies by exploiting the DNN estimation discussed in Section 5. To do so, we compute the distribution of the terminal profit and loss (PnL) of a market-maker as follows:

$$\text{PnL} = X_T + q_T(S_T - \phi q_T)$$

with the liquidation cost ϕ under the posting controls $\delta^{\pm}$, by generating the required sample paths.

A sample path of the terminal PnL can be generated by the following procedure. First, based on the thinning algorithm (Ogata, 1978), the sample paths of the intensity $\lambda_t^{\pm}$ and market order processes $M_t^{\pm}$ are obtained. Since our model contains additional synchrony factors compared with a typical exciting Hawkes model, we modify the original sampling

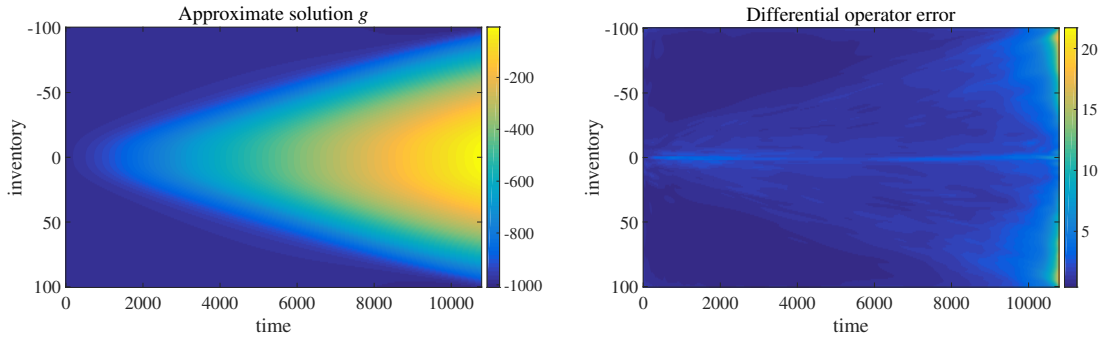

Figure 5: The DNN-approximate solution f for g (left) and the DNN-approximated value for the differential operator $\frac{\partial}{\partial t}g + \mathcal{L}g$ (right). Each value is obtained by taking the average of all the values (g and $\frac{\partial}{\partial t}g + \mathcal{L}g$) computed with 100,000 random samples of $\lambda^{\pm}$ and $c^{\pm}$.

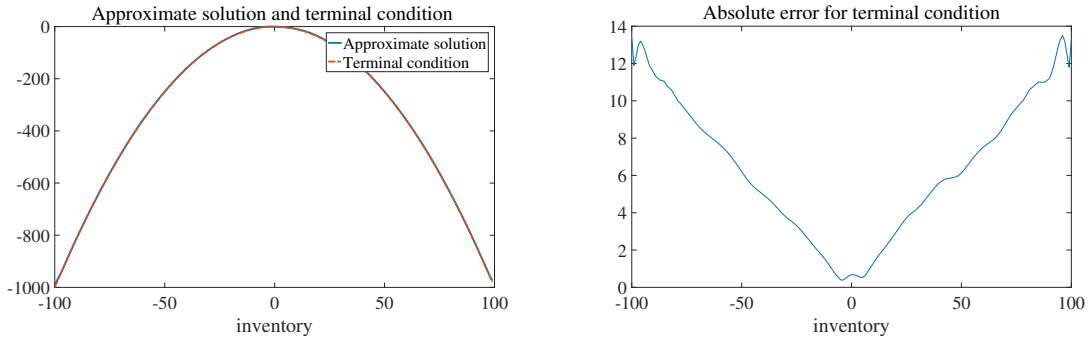

Figure 6: The terminal condition of g (blue line) and its DNN-approximated solution f (red-dot) by taking the average of the values with respect to $\lambda^{\pm}$ and $c^{\pm}$ with 100,000 random sampling for each (left) and its absolute difference (right).

algorithm that specified in Algorithm 2 (presented in Appendix A.4). Accordingly, the sample paths of the depth of the limit order book $c_t^{\pm}$ with $M_t^{\pm}$ are generated. Next, a sample of $N_t^{\pm}$ is filtered from $M_t^{\pm}$ using a Bernoulli variable generator under the fill probability $h(\delta_t^{\pm}, c_t^{\pm}) = \mathrm{e}^{-\delta_t^{\pm} c_t^{\pm}}$ with the estimated (or given) $\delta_t^{\pm}$. Lastly, with the generated stock price path S_t, X_t and q_t are computed. This procedure runs until the accumulated interarrival time of either N_t^+ or N_t^- reaches T. Initial wealth starts at zero and we obtain K samples of the terminal PnL by repeating the procedure. We generate K paths of the stock price S_t, market orders $M_t^{\pm}$, and depth of limit order book $c_t^{\pm}$. We then obtain the terminal PnL by activating market-making trades with the limit orders for which the execution is determined by the fill probability of the arrival of the market orders.

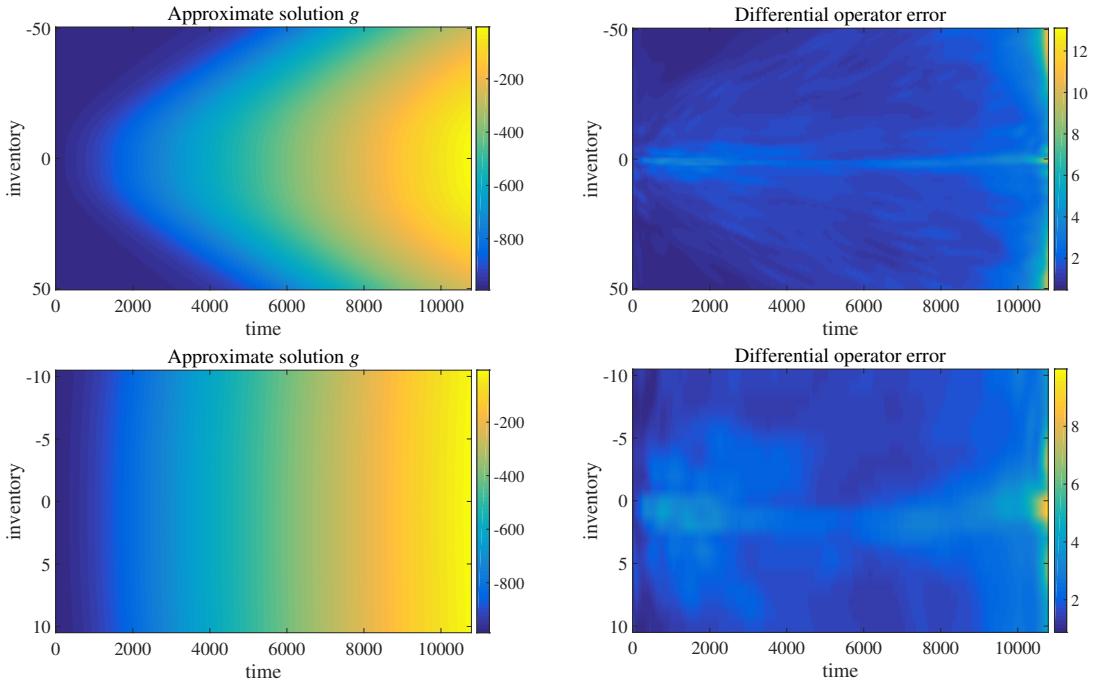

Figure 7: A closer look at the DNN-approximated solution f for g (left) and the DNN-approximated value for the differential operator $\frac{\partial}{\partial t}g + \mathcal{L}g$ (right) shown in Figure 5 on the partial domains of $|q| \leq 50$ (top) and $|q| \leq 10$.

6.1. Optimal or simple strategy?

This section considers several practically feasible market-making strategies and implements the corresponding market-making PnLs by comparing it with the optimal result.

First, we compute the performance of the optimal strategy and compare it with a rough constant strategy and a more elaborate strategy. The detailed strategies are given as follows: (i) the controls are constant (i.e. $\delta_t^- = \delta_t^+ = d_0$); (ii) the controls are chosen asymmetrically and are linearly dependent on inventory q_t, that is,

$$\delta_t^-(q_t) = \begin{cases} 2q_t + \frac{1}{2} & \text{if } q_t \geq 0 \\ -\frac{1}{2}q_t + \frac{1}{2} & \text{if } q_t < 0 \end{cases}, \quad \delta_t^+(q_t) = \begin{cases} \frac{1}{2}q_t + \frac{1}{2} & \text{if } q_t \geq 0 \\ -2q_t + \frac{1}{2} & \text{if } q_t < 0 \end{cases}; \tag{18}$$

and (iii) the optimal controls δ^* in Eq.(11) are taken.

Strategy (ii) is inspired by the estimation results of the optimal δ^* that has an asymmetric pattern depending on the current inventory size. Although the pattern is nonlinear, we consider the strategy with a linear but asymmetric relation with respect to the amount of inventory, which has zero computational costs. The reason for choosing a factor of two in this strategy is that it performs better than all other alternatives.

For this simulation, we adopt the market parameters based on the calibration results from the market data shown in Section 4, which are also applied to the numerical tests in Section 5. The parameters are chosen as follows: trading time $T = 10,800$; liquidation cost $\phi = 0.1$; volatility of the stock mid-price $\sigma = 0.0001$[9]; market order intensities are determined under $\beta = 100, \theta^- = 0.5, \theta^+ = 0.4, \kappa = 0.2, \eta = 20, \nu = 15$; and limit order intensities are determined under $\xi = 40, \alpha = 0.3, \kappa_c = 0.1, \eta_c = 8, \nu_c = 5$. With this selection, we generate 50,000 samples of the terminal PnL.

Figures 8 and 9 illustrate the histograms of the terminal wealth PnL achieved by the three strategies of posting $\delta^{\pm}$. Table 4 presents the terminal wealth PnL achieved under the three strategies of posting $\delta^{\pm}$ with the descriptive statistics (i.e. mean, standard deviation, minimum, and maximum) of the constant, asymmetric, and optimal posting strategies. For the expected terminal PnL, posting at the optimal distance performs better than the other cases. In addition, the optimal posting case significantly outperforms the others in all the samples, considering the minimum PnL of the optimal strategy and maximum PnLs of all the other constant cases. The PnL's diversity is considerably smaller than that of the others, which implies that the optimal strategy can achieve the highest expected profit with a comparatively lower risk than the others.

Table 3: Descriptive statistics of the terminal PnLs using three strategies (constant, asymmetric, and optimal) when trading lasts 10800 seconds.

Market-making strategy		Mean	SD	Min	Max
	$d_0 = 1$	-26,701	10,472	-87,524	1,937
	$d_0 = 2$	-6,242	6,174	-46,834	9,746
	$d_0 = 3$	2,969	3,667	-17,546	11,858
(i) Constant	$d_0 = 4$	6,629	2,226	-6,563	11,648
	$d_0 = 5$	7,578	1,346	-861	10,509
	$d_0 = 6$	7,253	832	1,452	9,230
	$d_0 = 7$	6,411	532	2,824	7,786
(ii) Asymmetric $\delta^{\pm}$ in Eq.(18)		8,902	**146**	8,351	9,582
(iii) Optimal δ^*		**12,333**	183	**11,583**	**13,349**

We next show the simulated path of the cash process X_t and inventory process q_t over time. Figure 10 illustrates the sample paths of X_t and q_t when using the optimal δ_t^* strategy (left) and a constant strategy with $d_0 = 5$ (right). Posting an optimal distance tends to

[9]This volatility is measured to 'second', which is transformed to the daily volatility of 1.47% ($= 0.0001 \times \sqrt{3600 \times 6}$ assuming the trading time in a day is 6 hours long.

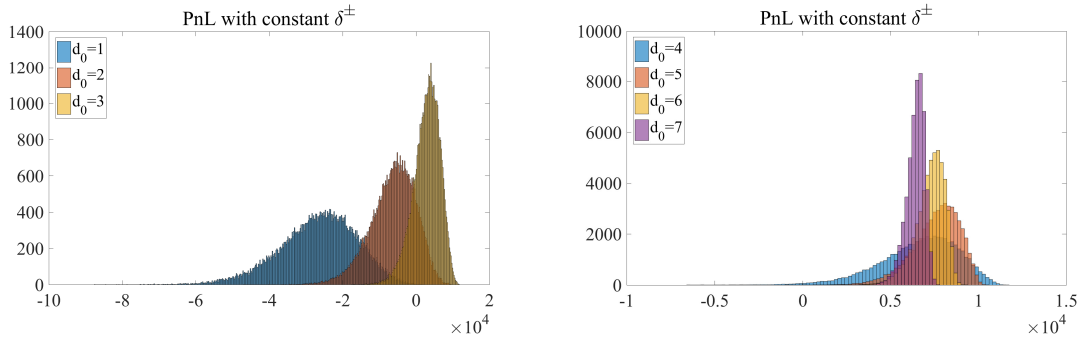

Figure 8: Distribution of the terminal PnL for a market-maker with constant $d_0 = 1, 2, 3$ (left) and $d_0 = 4, 5, 6, 7$ (right) when trading lasts 10,800 seconds.

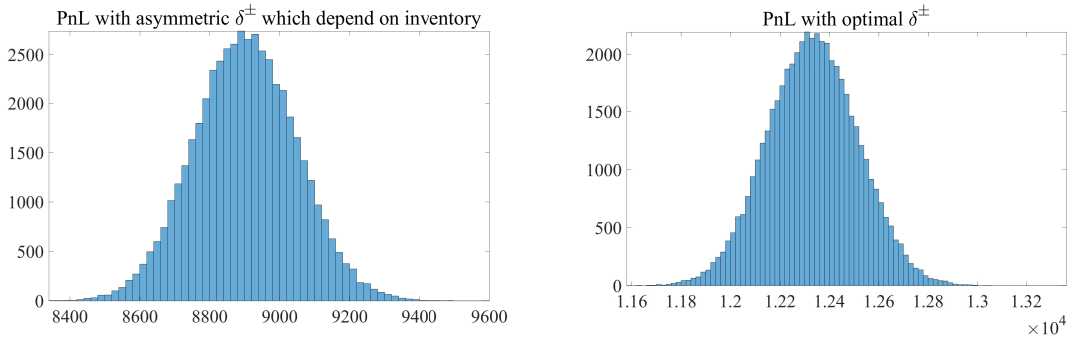

Figure 9: Distribution of the terminal PnL for a market-maker with asymmetric δ (left) and with optimal δ^* (right) which depend on inventory when trading lasts 10,800 seconds.

increase profit gradually over time, while posting a naive constant can result in a cumulative loss. In terms of inventory management, the inventory process under an optimal execution tends to move within a small range between ± 7 points at the most, whereas the inventory process under the constant execution increases over time and reaches over 150 points on this path.

Such inventory management within the tight bounds of optimal posting can be understood as a consequence of the objective function that aims to maximise final wealth by penalising large exposure to inventories. In practical terms, the capital amount available to market-makers is strictly constrained by their internal rules, because firms or regulators penalise activities that take large exposures. Another reason for the need for tight inventory management is that the piling (long or short) positions during market-making should be liquidated at the end of the trading day by paying transaction costs. Given that strict control of exposure to inventories is required for market-making, market-makers might prefer to adopt a high penalising rate ψ when their risk appetite reduces.

Second, we compute the difference in the market-maker's PnLs between the optimal and

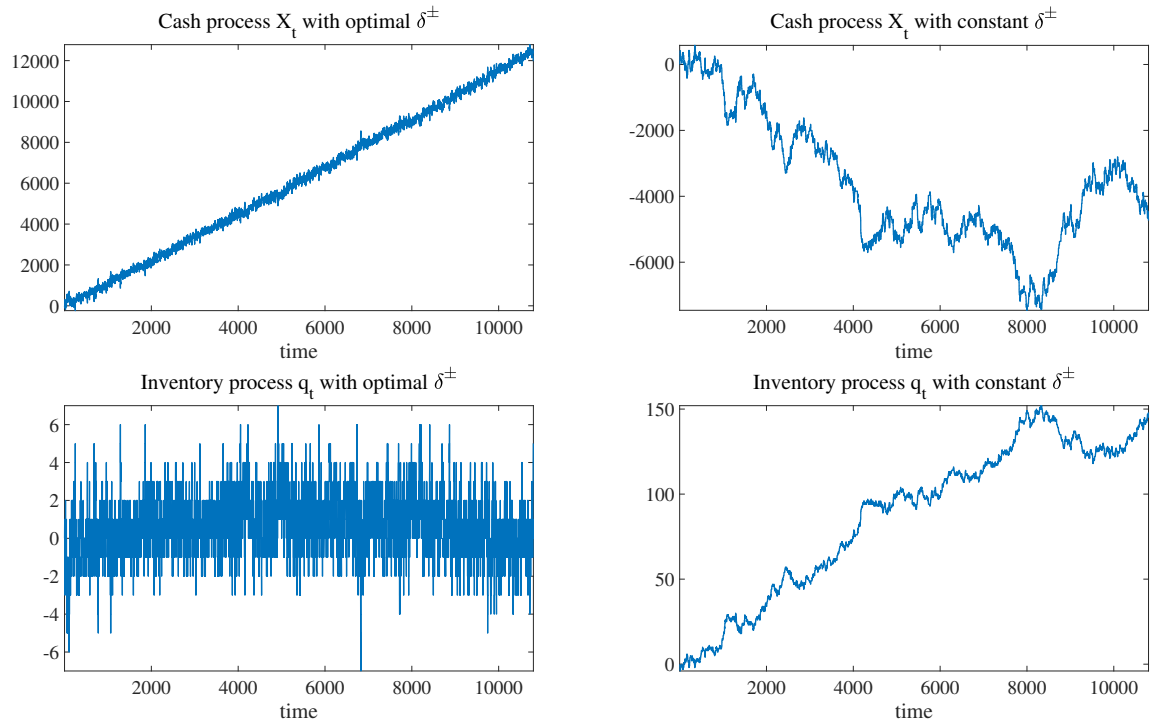

Figure 10: Sample paths of the cash process X_t (upper) and inventory process q_t (lower) for the optimal posting (left) and a constant posting (right) for 10,800 seconds trading. The values of terminal wealth PnL are 12,085 (optimal $\delta^{\pm *}$; upper left) and 7,477 (constant $d_0 = 5$; lower right).

suboptimal strategies. As a suboptimal strategy, we consider the optimal posting strategy under the no (or half) synchrony model that applies to scenarios with the synchronising factor for the PnL simulation. We investigate how much lower the performance of this suboptimal strategy than the fully optimal strategy in the market-maker's PnLs.

The optimal strategies are estimated with the same DNN architecture proposed in Eq.(16). This simulation chooses the following parameters[10]: $T = 1800$; for market orders $\beta = 50, \theta^- = 0.5, \theta^+ = 0.4, \eta = 4, \nu = 0.1$; and for limit orders $\xi = 40, \alpha = 0.3, \eta_c = 3, \nu_c = 0.1$. In addition, to investigate the effect of κ differentiated from κ_c on the market maker's PnL, we consider the following cases: (i) full synchrony with $\kappa = \kappa_c = 0.9$; (ii) half synchrony with $\kappa = 0$ and $\kappa_c = 0.9$; and (iii) no synchrony, that is, $\kappa = \kappa_c = 0$. Then, we generate 30,000 synchronising samples and apply the strategies estimated in each case.

Table 4 presents the terminal PnL's statistics under the optimal strategies with the full,

[10]To examine the impact of synchrony in the simulation, we decrease the mutually-exciting factor and increase the synchronising factor compared with the previous simulations.

half, and no synchrony models when the synchronising tendency exists in the market. The expected value of the terminal PnL when market-makers ignore both κ and κ_c (or, only κ) becomes significantly smaller than when they do not. On average, the market-maker's PnL of case (i) outperforms those of cases (ii) and (iii) by 9.2% and 42.2%, respectively. Indeed, the worst scenario of case (i) is better than the best of case (iii). If the synchrony effect decreases (or the mutually-exciting effect increases), the performance differences between the two cases can reduce.

Table 4: Descriptive statistics of the terminal PnLs using the optimal strategies of full, half, and no synchrony under the synchronising scenarios for market and limit orders with the 1800-second trading time.

Order Arrival Model	Mean	SD	Min	Max
(i) Full synchrony $\kappa = \kappa_c = 0.9$	801.9	9.2	768.3	839.5
(ii) Half synchrony $\kappa = 0$ and $\kappa_c = 0.9$	734.5	9.8	699.1	772.9
(iii) No synchrony $\kappa = \kappa_c = 0$	564.0	7.2	537.6	597.1

6.2. How good is the full synchrony model to choose?

In this section, we compare the terminal PnLs when posting the optimal strategy and posting a strategy without the synchronising factor when the synchrony tendency between buy and sell orders exists. In addition, we observe how the instability of the high-frequency market could affect the market-maker's PnLs. These tests aim to investigate how the posting strategy under various order book models and order arrival markets may produce differences in the market-maker's PnLs.

In the first part of simulations, we demonstrate the following models of λ_t and c_t: the synchrony model with $\kappa \neq 0$ and $\kappa_c \neq 0$, the no synchrony model with $\kappa = \kappa_c = 0$. To do so we calibrate each model from the high-frequency market order arrival data in the year of 2018. Table 5 exhibits the MLE results of all the required parameters in the models with/without the synchronising factor for the six stocks in 2018. Each value is computed by taking the average on all the daily estimations. This result shows that the no synchrony model leads to sharp increases in not only the self- and mutually-exciting effects but also the mean-reverting speed compared with the synchrony model. On the contrary, there is little difference in the constant mean-reverting speeds of $\theta^{\pm}$ between both models. From a stability perspective, the long-term moving patterns under the two models do not change considerably.

To investigate the respective contributions of κ and κ_c, we take the half synchrony model with $\kappa = 0$ and $\kappa \neq 0$, adding to the two models earlier mentioned. For this simulation, the λ_t parameters are adopted based on the scales of the true estimations (e.g.

Table 5: MLE results of the market order arrival model with the synchronising factor (top) and without the synchronising factor (bottom). Each estimated parameter is reported by taking the mean of all the daily results in 2018

Model		IBM	Chevron	Apple	Amazon	JP Morgan	Microsoft
Synchrony	θ^+	0.071	0.080	0.477	0.257	0.189	0.375
	θ^-	0.075	0.084	0.494	0.246	0.199	0.383
	η	224	215	3447	3089	156	4849
	ν	27.4	29.7	435	181	30.1	723
	β	749	767	7030	5349	684	9146
	$\kappa(\%)$	7.03	6.00	1.79	0.21	8.48	0.53
No Synchrony	θ^+	0.073	0.087	0.488	0.259	0.191	0.375
	θ^-	0.078	0.092	0.506	0.248	0.201	0.383
	η	828	938	3181	3169	1377	3944
	ν	140	156	445	186	188	615
	β	1916	2240	6455	5281	3055	7532

IBM, Chevron, and JP Morgan in Table 5) for η, ν, β, and κ, while the θ's are chosen as the average size of the samples because of high diversity in their values. Table 6 describes the considered models and their selected values. For the parameters of c_t, we consider the scales of the parameters estimated in λ_t and those employed in the numerical examples of Cartea et al. (2014). Under each parameter assumption in Table 6, we generate sample paths of λ^- and λ^+, and Figure 11 illustrates the generated sample paths of λ under each market order model. Note that, comparing the measure of the true intensity values of IBM stock in Figure 1, we can observe that the synchrony model generates λ's samples that have a more feasible scale than the no synchrony model does.

Table 6: The chosen order arrival model parameters for λ and c for testing the market-maker's PnL.

Order Arrival Model	Parameters for λ					Parameters for c				
	$\theta^{\pm}$	η	ν	β	κ	α	η_c	ν_c	ξ	κ_c
(i) Full Synchrony	0.3	200	20	700	0.1	0.3	100	20	400	0.1
(ii) Half Synchrony	0.3	800	130	2000	-	0.3	100	20	400	0.1
(iii) No Synchrony	0.3	800	130	2000	-	0.3	400	130	1200	-

Next, we find the optimal posting δ^* using the same DNN architecture proposed in Eq.(17) with a 300-second trading time for each case. Once solving the optimal control in each case separately, we repeat the PnL simulations obtained by the optimal controls

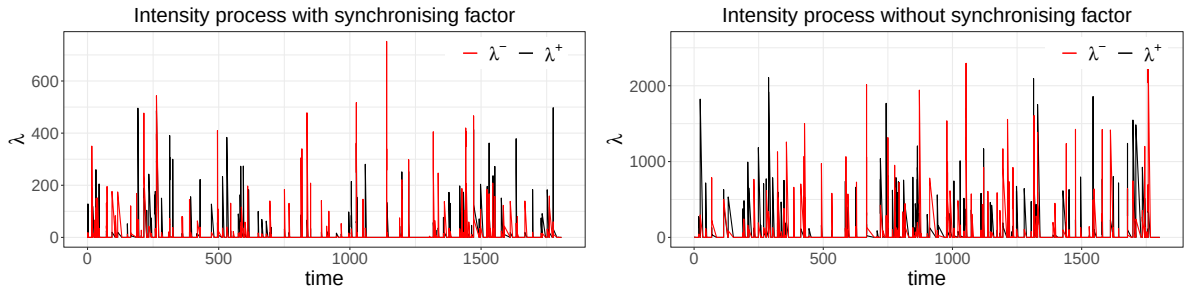

Figure 11: The generated samples of $\lambda^{\pm}$ under the market order model with the synchronising (left) and without the synchronising factor κ (right) using the calibrated parameter in Table 6, respectively.

using the same set of 30,000 full synchronising simulated samples. Table 7 presents the detailed statistics of market-making PnLs under the optimal strategies derived from each model when the synchronising tendency exists in the high-frequency market. Figure 12 (left) illustrates the distributions of the PnLs derived in the three cases.

The optimal strategy from the full synchrony model produces the highest expected profit, whereas it leads to the least risk (i.e. standard deviation), compared with the other models. The full synchrony model creates +2.9 (+1.4%) and +4.7 (+2.3%) higher returns than the half and no synchrony models, respectively, whereas it takes -4.97 (-15.9%) and -3.89 (-12.9%) less risk than the half and no synchrony models, respectively[11]. In addition, Figure 12 shows that the PnL in the full synchrony model has the shortest left tail of all the models; moreover, there is a remarkably high chance of earning a profit in the range of the expected return ± standard deviation in this model.

This simulation implies that it is beneficial to reflect the synchronising effect in the market and limit order dynamics in terms of wealth management from market-making in the high-frequency market. As discussed in Section 4, the synchronising tendency has appeared in the real market between buy and sell order arrivals and the tendency level has significantly increased in the past ten years. Market-makers may consider the synchronising factor in their order book models that enable them to expect higher and more stable performance than existing models.

In the second part, we conduct simulation to assess how a market-maker's PnLs can be different with respect to the stability of the market order arrivals. For the test, we consider the stability of market order arrivals that recorded the highest, middle, and lowest levels in 2018 among the six stocks under the synchrony model. This test compares the

[11]This difference in the expected profits between two model assumptions can be widened if the synchronising level increases or the exciting levels decrease.

Table 7: Terminal PnLs trading with the optimal strategy derived under the full, half, and no synchrony models when trading lasts 300 seconds.

Order Arrival Model		Mean	SD	Min	Max
(i) Full Synchrony		209.8	26.23	120.00	314.1
(ii) Half synchrony		206.8	30.13	50.44	314.2
(iii) No synchrony		205.1	31.20	55.02	346.3
Difference	(i)-(ii)	+2.9	-3.89	+69.55	-0.1
	(i)-(iii)	+4.7	-4.97	+67.98	-31.2

market-maker's optimal PnLs under the three different cases of stability under the same order arrival model. For the parameters of c, we employ the same values specified in the synchrony model (Table 6) to the three cases.

To ensure fair comparison of the stability indicator, we use $\theta = 0.3$ and $\kappa = 0.1$ in all cases. Based on Table 6 employed in the previous comparison test, we set the similar scales of the other synchrony model parameters (η, ν, and β), where Table 8 (left) presents the details. Using the same procedure, we estimate the optimal strategy through DNN training for a 300-second time horizon and conduct the market-making simulation under 30,000 scenarios generated in each parameter case; and we finally obtain the terminal PnLs. Table 8 (right) and Figure 12 (right) illustrate the descriptive statistics and the distributions of the market-maker's PnLs depending on the market stability level, respectively. This result shows that a higher expected return is more likely to occur in the optimal market-making as market instability rises. The reason may be that market-makers in an unstable market are able to explore more opportunities to earn profit as market orders arrive more frequently and irregularly. This finding may be useful to market practitioners because there have been apparent increases in instability of high-frequency market order arrivals in the past ten years for the particular stocks such as IBM, Chevron, and Amazon.

Table 8: The chosen market order arrival model parameter sets with $\theta = 0.3$ and $\kappa = 0.1$ which lead to different stability levels (left); and the descriptive statistics of the optimal market-maker's PnLs for 300-seconds trading time (right).

Stability		Model parameters			Market-maker's PnL			
		η	ν	β	Mean	SD	Min	Max
(i) High	3.5	150	30	700	204.7	25.90	92.34	313.3
(ii) Middle	2.9	200	20	700	209.8	26.23	120.0	314.1
(iii) Low	1.6	200	20	400	214.1	28.17	121.3	356.4

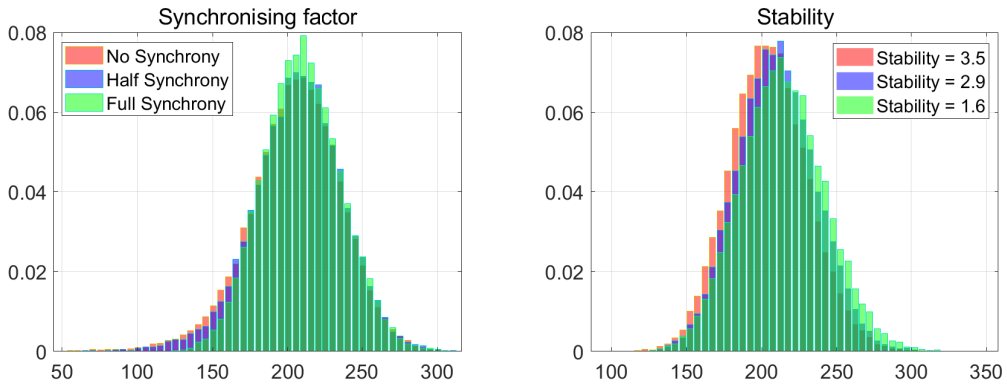

Figure 12: The market-maker's PnL distributions under the different assumptions. Comparison between the order book models with the full, half, and no synchronising factor (left); and comparison of the synchrony models with different stability assumptions – high (3.5), middle (2.9), and low (1.6) (right).

7. Conclusion

This study investigates the optimal execution strategy of market-making for the market and limit order arrival dynamics using a novel model with a synchronising factor between buy and sell orders. To examine the synchrony tendency in the market, we estimate the relevant model parameters using the MLE algorithm modified for the proposed model. The empirical results confirm that the arrival frequencies of market buy orders have the propensity to follow sell orders' long-term mean level and vice versa. This is presumably caused by the huge increase in the influence of HFT activities in the markets, and such a phenomenon may continue to appear or become stronger with further innovation in technologies.

To solve the high-dimensional HJB equation derived as the optimal decision problem, we propose the DNN approximation instead of mesh-based methods. We verify the existence of a DNN structure that guarantees that the loss function is sufficiently small on the restricted compact domain. We then estimate the solution of the HJB equation and optimal placement strategy within an acceptable range of numerical errors. Finally, we simulate the performance of the end profit from market-making under a variety of strategies and compare their statistical distributions. As a complementary test, we observe that the optimal controls estimated by the DNN method and existing method (Cartea et al., 2014) produce statistically non-distinguishable distributions for the market-maker's profits (Appendix C).

From the simulations, we find that the estimated optimal strategy outperforms the other deterministic and suboptimal strategies significantly. It enables us to accomplish a profit

profile with stable and steady accumulation over time under tight inventory management. In addition, we show the importance of reflecting the synchronising effects in both market and limit order dynamics to improve the wealth management (higher expected profit and less risk) of high-frequency market-making. In particular, we show that more unstable the high-frequency market provides the more likely that market-makers are to obtain higher expected returns.

Acknowledgements

We thank the editors and anonymous reviewers for their valuable comments and suggestions which have helped us to significantly improve the quality of our paper of the previous three versions. This work was partially supported by the National Research Foundation (NRF) of Korea grant (no.2020R1F1A1A01066238). The last author was partially supported by the Engineering and Physical Sciences Research Council (EPSRC) of UK Standard Research Grant (EP/V008331/1).

References

Aït-Sahalia, Y., Cacho-Diaz, J., Laeven, R.J., 2015. Modeling financial contagion using mutually exciting jump processes. Journal of Financial Economics 117, 585–606. doi:10.1016/j.jfineco.2015.03.002.

Avellaneda, M., Stoikov, S., 2008. High-frequency trading in a limit order book. Quantitative Finance 8, 217–224.

Bacry, E., Dayri, K., Muzy, J.F., 2012. Non-parametric kernel estimation for symmetric Hawkes processes. Application to high frequency financial data. The European Physical Journal B 85, 157.

Bacry, E., Jaisson, T., Muzy, J.F., 2016. Estimation of slowly decreasing Hawkes kernels: application to high-frequency order book dynamics. Quantitative Finance 16, 1179–1201.

Bacry, E., Muzy, J.F., 2014. Hawkes model for price and trades high-frequency dynamics. Quantitative Finance 14, 1147–1166.

Baesens, B., Setiono, R., Mues, C., Vanthienen, J., 2003. Solving high-dimensional partial differential equations using deep learning. Management Science 46, 312–329.

Baillie, R., Bollerslev, T., 1989. The message in daily exchange rates: A conditional-variance tale. Journal of Business & Economic Statistics 7, 197–305.

Berg, J., Nystrom, K., 2018. A unified deep artificial neural network approach to partial differential equations in complex geometries. Neurocomputing 317, 28–41.

Bowsher, C.G., 2007. Modelling security market events in continuous time: Intensity based, multivariate point process models. Journal of Econometrics 141, 876–912.

Cartea, A., Jaimungal, S., 2013. Modelling asset prices for algorithmic and high-frequency trading. Appl Math Fin 20, 512–547.

Cartea, A., Jaimungal, S., Ricci, J., 2014. Buy low, sell high: A high frequency trading perspective. SIAM Journal on Financial Mathematics 60, 415–444.

Cerchi, M., Havenner, A., 1988. Cointegration and stock prices: The random walk and wall street revisited. Journal of Economic Dynamics and Control 12, 333–346.

Chiu, M.C., Wong, H.Y., Zhao, J., 2015. Commodity derivatives pricing with cointegration and stochastic covariances. European Journal of Operational Research 246, 476–486.

Czarnecki, W., Osindero, S., Jaderberg, M., Swirszcz, G., Pascanu, R., 2017. Sobolev training for neural networks. Advances in Neural Information Processing Systems -, 4281–4290.

Da Fonseca, J., Zaatour, R., 2014. Hawkes process: Fast calibration, application to trade clustering, and diffusive limit. Journal of Futures Market 34, 548–579.

Danielsson, J., Shin, H.S., Zigrand, J.P., 2012. Endogenous and systemic risk, in: Quantifying systemic risk. University of Chicago Press, pp. 73–94.

De Spiegeleer, J., Madan, D., Reyners, S., Schoutens, W., 2018. Machine learning for quantitative finance: fast derivatives pricing, hedging and fitting. Quantitative Finance 18, 1635–1643.

Dixon, M., Klabjan, D., Bang, J., 2019. Classification-based financial markets prediction using deep neural networks. Working paper -, –.

Duan, J.C., Pliska, S., 2004. Option valuation with co-integrated asset prices. Journal of Economic Dynamics & Control 28, 727–754.

E, W., Han, J., Jentzen, A., 2017. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. Commun. Math. Stat. 5, 349–380.

Engle, R.F., Granger, C.W.J., 1987. Co-integration and error correction: representation, estimation, and testing. Econometrica: journal of the Econometric Society , 251–276.

Filimonov, V., Sornette, D., 2012. Quantifying reflexivity in financial markets: Toward a prediction of flash crashes. Physical Review E 85, 56108.

Fitzpatrick, T., Mues, C., 2016. An empirical comparison of classification algorithms for mortgage default prediction: evidence from a distressed mortgage market. European Journal of Operational Research 249, 427–439.

Fletcher, R., 1987. Practical methods of optimization . Wiley 2.

Fujii, M., Takahashi, A., Takahashi, M., 2019. Asymptotic expansion as prior knowledge in deep learning method for high dimensional bsdes. Asia-Pacific Financial Markets 26, 391–408.

Gerig, A., 2015. High-frequency trading synchronizes prices in financial markets. US Securities and Exchange Commission January, –. URL: `https://www.sec.gov/files/dera-wp-hft-synchronizes.pdf`.

Gramacy, R., Ludkovski, M., 2015. Sequential design for optimal stopping problems. SIAM Journal of Financial Mathematics 6, 748–745.

Granger, C.W.J., 1981. Some properties of time series data and their use in econometric model specification. Journal of econometrics 16, 121–130.

Gueant, O., Lehalle, C.A., Fernandez-Tapia, J., 2013. Dealing with the inventory risk: a solution to the market making problem. Math Finan Econ 8, 477–507.

Guilbaud, F., Pham, H., 2013. Optimal high-frequency trading with limit and market orders. Quantitative Finance 13, 79–94.

Guo, X., Larrard, A., Ruan, Z., 2017. Optimal placement in a limit order book: an analytical approach. Math Finan Econ 11, 189–213.

Ha, S.Y., Kim, K.K., Lee, K., 2015. A mathematical model for multi-name credit based on community flocking. Quantitative Finance 15, 841–851.

Han, J., Jentzen, A., E, W., 2018. Solving high-dimensional partial differential equations using deep learning. PNAS 115, 8505–8510.

Hardiman, S.J., Bercot, N., Bouchaud, J.P., 2013. Critical reflexivity in financial markets: a Hawkes process analysis. The European Physical Journal B 86, 442.

Hawkes, A.G., 1971. Point spectra of some mutually exciting point processes. Journal of the Royal Statistical Society. Series B (Methodological) 33, 438–443.

Hawkes, A.G., Oakes, D., 1974. A cluster process representation of a self-exciting process. Journal of Applied Probability 11, 493–503.

Henningsen, A., Toomet, O., 2011. maxLik: A package for maximum likelihood estimation in R. Computational Statistics 26, 443–458.

Ho, T., Stoll, H.R., 1981. Optimal dealer pricing under transactions and return uncertainty. Journal of Financial Economics 9, 47–73.

Hornik, K., 1991. Approximation capabilities of multilayer feedforward networks. Neural Networks 4, 251–257. doi:10.1016/0893-6080(91)90009-T.

Hornik, K., Stinchcombe, M., White, H., 1989. Multilayer feedforward networks are universal approximators. Neural Networks 2, 359–366. doi:10.1016/0893-6080(89)90020-8.

Huepe, C., Aldana, M., 2008. New tools for characterizing swarming systems: A comparison of minimal models. Physic A: Statistical Mechanics and its Applications 387, 2809–2822.

Hutchinson, J., Lo, A., Possio, T., 1994. A nonparametric approach to pricing and hedging derivative securities via learning networks. Journal of Finance 49, 851–889.

Jang, H., Lee, K., Lee, K., 2020. Systemic risk in market microstructure of crude oil and gasoline futures prices: A hawkes flocking model approach. Journal of Futures Markets 40, 247–275.

Kellard, N., Dunis, C., Sarantis, N., 2010. Foreign exchange, fractional cointegration and the implied–realized volatility relation. Journal of Banking & Finance 34, 882–891.

Khandani, A., Kim, A., Lo, A., 2010. Consumer credit-risk models via machine-learning algorithms. International Journal of Forecasting 34, 2767–2787.

Kingma, D.P., Ba, J., 2014. Adam: A Method for Stochastic Optimization arXiv:1412.6980.

Kruskal, W., Wallis, W., 1952. Use of ranks in one-criterion variance analysis. Journal of the American Statistical Association 47, 538–621.

Large, J., 2007. Measuring the resiliency of an electronic limit order book. Journal of Financial Market 10, 1–25.

Lee, C.M.C., Ready, M.J., 1991. Inferring trade direction from intraday data. The Journal of Finance 46, 733–746.

Lee, K., Seo, B.K., 2017. Modeling microstructure price dynamics with symmetric Hawkes and diffusion model using ultra-high-frequency stock data. Journal of Economic Dynamics and Control 79, 154–183.

Loterman, G., Brown, I., Martens, D., Mues, C., Baesens, B., 2012. Benchmarking regression algorithms for loss given default modeling. International Journal of Forecasting 28, 161–170.

Lu, X., Abergel, F., 2018. High dimensional hawkes processes for limit order books. Quantitative Finance 18, 249–264.

Matthias, K., Diepold, K., 2019. Sobolev training with approximated derivatives for black-box function regression with neural networks. Joint European Conference on Machine Learning and Knowledge Discovery in Databases -, –.

Miller, R.S., Shorter, G., 2016. High frequency trading: Overview of recent developments. Congressional Research Service .

Ng, V.K., Pirrong, S.C., 1994. Fundamentals and volatility: Storage, spreads, and the dynamics of metals prices. The Journal of Business 67, 203–230.

Ogata, Y., 1978. The asymptotic behaviour of maximum likelihood estimators for stationary point processes. Annals of the Institute of Statistical Mathematics 30, 243–261. doi:10.1007/BF02480216.

Rauch, E.M., Millonas, M.M., Chialvo, D.R., 1995. Pattern formation and functionality in swarm models. Physics Letters A 207, 185–193.

Reynolds, C.W., 1987. Flocks, herds and schools: A distributed behavioral model. ACM SIGGRAPH Computer Graphics 21, 25–34.

Rosu, I., 2009. A dynamic model of the limit order book. The Review of Financial Studies 22, 4601–4641.

SEC, 2010. Concept Release on Equity Market Structure. US Securities and Exchange Commission 17 CFR Part 242, 34–61358. URL: https://www.sec.gov/rules/concept/2010/34-61358.pdf.

SEC, 2014. Equity market structure literature review part ii: High frequency trading. US Securities and Exchange Commission March, –. URL: https://www.sec.gov/marketstructure/research/hft_lit_review_march_2014.pdf.

Sirignano, J., 2019. Deep learning for limit order books. Quantitative Finance 19, 549–570.

Sirignano, J., Spiliopoulos, K., 2018. DGM: A deep learning algorithm for solving partial differential equations. Journal of Computational Physics 375, 1339–1364. URL: https://linkinghub.elsevier.com/retrieve/pii/S0021999118305527, doi:10.1016/j.jcp.2018.08.029.

Veraart, L.A.M., 2010. Optimal market making in the foreign exchange market. Appl Math Fin 17, 359–372.

Appendix A. Proofs

Appendix A.1. Lemma 1

We derive the conditions that guarantee the intensity processes for market orders be stable in Lemma 1 with the mean future rate $m_t^{\pm}(u) = \mathbb{E}[\lambda_u^{\pm}|\mathcal{F}_t]$ for $u \geq t$.

Lemma 1. *The mean future rates $m_t^{\pm}(u) = \mathbb{E}[\lambda_u^{\pm}|\mathcal{F}_t]$ are bounded for all $u \geq t$ if and only if $(1-\kappa)\beta > \eta + \nu$. Furthermore,*

$$\lim_{u\to\infty} m_t^{\pm}(u) = D^{-1}\pi, \quad \textit{where} \quad D = \begin{bmatrix} \beta-\eta & -\kappa\beta-\nu \\ -\kappa\beta-\nu & \beta-\eta \end{bmatrix} \quad \textit{and} \quad \pi = \begin{bmatrix} \beta\theta^- \\ \beta\theta^+ \end{bmatrix}.$$

Proof. Similar to the scheme of Cartea et al. (2014), taking the integral of both sides of Eq.(3), the conditional expectation $\mathbb{E}[\cdot|\mathcal{F}_t]$ by applying Fubini's theorem, and the derivative gives the following system of ordinary differential equations (ODEs) for $m_t^{\pm}(u)$, then we have

$$\frac{d}{du}\begin{bmatrix} m_t^-(u) \\ m_t^+(u) \end{bmatrix} = \begin{bmatrix} \beta\theta^- \\ \beta\theta^+ \end{bmatrix} + \begin{bmatrix} -\beta+\eta & \kappa\beta+\nu \\ \kappa\beta+\nu & -\beta+\eta \end{bmatrix}\begin{bmatrix} m_t^-(u) \\ m_t^+(u) \end{bmatrix} \tag{A.1}$$

with initial values $m_t^-(t) = \lambda_t^-, m_t^+(t) = \lambda_t^+$. If D has no zero eigenvalues, it takes a unique solution such that

$$\begin{bmatrix} m_t^-(u) \\ m_t^+(u) \end{bmatrix} = e^{-D(u-t)}\left(\begin{bmatrix} \lambda_t^- \\ \lambda_t^+ \end{bmatrix} - D^{-1}\pi\right) + D^{-1}\pi. \tag{A.2}$$

Since the eigenvalues of D are $(1\pm\kappa)\beta - (\eta \mp \nu)$, $\lim_{u\to\infty} m_t(u)$ converges to $D^{-1}\pi$ if and only if $(1\pm\kappa)\beta > \eta \mp \nu$. It implies $(1-\kappa)\beta > \eta + \nu$ since $\beta, \eta, \nu, \kappa > 0$. □

We also obtain the conditional expectation $n_t^{\pm}(u) = \mathbb{E}[c_u^{\pm}|\mathcal{F}_t]$ by solving a system of ODEs derived by taking the integral, conditional expectation, and derivative in Eq.(4).

Remark 2. *The mean future rate $n_t^{\pm}(u)$ is given by, for all $u \geq t$,*

$$\begin{bmatrix} n_t^-(u) \\ n_t^+(u) \end{bmatrix} = e^{-\tilde{D}(u-t)} \int_t^u e^{\tilde{D}(s-t)} G(s) ds + e^{-\tilde{D}(u-t)} \begin{bmatrix} n_t^- \\ n_t^+ \end{bmatrix}, \tag{A.3}$$

where

$$\tilde{D} = \begin{bmatrix} \xi & -\kappa_c \xi \\ -\kappa_c \xi & \xi \end{bmatrix} \quad \textit{and} \quad G(s) = \begin{bmatrix} \xi\alpha \\ \xi\alpha \end{bmatrix} + \eta_k \begin{bmatrix} m_t^-(s) \\ m_t^+(s) \end{bmatrix} + \nu_k \begin{bmatrix} m_t^+(s) \\ m_t^-(s) \end{bmatrix} \tag{A.4}$$

with the initial condition $n_t^{\pm} = n_t^{\pm}(t)$ and $m_t^{\pm}(s)$ in Eq.(A.2).

Appendix A.2. Solution for Eq. (3)

In Section 3, we consider the intensity processes $(\lambda_t^-, \lambda_t^+)_{t\geq 0}$ for market sell and buy orders $(M_t^-, M_t^+)_{t\geq 0}$ such that

$$\begin{aligned} \mathrm{d}\lambda_t^- &= \beta(\theta^- - \lambda_t^- + \kappa\lambda_t^+)\mathrm{d}t + \eta \mathrm{d}M_t^- + \nu \mathrm{d}M_t^+ \\ \mathrm{d}\lambda_t^+ &= \beta(\theta^+ - \lambda_t^+ + \kappa\lambda_t^-)\mathrm{d}t + \eta \mathrm{d}M_t^+ + \nu \mathrm{d}M_t^- \end{aligned} \tag{A.5}$$

with all non-negative coefficients.

For the goodness-of-fit test of the proposed intensity model using the maximum likelihood estimator, the explicit form of $(\lambda_t^-, \lambda_t^+)_{t\geq 0}$ must have a likelihood function. Since the synchronising terms are $\mathrm{d}\lambda_t^+$ and $\mathrm{d}\lambda_t^-$, the system of the stochastic differential equations (SDEs) is not solved straightforward. Following Ha et al. (2015), we split each process into an average process and each fluctuation process, which are the average process $\bar{\lambda}_t = \frac{1}{2}(\lambda_t^+ + \lambda_t^-)$, and the fluctuation processes $\tilde{\lambda_t^-} = \lambda_t^- - \bar{\lambda}_t$ and $\tilde{\lambda_t^+} = \lambda_t^+ - \bar{\lambda}_t$. From this, $\tilde{\lambda_t^-} + \tilde{\lambda_t^+} = 0$ which is applied to derive the solution $\tilde{\lambda}_t$.

First, we have the average process such as

$$\mathrm{d}\bar{\lambda}_t = \frac{\beta}{2}\left(\theta + 2(\kappa - 1)\bar{\lambda}_t\right) dt + \zeta \mathrm{d}M_t^- + \zeta \mathrm{d}M_t^+ \tag{A.6}$$

where $\theta = \theta^- + \theta^+$ and $\zeta = \frac{1}{2}(\eta + \nu)$. By letting $A = \beta(\kappa - 1)$, a (strong) solution of the linear SDE in Eq.(A.6) is given by

$$\bar{\lambda}_t = e^{At}\bar{\lambda}_0 + \frac{\theta}{2(\kappa - 1)}(e^{At} - 1) + \zeta e^{At}\left(\int_0^t e^{-As} dM_s^- + e^{-As} dM_s^+\right). \tag{A.7}$$

Second, we have the SDE for each fluctuation process $\tilde{\lambda_t^+}, \tilde{\lambda_t^-}$ such as

$$\begin{aligned} \mathrm{d}\tilde{\lambda_t^-} &= \frac{\beta}{2}\left(\tilde{\theta} - 2(\kappa+1)\tilde{\lambda_t^-}\right)dt + \tilde{\zeta}\mathrm{d}M_t^- - \tilde{\zeta}\mathrm{d}M_t^+ \\ \mathrm{d}\tilde{\lambda_t^+} &= -\frac{\beta}{2}\left(\tilde{\theta} + 2(\kappa+1)\tilde{\lambda_t^+}\right)dt - \tilde{\zeta}\mathrm{d}M_t^- + \tilde{\zeta}\mathrm{d}M_t^+ \end{aligned} \tag{A.8}$$

where $\tilde{\theta} = \theta^- - \theta^+$ and $\tilde{\zeta} = \frac{1}{2}(\eta - \nu)$. By letting $\tilde{A} = \beta(\kappa+1)$, the (strong) solutions of Eq.(A.8) are given by

$$\begin{aligned} \tilde{\lambda_t^-} &= e^{-\tilde{A}t}\tilde{\lambda_0^-} + \frac{\tilde{\theta}}{2(\kappa+1)}(1-e^{-\tilde{A}t}) + \tilde{\zeta}e^{-\tilde{A}t}\left(\int_0^t e^{-\tilde{A}s}dM_s^- - e^{-\tilde{A}s}dM_s^+\right) \\ \tilde{\lambda_t^+} &= e^{-\tilde{A}t}\tilde{\lambda_0^+} - \frac{\tilde{\theta}}{2(\kappa+1)}(1-e^{-\tilde{A}t}) - \tilde{\zeta}e^{-\tilde{A}t}\left(\int_0^t e^{-\tilde{A}s}dM_s^- - e^{-\tilde{A}s}dM_s^+\right), \end{aligned} \tag{A.9}$$

respectively. From Eqs.(A.7) and (A.9), we have $\lambda_t^- = \bar{\lambda}_t + \tilde{\lambda_t^-}$ and $\lambda_t^+ = \bar{\lambda}_t + \tilde{\lambda_t^+}$.

Appendix A.3. Convergence theorem for the loss function $L(f)$

Theorem 3. *Denote the set of all functions implemented by one hidden layer with n hidden units and one output unit as $\mathfrak{N}_k^n(\varphi) = \left\{ g : \mathbb{R}^k \to \mathbb{R} \middle| g(x) = \sum_{j=1}^n b_j \varphi\left(a_j' x + c_j\right) \right\}$ where φ is the activation function of the hidden units. When a function is implemented by a network with an arbitrarily large number of hidden units, the function is contained in $\mathfrak{N}_k(\varphi) = \bigcup_{n=1}^\infty \mathfrak{N}_k^n(\varphi)$. For every $\epsilon > 0$, there exists a function $f \in \mathfrak{N}_k(\varphi)$ such that $L(f) \le \epsilon$ for the nonlinear operator $\mathcal{L}$ defined by*

$$\begin{aligned} \mathcal{L}u(t,q,\boldsymbol{\lambda},\boldsymbol{c}) = {} & \beta(\theta^- - \lambda^- + \kappa\lambda^+)\frac{\partial u}{\partial \lambda^-} + \beta(\theta^+ - \lambda^+ + \kappa\lambda^-)\frac{\partial u}{\partial \lambda^+} \\ & + \xi(\alpha - c^- + \kappa_c c^+)\frac{\partial u}{\partial c^-} + \xi(\alpha - c^+ + \kappa_c c^-)\frac{\partial u}{\partial c^+} \\ & + \lambda^-(\Delta_{q,\lambda,c}^- u - u)\mathbb{1}_{\{c^-(\Delta_{q,\lambda,c}^- u - \Delta_{\lambda,c}^- u) \ge 1\}} + \lambda^+(\Delta_{q,\lambda,c}^+ u - u)\mathbb{1}_{\{c^+(\Delta_{q,\lambda,c}^+ u - \Delta_{\lambda,c}^+ u) \ge 1\}} \\ & + \lambda^-\left(\frac{\mathrm{e}^{c^-(\Delta_{q,\lambda,c}^- u - \Delta_{\lambda,c}^- u)}}{\mathrm{e}c^-} + \Delta_{\lambda,c}^- u - u\right)\mathbb{1}_{\{c^-(\Delta_{q,\lambda,c}^- u - \Delta_{\lambda,c}^- u) < 1\}} \\ & + \lambda^+\left(\frac{\mathrm{e}^{c^+(\Delta_{q,\lambda,c}^+ u - \Delta_{\lambda,c}^+ u)}}{\mathrm{e}c^+} + \Delta_{\lambda,c}^+ u - u\right)\mathbb{1}_{\{c^+(\Delta_{q,\lambda,c}^+ u - \Delta_{\lambda,c}^+ u) < 1\}} - \psi q^2. \end{aligned} \tag{A.10}$$

Proof. By the universal approximation theorem proposed by Hornik (1991), for any $\epsilon > 0$, there exists a function $f \in \mathfrak{N}_6(\varphi)$ such that

$$\max_{|\alpha| \le 1} \sup_{\tilde{\mathcal{D}}_T} |D^\alpha(f(t,q,\boldsymbol{\lambda},\boldsymbol{c};\Theta) - u(t,q,\boldsymbol{\lambda},\boldsymbol{c}))| < \epsilon, \tag{A.11}$$

where D^{α} is the αth order derivative operator.

The second term of loss function $L(f)$ in Eq.(15) is

$$\left\| f(T, q, \boldsymbol{\lambda}, \boldsymbol{c}; \Theta) - \phi q^2 \right\|^2_{\tilde{\mathcal{D}}, \mu_2} = \int_{\tilde{\mathcal{D}}} |f - u|^2 \, d\mu_2. \tag{A.12}$$

For the first term of Eq.(15), we can write that

$$\begin{aligned} &\left\| \frac{\partial f}{\partial t}(t, q, \boldsymbol{\lambda}, \boldsymbol{c}; \Theta) + \mathcal{L} f(t, q, \boldsymbol{\lambda}, \boldsymbol{c}; \Theta) \right\|_{\tilde{\mathcal{D}}_T, \mu_1} \\ &= \left\| \frac{\partial f}{\partial t}(t, q, \boldsymbol{\lambda}, \boldsymbol{c}; \Theta) - \frac{\partial u}{\partial t}(t, q, \boldsymbol{\lambda}, \boldsymbol{c}) + \mathcal{L} f(t, q, \boldsymbol{\lambda}, \boldsymbol{c}; \Theta) - \mathcal{L} u(t, q, \boldsymbol{\lambda}, \boldsymbol{c}) \right\|_{\tilde{\mathcal{D}}_T, \mu_1}. \end{aligned} \tag{A.13}$$

Let $A_h^{\pm} = \{(t, q, \boldsymbol{\lambda}, \boldsymbol{c}) \in \tilde{\mathcal{D}}_T \,\big|\, c^{\pm}(\Delta^{\pm}_{q,\lambda,c} h - \Delta^{\pm}_{\lambda,c} h) < 1\}$. By the Minkowski inequality,

Eq.(A.13) is less than the following:

$$
\begin{aligned}
&\left(\int_{\tilde{\mathcal{D}}_T}\left|\frac{\partial f}{\partial t}-\frac{\partial u}{\partial t}\right|^2 d\mu_1\right)^{1/2}+\left(\int_{\tilde{\mathcal{D}}_T}\left|\beta(\Theta^- -\lambda^- +\kappa\lambda^+)\left(\frac{\partial f}{\partial \lambda^-}-\frac{\partial u}{\partial \lambda^-}\right)\right|^2 d\mu_1\right)^{1/2}\\
&+\left(\int_{\tilde{\mathcal{D}}_T}\left|\beta(\Theta^+ -\lambda^+ +\kappa\lambda^-)\left(\frac{\partial f}{\partial \lambda^+}-\frac{\partial u}{\partial \lambda^+}\right)\right|^2 d\mu_1\right)^{1/2}+\left(\int_{\tilde{\mathcal{D}}_T}\left|\xi(\alpha-c^- +\kappa_c c^+)\left(\frac{\partial f}{\partial c^-}-\frac{\partial u}{\partial c^-}\right)\right|^2 d\mu_1\right)^{1/2}\\
&+\left(\int_{\tilde{\mathcal{D}}_T}\left|\xi(\alpha-c^+ +\kappa_c c^-)\left(\frac{\partial f}{\partial c^+}-\frac{\partial u}{\partial c^+}\right)\right|^2 d\mu_1\right)^{1/2}+\left(\int_{\tilde{\mathcal{D}}_T}\left|(\lambda^- +\lambda^+)(f-u)\right|^2 d\mu_1\right)^{1/2}\\
&+\left(\int_{\tilde{\mathcal{D}}_T\setminus(A_f^-\cup A_u^-)}\left|\lambda^-(\Delta^-_{q,\lambda,c}f-\Delta^-_{q,\lambda,c}u)\right|^2 d\mu_1\right)^{1/2}+\left(\int_{\tilde{\mathcal{D}}_T\setminus(A_f^+\cup A_u^+)}\left|\lambda^+(\Delta^+_{q,\lambda,c}f-\Delta^+_{q,\lambda,c}u)\right|^2 d\mu_1\right)^{1/2}\\
&+\left(\int_{A_f^-\cap\overline{A_u^-}}\left|\lambda^-\left(\frac{\mathrm{e}^{c^-(\Delta^-_{q,\lambda,c}f-\Delta^-_{\lambda,c}f)}}{\mathrm{e}c^-}-\frac{\mathrm{e}^{c^-(\Delta^-_{q,\lambda,c}u-\Delta^-_{\lambda,c}u)}}{\mathrm{e}c^-}\right)\right|^2 d\mu_1\right)^{1/2} \qquad (*)\\
&+\left(\int_{A_f^+\cap\overline{A_u^+}}\left|\lambda^+\left(\frac{\mathrm{e}^{c^+(\Delta^+_{q,\lambda,c}f-\Delta^+_{\lambda,c}f)}}{\mathrm{e}c^+}-\frac{\mathrm{e}^{c^+(\Delta^+_{q,\lambda,c}u-\Delta^+_{\lambda,c}u)}}{\mathrm{e}c^+}\right)\right|^2 d\mu_1\right)^{1/2} \qquad (*)\\
&+\left(\int_{A_f^-\cap A_u^-}\left|\lambda^-(\Delta^-_{\lambda,c}f-\Delta^-_{\lambda,c}u)\right|^2 d\mu_1\right)^{1/2}+\left(\int_{A_f^+\cap A_u^+}\left|\lambda^+(\Delta^+_{\lambda,c}f-\Delta^+_{\lambda,c}u)\right|^2 d\mu_1\right)^{1/2}\\
&+\left(\int_{A_u^-\setminus A_f^-}\left|\lambda^-\left(\Delta^-_{q,\lambda,c}f-\Delta^-_{\lambda,c}u-\frac{\mathrm{e}^{c^-(\Delta^-_{q,\lambda,c}u-\Delta^-_{\lambda,c}u)}}{\mathrm{e}c^-}\right)\right|^2 d\mu_1\right)^{1/2} \qquad (**)\\
&+\left(\int_{A_f^-\setminus\overline{A_u^-}}\left|\lambda^-\left(\Delta^-_{q,\lambda,c}u-\Delta^-_{\lambda,c}f-\frac{\mathrm{e}^{c^-(\Delta^-_{q,\lambda,c}f-\Delta^-_{\lambda,c}f)}}{\mathrm{e}c^-}\right)\right|^2 d\mu_1\right)^{1/2} \qquad (\dagger)\\
&+\left(\int_{A_u^+\setminus A_f^+}\left|\lambda^+\left(\Delta^+_{q,\lambda,c}f-\Delta^+_{\lambda,c}u-\frac{\mathrm{e}^{c^+(\Delta^+_{q,\lambda,c}u-\Delta^+_{\lambda,c}u)}}{\mathrm{e}c^+}\right)\right|^2 d\mu_1\right)^{1/2} \qquad (**)\\
&+\left(\int_{A_f^+\setminus\overline{A_u^+}}\left|\lambda^+\left(\Delta^+_{q,\lambda,c}u-\Delta^+_{\lambda,c}f-\frac{\mathrm{e}^{c^+(\Delta^+_{q,\lambda,c}f-\Delta^+_{\lambda,c}f)}}{\mathrm{e}c^+}\right)\right|^2 d\mu_1\right)^{1/2}. \qquad (\dagger)
\end{aligned}
$$

For the integrals $(*)$, the exponential functions are uniformly continuous because

$$c^{\pm}\left(\Delta^{\pm}_{q,\lambda,c}f-\Delta^{\pm}_{\lambda,c}f\right)<1 \ \text{ and } \ c^{\pm}\left(\Delta^{\pm}_{q,\lambda,c}u-\Delta^{\pm}_{\lambda,c}u\right)\le 1.$$

We can make the integrand sufficiently small by choosing $f\in\mathfrak{N}_k(\varphi)$ to be sufficiently close

to u. For the integrals $(**)$,

$$\left(\Delta^{\pm}_{q,\lambda,c}u - \Delta^{\pm}_{\lambda,c}u\right) < \frac{1}{c^{\pm}} \leq \left(\Delta^{\pm}_{q,\lambda,c}f - \Delta^{\pm}_{\lambda,c}f\right)$$

and with Eq.(A.11), $A^{\pm}_u \setminus A^{\pm}_f \subset \{(t,q,\boldsymbol{\lambda},\boldsymbol{c}) \in \tilde{\mathcal{D}}_T \,|\, \Delta^{\pm}_{q,\lambda,c}u - \Delta^{\pm}_{\lambda,c}u \in \left(\frac{1}{c^{\pm}} - 2\epsilon, \frac{1}{c^{\pm}}\right)\}$. Thus, $\mu_1(A^{\pm}_u \setminus A^{\pm}_f)$ decreases to zero if ϵ becomes arbitrarily small. The integrals $(\dagger)$ can be treated similarly. The other integrals and Eq.(A.12) are bounded by ϵ times some constant related to the size of domain $\mathcal{D}_T$, therefore we can make the loss $L(f)$ arbitrarily small.

□

Appendix A.4. Thinning algorithm for the proposed model

Let $t_1, t_2, \ldots$ with $0 < t_1 < t_2 < \ldots$ denote the random arrival times at which the counting processes M^-_t or M^+_t jump. Algorithm 2 describes a thinning algorithm to generate the random arrival times.

Appendix B. Tables related to the MLE results

Appendix C. PnL comparison with the Cartea et al. (2014)'s formula

This test compares the market-maker's PnLs with the optimal strategies using the DNN-estimation and the asymptotic formula of Cartea et al. (2014), when $\kappa = \kappa_c = 0$. Table C.10 shows the selected parameters of λ_t and c_t; and $\psi = 0.00001$, $\phi = 0$[12], $T = 180$, and $\sigma = 0.0001$ are adopted for the comparison. Figure C.13 illustrates the PnL distributions obtained from the two different schemes for the four cases, which shows that the expected values of the PnLs achieved under the both methods are reasonably close to each other.

To verify the distributional similarity rather thoroughly, we apply the Kruskal–Wallis (KW) test, which validates whether two independent samples are selected from the same distribution (Kruskal and Wallis, 1952). At a x%-significance level, the KW test rejects the null hypothesis H_0 that the two samples originate from the same distribution, if the p-value is less than x%; otherwise, it does not reject it indicating that the two samples are from the same distribution. The KW test with a 5%-significance level is taken for the two PnL distributions from the DNN-estimation and asymptotic formula. Table C.11 displays the expected terminal PnLs with their absolute difference and the p-value for each case. We can see that H_0 cannot be rejected for all the cases, as the p-values are greater than 0.05, indicating that the two PnL values obtained using the two approaches can be statistcally non-distinguishable.

[12]The asymptotic optimal controls contains an expansion error that depends on the inventory penalty ψ; and this optimisation model considers no liquidating cost.

Algorithm 2 A thinning algorithm to simulate the processes $M_t^{\pm}$.

Initialize $t_0 = 0$ and $n = 0$.
while $t_n < T$ **do**
  Set $s^- = 0$ and $s^+ = 0$.
  while true do
    Draw $E^- \sim \exp(\lambda^-_{t_n})$ and $U^- \sim \mathrm{Unif}(0, \lambda^-_{t_n})$.
    Update $s^- = s^- + E^-$.
    Set $\lambda^-_{new} = \left(\frac{\lambda^-_{t_n}+\lambda^+_{t_n}}{2} - \frac{\theta^-+\theta^+}{2(1-\kappa)}\right) \mathrm{e}^{-\beta(1-\kappa)s^-} + \left(\frac{\lambda^-_{t_n}-\lambda^+_{t_n}}{2} - \frac{\theta^--\theta^+}{2(1+\kappa)}\right) \mathrm{e}^{-\beta(1+\kappa)s^-} + \frac{\theta^-+\theta^+}{2(1-\kappa)} + \frac{\theta^--\theta^+}{2(1+\kappa)}$.
    if $U^- \le \lambda^-_{new}$ **then**
      break
    end if
  end while
  while true do
    Draw $E^+ \sim \exp(\lambda^+_{t_n})$ and $U^+ \sim \mathrm{Unif}(0, \lambda^+_{t_n})$.
    Update $s^+ = s^+ + E^+$.
    Set $\lambda^+_{new} = \left(\frac{\lambda^-_{t_n}+\lambda^+_{t_n}}{2} - \frac{\theta^-+\theta^+}{2(1-\kappa)}\right) \mathrm{e}^{-\beta(1-\kappa)s^+} - \left(\frac{\lambda^-_{t_n}-\lambda^+_{t_n}}{2} - \frac{\theta^--\theta^+}{2(1+\kappa)}\right) \mathrm{e}^{-\beta(1+\kappa)s^+} + \frac{\theta^-+\theta^+}{2(1-\kappa)} - \frac{\theta^--\theta^+}{2(1+\kappa)}$.
    if $U^+ \le \lambda^+_{new}$ **then**
      break
    end if
  end while
  Set $t_{n+1} = t_n + \min(s^-, s^+)$.
  if $s^- \le s^+$ **then**
    Set $\lambda^-_{t_{n+1}} = \lambda^-_{new} + \eta$ and $\lambda^+_{t_{n+1}} = \frac{\theta^-+\theta^+}{2(1-\kappa)} - \frac{\theta^--\theta^+}{2(1+\kappa)} + \left(\frac{\lambda^-_{t_n}+\lambda^+_{t_n}}{2} - \frac{\theta^-+\theta^+}{2(1-\kappa)}\right) \mathrm{e}^{-\beta(1-\kappa)s^-} - \left(\frac{\lambda^-_{t_n}-\lambda^+_{t_n}}{2} - \frac{\theta^--\theta^+}{2(1+\kappa)}\right) \mathrm{e}^{-\beta(1+\kappa)s^-} + \nu$.
  else
    Set $\lambda^+_{t_{n+1}} = \lambda^+_{new} + \eta$ and $\lambda^-_{t_{n+1}} = \frac{\theta^-+\theta^+}{2(1-\kappa)} + \frac{\theta^--\theta^+}{2(1+\kappa)} + \left(\frac{\lambda^-_{t_n}+\lambda^+_{t_n}}{2} - \frac{\theta^-+\theta^+}{2(1-\kappa)}\right) \mathrm{e}^{-\beta(1-\kappa)s^+} + \left(\frac{\lambda^-_{t_n}-\lambda^+_{t_n}}{2} - \frac{\theta^--\theta^+}{2(1+\kappa)}\right) \mathrm{e}^{-\beta(1+\kappa)s^+} + \nu$.
  end if
  Update $n = n + 1$.
end while
return $\{t_k\}_{1 \le k \le n-1}$.

Table B.9: Estimation results for $\theta^+, \theta^-, \eta, \nu, \beta$: mean(1), standard deviation(2), skewness(3), kurtosis(4)

		2008					2018				
		θ^+	θ^-	η	ν	β	θ^+	θ^-	η	ν	β
IBM	1	0.245	0.246	304.0	119.3	1287	0.071	0.075	224.4	27.39	749.0
	2	0.202	0.189	106.5	48.15	345.7	0.068	0.075	115.4	10.53	277.8
	3	0.978	1.176	0.344	0.506	1.528	2.027	2.479	1.990	0.905	1.175
	4	4.876	5.388	2.381	2.333	4.572	8.076	12.52	8.451	4.689	4.935
Chevron	1	0.375	0.381	306.9	117.4	1322	0.083	0.084	214.6	29.69	767.4
	2	0.255	0.247	143.1	48.15	394.5	0.069	0.060	99.95	12.05	270.6
	3	0.608	0.680	1.096	0.374	1.553	2.790	1.243	0.893	1.088	0.644
	4	2.699	2.804	4.410	3.092	4.690	19.59	5.908	3.594	5.401	3.072
Apple	1	1.266	1.253	859.9	179.9	1984	0.477	0.493	3447	434.5	7030
	2	0.787	0.767	335.1	127.1	794.4	0.417	0.389	1796	354.6	4261
	3	0.146	-0.060	-0.420	0.218	-0.453	1.005	0.661	4.262	5.244	7.452
	4	3.945	3.452	3.060	1.439	2.583	3.921	2.857	33.31	44.38	85.26
Amazon	1	0.373	0.372	510.8	154.4	1267	0.257	0.246	3089	181.1	5349
	2	0.434	0.428	318.4	126.9	815.6	0.274	0.252	1311	104.3	2077
	3	0.772	0.785	-0.341	-0.084	-0.223	1.429	1.818	0.933	1.495	0.687
	4	3.419	3.410	1.970	1.247	1.987	5.224	7.994	3.835	5.908	3.266
JP Morgan	1	0.688	0.680	2941	674.0	8025	0.189	0.200	156.0	30.13	684.4
	2	0.572	0.570	4176	2003	11040	0.159	0.170	53.14	11.23	193.9
	3	1.316	1.330	4.659	8.135	5.106	1.465	1.564	1.454	2.505	1.152
	4	5.791	5.831	31.18	84.65	40.18	5.176	6.166	7.062	17.26	5.192
Microsoft	1	0.497	0.501	11191	878.4	12239	0.375	0.383	4849	722.8	9146
	2	1.081	1.080	14280	3533	13938	0.465	0.485	2895	377.4	4884
	3	1.791	1.797	3.610	11.36	3.198	1.569	1.5798	2.964	1.029	2.612
	4	7.138	7.193	19.34	154.1	14.86	5.532	5.750	16.04	4.500	14.22
Average		0.574	0.572	2703	356.8	4417	0.242	0.247	1998	237.3	3957

Table C.10: The chosen order arrival model parameters for λ and c for comparison test with the asymptotic solution.

	Parameters for λ					Parameters for c				
	$\theta^{\pm}$	η	ν	β	κ	α	η_c	ν_c	ξ	κ_c
Case (i)	1	40	10	100	-	1	10	25	45	-
Case (ii)	1	20	5	60	-	1	5	17.5	45	-
Case (iii)	2	20	5	60	-	2	5	17.5	45	-
Case (iv)	1	60	15	150	-	1	10	25	45	-

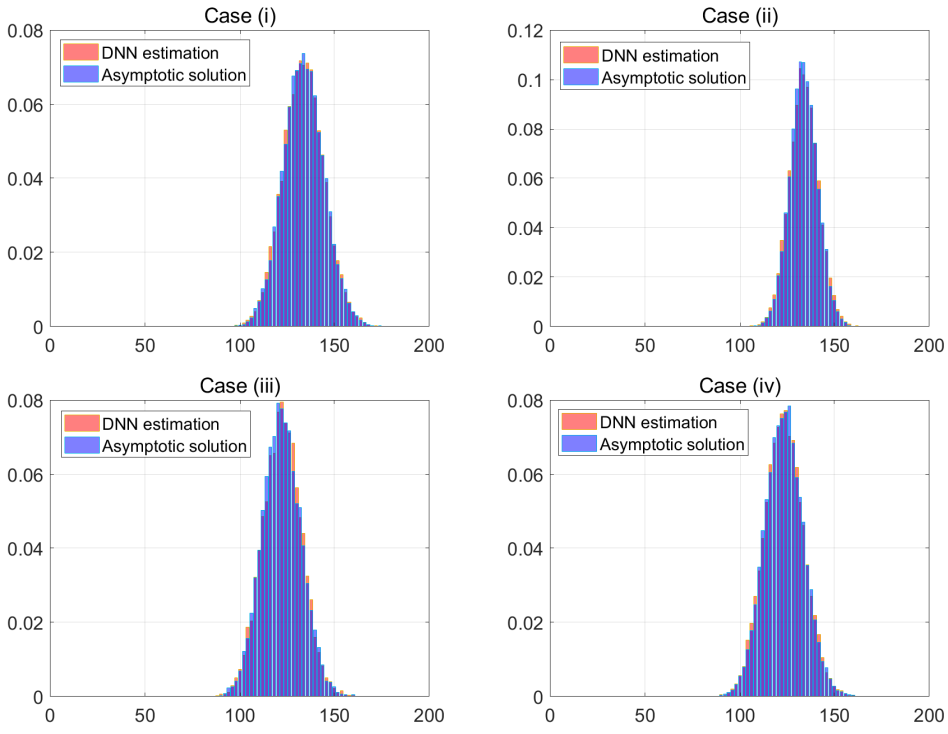

Figure C.13: The market-maker's PnL distributions under taking the optimal strategies with the DNN-estimation and the Cartea's asymptotic formula for each parameter case specified in Table C.10.

Table C.11: The expected PnLs using the optimal strategy derived with the DNN-estimation and Cartea's asymptotic solution; and their absolute difference (top). The KW test p-values under the null hypothesis H_0 that the two samples are from the same distribution, and the test result of rejecting H_0 at the 5% significance level (bottom).

Expected PnL	Case (i)	Case (ii)	Case (iii)	Case (iv)
DNN-estimation	133.6656	133.8209	122.2420	122.3384
Asymptotic solution	133.7425	133.7761	122.1516	122.4094
Difference	0.0769	0.0448	0.0904	0.0710
KW test at a 5%				
p-value	0.5391	0.5729	0.2390	0.4963
H_0	Not reject	Not reject	Not reject	Not reject