

Latency and Liquidity Risk

Álvaro Cartea^{a,b}, Sebastian Jaimungal^c, Leandro Sánchez-Betancourt^a

^a*Mathematical Institute, University of Oxford*

^b*Oxford-Man Institute of Quantitative Finance*

^c*Department of Statistical Sciences, University of Toronto*

Abstract

Latency (i.e., time delay) in electronic markets affects the efficacy of liquidity taking strategies. During the time liquidity takers process information and send marketable limit orders (MLOs) to the exchange, the limit order book (LOB) might undergo updates, so there is no guarantee that MLOs are filled. We develop a latency-optimal trading strategy that improves the marksmanship of liquidity takers. The interaction between the LOB and MLOs is modelled as a marked point process. Each MLO specifies a price limit so the order can receive worse prices and quantities than those the liquidity taker targets if the updates in the LOB are against the interest of the trader. In our model, the liquidity taker balances the tradeoff between the costs of missing trades and the costs of walking the book. In particular, we show how to build cost-neutral strategies, that on average, trade price improvements for fewer misses. We employ techniques of variational analysis to obtain the price limit of each MLO the agent sends. The price limit of a MLO is characterized as the solution to a class of forward-backward stochastic differential equations (FBSDEs) driven by random measures. We prove the existence and uniqueness of the solution to the FBSDE and numerically solve it to illustrate the performance of the latency-optimal strategies.

Keywords: Marked point processes, high-frequency trading, algorithmic trading, latency, forward-backward stochastic differential equations.

1. Introduction

Speed to make decisions and to access the market is a key element in the success of trading strategies in electronic markets. Liquidity providers monitor and update their limit orders (LOs) resting in the limit order book (LOB), and liquidity takers send orders that target LOs. The efficacy of the strategies of the makers and takers of liquidity depends on their latency in the marketplace. Latency is the time delay between an exchange streaming market data to a trader, the trader processing information and making a decision, and the exchange receiving the instruction from the trader. Thus, due to latency, there is no guarantee that liquidity providers can place a LO in

*SJ would like to acknowledge support from the Natural Sciences and Engineering Research Council of Canada (grants RGPIN-2018-05705 and RGPAS-2018-522715). LSB acknowledges support from CONACyT, México and the Mathematical Institute, University of Oxford.

Email addresses: `alvaro.cartea@maths.ox.ac.uk` (Álvaro Cartea), `sebastian.jaimungal@utoronto.ca` (Sebastian Jaimungal), `leandro.sanchezbetancourt@maths.ox.ac.uk` (Leandro Sánchez-Betancourt)

a desired queue position in the book or withdraw a stale quote before it is picked off by another trader.

Furthermore, there are no assurances that marketable limit orders (MLOs) from liquidity takers, which aim at a quantity and price they observed in the LOB, hit the desired target. A MLO is a liquidity taking order for immediate execution against the LOs resting in the book, and each MLO specifies the quantity of the security (e.g., equity, currency pairs, futures, etc.) and a price limit to execute against LOs. A MLO and a market order differ in that the MLO walks the LOB until it reaches the limit price specified by the trader, while a market order walks the LOB until it is filled in full. Due to latency, by the time the exchange processes a MLO, prices and quantities could have improved, so the order is filled at a better price, or prices and quantities could have worsened, so the order is filled if the limit price allows, otherwise the order is rejected.

In this paper, we focus on how latency affects the marksmanship of liquidity takers and we develop a latency-optimal trading strategy that accounts for the time delays in the marketplace. We frame the interaction between the LOB and MLOs as a marked point process (MPP). In our model, the random times of the MPP are the times at which the trade attempts are processed by the exchange, and the marks of the MPP encode the quantity and price updates that LOB undergoes during the latency period – these updates can be in favour or against the interest of the trader. We assume the agent sends fill-or-kill (FoK) MLOs, that is, the orders are either filled in full or rejected.¹ The price limit of the MLO consists of the best quote the agent observes at the time she decides to trade and a discretion to walk the LOB.

The LOB is a moving target, so liquidity takers hit or miss the LOs they are attempting to execute. Everything else being the same, the chances of filling a MLO increase if the agent is willing to receive prices for FoK that are worse than those of the best quotes the agent observes in the LOB when she decides to trade. If the discretion to walk the book is unlimited, the MLO will be filled, but potentially at much worse prices than those of the best quotes the agent observed – i.e., the MLO becomes a market order. On the other hand, if the updates in the LOB are in the interest of the agent, the MLO will be filled at better prices than those of the LOs that the agent targeted.

Traders internalise the costs of missed trades. Missing a trade is costly for several reasons. We provide two examples. i) A broker sends a liquidity taking order to externalise a position. If the trade misses the target price, the broker internalises the position and bears the costs of holding and unwinding the position at a later date. ii) The trade is one leg of a trading strategy that requires multiple executions, so the trader incurs a cost if one of the trades in the strategy is not completed.

In our model, the agent chooses the optimal discretion of the MLOs to minimize the cost of multiple trades over a fixed time horizon (e.g., minutes, hours, days, etc.). The costs are those that arise from walking the LOB and from missed trades. We employ techniques of variational analysis to obtain the optimal discretion for each MLO the agent sends, which we characterize as the solution to a forward-backward stochastic differential equation (FBSDE). We show existence and uniqueness of the solution to the forward and backward parts of the FBSDE and show existence and uniqueness of the solution to the full FBSDE. To the best of our knowledge, existence and uniqueness of the resulting random-measure driven FBSDE is not covered in the extant literature.

¹This is in contrast to an immediate-or-cancel order, which has the property that the order can be partially filled if there is liquidity in the LOB that meets the requirements of the MLO. The unfilled portion of the order is rejected.

We obtain the optimal strategy in closed-form when the cost for missing trades is linear in the number of missed trades (the latency-optimal strategy consists of sending all MLOs with a fixed discretion), and when the cost is quadratic in the number of missed trades, we solve the FBSDE numerically. We illustrate the performance of the latency-optimal strategies for a range of model parameters and examine the tradeoff between costs from walking the book and costs from missed trades. Finally, we discuss strategies that are cost-neutral to the agent. That is, the latency-optimal strategy is devised so the expected costs from walking the book to fill MLOs when the LOB moves against the agent's interests is the same as the expected benefits (i.e., negative costs) from executing trades at better prices than the ones the agent targets.

In the extant literature, there are lines of research that focus on various aspects of latency in electronic markets. In [Moallemi and Sağlam \(2013\)](#), the objective is to quantify the cost of latency on transaction costs, while our paper develops a latency-optimal trading strategy that improves the marksmanship of liquidity takers. In another strand of the literature, the work of [Stoikov and Waeber \(2016\)](#) shows how to execute a large order in electronic markets by employing the volume imbalance of the LOB to predict price changes and study the effect of latency in the efficacy of the execution strategy. [Lehalle and Mounjid \(2017\)](#) employ data from Nasdaq-Omx and also find that as latency increases, the informational content in the volumes of the LOB diminishes.

[Cartea and Sánchez-Betancourt \(2021\)](#) employ proprietary foreign exchange data to show how latency and volatility of the midprice of the security affect the fill ratio of liquidity taking strategies. The authors show how traders could employ latency-optimal strategies to improve fill ratios, while minimizing costs, and they show how to compute the shadow price of latency in foreign exchange markets. [Gao and Wang \(2018\)](#) use Markov decision processes to model the problem of a market maker with latency who trades in a LOB, where the size of the quoted spread is always one tick. The authors find that as latency increases, the profits from making markets decrease.

Recent literature on high-frequency trading and algorithmic trading discusses various characteristics of trading and how traders use speed to obtain informational advantages, see e.g., [Lehalle and Laruelle \(2013\)](#). Other strands of the literature discuss the relationship of market quality, the speed of market participants, and stochastic liquidity, see for example [Almgren \(2012\)](#) and [Guéant \(2016\)](#) for trading in illiquid markets. [Barger and Lorig \(2019\)](#) model the rapid updates of the best quotes in the LOB to propose a model of stochastic price impact.

In what follows, the problem is formulated from the point of view of a trader who wishes to exchange price improvements for fewer misses. Our framework accounts for both buy and sell orders in a unified manner. We disregard the motives that led the trader to send the trade attempt in the first place and focus on the tradeoff between the extra cost to obtain a trade and the number of missed attempts. The problem we study can also be thought from the point of view of a firm-liquidity exchange, or a liquidity provider. Here, the exchange (or liquidity provider) receives trade attempts from a trader at random times, and devises the latency-optimal discretion (or tolerance) to attach to the trader's orders to exchange the trader's price improvements for fewer misses. From the point of view of the exchange, the number of missed trades and the extra cost of the strategy are the two variables of interest.

The remainder of the paper proceeds as follows. Section 2 proposes the agent's performance criterion and characterizes the latency-optimal strategy as the solution to a FBSDE. Section 3 shows existence and uniqueness of the solution of the forward and backward part of the FBSDE, and existence and uniqueness of the solution to the full FBSDE. Section 4 discusses the performance of the strategy for various scenarios. We conclude in Section 5 and collect some proofs in the

Appendix.

2. Optimal discretion to walk the book

2.1. MLOs: revealed preferences

Most electronic exchanges have order types that are designed to protect liquidity takers and liquidity makers against the frictions that stem from latency. MLOs protect liquidity takers from adverse price movements that occur between the time the trader makes a decision to trade and the time the exchange matches the order with a limit order in the book (if possible). If latency were zero for all liquidity takers, MLOs would not be as useful because traders would take liquidity at the price and quantity they observe in the book.

There are, however, no traders with zero latency in the marketplace. Indeed, each market participant is exposed to their own levels of latency, which are random. There are two components that affect the latency of traders: (i) the idiosyncrasy of each trader (hardware, software, co-location) is a key determinant of the baseline latency that each trader experiences in the market; and (ii) the matching engine becomes overcrowded at random times (the capacity of the exchanges to process messages is limited), and the level of quality and the speed of communication between traders and the exchange are also random. Therefore, the time it takes traders to process information and the time it takes the exchange to process instructions are stochastic and constitute additional layers of random delays that affect the latency of each trader in the marketplace; see [Cartea and Sánchez-Betancourt \(2021\)](#) for empirical evidence in the foreign exchange markets.

Traders reveal their preferences when they choose a type of liquidity taking order for a trade or a sequence of trades whose outcome is contingent on latency. These choices demonstrate that traders balance the cost of completing a trade and the costs of price protection. A trader who must complete a trade without delay will choose a market order to guarantee execution in full – and expose the trade to price movements over the latency period. On the other hand, a trader who can afford to miss the trade if the price is ‘not right’ will choose a MLO; thus, in exchange for price protection, the trader concedes that the trade might not get filled.

Everything else being equal, the higher (lower) is the value of the discretion of the MLO to walk the book, the lower (higher) is the probability that the trade is missed and the order is less (more) protected against adverse price moves. Therefore, the price limit of a MLO determines how a trader balances the costs of missing a trade and the potential adverse move in prices to complete the trade. Missing a trade is costly for a number of reasons: i) The trade may be a leg of a trading strategy with multiple executions, so its success depends on the completion of all or most legs of the trading strategy. ii) A trader (e.g., a broker) needs to hedge a position or internalise it. There is a tradeoff between hedging and unwinding the position at an acceptable price or keeping the position in their own books and unwind it at a later date. iii) If a missed trade is ‘quickly’ tried again, its eventual execution price is likely to be worse than the original target price – a trade is missed because the price moved against the interest of the trader – see Tables 3 to 8 in [Cartea and Sánchez-Betancourt \(2021\)](#), where the authors compute the recovery cost of missed trades in foreign exchange.

In asset classes such as equity, fixed income, foreign exchange, and commodities, the MLOs we describe in this paper are available in exchanges that trade in a limit order book.² Each exchange

²For example, see order types in <https://markets.cboe.com/>, <https://www.londonstockexchange.com/>, www.

has its own rules, but all acknowledge that liquidity takers seek to protect their trades from the risks that arise from latency, so they offer order types where the liquidity taker specifies the price limit and the time-in-force of the order. The time-in-force refers to how long the order is active in the market. The shortest time-in-force is immediate execution and the longest time-in-force is typically for the rest of the trading day – the time-in-force for the MLOs we consider in this paper is immediate execution.

Therefore, the mere existence and the widespread use of MLOs in exchanges shows that liquidity takers balance the costs of missing trades and the costs of walking the book to complete trades.

2.1.1. Empirical evidence: price limits

Here, we employ foreign exchange data to examine the strategy of traders who employ MLOs to take liquidity in LMAX and who trade very frequently. We analyse the trading patterns of various market participants and find that over long periods of times (days, weeks, months), traders do not change the value of the discretion to walk the LOB they specify in their MLOs. In the interest of space, below we discuss the activity of one such trader who sent MLOs with a estimated discretion of 10 ticks for a period of four months, and an estimated discretion of 5 ticks for a period of three months.

Our data set contains a number of features of the trading activity in the LOB of the exchange, including: the time-stamp of when the order was processed by the exchange, the direction of the trade (i.e., buy or sell), and the maximum rate willing to pay for a buy order or the minimum rate willing to receive for a sell order. Our data set does not contain the time-stamp of when the trader sent the MLO to the exchange, nor does it point to the best quote the trader observed when she decided to trade; hence, we cannot compute the discretion specified by the trader in the MLO to walk the book. Instead, we use the slack of the MLO as an estimate of the discretion specified by the trader. Here, the slack of a buy MLO is the difference between the limit rate specified in the MLO and the average rate the order would have payed if the order was filled in full. Similarly, the slack of a sell MLO is the difference between the average rate the order would have received if filled in full and the limit sell rate specified in the MLO. Therefore, if slack is non-negative, the MLO was filled in full, and if slack is negative, the order was either fully or partially rejected by the exchange.

Figure 1 shows the liquidity taking activity of the trader in the currency pair USD/JPY between 5 September 2016 and 30 June 2017. During this period, the trader sent 81,641 buy MLOs and 78,086 sell MLOs, of which 491 buys and 463 sells were fully or partially rejected by the exchange because the best quotes had moved beyond the acceptable limit rate specified in the MLO.

In the figure, the unit of slack is ticks (one tick in the pair USD/JPY is 10^{-3} JPY), and the circles represent the slack of each MLO. The left-hand panel shows the slack of each order in our sample and the solid line represents a moving average of the slack of the last 500 MLOs. Observe that between December 2016 and March 2017 the moving average of slacks strongly suggests that the trader sent MLOs with discretion of 10 ticks. Similarly, from April 2017 to June 2017, the moving average of slacks strongly suggests that the discretion to walk the LOB was 5 ticks. Thus, for extended periods of times, and for a large number of liquidity taking orders, the data strongly suggest that the trader sent all MLOs with a constant discretion to walk the book, i.e., 10 ticks during approximately four months and 5 ticks during approximately three months.

lmax.com/, www.cmegroup.com, www.nasdaq.com. Note that FoK, IoC orders have immediate execution in the time-in-force qualifier of the order.

The middle panel shows the slack of the MLOs sent by the trader on 18 June 2017. During this day, all MLOs sent by the trader (353 buys and 297 sells) were filled in full. We observe that there are three clusters of trades (buys and sells) in which the slacks of the MLOs are more volatile than in the other periods. In the right-hand panel of the figure we focus on one of these clusters and show the best ask rate and the MLOs sent by the trader between 15:59:50 and 16:00:30 – during this time window, all orders were buy MLOs and all were filled. The squares in the panel show the limit rate the trader is willing to accept, i.e., the price they observed when deciding to send the buy MLO plus the discretion to walk the LOB.

Clearly, slack is a noisy estimate of the discretion specified in the MLO because during the latency period the best quote in the LOB might change. For example, at time 15:59:56.590 the exchange processed an MLO with slack of 18 ticks, which is the result of an advantageous (for the trader) drop in the rate and of the discretion of the MLO. We note that during the 30 milliseconds before the buy MLO was processed, the USD/JPY exchange rate dropped by approximately 8 ticks. This drop in the rate plus the discretion of 10 ticks, which is our estimate of the discretion of the MLO, accounts for the 18 ticks of slack we observe.³

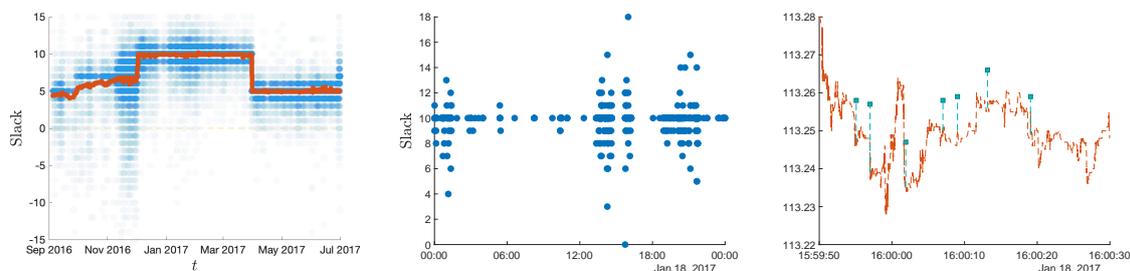


Figure 1: MLOs for a trader in the USD/JPY currency pair. Period: 5 September 2016 to 30 June 2017. The blue circles denote slack of MLOs. The continuous line in left-hand panel is moving average of slack over last 500 trades. The squares in the right-hand panel denote the price limit of the buy MLO and the line denotes the best ask rate in the LOB of the exchange.

2.2. LOB: moving target

Liquidity takers in electronic markets face a moving target problem. Traders send orders that target a price and quantity they observe in the LOB, but due to latency, when the order arrives in the exchange, the target could have moved. If prices and quantities worsen, the agent's order is rejected, and if prices and quantities improve or do not worsen, the order is filled.

We frame the moving target problem as a finite activity MPP $\mathcal{N} = \{(T_n, Z_n)\}_{(n \geq 1)}$ in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Here, (T_n) is an increasing sequence of random points in $(0, T] \cup \{\infty\}$, which represent the times when the agent sends MLOs to the exchange, and (Z_n) is a sequence of marks in \mathbb{R} , which represent the shock to the average price per share due to changes in prices, quantities, and temporary price impact.

We assume that each order is for one unit of the security or for a lot of securities, where the lots have a fixed size throughout the trading horizon. When the volume of the MLO is in lots of the

³In [Cartea and Sánchez-Betancourt \(2021\)](#), the authors analysed in detail one foreign exchange trader and found that their mean latency in the market place is approximately 20.02 milliseconds with 11.31 milliseconds of standard deviation.

security, the mark Z represents a shock to the LOB commensurate with the volume of the MLO.⁴

Let $T \in (0, \infty)$ denote a fixed time horizon. The filtration $(\mathcal{F}_t)_{t \geq 0}$ is generated by \mathcal{N} and is the smallest filtration such that for each n , the point T_n is a stopping time and the mark Z_n is \mathcal{F}_{T_n} -measurable.

The random measure associated with \mathcal{N} is

$$p(dz, dt) = \sum_{n \geq 1} \mathfrak{D}_{(Z_n, T_n)}(dz, dt),$$

where \mathfrak{D} denotes the Dirac measure, and we assume that

$$\mathbb{E} \left[(p(\mathbb{R}, [0, T]))^2 \right] < \infty \quad \text{and} \quad \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} |z| p(dz, dt) \right] < \infty. \quad (1)$$

We denote by \tilde{p} the predictable compensator of the random measure p with decomposition

$$\tilde{p}(dz, dt) = \Phi_t(dz) dA_t, \quad (2)$$

see Definition 8.8.2.1 in [Jeanblanc et al. \(2009\)](#). Here, the compensator has the property that for $q := p - \tilde{p}$, the stochastic integral $(H \star q)_t = \int_0^t \int_{\mathbb{R}} H_s q(dz, ds)$ is a martingale for any integrable and predictable process H . In (2), the predictable process $(A_t)_{t \in \mathfrak{T}}$, where $\mathfrak{T} := [0, T]$, is the compensator of the counting process of the MLOs, which we denote by $N_t := p(\mathbb{R} \times [0, t])$.

We make the following two assumptions throughout the analysis.

Assumption 1. *The process A admits a bounded stochastic intensity so that we may write $A_t = \int_0^t \lambda_u du$ for a predictable process $(\lambda_t)_{t \in \mathfrak{T}}$ and $\exists \bar{\lambda} \in \mathbb{R}$, such that $\forall (t, \omega) \in [0, T] \times \Omega$, $0 < \lambda_t(\omega) \leq \bar{\lambda}$.*

Assumption 2. *The density function ϕ_t of the marks has support in \mathbb{R} , is bounded, and its cumulative distribution function Φ_t is uniformly Lipschitz in $[0, T] \times \Omega$ with Lipschitz constant k .*

Let $(\delta_t)_{t \in \mathfrak{T}}$ be a predictable process that specifies the cash per unit of the security (or lots of the security) the agent is willing to walk the LOB to increase the chances of filling her liquidity taking order, i.e., δ is the discretion of the MLO. For example, in equity markets, if the agent sends a buy order to lift the offer at the best ask a_t , the discretionary amount δ_t is the extra cash per share the order may walk the book, i.e., $a_t + \delta_t$ is the highest price the agent is willing to pay for one share of equity. Similarly, if the agent sends a sell order to hit the best bid b_t , the amount δ_t is the cash discount per share the order may walk the book, i.e., $b_t - \delta_t$ is the lowest price the agent is willing to accept to sell one share of equity.

In the examples above, the best bid and best ask prices (b_t and a_t) refer to those the agent ‘observes’ when she decides to trade, but due to latency, these prices could be stale. In addition,

⁴Note that modelling these shocks is much more effective than modelling the best bid/ask prices and volumes. If one models the latter, then, the framework gets more involved and we do not gain any further insights. Let us consider this point in a little more detail. Suppose that a trader makes a decision to buy (sell) at τ , and the average price of the transaction at time τ is a_τ (b_τ), after hypothetically walking the LOB if necessary. Let $\tilde{\tau} > \tau$ be the time at which the trade attempt is processed by the exchange, and $a_{\tilde{\tau}}$ ($b_{\tilde{\tau}}$) be the average price of the transaction at time $\tilde{\tau}$ if it were to be executed in full. In our model, the shock Z is $a_{\tilde{\tau}} - a_\tau$ for a sell attempt ($b_\tau - b_{\tilde{\tau}}$ for a buy attempt). Note that Z encodes information of how prices and volumes improved/deteriorated from τ to $\tilde{\tau}$ without modeling the best quotes and volumes. This helps to keep the framework tractable, and to treat buy and sell attempts in a unified manner.

by the time the exchange processes the order of the agent, prices and quantities in the LOB could have borne further updates. Price changes could be against or in favour of the agent's interest. When the price per unit of the security moves against the interest of the agent, the order is filled only if the discretion δ of the MLO is enough to cover the adverse change in price and quantity; we refer to this as a price deterioration. On the other hand, if the price per unit of the security moves in favour of the agent's trade interest, the order is filled at a better price; we refer to this as a price improvement.

2.3. Tradeoff: cost of walking the LOB and cost of missed trades

2.3.1. Cost I: discretion to walk the book

Here, we discuss the extra costs incurred by the trader when the MLOs walk the book. We assume that the volume of each MLO is equal to one unit of the security – the costs for MLOs where volume is in lots of the security are computed in a similar way. For buy orders, the contribution to the cost of the strategy is the cash the agent pays for the security minus the price on the offer side of the LOB that the agent targets. Similarly, for sell orders, the contribution to the cost of the strategy is the target price in the bid side of the LOB minus the cash received for the security. That is, for a purchase (sell) order, the contribution to the cost of the strategy is the cash paid (not received) to walk the LOB, which is zero if the order is not executed. We denote the controlled cost process that arises from walking the book by $C^\delta = (C_t^\delta)_{t \in \mathfrak{T}}$, which is given by

$$C_t^\delta = \int_0^t \int_{\mathbb{R}} z \hat{G}(\delta_s - z) p(dz, ds). \quad (3)$$

For a FoK, the function $\hat{G}(x)$ is the step function: $\hat{G}(x) = 1$ if $x \geq 0$ and $\hat{G}(x) = 0$ otherwise. For an IoC, the function $\hat{G}(x)$ is a smooth version of the step function. From this point forward, we assume that the MLOs are FoK orders. All proofs in this paper work for the smooth version of \hat{G} .

The walk-the-book cost for each filled trade is $z \hat{G}(\delta_s - z)$, which can be negative (price improvement), positive (price deterioration), or zero. This cost is negative when the shock to the LOB is negative ($z < 0$), in which case the order is filled at a better price than that targeted by the agent – the price improvement is $|z|$. On the other hand, this cost is positive when the shock to the LOB is positive ($z > 0$), in which case the order is filled (because $\delta \geq z$) at a worse price than that targeted by the agent – the price deterioration is z . Finally, when the shock to the LOB is zero ($z = 0$) or the trade is missed, the cost is zero.

2.3.2. Cost II: missed trades

We denote by $D^\delta = (D_t^\delta)_{t \in \mathfrak{T}}$ the controlled number of missed trades, which is given by

$$D_t^\delta = \int_0^t \int_{\mathbb{R}} G(\delta_s - z) p(dz, ds), \quad (4)$$

and where $G(x) = 1 - \hat{G}(x)$. Recall that the MLO is for one unit of the security or for lots of the security, which are of fixed size throughout the trading horizon. In the latter case, the number of misses is in lots of the security.

The cost of missing D_t^δ trades is given by the quadratic cost function:

$$\alpha D_t^\delta + \gamma \left(D_t^\delta\right)^2, \quad (5)$$

where $\alpha \geq 0$ and $\gamma \geq 0$ are cost parameters that are specific to each trader.

In our framework, we can employ any cost function that is increasing, differentiable, and convex in the number of missed trades. However, these alternative cost functions would make the control problem less tractable and we would not gain further insights into the trader's strategy.

2.4. Performance criterion

The agent's performance criterion is

$$J(\delta) = \mathbb{E} \left[C_T^\delta + \alpha D_T^\delta + \gamma \left(D_T^\delta \right)^2 \right], \quad (6)$$

where T is the terminal date of the trading horizon and the set of admissible strategies is the reflexive Banach space

$$\mathcal{A} := \left\{ \delta = (\delta_t)_{t \in \mathfrak{T}} \mid \delta \text{ is } \mathcal{F}\text{-predictable \& } \mathbb{E} \left[\int_0^T (\delta_t)^2 dt \right] < \infty \right\}, \quad (7)$$

equipped with the norm

$$\|\delta\|^2 = \mathbb{E} \left[\int_0^T (\delta_t)^2 dt \right].$$

Throughout the paper, all spaces are vector spaces because we employ the quotient space with respect to the kernel of the given norm.

The agent wishes to find a control $\delta^* \in \mathcal{A}$ that minimizes the performance criterion (6), that is, the agent solves the problem

$$\delta^* = \underset{\delta \in \mathcal{A}}{\operatorname{argmin}} J(\delta).$$

Note that $J(\delta) < \infty$ because $G \leq 1$ and (1) holds. We choose the units of the cost parameters α, γ , so that the performance criterion has the same units as those of the costs C .

In the performance criterion, everything else being equal, an increase in the value of the cost parameters makes the strategy post orders with higher discretion to walk the LOB. In the extreme case where one of the cost parameters is arbitrarily large, the optimal strategy is to post orders with discretion to walk the LOB as deep as necessary to fill the trades, i.e., the MLO with infinite discretion is a market order.

2.5. Variational Analysis Approach

We employ techniques of variational analysis to obtain the optimal discretion strategy. For ease of presentation, we write

$$J(\delta) = J^C(\delta) + \alpha J^{\text{LP}}(\delta) + \gamma J^{\text{QP}}(\delta), \quad (8)$$

where $J^C(\delta) = \mathbb{E} [C_T^\delta]$, $J^{\text{LP}}(\delta) = \mathbb{E} [D_T^\delta]$, and $J^{\text{QP}}(\delta) = \mathbb{E} \left[(D_T^\delta)^2 \right]$.

Proposition 1. *The following equation holds*

$$J^C(\delta) = \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} z \hat{G}(\delta_t - z) p(dz, ds) \right] = \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} z \hat{G}(\delta_t - z) \Phi_t(dz) dA_t \right]. \quad (9)$$

Proof. Follows from the definition of the predictable compensator. ■

The next proposition provides expressions for $J^{LP}(\delta)$ and $J^{QP}(\delta)$.

Proposition 2. *The following equations hold*

$$J^{LP}(\delta) = \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} G(\delta_t - z) \Phi_t(dz) dA_t \right], \quad \text{and} \quad (10)$$

$$J^{QP}(\delta) = \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \left(2 D_{t-}^{\delta} G(\delta_t - z) + G(\delta_t - z) \right) \Phi_t(dz) dA_t \right]. \quad (11)$$

Proof. Equation (10) follows from the predictability of the integrand. Next, we show (11). The number of missed trades D_t^{δ} satisfy the SDE

$$dD_t^{\delta} = \int_{\mathbb{R}} G(\delta_t - z) p(dz, dt).$$

Let $h(x) = x^2$ and use an integration formula (see Jeanblanc et al. (2009)) to write

$$dh(D_t^{\delta}) = \int_{\mathbb{R}} \left(h \left(D_{t-}^{\delta} + G(\delta_t - z) \right) - h \left(D_{t-}^{\delta} \right) \right) p(dz, dt).$$

Then,

$$\begin{aligned} d \left(D_t^{\delta} \right)^2 &= \int_{\mathbb{R}} \left(2 D_{t-}^{\delta} G(\delta_t - z) + G^2(\delta_t - z) \right) p(dz, dt) \\ &= \int_{\mathbb{R}} \left(2 D_{t-}^{\delta} G(\delta_t - z) + G(\delta_t - z) \right) p(dz, dt), \end{aligned}$$

where the second equality holds because $G^2 = G$. Integrate from zero to T , take expectations, and because the integrand $2 D_{t-}^{\delta} G(\delta_t - z) + G(\delta_t - z)$ is predictable, obtain

$$\begin{aligned} \mathbb{E} \left[\left(D_T^{\delta} \right)^2 \right] &= \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \left(2 D_{t-}^{\delta} G(\delta_t - z) + G(\delta_t - z) \right) p(dz, dt) \right] \\ &= \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \left(2 D_{t-}^{\delta} G(\delta_t - z) + G(\delta_t - z) \right) \tilde{p}(dz, dt) \right]. \end{aligned}$$

■

2.5.1. Optimal discretion to walk the LOB

We employ Gâteaux derivatives to obtain the latency-optimal strategy that minimizes the performance criterion of the agent. Let $w, \delta \in \mathcal{A}$. The directional derivative of J at δ in the direction of w is given by

$$\langle \mathcal{D} J(\delta), w \rangle = \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} [J(\delta + \epsilon w) - J(\delta)], \quad (12)$$

when the limit exists. Now, let \mathcal{A}' be the dual space of \mathcal{A} . If there is $A' \in \mathcal{A}'$ such that $\langle \mathcal{D} J(\delta), w \rangle = A'(w)$ for all $w \in \mathcal{A}$, then A' is called the Gâteaux derivative of J at δ . In this paper, the directional derivatives are elements of the dual of \mathcal{A} , hence we refer to the directional derivatives as Gâteaux derivatives.

Lemma 1. *The Gâteaux derivative at $\delta \in \mathcal{A}$ in the direction $w \in \mathcal{A}$ of the:*

(a) *cost functional J^C is*

$$\langle \mathcal{D} J^C(\delta), w \rangle = \mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) \delta_t \, dA_t \right];$$

(b) *linear cost functional J^{LP} is*

$$\langle \mathcal{D} J^{LP}(\delta), w \rangle = -\mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) \, dA_t \right];$$

(c) *quadratic cost functional J^{QP} is*

$$\begin{aligned} \langle \mathcal{D} J^{QP}(\delta), w \rangle &= -2 \mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) \mathbb{E}_{t-} \left[\int_t^T \int_{\mathbb{R}} G(\delta_s - z') \tilde{p}(dz', ds) \right] \, dA_t \right] \\ &\quad - 2 \mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) D_{t-}^{\delta} \, dA_t \right] - \mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) \, dA_t \right], \end{aligned}$$

where \mathbb{E}_{t-} is the conditional expectation with respect to the predictable σ -algebra at time t .

Proof. See [Appendix A](#). ■

The next theorem provides the Gâteaux derivative of the performance criterion of the agent and provides a characterization of the optimal discretion to walk the LOB.

Theorem 1. *The Gâteaux derivative of the functional J at $\delta \in \mathcal{A}$ in the direction of $w \in \mathcal{A}$ is*

$$\langle \mathcal{D} J(\delta), w \rangle = \mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) \left(\delta_t - 2\gamma D_{t-}^{\delta} - \gamma - \alpha - 2\gamma \mathbb{E}_{t-} \left[\int_t^T \int_{\mathbb{R}} G(\delta_s - z') \tilde{p}(dz', ds) \right] \right) \, dA_t \right],$$

and vanishes in every direction $w \in \mathcal{A}$ if and only if there is a process $\delta^* \in \mathcal{A}$ such that

$$\delta_t^* = 2\gamma \mathbb{E}_{t-} \left[D_T^{\delta^*} \right] + \gamma + \alpha, \quad (13)$$

almost everywhere in $\mathfrak{T} \times \Omega$.

Proof. By Lemma 1 and the performance criterion (8), the Gâteaux derivative of J vanishes at

$$\begin{aligned} \delta_t^* &= 2\gamma \mathbb{E}_{t-} \left[\int_t^T \int_{\mathbb{R}} G(\delta_s^* - z') \tilde{p}(dz', ds) \right] + 2\gamma \left(D_{t-}^{\delta^*} + \frac{1}{2} \right) + \alpha \\ &= 2\gamma \mathbb{E}_{t-} \left[\int_t^T \int_{\mathbb{R}} G(\delta_s^* - z') p(dz', ds) \right] + 2\gamma \left(D_{t-}^{\delta^*} + \frac{1}{2} \right) + \alpha \\ &= 2\gamma \mathbb{E}_{t-} \left[D_T^{\delta^*} - D_{t-}^{\delta^*} \right] + 2\gamma \left(D_{t-}^{\delta^*} + \frac{1}{2} \right) + \alpha \\ &= 2\gamma \mathbb{E}_{t-} \left[D_T^{\delta^*} \right] + \gamma + \alpha. \end{aligned} \quad (14)$$

Now we show that if the Gâteaux derivative at δ vanishes in every direction w , the control δ satisfies (14). We proceed by contradiction. Suppose there exists $\hat{\delta} \in \mathcal{A}$ such that $\langle \mathcal{D} J(\hat{\delta}), w \rangle = 0$ for all

$w \in \mathcal{A}$ and there is $(\mathbb{T}, \mathfrak{D}) \in \mathcal{B}(\mathfrak{T}) \times \mathcal{F}_T$ with $\mathbb{L}(\mathbb{T}) \mathbb{P}(\mathfrak{D}) > 0$ such that $\hat{\delta}_t(\omega) \neq \delta_t^*(\omega)$ for $t \in \mathbb{T}$, and $\omega \in \mathfrak{D}$, where $\mathbb{L}(\mathbb{T})$ denote the Lebesgue measure of $\mathbb{T} \in \mathcal{B}(\mathfrak{T})$, and $\mathcal{B}(\mathfrak{T})$ is the Borel σ -algebra of \mathfrak{T} . Thus, on $\mathbb{T} \times \mathfrak{D}$ we have

$$\hat{\delta}_t(\omega) - 2\gamma \mathbb{E}_{t-} \left[D_T^{\hat{\delta}} \right] (\omega) - \gamma - \alpha \neq 0.$$

Hence, $w_t = \hat{\delta}_t - 2\gamma \mathbb{E}_{t-} [D_T^{\hat{\delta}}] - \gamma - \alpha$ is predictable and $\mathbb{E}[\sup_{t \in \mathfrak{T}} (w_t)^2] < \infty$. Furthermore, the Gâteaux derivative of $\hat{\delta}$ in the direction of w satisfies the inequality $\langle \mathcal{D}J(\hat{\delta}), w \rangle > 0$, which is a contradiction. Therefore, there is no $(\mathbb{T}, \mathfrak{D}) \in \mathcal{B}(\mathfrak{T}) \times \mathcal{F}$ with $\mathbb{L}(\mathbb{T}) \mathbb{P}(\mathfrak{D}) > 0$ such that $\hat{\delta}_t(\omega) \neq \delta_t^*(\omega)$ for $t \in \mathbb{T}$ and $\omega \in \mathfrak{D}$. ■

Theorem 2. *The control δ^* in (13) is a local minimizer of the agent's performance criterion $J(\delta)$.*

Proof. Recall that the Gâteaux derivative $\langle \mathcal{D}J(\delta^*), w \rangle$ vanishes in every direction $w \in \mathcal{A}$. The second Gâteaux derivative⁵ at $\delta^* \in \mathcal{A}$ in the direction $\nu \in \mathcal{A}$ is

$$\langle \mathcal{D}^2 J(\delta^*), \nu, \nu \rangle = \mathbb{E} \left[\int_0^T \nu_t^2 \phi_t(\delta_t^*) dA_s \right] + \mathbb{E} \left[\left(\int_0^T \nu_t \phi_t(\delta_t^*) dA_s \right)^2 \right]. \quad (15a)$$

This Gâteaux derivative is non-negative. Therefore, δ^* is a local minimum. ■

If the value of the quadratic cost parameter γ is zero, the candidate optimal control in (13) has the simple closed-form expression

$$\delta_t^* = \alpha, \quad (16)$$

which is independent of the number of missed trades.

In our framework, the trading pattern of the trader we analysed in subsection 2.1.1 would correspond to a performance criterion as that in (6) with cost parameters $\gamma = 0$ and $\alpha = 10$ between December 2016 and March 2017, and with cost parameters $\gamma = 0$ and $\alpha = 5$ between April 2017 and June 2017.

One expects, everything else being equal, that the discretion to walk the book is increasing in the number of missed trades. This is true when the value of the cost parameter γ is greater than zero because the marginal cost of missing a trade is positive. Although intuitive, this result is not trivial to show because to determine the optimal discretion one requires fixed point arguments. Note that the expected number of missed trades appears on the right-hand side of (13), which affects the optimal discretion, i.e., determines the left-hand side of (13), which in turn affects the expected number of missed trades that appears on the right-hand side of (13), and so on. In Section 4 we return to this point when we illustrate the performance of the strategy.

When, however, the cost of missing trades is linear (i.e., $\gamma = 0$ and $\alpha \geq 0$) the optimal discretion is independent of the number of missed trades, as shown in (16). This corresponds to the behaviour of traders described in the empirical analysis of subsection 2.1.1. Finally, as discussed above, the value of the cost parameters α and γ is specified by each trader. One can estimate these values from the discretion included in MLOs sent by the trader to the exchange. When traders send MLOs with a fixed discretion to walk the book it is straightforward to estimate the cost parameter α .

⁵See Appendix C for details of the second Gâteaux derivative.

3. Existence and Uniqueness of the FBSDE

To the best of our knowledge, the FBSDE

$$D_t^\delta = \int_0^t \int_{\mathbb{R}} G(\delta_s - z) p(dz, ds), \quad D_0^\delta = 0,$$

$$\delta_t = 2\gamma \mathbb{E}_{t-} \left[D_T^\delta \right] + \gamma + \alpha, \quad \delta_T = 2\gamma D_{T-}^\delta + \gamma + \alpha,$$

has no existence and uniqueness results in the extant literature. By simple inspection we observe that its representation is non-standard – see [Ma and Yong \(2007\)](#); [Carmona \(2016\)](#) for an introduction to the topic. We also note that our FBSDE is driven by random measures, with the forward part being adapted and the backward part being predictable. Finally, the function $G(x)$ –which is one if $x < 0$ or zero otherwise, is not Lipschitz and it is not continuous. Furthermore, although the terminal condition for δ is square integrable, we remark that it is not an input to the FBSDE because the terminal condition is controlled. Hence, we cannot employ the results in [Tang and Li \(1994\)](#) or those in [Quenez and Sulem \(2013\)](#) to study the FBSDE that we derive in this paper. The mathematical tools we use are different. For example, Theorem 2.3 (existence and uniqueness) in [Quenez and Sulem \(2013\)](#) is proved in [Tang and Li \(1994\)](#). Their proof follows from a Martingale representation – see Lemma 2.3 in [Tang and Li \(1994\)](#), whereas our proof follows from fixed point arguments in Banach spaces. Recall that, as stated in [Peng and Wu \(1999\)](#), there are two known methods to study FBSDEs: (i) The probabilistic approach that uses contraction mappings, which may be regarded as an extension of Picard’s iterations, and (ii) the four-step scheme that relies on partial differential equations and stochastic optimal control. Here, we employ the first approach. In this section we prove existence and uniqueness of the solution of the FBSDE. For FBSDEs in a semimartingale setting see [Antonelli \(1993\)](#), whose work was inspired by [Duffie and Epstein \(1992\)](#) and [Jeanblanc-Picque and Pontier \(1990\)](#). In [Duffie and Epstein \(1992\)](#) the authors construct a stochastic utility function to find optimal portfolios. The BSDE they study shares similarities with the one we analyze in this paper.⁶ For fully coupled FBSDEs in the Brownian motion case see [Peng and Wu \(1999\)](#). For an account of Brownian motion and Poisson processes in FBSDEs, see [Zhen \(1999\)](#). [Xia \(2000\)](#), [Confortola and Fuhrman \(2013\)](#), [Confortola et al. \(2016\)](#), and [Bandini \(2016\)](#) study the framework of BSDEs and MPPs. Finally, for the study of FBSDEs that arise from vanishing Gâteaux derivatives in stochastic games stemming from algorithmic trading problems, see [Casgrain and Jaimungal \(2018\)](#) and [Casgrain and Jaimungal \(2020\)](#).

To streamline the results in this section, we define the spaces we use to prove existence and uniqueness of the solution to the FBSDE (13).

Define

$$\mathcal{C}^1 := \left\{ U = (U_t)_{t \in \mathfrak{T}} \mid U \text{ is } \mathcal{F}\text{-adapted and } \mathbb{E} \left[\int_0^T |U_t| dt \right] < \infty \right\},$$

$$\mathcal{C}^\infty := \left\{ U = (U_t)_{t \in \mathfrak{T}} \mid U \text{ is } \mathcal{F}\text{-adapted and } \mathbb{E} \left[\sup_{t \in \mathfrak{T}} |U_t| \right] < \infty \right\},$$

then, the spaces $(\mathcal{C}^\infty, \|\cdot\|_\infty)$, and $(\mathcal{C}^1, \|\cdot\|_1)$ are Banach spaces, where

$$\|\delta\|_\infty = \mathbb{E} \left[\sup_{t \in \mathfrak{T}} |\delta_t| \right] \quad \text{and} \quad \|\delta\|_1 = \mathbb{E} \left[\int_0^T |\delta_t| dt \right].$$

⁶In [Appendix B](#) we return to the paper by [Duffie and Epstein \(1992\)](#) when we prove the existence and uniqueness of the BSDE part of the FBSDE we study.

This is direct consequence of classical results such as: (i) the L^p spaces defined on a measurable spaces are Banach spaces for $1 \leq p < \infty$, see Volume III and IV of [Stein and Shakarchi \(2003\)](#), (ii) the space $(\Omega \times [0, T], \mathcal{F}_T \otimes \mathcal{B}([0, T]), \mathbb{P} \times \mathbb{L})$ is a measurable space, where $\mathcal{B}([0, T])$ is the Borel σ -algebra of $[0, T]$ and \mathbb{L} is the Lebesgue measure, see Chapter 1 in [Cohen and Elliott \(2015\)](#), and (iii) any left (or right) continuous process is progressively measurable, see page 42 in [Revuz and Yor \(2013\)](#). We remark that [Duffie and Epstein \(1992\)](#) use the Banach space $(\mathcal{C}^\infty, \|\cdot\|_\infty)$ in Appendix A to prove existence and uniqueness of a BSDE via the Banach fixed point theorem.

An immediate consequence of the previous statement is that the space $\mathcal{C}^1 \times \mathcal{C}^1$ with norm

$$\|(U, V)\|_{\mathcal{C}^1 \times \mathcal{C}^1} = \|U\|_1 + \|V\|_1$$

is a Banach space.

Next, we transform the FBSDE in (13), and then we proceed to prove the three uniqueness and existence results of this section. By the change of variable $\delta_t = \hat{\delta}_{t-}$, we assert that a solution to

$$\begin{aligned} \delta_t &= 2\gamma \mathbb{E}_{t-} \left[D_T^\delta \right] + \gamma + \alpha, \\ D_t^\delta &= \int_0^t \int_{\mathbb{R}} G(\delta_s - z) p(dz, ds), \quad D_0^\delta = 0, \end{aligned} \tag{17}$$

with δ being \mathcal{F} -predictable and D^δ being \mathcal{F} -adapted, exists and is unique, if a solution to

$$\begin{aligned} \hat{\delta}_t &= 2\gamma \mathbb{E}_t \left[D_T^{\hat{\delta}} \right] + \gamma + \alpha, \\ D_t^{\hat{\delta}} &= \int_0^t \int_{\mathbb{R}} G(\hat{\delta}_{s-} - z) p(dz, ds), \quad D_0^{\hat{\delta}} = 0, \end{aligned} \tag{18}$$

exists and is unique, with $\hat{\delta}$ and $D^{\hat{\delta}}$ \mathcal{F} -adapted. Finally, by a further change of variables, $\tilde{\delta}_t = \hat{\delta}_t - 2\gamma D_t^{\hat{\delta}}$, a solution to (17) exists and is unique if a solution to

$$\begin{aligned} \tilde{\delta}_t &= 2\gamma \mathbb{E}_t \left[\int_t^T \int_{\mathbb{R}} G(\tilde{\delta}_{s-} + 2\gamma D_{s-}^{\tilde{\delta}} - z) p(dz, ds) \right] + \gamma + \alpha, \\ D_t^{\tilde{\delta}} &= \int_0^t \int_{\mathbb{R}} G(\tilde{\delta}_{s-} + 2\gamma D_{s-}^{\tilde{\delta}} - z) p(dz, ds), \quad D_0^{\tilde{\delta}} = 0, \end{aligned} \tag{19}$$

exists and is unique, with $\tilde{\delta}$ and $D^{\tilde{\delta}}$ adapted.

To analyse solutions to the FBSDE (19), we study the fixed points of the functional

$$\Upsilon(U, V)_t = \begin{pmatrix} H(U, V)_t \\ I(U, V)_t \end{pmatrix} = \begin{pmatrix} 2\gamma \mathbb{E}_t \left[\int_t^T \int_{\mathbb{R}} G(U_{s-} + 2\gamma V_{s-} - z) p(dz, ds) \right] + \gamma + \alpha \\ \int_0^t \int_{\mathbb{R}} G(U_{s-} + 2\gamma V_{s-} - z) p(dz, ds) \end{pmatrix}, \tag{20}$$

and, for completeness, prove existence and uniqueness of the solution of: (i) the backward part of the FBSDE; (ii) the forward part of the FBSDE; and (iii) the full FBSDE – a result which we derive independently from the existence of the backward and forward parts of the FBSDE. For the proofs of (i) and (ii) we refer the reader to [Appendix B](#).

The next theorem shows the existence and uniqueness of the solution to the FBSDE (19).

Theorem 3. *Assumptions 1 and 2 hold. Then, if the parameters $k, \bar{\lambda}, T, \gamma$ are such that*

$$k T \bar{\lambda} (\max\{1, 2\gamma\})^2 < 1, \quad (21)$$

there exists a unique solution to the FBSDE

$$\tilde{\delta}_t = 2\gamma \mathbb{E}_t \left[\int_t^T \int_{\mathbb{R}} G(\tilde{\delta}_{s^-} + 2\gamma D_{s^-}^{\tilde{\delta}} - z) p(dz, ds) \right] + \gamma + \alpha, \quad (22a)$$

$$D_t^{\tilde{\delta}} = \int_0^t \int_{\mathbb{R}} G(\tilde{\delta}_{s^-} + 2\gamma D_{s^-}^{\tilde{\delta}} - z) p(dz, ds), \quad D_0^{\tilde{\delta}} = 0. \quad (22b)$$

Proof. Consider the functional $\Upsilon : \mathcal{C}^1 \times \mathcal{C}^1 \rightarrow \mathcal{C}^1 \times \mathcal{C}^1$ defined in (20). We know $\mathcal{C}^1 \times \mathcal{C}^1$ is a Banach space when equipped with the norm

$$\|\Upsilon(U, V)\|_{\mathcal{C}^1 \times \mathcal{C}^1} = \|H(U, V)\|_1 + \|I(U, V)\|_1, \quad \text{where} \quad \|U\|_1 = \mathbb{E} \left[\int_0^T |U_s| ds \right].$$

Let (U, V) and (X, Y) be in $\mathcal{C}^1 \times \mathcal{C}^1$ and write

$$\begin{aligned} \|\Upsilon(U, V) - \Upsilon(X, Y)\|_{\mathcal{C}^1 \times \mathcal{C}^1} &= \mathbb{E} \left[\int_0^T |H(U, V)_t - H(X, Y)_t| dt \right] \\ &\quad + \mathbb{E} \left[\int_0^T |I(U, V)_t - I(X, Y)_t| dt \right]. \end{aligned} \quad (23)$$

The first term on the right-hand side of (23) satisfies the bound

$$\begin{aligned} &\mathbb{E} \left[\int_0^T |H(U, V)_t - H(X, Y)_t| dt \right] \\ &\leq \mathbb{E} \left[\int_0^T 2\gamma \mathbb{E}_t \left[\int_t^T |\Phi(U_{s^-} + 2\gamma V_{s^-}) - \Phi(X_{s^-} + 2\gamma Y_{s^-})| dA_s \right] dt \right] \\ &\leq 2k\gamma\bar{\lambda} \int_0^T \mathbb{E} \left[\int_t^T |U_{s^-} + 2\gamma V_{s^-} - X_{s^-} - 2\gamma Y_{s^-}| ds \right] dt. \end{aligned}$$

The second term on the right-hand side of (23) satisfies the bound

$$\begin{aligned} &\mathbb{E} \left[\int_0^T |I(U, V)_t - I(X, Y)_t| dt \right] \\ &\leq \int_0^T \mathbb{E} \left[\int_0^t \int_{\mathbb{R}} |G(U_{s^-} + 2\gamma V_{s^-} - z) - G(X_{s^-} + 2\gamma Y_{s^-} - z)| \tilde{p}(dz, ds) \right] dt \\ &= \int_0^T \mathbb{E} \left[\int_0^t |\Phi(U_{s^-} + 2\gamma V_{s^-}) - \Phi(X_{s^-} + 2\gamma Y_{s^-})| \lambda_s ds \right] dt \\ &\leq k\bar{\lambda} \int_0^T \mathbb{E} \left[\int_0^t |U_{s^-} + 2\gamma V_{s^-} - X_{s^-} - 2\gamma Y_{s^-}| ds \right] dt. \end{aligned}$$

Now, let $k_1 = k \bar{\lambda} \max\{2\gamma, 1\}$ and $k_2 = k_1 \max\{2\gamma, 1\}$, and write

$$\begin{aligned} \|\Upsilon(U, V) - \Upsilon(X, Y)\|_{\mathcal{C}^1 \times \mathcal{C}^1} &\leq k_1 \int_0^T \mathbb{E} \left[\int_0^T |U_{s^-} + 2\gamma V_{s^-} - X_{s^-} - 2\gamma Y_{s^-}| ds \right] dt \\ &\leq k_1 T \mathbb{E} \left[\int_0^T |U_{t^-} + 2\gamma V_{t^-} - X_{t^-} - 2\gamma Y_{t^-}| dt \right] \\ &\leq k_2 T \mathbb{E} \left[\int_0^T |U_{t^-} - X_{t^-}| dt \right] + k_2 T \mathbb{E} \left[\int_0^T |V_{t^-} - Y_{t^-}| dt \right] \\ &< \|(U, V) - (X, Y)\|_{\mathcal{C}^1 \times \mathcal{C}^1} . \end{aligned}$$

Thus, Υ is a contraction mapping in the Banach space $\mathcal{C}^1 \times \mathcal{C}^1$, so there exists a unique pair of processes U^* and V^* such that $\Upsilon(U^*, V^*) = (U^*, V^*)$. ■

Note that the inequality in (21) involves the upper bound $\bar{\lambda}$ (see Assumption 1) and the jump-size parameter k . Then, for a given cost parameter γ , there are two degrees of freedom to choose the units of $\bar{\lambda}$ and of k , so that the inequality is satisfied.

In all, we have shown that the candidate optimal control in (13) exists and is unique, this was proved in Theorem 3. It is straightforward to see that $\delta^* \in \mathcal{A}$. By definition, the control δ^* is predictable. A short calculation shows

$$\begin{aligned} \mathbb{E} \left[\int_0^T (\delta_t^*)^2 dt \right] &\leq T \mathbb{E} \left[\sup_{0 \leq t \leq T} (2\gamma \mathbb{E}_{t^-} [N_T] + \gamma + \alpha)^2 \right] \\ &= T \mathbb{E} \left[\sup_{0 \leq t \leq T} 2(2\gamma)^2 \left(N_t + \mathbb{E}_{t^-} \left[\int_t^T \lambda_s ds \right] \right)^2 + 2(\gamma + \alpha)^2 \right] \\ &\leq 16T \gamma^2 \mathbb{E} [N_T^2] + 16\gamma^2 \bar{\lambda}^2 T^3 + 2T(\gamma + \alpha)^2 < \infty , \end{aligned}$$

where the last inequality follows from (1). Therefore the control δ^* that satisfies (13) is an element of \mathcal{A} . Finally, uniqueness is given in the sense of the norm $\|\cdot\|_1$, meaning that if we have two processes $\delta^*, \delta^* \in \mathcal{A}$ satisfying (17), then $\|\delta^* - \delta^*\|_1 = 0$.

4. Performance of strategy

The expectation that appears in (13) is conditional on the information \mathcal{F}_{t^-} . Here, we study a slight variation of the FBSDE in (13) and derive a partial-integro differential equation for the optimal control.

To this end, fix the optimal control $\delta^* \in \mathcal{A}$ and define the process $(\check{\delta}_t)_{t \in \mathfrak{T}}$, where

$$\check{\delta}_t = 2\gamma \mathbb{E}_t \left[D_T^{\delta^*} \right] + \gamma + \alpha .$$

Observe that δ^* in (13) is the càglàd (LCRL) version of the càdlàg (RCLL) process $\check{\delta}$, and $\delta_t^* = \check{\delta}_{t^-}$. Define the dynamics of the missed trades D^{δ^*} as a function of the process $\check{\delta}$:

$$D_t^{\check{\delta}} = \int_0^t \int_{\mathbb{R}} G(\check{\delta}_{s^-} - z) \tilde{p}(dz, ds) + \int_0^t \int_{\mathbb{R}} G(\check{\delta}_{s^-} - z) q(dz, ds) ,$$

and recall that $q = p - \tilde{p}$ is the compensated random measure of \mathcal{N} .

Assumption 3. *The stochastic intensity $(\lambda_t)_{t \in \mathcal{T}}$ has the Markov property, and the quadratic co-variation between the processes λ and $D^{\check{\delta}}$ is zero.*

By Assumption 3, we derive the Markov property of $\check{\delta}$, which we use to write $\check{\delta}_t = h(t, D_t^{\check{\delta}}, \lambda_t)$ for a differentiable function h with respect to the first argument. Then the process $D^{\check{\delta}}$ is given by

$$D_t^{\check{\delta}} = \int_0^t \int_{\mathbb{R}} G(h(s, D_{s-}^{\check{\delta}}, \lambda_{s-}) - z) \tilde{p}(dz, ds) + \int_0^t \int_{\mathbb{R}} G(h(s, D_{s-}^{\check{\delta}}, \lambda_{s-}) - z) q(dz, ds),$$

and because $\check{\delta}$ is a martingale, the function h is the solution of a PIDE that we characterize in the following theorem.

Theorem 4. *Let $\check{\delta}_t = h(t, D_t^{\check{\delta}}, \lambda_t)$. When assumptions 1 and 3 hold, and h is once differentiable in the first argument and twice differentiable in the third argument, the function h satisfies the PIDE*

$$0 = \partial_t h(t, D, \lambda) + \mathcal{L}_t^\lambda h(t, D, \lambda) + \left(\int_{h(t, D, \lambda)}^{\infty} \lambda \phi_t(z) dz \right) (h(t, D + 1, \lambda) - h(t, D, \lambda)), \quad (24)$$

with boundary and terminal conditions

$$\lim_{D \rightarrow \infty} h(t, D, \lambda) = \infty \quad \text{and} \quad h(T, D, \lambda) = 2\gamma D + \gamma + \alpha.$$

Here, $\mathcal{L}_t^\lambda h(t, D, \lambda)$ is the infinitesimal generator of the arrival intensity process λ acting on the function h .

Proof. Apply Itô's formula to $\check{\delta}_t = h(t, D_t^{\check{\delta}}, \lambda_t)$ and note that the drift term (i.e., the dt -term) vanishes because $\check{\delta}$ is a martingale. ■

The above result provides us with h and hence $\check{\delta}$. To obtain δ^* , note that h is continuous in time by the following observation

$$h(t, D, \lambda) = 2\gamma D + \gamma + \alpha + \int_t^T \mathbb{E}[(1 - \Phi(h(u, D_u, \lambda_u))) \mid D_t = D, \lambda_t = \lambda] du.$$

Recall $\delta_t^* = \check{\delta}_{t-}$, hence, from continuity in time, we have $\delta_t^* = h(t, D_{t-}^{\check{\delta}}, \lambda_{t-})$.

4.1. Poisson arrival of trades

We solve the PIDE in (24) numerically to illustrate the performance of the latency-optimal strategy. We assume the agent sends MLOs according to a homogeneous Poisson process with intensity $\lambda = 100$, the linear cost parameter is $\alpha = 0$, the quadratic cost parameter γ takes values in $\{0.01, 0.05, 0.15\}$, the marks $Z_n \sim \mathbf{N}(\mu, \sigma^2)$ (price and quantity shocks to the LOB) are iid normal with parameters μ and σ^2 , $n = 1, 2, \dots$, and the trading horizon is $T = 1$ day.

We employ a subset of the data in subsection 2.1.1 to estimate the parameters of the distribution of the marks in the MPP. We focus on the period when the trader sends MLOs with 10 ticks of discretion to walk the LO; i.e., from December 2016 to March 2017. The mean and the standard deviation of slacks are -0.099 and 0.977 , respectively – below, in the numerical studies, we assume that the mean value of the shock $\mu = 0.1$, and the standard deviation of the shock $\sigma = 1$, i.e., we

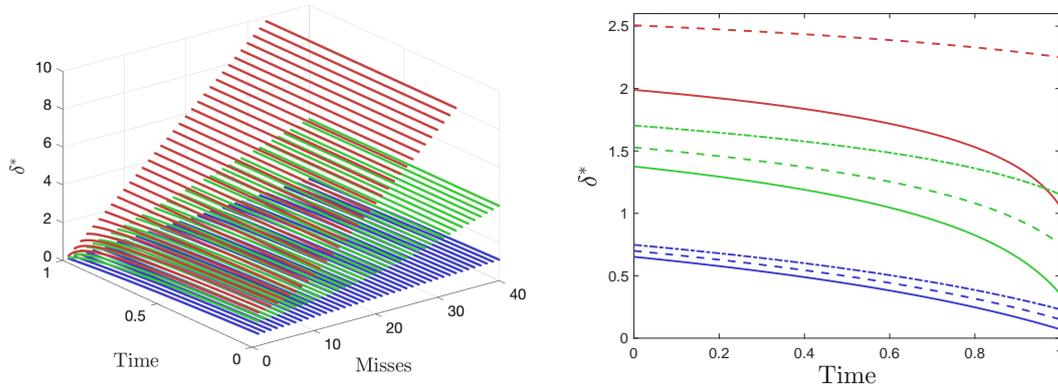


Figure 2: Left panel: Optimal strategy δ^* as a function of time and the number of missed trades for $\gamma = 0.01$ (bottom surface), $\gamma = 0.05$ (middle surface), and $\gamma = 0.15$ (top surface). The remaining parameters are: $\lambda = 100$, $\alpha = 0$, and $Z_n \sim \mathbf{N}(0.1, 1)$ for every n . Right panel: Optimal strategy for various values of missed trades; blue curves are for $D = 4$, green curves are for $D = 8$, and red curves are for $D = 12$.

assume that $Z_n \sim \mathbf{N}(0.1, 1)$. In Section 4.1.1 we explore the results for other values of the cost parameters α and γ .

Figure 2 shows the optimal discretion δ^* as a function of the number of missed trades. The left panel shows three surfaces, one for each value of the quadratic cost parameter γ . The higher the value of the quadratic cost parameter for missing trades, the higher is the optimal discretion employed in the strategy. The right panel shows the optimal discretion when the number of missed trades is $D^{\delta^*} \in \{4, 8, 12\}$, and the quadratic cost parameter is $\gamma \in \{0.01, 0.05, 0.15\}$. Blue denotes cases with $\gamma = 0.01$, green for $\gamma = 0.05$, and red for $\gamma = 0.15$. Solid lines are for $D^{\delta^*} = 4$, dashed lines are for $D^{\delta^*} = 8$, and dash-dotted lines are for $D^{\delta^*} = 12$. Observe that, everything else being equal, as time approaches T , the optimal discretion to walk the LOB decreases because the conditional expectation of the number of misses D_T decreases as there is less time left for new trade attempts.

We perform 100,000 simulations of the agent's trading activity and Figure 3 shows three sample paths. The top panel shows the optimal discretion of the agent's orders and the cumulative costs accrued from walking the book and from receiving price improvements. The bottom panel shows the number of missed trades and the number of trade attempts. Clearly, as the number of missed trades increases (decreases), the optimal strategy is to increase (decrease) the discretion of the MLOs to walk the LOB.

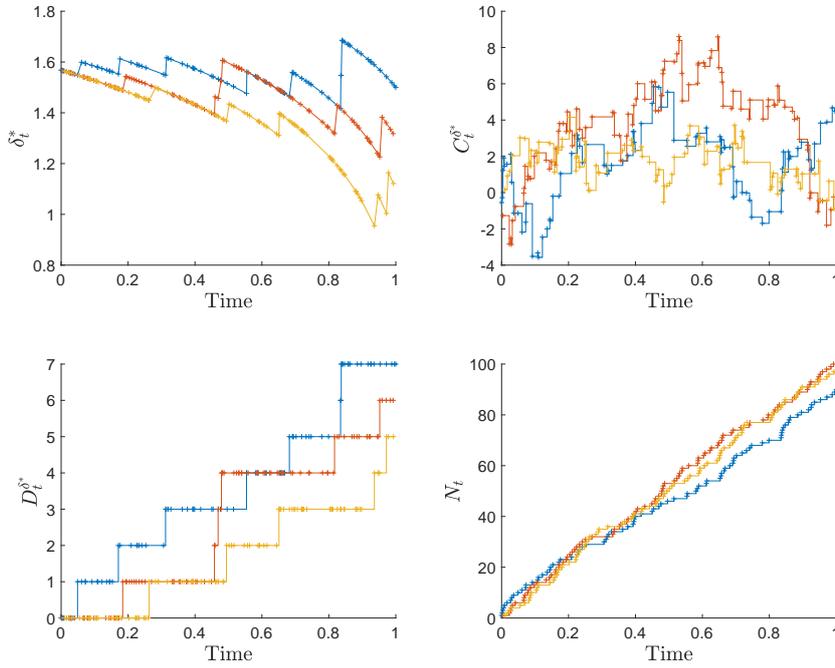


Figure 3: Sample paths for the optimal discretion δ^* (top left panel), number of missed trades D^{δ^*} (lower left panel), cost of strategy C^{δ^*} (top right panel), and number of trade attempts N (lower right panel) for three simulations of the MPP. Parameters: $\alpha = 0$, $\gamma = 0.1$, $\lambda = 100$, $T = 1$.

Figure 4 reports various cost metrics of the optimal strategy for three values of the quadratic cost parameter γ . The top panel shows histograms of the cost incurred by the strategy to fill trades, i.e., $C_T^{\delta^*}$, and the average cost of walking the LOB to fill trades, i.e., $C_T^{\delta^*} / (N_T - D_T^{\delta^*})$. Recall that the cost is negative (positive) when the trade is executed with price improvement (deterioration). The Figure shows that as the value of the quadratic cost parameter increases: (i) the average cost of walking the book to fill trades increases, the total cost increases, and the average number of misses decreases, see bottom panels; (ii) the costs of walking the LOB increase because the strategy fills more orders (i.e., misses fewer trades), see the bottom-left panel. The bottom-right panel shows that the average ratio of missed trades to trade attempts decreases when the cost for missing trades increases.

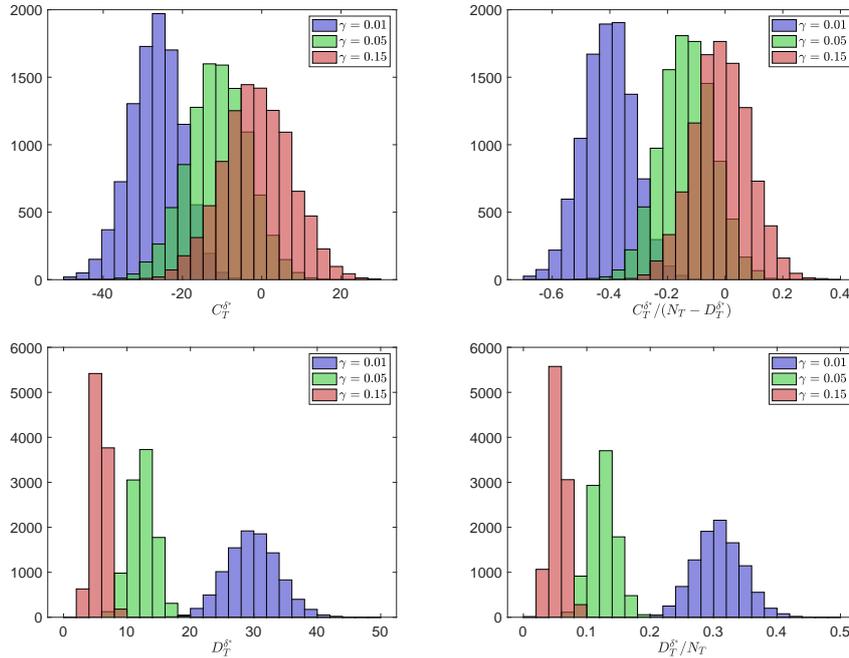


Figure 4: Top left panel: Histogram of the cost $C_T^{\delta^*}$ of the strategy. Top right panel: Histogram of the extra cost per filled trade $C_T^{\delta^*}/(N_T - D_T^{\delta^*})$. Bottom left panel: Histogram of the number of misses $D_T^{\delta^*}$. Bottom right panel: Histogram of percentage of misses $D_T^{\delta^*}/N_T$.

The tradeoff between higher fill ratios and costs of walking the book are clear. An agent who seeks very high fill ratios, i.e., high values of $(N_T - D_T^{\delta^*})/N_T$, employs very high values of the cost parameters in the performance criterion. Other agents may prefer to swap price improvements for price deteriorations in their overall trading strategy. For example, in the 100,000 simulations we discuss, when $\gamma \approx 0.1801$ the average cost of filled trades, $C_T^{\delta^*}/(N_T - D_T^{\delta^*})$, is zero and the average rate of missed trades, $D_T^{\delta^*}/N_T$ is 0.0471.

Finally, a naive strategy employed by liquidity takers is to send MLOs with no discretion to walk the LOB, see [Cartea and Sánchez-Betancourt \(2021\)](#). Here, the expected ratio of missed trades to number of attempts and the expected cost of the strategy for an agent who sends all MLOs with no discretion to walk the LOB is $\mathbb{E}[D_T^0/N_T] = 0.5404$ and $\mathbb{E}[C_T^0] = -33.45$, respectively. The expected cost is negative because the strategy does not accrue costs from walking the book, but may receive price improvements.

4.1.1. Variable-discretion vs fixed-discretion to walk the LOB

Here, we compare the performance of cost-neutral latency-optimal strategies δ^* that send MLOs with fixed discretion (i.e., $\gamma = 0$, so that $\delta^* = \alpha$) or with variable discretion (i.e., $\gamma > 0$, $\alpha \geq 0$) to walk the limit order book – a cost-neutral strategy δ^* is one for which $\mathbb{E}[C_T^{\delta^*}] = 0$.

For each cost-neutral strategy δ^* we use simulations to estimate two quantities that are used by the agent to assess the performance of the strategy. Specifically, the agent estimates: (i) the expected final number of misses $\mathbb{E}[D_T^{\delta^*}]$, and (ii) the probability that the final number of misses $D_T^{\delta^*}$ is less than a percentage of the trade attempts N_T – i.e., estimate $\mathbb{P}[D_T^{\delta^*} < \tau N_T]$, where the parameter $\tau \in [0, 1]$ denotes the agent’s tolerance level to missed trades.

We proceed with the analysis of various cost-neutral strategies that employ fixed and variable discretion in their MLOs.

The top panels in Figure 5 show the probability that the number of missed trades is less than $\tau = 10\%$ of trade attempts, i.e., $\mathbb{P}[D_T^{\delta^*} < 0.1 N_T]$, and show the expected cost of the strategy, i.e., $\mathbb{E}[C_T^{\delta^*}]$, when the agent sends orders with a fixed discretion to walk the LOB, i.e., $\gamma = 0$ and $\alpha \in [0, 2.5]$. Similarly, the bottom panels show the probability that the number of missed trades is less than 10% of trade attempts, i.e., $\mathbb{P}[D_T^{\delta^*} < 0.1 N_T]$ and the expected cost of the strategy, i.e., $\mathbb{E}[C_T^{\delta^*}]$ for $\gamma \in [0.02, 0.2]$ and $\alpha = 0$. The orange dot in each picture shows the strategy for which $\mathbb{E}[C_T^{\delta^*}] = 0$.

We study two cost-neutral strategies in detail: (i) fixed discretion: $\gamma = 0$ and $\alpha^* = 1.79$ (orange dot in the top panels), and (ii) variable discretion: $\alpha = 0$ and $\gamma^* = 0.18$ (orange dot in the bottom panels); i.e., the discretion of the MLOs is variable because it depends on the number of missed trades. When the agent uses the cost-neutral fixed-discretion strategy, the expected number of misses is $\mathbb{E}[D_T^{\delta^*}] = 4.32$, and when the agent uses the cost-neutral variable-discretion strategy, the expected number of misses is $\mathbb{E}[D_T^{\delta^*}] = 4.41$. Therefore, an agent who employs cost-neutral strategies and who also wishes to bear a low expected number of misses, will prefer MLOs with the discretion of the cost-neutral strategy in (i) than that cost-neutral strategy in (ii). Note that for each $\alpha \in [0, \alpha^*]$, there exists $\gamma^\alpha \in [0, \gamma^*]$ such that the latency-optimal strategy in (13) with parameters α and γ^α is cost-neutral.

With the cost-neutral strategy in (i), the agent misses more than 10% of her trade attempts in 1,294 out of 100,000 simulations; i.e., $\mathbb{P}[D_T^{\delta^*} < 0.1 N_T] = 98.706\%$. Similarly, with the cost-neutral strategy in (ii), the agent misses more than 10% of her trade attempts in 3 out of 100,000 simulations; i.e., $\mathbb{P}[D_T^{\delta^*} < 0.1 N_T] = 99.997\%$.

Thus, for the specific choice of parameters in this example, an agent who employs cost-neutral strategies and who also expects to fill more than 90% of the trades for each trading day, will prefer to send MLOs with the discretion of the cost-neutral strategy in (ii) than those of the cost-neutral strategy in (i).

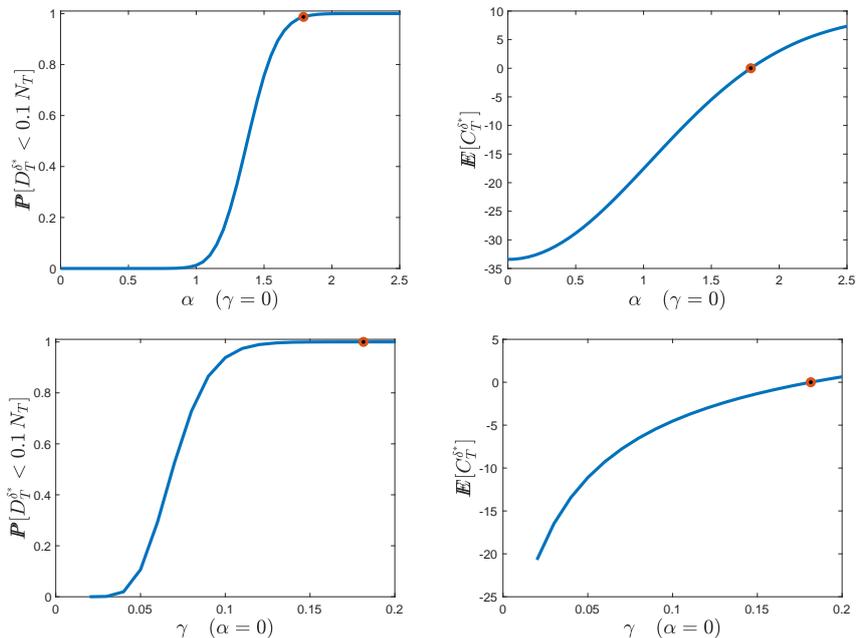


Figure 5: Top panel shows $\mathbb{P}(D_T^{\delta^*} < 0.1 N_T)$ and $\mathbb{E}[C_T^{\delta^*}]$ when $\gamma = 0$ and for $\alpha \in [0, 2.5]$, recall that $\delta^* = \alpha$ when $\gamma = 0$, see (16). Similarly, bottom panel shows $\mathbb{P}(D_T^{\delta^*} < 0.1 N_T)$ and $\mathbb{E}[C_T^{\delta^*}]$ when $\alpha = 0$ and $\gamma \in [0.02, 0.2]$. In all pictures, the orange dot marks the lowest value of $\mathbb{E}[C_T]$ when $\mathbb{P}(D_T < 0.1 N_T) \geq 0.99$. Other model parameters: $\lambda = 100$, $\alpha = 0$, and $Z \sim \mathbf{N}(0.1, 1)$ for all trades.

In general, the expected number of misses and the probability that the final number of misses is less than a percentage of the trade attempts depend on the value of the model parameters. For example, if the shock to the average price per share due to changes in prices is distributed as $Z_n \sim \mathbf{N}(\mu, 1)$ with $\mu \in [0.01, 0.185]$, we observe that: (a) the expected number of misses of the cost-neutral fixed-discretion strategy is lower than the expected number of misses of the cost-neutral variable-discretion strategy, and (b) the probability that the agent misses less than 10% of her trades is larger for the cost-neutral variable-discretion strategy. However, if the value of the parameter $\mu > 0.185$, then (a) holds, but (b) does not.

The results also depend on the agent's tolerance to missed trades. For example, if the value of the tolerance level parameter τ is greater than 0.05, the probability that the agent misses less than 100 $\tau\%$ of her trades is larger for the cost-neutral variable-discretion strategy. However, if $\tau \in (0, 0.05]$, the probability that the agent misses less than 100 $\tau\%$ of her trades is larger for the cost-neutral fixed-discretion strategy.

4.2. Pinned arrival rates

In this section, we assume the arrival intensity of the agent's MLOs is

$$\lambda_t^* = \frac{M - N_t}{T - t + \epsilon}, \quad (25)$$

where $M > 0$ is a positive integer, $\epsilon > 0$ and recall that N_t denotes the number of trade attempts. The intensity λ_t^* is bounded by $\bar{\lambda} = M/\epsilon$, which is a condition we require in the latency-optimal

strategy we derived above, and if $\epsilon = 0$, the intensity guarantees that $N_T = M$, see [Conforti \(2016\)](#) and [Hoyle \(2010\)](#).

Now, use the Markov property of δ^* to write $\delta^* = h(t, D_{t-}, N_{t-})$, where the function h satisfies the PIDE

$$0 = \partial_t h(t, D, N) + \left(\int_{h(t,D,N)}^{\infty} \frac{M - N}{T - t + \epsilon} \phi_t(z) dz \right) (h(t, D + 1, N + 1) - h(t, D, N)) \\ + \left(\int_{-\infty}^{h(t,D,N)} \frac{M - N}{T - t + \epsilon} \phi_t(z) dz \right) (h(t, D, N + 1) - h(t, D, N)) ,$$

with

$$h(t, D, M) = 2\gamma D + \gamma + \alpha \quad \text{and} \quad h(T, D, N) = 2\gamma D + \gamma + \alpha .$$

Figure 6 shows the optimal discretion to walk the LOB for various values of missed trades and target number of trades $M = 100$. The interpretation is similar to that of Figure 2.

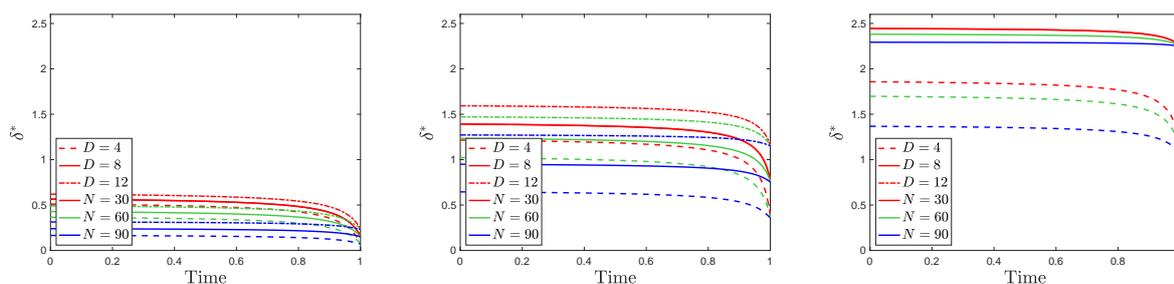


Figure 6: Optimal strategy δ^* for various values of γ , number of misses, and number of attempts. From left to right, cost parameter is $\gamma = 0.01$, $\gamma = 0.05$, and $\gamma = 0.15$. Dotted line $N_t = 30$, solid line $N_t = 60$, and dot-dash line $N_t = 90$. Blue lines $D_t^{\delta^*} = 4$, green lines $D_t^{\delta^*} = 8$, red lines $D_t^{\delta^*} = 12$. The remaining parameters are: $M = 100$, $\alpha = 0$, $\epsilon = 0.1$, $Z \sim \mathbf{N}(0.1, 1)$.

We perform 100,000 simulations with the same parameters as above and use the arrival rate of the MLOs as in (25) with $\epsilon = 0.1$. Figures 7 and 8 report the results, which have a similar interpretation to that of Figures 3 and 4, respectively.

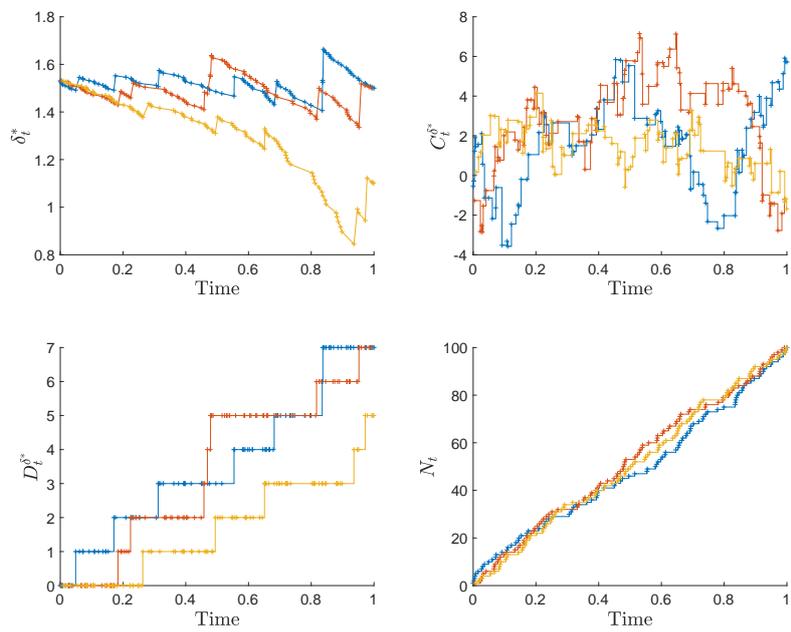


Figure 7: Sample paths for the optimal discretion δ^* (top left panel), number of missed trades D^{δ^*} (lower left panel), cost of strategy C^{δ^*} (top right panel), and number of trade attempts N (lower right panel) for three simulations of the MPP. Parameters: $\alpha = 0$, $\gamma = 0.1$, $\epsilon = 0.1$, $M = 100$, $T = 1$.

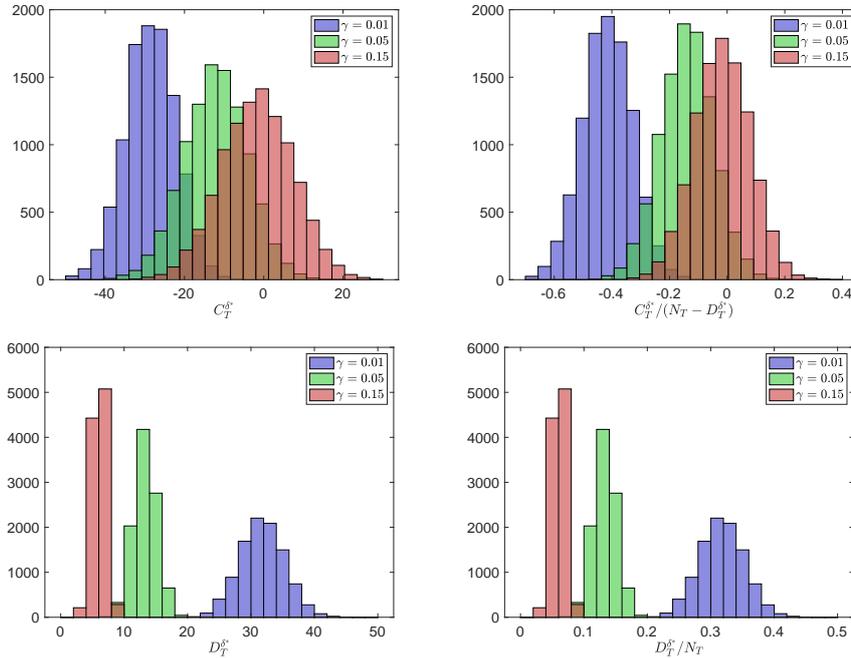


Figure 8: Top left panel: Histogram of the cost $C_T^{\delta^*}$ of the strategy. Top right panel: Histogram of the extra cost per filled trade $C_T^{\delta^*}/(N_T - D_T^{\delta^*})$. Bottom left panel: Histogram of the number of misses $D_T^{\delta^*}$. Bottom right panel: Histogram of percentage of misses $D_T^{\delta^*}/N_T$.

5. Conclusions

With few exceptions, the literature on algorithmic trading assumes that latency in the marketplace is zero. This is not accurate, and the effects of latency on the efficacy of liquidity making and taking strategies are economically significant. In this paper we proposed a model to improve the marksmanship of the orders sent by liquidity takers when, due to latency, the limit order book is a moving target.

We showed how a liquidity taker chooses the price limit of marketable orders when there is latency in the marketplace. The optimal strategy balances the tradeoff between the costs of walking the book and the number of missed trades over a trading horizon. We modelled the effects of latency as a marked point process that captures the interaction between liquidity taking orders and the limit orders resting in the book. We characterized the optimal price limit of marketable orders as a solution to a FBSDE, which, to the best of our knowledge, is new and, as the extant literature does not have uniqueness and existence results, we prove both.

The strategy developed here may be implemented as another layer of any liquidity taking strategy (especially those that follow a stochastic trading schedule) that incorrectly assumes zero latency. Our framework can be applied in other contexts too. In its most general form, we solve a problem in which the agent decides how much she is willing to pay to absorb a stochastic shock to achieve an objective or complete a task. For example, (i) market makers in foreign exchange markets with ‘last look’ can employ the framework developed in this paper. The last look feature allows liquidity makers to reject trades, so they are not picked off by ultra-fast traders, see [Oomen \(2017\)](#)

and [Cartea et al. \(2019\)](#). Specifically, a foreign exchange market maker can use our framework to derive the optimal threshold (to reject/accept trades) that maximizes the number of incoming marketable orders she is willing to fill while minimizing losses to the ultra-fast traders who snipe her stale quotes in the LOB. (ii) Firm liquidity venues in all asset classes can use our framework to design a cost-neutral strategy for liquidity takers who wish to improve the fill ratio of their trades. The venue uses the price improvements that the trader would have received to pay for the costs that stem from walking the LOB to increase fills, i.e., the exchange chooses a pair of values of the cost parameters α and γ such that the expected cost of the strategy is zero.

References

- Almgren, R. (2012). Optimal trading with stochastic liquidity and volatility. *SIAM Journal on Financial Mathematics*, 3(1):163–181.
- Antonelli, F. (1993). Backward-forward stochastic differential equations. *The Annals of Applied Probability*, 3(3):777–793.
- Bandini, E. (2016). Probabilistic representation of HJB equations for optimal control of jump processes, BSDEs and related stochastic calculus. *PhD Thesis*.
- Barger, W. and Lorig, M. (2019). Optimal liquidation under stochastic price impact. *International Journal of Theoretical and Applied Finance*, 22(02):1850059.
- Carmona, R. (2016). *Lectures on BSDEs, stochastic control, and stochastic differential games with financial applications*, volume 1. SIAM.
- Cartea, Á., Jaimungal, S., and Walton, J. (2019). Foreign exchange markets with last look. *Mathematics and Financial Economics*, 13(1):1–30.
- Cartea, Á. and Sánchez-Betancourt, L. (2021). The shadow price of latency: Improving intraday fill ratios in foreign exchange markets. *SIAM Journal on Financial Mathematics*, 12(1):254–294.
- Casgrain, P. and Jaimungal, S. (2018). Mean-field games with partial information for algorithmic trading. *arXiv preprint arXiv:1803.04094*.
- Casgrain, P. and Jaimungal, S. (2020). Mean-field games with differing beliefs for algorithmic trading. *Mathematical Finance*, 30(3):995–1034.
- Cohen, S. N. and Elliott, R. J. (2015). *Stochastic calculus and applications*, volume 2. Springer.
- Conforti, G. (2016). Bridges of Markov counting processes: quantitative estimates. *Electronic Communications in Probability*, 21.
- Confortola, F. and Fuhrman, M. (2013). Backward stochastic differential equations and optimal control of marked point processes. *SIAM Journal on Control and Optimization*, 51(5):3592–3623.
- Confortola, F., Fuhrman, M., and Jacod, J. (2016). Backward stochastic differential equation driven by a marked point process: An elementary approach with an application to optimal control. *The Annals of Applied Probability*, 26(3):1743–1773.
- Duffie, D. and Epstein, L. G. (1992). Stochastic differential utility. *Econometrica: Journal of the Econometric Society*, pages 353–394.
- Gao, X. and Wang, Y. (2018). Electronic market making and latency. *arXiv preprint arXiv:1806.05849*.
- Guéant, O. (2016). *The financial mathematics of market liquidity: From optimal execution to market making*, volume 33. CRC Press.
- Hoyle, E. (2010). Information-based models for finance and insurance. *PhD Thesis*.
- Jeanblanc, M., Yor, M., and Chesney, M. (2009). *Mathematical methods for financial markets*. Springer Science & Business Media.
- Jeanblanc-Picque, M. and Pontier, M. (1990). Optimal portfolio for a small investor in a market model with discontinuous prices. *Applied mathematics and optimization*, 22(1):287–310.
- Lehalle, C.-A. and Laruelle, S. (2013). *Market Microstructure in Practice*. World Scientific.
- Lehalle, C.-A. and Mounjid, O. (2017). Limit order strategic placement with adverse selection risk and the role of latency. *Market Microstructure and Liquidity*, 03(01):1750009.
- Ma, J. and Yong, J. (2007). *Forward-Backward Stochastic Differential Equations and their Applications*. Springer.
- Moallemi, C. C. and Sağlam, M. (2013). The cost of latency in high-frequency trading. *Operations Research*, 61(5):1070–1086.

- Oomen, R. (2017). Last Look. Quantitative Finance, 17:1057–1070.
- Peng, S. and Wu, Z. (1999). Fully coupled forward-backward stochastic differential equations and applications to optimal control. SIAM Journal on Control and Optimization, 37(3):825–843.
- Quenez, M.-C. and Sulem, A. (2013). BSDEs with jumps, optimization and applications to dynamic risk measures. Stochastic Processes and their Applications, 123(8):3328–3357.
- Revuz, D. and Yor, M. (2013). Continuous martingales and Brownian motion, volume 293. Springer Science & Business Media.
- Stein, E. M. and Shakarchi, R. (2003). Princeton lectures in analysis. Princeton University Press.
- Stoikov, S. and Waeber, R. (2016). Reducing transaction costs with low-latency trading algorithms. Quantitative Finance, 16(9):1445–1451.
- Tang, S. and Li, X. (1994). Necessary conditions for optimal control of stochastic systems with random jumps. SIAM Journal on Control and Optimization, 32(5):1447–1475.
- Xia, J. (2000). Backward stochastic differential equation with random measures. Acta Mathematicae Applicatae Sinica, 16(3):225–234.
- Zhen, W. (1999). Forward-backward stochastic differential equations with brownian motion and poisson process. Acta Mathematicae Applicatae Sinica, 15(4):433–443.

Appendix A. Proof of Lemma 1

We prove the lemma in three parts. First we work out the Gâteaux derivative of the cost function. We use (9) to write

$$\begin{aligned} \frac{1}{\epsilon} \left\{ \mathbb{E} [C_T^{\delta+\epsilon w}] - \mathbb{E} [C_T^\delta] \right\} &= \frac{1}{\epsilon} \left\{ \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} z \left(\hat{G}(\delta_t + \epsilon w_t - z) - \hat{G}(\delta_t - z) \right) \phi_t(dz) dA_t \right] \right\} \\ &= \frac{1}{\epsilon} \left\{ \mathbb{E} \left[\int_0^T \int_{\delta_t}^{\delta_t + \epsilon w_t} z \phi_t(dz) dA_t \right] \right\} \\ &= \mathbb{E} \left[\int_0^T \frac{1}{\epsilon} \left\{ \int_{\delta_t}^{\delta_t + \epsilon w_t} z \phi_t(dz) \right\} dA_t \right]. \end{aligned}$$

Then, by the dominated convergence theorem and the fundamental theorem of calculus, we have

$$\begin{aligned} \langle \mathcal{D} J^C(\delta), w \rangle &= \lim_{\epsilon \rightarrow 0} \frac{J^C(\delta + \epsilon w) - J^C(\delta)}{\epsilon} \\ &= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left\{ \mathbb{E} [C_T^{\delta+\epsilon w}] - \mathbb{E} [C_T^\delta] \right\} = \mathbb{E} \left[\int_0^T \delta_t w_t \phi_t(\delta_t) dA_t \right]. \end{aligned}$$

Next we work out the Gâteaux derivative of the linear penalty. Note that

$$\begin{aligned} \frac{1}{\epsilon} \left\{ \mathbb{E} [D_T^{\delta+\epsilon w}] - \mathbb{E} [D_T^\delta] \right\} &= \frac{1}{\epsilon} \left\{ \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} (G(\delta_t + \epsilon w_t - z) - G(\delta_t - z)) \tilde{p}(dz, dt) \right] \right\} \\ &= \mathbb{E} \left[\int_0^T \frac{1}{\epsilon} \left\{ \int_{\delta_t + \epsilon w_t}^{\delta_t} \phi_t(dz) \right\} dA_t \right], \end{aligned}$$

therefore, we have

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left\{ \mathbb{E} [D_T^{\delta+\epsilon w}] - \mathbb{E} [D_T^\delta] \right\} = -\mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) dA_t \right].$$

Finally, we work out the Gâteaux derivative of the quadratic penalty. We write

$$\begin{aligned} \frac{1}{\epsilon} \left\{ \mathbb{E} \left[(D_T^{\delta+\epsilon w})^2 \right] - \mathbb{E} \left[(D_T^\delta)^2 \right] \right\} &= \frac{2}{\epsilon} \left\{ \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} D_{t-}^{\delta+\epsilon w} G(\delta_t + \epsilon w_t - z) \tilde{p}(dz, dt) \right] \right. \\ &\quad \left. - \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} D_{t-}^\delta G(\delta_t - z) \tilde{p}(dz, dt) \right] \right\} \\ &\quad + \frac{1}{\epsilon} \left\{ \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} G(\delta_t + \epsilon w_t - z) - G(\delta_t - z) \tilde{p}(dz, dt) \right] \right\}. \end{aligned}$$

Subtract and add

$$\frac{2}{\epsilon} \left\{ \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} D_{t-}^{\delta+\epsilon w} G(\delta_t - z) \tilde{p}(dz, dt) \right] \right\}$$

to the right-hand side of the equation above and write

$$\begin{aligned}
& \frac{1}{\epsilon} \left\{ \mathbb{E} \left[(D_T^{\delta+\epsilon w})^2 \right] - \mathbb{E} \left[(D_T^\delta)^2 \right] \right\} \\
&= \frac{2}{\epsilon} \left\{ \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} (D_{t-}^{\delta+\epsilon w} G(\delta_t + \epsilon w_t - z) - D_{t-}^{\delta+\epsilon w} G(\delta_t - z)) \tilde{p}(dz, dt) \right] \right. \\
&\quad \left. + \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} (D_{t-}^{\delta+\epsilon w} G(\delta_t - z) - D_{t-}^\delta G(\delta_t - z)) \tilde{p}(dz, dt) \right] \right\} \\
&+ \frac{1}{\epsilon} \left\{ \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} (G(\delta_t + \epsilon w_t - z) - G(\delta_t - z)) \tilde{p}(dz, dt) \right] \right\} \\
&= 2 \left\{ \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \left(D_{t-}^{\delta+\epsilon w} \frac{G(\delta_t + \epsilon w_t - z) - G(\delta_t - z)}{\epsilon} \right) \tilde{p}(dz, dt) \right] \right. \tag{QP1} \\
&\quad \left. + \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \left(\frac{D_{t-}^{\delta+\epsilon w} - D_{t-}^\delta}{\epsilon} G(\delta_t - z) \right) \tilde{p}(dz, dt) \right] \right\} \tag{QP2} \\
&+ \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \left(\frac{G(\delta_t + \epsilon w_t - z) - G(\delta_t - z)}{\epsilon} \right) \tilde{p}(dz, dt) \right]. \tag{QP3}
\end{aligned}$$

Next, take the limit of QP1, QP2, and QP3 as ϵ approaches zero. The limit of QP1 is given by

$$\begin{aligned}
\lim_{\epsilon \rightarrow 0} \text{QP1} &= \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} D_{t-}^{\delta+\epsilon w} \frac{G(\delta_t + \epsilon w_t - z) - G(\delta_t - z)}{\epsilon} \tilde{p}(dz, dt) \right] \\
&= \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[\int_0^T D_{t-}^{\delta+\epsilon w} \frac{1}{\epsilon} \left\{ \int_{\delta_t + \epsilon w_t}^{\delta_t} \phi_t(dz) \right\} dA_t \right] \\
&= -\mathbb{E} \left[\int_0^T D_{t-}^\delta w_t \phi_t(\delta_t) dA_t \right].
\end{aligned}$$

The last equality follows from the dominated convergence theorem and because $\lim_{\epsilon \rightarrow 0} D_{t-}^{\delta+\epsilon w} = D_{t-}^\delta$ almost surely.

The limit of QP2 is given by

$$\begin{aligned}
& \lim_{\epsilon \rightarrow 0} \text{QP2} \\
&= \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \frac{D_{t-}^{\delta+\epsilon w} - D_{t-}^{\delta}}{\epsilon} G(\delta_t - z) \tilde{p}(dz, dt) \right] \\
&= \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} G(\delta_t - z) \left(\int_0^{t-} \int_{\mathbb{R}} \frac{G(\delta_s + \epsilon w_s - z') - G(\delta_s - z')}{\epsilon} p(dz', ds) \right) \tilde{p}(dz, dt) \right] \\
&= \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \frac{G(\delta_t + \epsilon w_t - z) - G(\delta_t - z)}{\epsilon} \left(\int_t^T \int_{\mathbb{R}} G(\delta_s - z') \tilde{p}(dz', ds) \right) p(dz, dt) \right] \\
&= \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \frac{G(\delta_t + \epsilon w_t - z) - G(\delta_t - z)}{\epsilon} \mathbb{E}_{t-} \left[\int_t^T \int_{\mathbb{R}} G(\delta_s - z') \tilde{p}(dz', ds) \right] p(dz, dt) \right] \\
&= \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \frac{G(\delta_t + \epsilon w_t - z) - G(\delta_t - z)}{\epsilon} \mathbb{E}_{t-} \left[\int_t^T \int_{\mathbb{R}} G(\delta_s - z') \tilde{p}(dz', ds) \right] \tilde{p}(dz, dt) \right] \\
&= \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[\int_0^T \frac{1}{\epsilon} \int_{\delta_t + \epsilon w_t}^{\delta_t} \phi_t(dz) \mathbb{E}_{t-} \left[\int_t^T \int_{\mathbb{R}} G(\delta_s - z') \tilde{p}(dz', ds) \right] dA_t \right] \\
&= -\mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) \mathbb{E}_{t-} \left[\int_t^T \int_{\mathbb{R}} G(\delta_s - z') \tilde{p}(dz', ds) \right] dA_t \right].
\end{aligned}$$

Finally, the limit of QP3 is given by

$$\begin{aligned}
\lim_{\epsilon \rightarrow 0} \text{QP3} &= \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \frac{G(\delta_t + \epsilon w_t - z) - G(\delta_t - z)}{\epsilon} \tilde{p}(dz, dt) \right] \\
&= \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[\int_0^T \frac{1}{\epsilon} \int_{\delta_t + \epsilon w_t}^{\delta_t} \phi_t(dz) dA_t \right] \\
&= -\mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) dA_t \right],
\end{aligned}$$

which concludes the proof.

Appendix B. Existence and Uniqueness of forward and backward parts of the FBSDE

The following theorem shows the existence and uniqueness of the solution to the backward part of the FBSDE (19).

Theorem 5. Fix $V \in \mathcal{C}^\infty$. Assumptions 1 and 2 hold. Then, the functional $\Psi : \mathcal{C}^\infty \rightarrow \mathcal{C}^\infty$ given by

$$\Psi(U)_t = 2\gamma \mathbb{E}_t \left[\int_t^T \int_{\mathbb{R}} G(U_{s-} + 2\gamma V_{s-} - z) p(dz, ds) \right] + \gamma + \alpha, \quad V \in \mathcal{C}^\infty,$$

has a unique fixed point.

Proof. We proceed as in Proposition A1 in [Duffie and Epstein \(1992\)](#). Define $Z = \sup_{t \in \mathfrak{T}} |X_t - Y_t|$ for any X and Y in \mathcal{C}^∞ . Let $\Psi^{(1)} = \Psi$ and $\Psi^{(n)} = \Psi(\Psi^{(n-1)})$. Then

$$\begin{aligned} |\Psi(X)_t - \Psi(Y)_t| &= 2\gamma \left| \mathbb{E}_t \left[\int_t^T (\Phi(Y_{s-} + 2\gamma V_{s-}) - \Phi(X_{s-} + 2\gamma V_{s-})) \lambda_s ds \right] \right| \\ &\leq 2\gamma k \mathbb{E}_t \left[\int_t^T |X_{s-} - Y_{s-}| \lambda_s ds \right] \leq 2\gamma k \bar{\lambda} (T-t) \mathbb{E}_t [Z]. \end{aligned}$$

Use Fubini's theorem for conditional expectations to write

$$\begin{aligned} |\Psi^{(2)}(X)_t - \Psi^{(2)}(Y)_t| &\leq 2\gamma k \bar{\lambda} \mathbb{E}_t \left[\int_t^T |\Psi(X_s) - \Psi(Y_s)| ds \right] \leq 2\gamma k \bar{\lambda} \mathbb{E}_t \left[\int_t^T 2\gamma k \bar{\lambda} (T-s) \mathbb{E}_s [Z] ds \right] \\ &\leq (2\gamma k \bar{\lambda})^2 \mathbb{E}_t \left[\int_t^T (T-s) \mathbb{E}_s [Z] ds \right] \leq (2\gamma k \bar{\lambda})^2 \frac{(T-t)^2}{2!} \mathbb{E}_t [Z], \end{aligned}$$

which after n iterations becomes

$$|\Psi^{(n)}(X)_t - \Psi^{(n)}(Y)_t| \leq (2\gamma k \bar{\lambda})^n \frac{(T-t)^n}{n!} \mathbb{E}_t [Z].$$

Finally,

$$\begin{aligned} \mathbb{E} \left[\sup_{t \in \mathfrak{T}} |\Psi^{(n)}(X)_t - \Psi^{(n)}(Y)_t| \right] &\leq \frac{(2\gamma k \bar{\lambda} T)^n}{n!} \mathbb{E} \left[\sup_{t \in \mathfrak{T}} \mathbb{E}_t [Z] \right] \\ &\leq 4 \frac{(2\gamma k \bar{\lambda} T)^n}{n!} \mathbb{E} \left[\sup_{t \in \mathfrak{T}} |X_t - Y_t| \right]. \end{aligned}$$

Therefore, for n sufficiently large, the function $\Psi^{(n)}$ is a contraction mapping in the Banach space \mathcal{C}^∞ equipped with the supremum norm $(\mathcal{C}^\infty, \|\cdot\|_\infty)$. Thus, there exists a unique⁷ process $U \in \mathcal{C}^\infty$ such that $\Psi^{(n)}(U) = U$ and because $\Psi^{(n)}(\Psi(U)) = \Psi(\Psi^{(n)}(U)) = \Psi(U)$ and by uniqueness of the fixed point, we have $\Psi(U) = U$, which proves the existence of the fixed point for Ψ . Uniqueness of this fixed point for Ψ follows from uniqueness of the fixed point in $\Psi^{(n)}$, which concludes the proof. ■

The next theorem shows the existence and uniqueness of the solution to the forward part of the FBSDE [\(19\)](#).

Theorem 6. Fix $U \in \mathcal{C}^\infty$. Assumptions [1](#) and [2](#) hold. Then, the functional $\Theta : \mathcal{C}^\infty \rightarrow \mathcal{C}^\infty$ given by

$$\Theta(V)_t = \int_0^t \int_{\mathbb{R}} G(U_{s-} + 2\gamma V_{s-} - z) p(dz, ds), \quad U \in \mathcal{C}^\infty,$$

has a unique fixed point.

Proof. First we prove that Θ is a functional from \mathcal{C}^∞ to \mathcal{C}^∞ . Let $U, V \in \mathcal{C}^\infty$. By definition, the function $\Theta(V)$ is adapted and because $G \leq 1$ we have

$$\mathbb{E} \left[\sup_{t \in \mathfrak{T}} |\Theta(V)_t| \right] \leq \mathbb{E} [p([0, T], \mathbb{R})] < \infty.$$

⁷Unique in the sense of indistinguishability.

Thus, $\Theta(V) \in \mathcal{C}^\infty$. Next, denote $\Theta^n = \Theta(\Theta^{n-1})$ with $\Theta^0 = \Theta(0)$ and define $h_n : [0, T] \rightarrow \mathbb{R}$ as

$$h_n(t) = \mathbb{E} \left[\sup_{0 \leq s \leq t} \left| \Theta_s^{(n+1)} - \Theta_s^{(n)} \right| \right].$$

We find an upper bound for $h_n(t)$ as follows:

$$\begin{aligned} h_n(t) &= \mathbb{E} \left[\sup_{0 \leq s \leq t} \left| \Theta_s^{(n+1)} - \Theta_s^{(n)} \right| \right] \\ &= \mathbb{E} \left[\sup_{0 \leq s \leq t} \left| \int_0^s \int_{\mathbb{R}} \left(G \left(U_{u^-} + 2\gamma \Theta_{u^-}^{(n)} - z \right) - G \left(U_{u^-} + 2\gamma \Theta_{u^-}^{(n-1)} - z \right) \right) p(dz, du) \right| \right] \\ &\leq \mathbb{E} \left[\int_0^t \int_{\mathbb{R}} \left| G \left(U_{u^-} + 2\gamma \Theta_{u^-}^{(n)} - z \right) - G \left(U_{u^-} + 2\gamma \Theta_{u^-}^{(n-1)} - z \right) \right| \tilde{p}(dz, du) \right] \\ &= \mathbb{E} \left[\int_0^t \left| \Phi \left(U_{u^-} + 2\gamma \Theta_{u^-}^{(n)} \right) - \Phi \left(U_{u^-} + 2\gamma \Theta_{u^-}^{(n-1)} \right) \right| \lambda_t dt \right] \\ &\leq 2\gamma k \bar{\lambda} \mathbb{E} \left[\int_0^t \left| \Theta_{u^-}^{(n)} - \Theta_{u^-}^{(n-1)} \right| dt \right] \\ &\leq 2\gamma k \bar{\lambda} \int_0^t h_n(s) ds. \end{aligned}$$

The above inequality, together with the observation that $h_0(T) = M < \infty$, implies

$$0 \leq h_n(T) \leq \frac{M (2\gamma k \bar{\lambda})^n T^n}{n!},$$

and use Markov's inequality to obtain the bound:

$$\mathbb{P} \left(\sup_{t \in \mathfrak{T}} \left| \Theta_t^{(n+1)} - \Theta_t^{(n)} \right| \geq 2^{-n} \right) \leq \frac{M (2\gamma k \bar{\lambda})^n T^n 2^{2n}}{n!} \xrightarrow{n \rightarrow \infty} 0.$$

By Borel-Cantelli arguments, there is $\mathfrak{D} \subset \Omega$ such that for all $\omega \in \mathfrak{D}$ the functions $t \rightarrow \Theta_t^{(n)}(\omega)$ form a Cauchy sequence in the supremum norm of \mathcal{C}^∞ with probability one. Thus, $\forall \omega \in \mathfrak{D}$ there is a function $\Theta_t^*(\omega)$ such that $\Theta_t^{(n)}(\omega)$ converges uniformly to $\Theta_t^*(\omega)$ in \mathfrak{T} . Furthermore, there is an adapted modification of Θ^* in Ω .

Thus, the process Θ^* is a fixed point of the mapping defined by Θ , and therefore satisfies the forward part of the FBSDE. ■

Appendix C. Second Gâteaux derivative

The first Gâteaux derivative of the functional J is given by

$$\begin{aligned} \langle \mathcal{D} J(\delta), w \rangle &= \mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) \left(\delta_t - 2\gamma \left(\int_t^T \int_{\mathbb{R}} G(\delta_s - z') \tilde{p}(dz', ds) \right) - 2\gamma D_{t^-}^\delta - (\gamma + \alpha) \right) dA_t \right] \\ &= \mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) \left(\delta_t - 2\gamma \mathbb{E}_{t^-} [D_T^\delta] - (\gamma + \alpha) \right) dA_t \right]. \end{aligned}$$

Let $\delta, w, \nu \in \mathcal{A}$. The second Gâteaux derivative of $J(\delta)$ in the directions w and ν , is defined as

$$\langle \mathcal{D}^2 J(\delta), w, \nu \rangle = \lim_{\epsilon \rightarrow 0} \frac{\langle \mathcal{D} J(\delta + \epsilon \nu), w \rangle - \langle \mathcal{D} J(\delta), w \rangle}{\epsilon},$$

which converges to

$$\begin{aligned} \langle \mathcal{D}^2 J(\delta), w, \nu \rangle &= \mathbb{E} \left[\int_0^T w_t \nu_t \phi'_t(\delta_t) (\delta_t - 2\gamma \mathbb{E}_{t-} [D_T^\delta] - \gamma - \alpha) \, dA_t \right] \\ &+ \mathbb{E} \left[\int_0^T w_t \phi_t(\delta_t) \left(\nu_t + 2\gamma \mathbb{E}_{t-} \left[\int_0^T \phi_s(\delta_s) \nu_s \, dA_s \right] \right) \, dA_t \right]. \end{aligned}$$