

Uncertain Execution in Order-Driven Markets



Leandro Sánchez-Betancourt
Oriol College
University of Oxford

Thesis submitted for the degree of
DPhil in Mathematics

Trinity 2021

To my mother and my father

María Antonia Betancourt-Betancourt
&
Leonardo Sánchez-Milanés

Acknowledgements

First and foremost, I would like to express how fortunate I feel to have had Professor Álvaro Cartea as my supervisor. I thank him for our frequent conversations, for his astute criticism, and his unflagging support. I hope that we continue to collaborate for many years to come.

My profound gratitude also goes to Professor Lane Hughston and Professor Sebastian Jaimungal. They have played a key role in my growth as a researcher. I cherish each and every one of our mathematical discussions.

I would also like to thank my colleagues at LMAX Exchange. In particular, I am indebted to Andrew Stewart for our weekly meetings and his insightful comments.

I feel fortunate for the valuable feedback received from Professor Samuel N. Cohen, Professor Rama Cont, Professor Ben Hambly, Professor Sam Howison, Professor Michael Monoyios, Professor Johannes Muhle-Karbe, and Professor José Penalva. I have also benefited from mathematical discussions with collaborators, and friends: Theerawat Bhudisaksang, Georgios Bouzianis, Dr Ryan Donnelly, Gerardo Durán-Martín, Dr Maria Flora, Dr Martin Forde, Dr Mauricio Labadie, Anthony Marriot, Dr José Manuel Pedraza-Ramírez, Dr Imanol Pérez-Arribas, and Yufei Zhang.

I gratefully acknowledge financial support received from (a) the Mathematical Institute, Oxford, (b) Consejo Nacional de Ciencia y Tecnología (CONACyT), Ciudad de México, (c) LMAX Exchange, London, and (d) Oriel College, Oxford.

Lastly, I thank my friends and family for their constant encouragement.

Abstract

So-called ‘latency’ refers to the various small but significant time delays that occur in the course of the communications between a trader and a market. Such delays happen between the time an exchange streams market data to a trader, the time at which the trader processes the information and decides to trade, and the time at which the exchange receives and processes the order from the trader. Latency is a challenge faced by all traders and is of great importance in modern financial markets. In the present work, we develop mathematical models to solve a variety of problems faced by liquidity takers regarding uncertainty in executions.

Firstly, we devise a model for computing the price that traders are willing to pay to reduce their latency. This latency-optimal strategy balances the tradeoff, over a period of time, between the costs of walking the limit order book and the percentage of orders filled. This work may lead to social benefits, since it offers a way to stop the arms race to being faster in the marketplace.

Secondly, we develop a latency-optimal trading strategy that improves the marksmanship of liquidity takers. We make use of the techniques of variational analysis to obtain the optimal price limit of each marketable limit order (MLO) that the trader sends. The price limit of each MLO is characterized as the solution to a new class of forward-backward stochastic differential equations (FBSDEs) driven by random measures. We prove the existence and uniqueness of the FBSDE solution and solve the FBSDE numerically to illustrate the performance of the strategies.

Finally, we show how traders can optimally liquidate a position over a trading window when there is latency in the marketplace. We frame our model as an impulse control problem with stochastic delay – this work contributes to the stochastic control literature by allowing one to have random delays before the impulses take place. We show that impatient liquidity takers submit MLOs that may walk the book (capped by the limit price) to increase the probability of filling the trades. Patient traders who are fast do not use their speed to hit the quotes they observe, nor to finish the execution programme early: they use speed to complete the execution with as many speculative MLOs as possible. We use foreign exchange data to implement the random-latency-optimal strategy and to compare it with various benchmarks. We find that for patient traders, the random-latency optimal strategy outperforms the benchmarks that do not account for latency by a quantity that is greater than the transaction costs paid by liquidity takers. Around news announcements, the value of the outperformance significantly increases. The superiority of the latency-optimal strategies is due to both the speculative MLOs that are filled and the price protection of the MLOs.

Contents

1	Introduction	1
2	Background material	4
2.1	Limit order books	4
2.2	Order types	5
2.3	Definition of latency	7
2.4	Incentives to use MLOs	7
2.5	Previous work on latency	10
3	The shadow price of latency	12
3.1	Latency and fill ratios	15
3.1.1	Latency	15
3.1.2	Fill ratios and post-trade analysis	16
3.2	Empirical analysis	18
3.2.1	Data	18
3.2.2	Fill ratios and volatility	18
3.3	Model	19
3.3.1	Model uncertainty	22
3.3.2	Value function	23
3.3.2.1	Infinite-horizon latency-optimal strategy	26
3.4	Estimation of model parameters	26
3.5	Performance of strategy	28
3.5.1	Stylized features of latency-optimal strategy	28
3.5.2	Performance of strategy for traders T1 and T2	30
3.5.2.1	Performance of latency-optimal strategy for T1	30
3.5.2.2	Performance of optimal strategy: T2	33
3.5.3	Infinite-horizon strategy and robustness checks	35
3.6	The cost of latency	35
3.6.1	Latency and fill ratios	36
3.6.2	Shadow price of latency	39
3.7	Conclusions	39

4	Tradeoff between costs and misses	41
4.1	Optimal discretion to walk the book	43
4.1.1	MLOs: revealed preferences	43
4.1.1.1	Empirical evidence: price limits	45
4.1.2	LOB: moving target	46
4.1.3	Tradeoff: cost of walking the LOB and cost of missed trades	48
4.1.3.1	Cost I: discretion to walk the book	48
4.1.3.2	Cost II: missed trades	49
4.1.4	Performance criterion	50
4.1.5	Variational analysis approach	50
4.1.5.1	Optimal discretion to walk the LOB	51
4.2	Existence and uniqueness of the FBSDE	53
4.3	Optimality	57
4.4	Performance of strategy	62
4.4.1	Poisson arrival of trades	63
4.4.1.1	Variable-discretion vs fixed-discretion to walk the LOB	66
4.4.2	Pinned arrival rates	68
4.5	Conclusions	70
5	Optimal trading with delay	72
5.1	Order types and data	75
5.1.1	Orders and revealed preferences	75
5.1.2	Data	77
5.1.3	Flickering of best rates in the LOB	78
5.1.4	Distribution of flickers: hit and miss	79
5.1.5	Liquidity: make and take	81
5.2	Optimal execution with random latency	82
5.2.1	Trader's clock	83
5.2.2	Exchange rate dynamics and execution of MLOs	83
5.2.3	Outcome of trade attempts: miss or fill	85
5.2.4	Admissible strategies	87
5.2.5	System dynamics	87
5.2.6	Liquidity taking with stochastic delay	90
5.2.7	Latency-optimal strategy: price protection and speculation	92
5.3	Benchmarks	93
5.3.1	Optimal execution with zero latency	93
5.3.2	Optimal execution with deterministic latency	94
5.4	Performance of execution strategy	96
5.4.1	Model parameters	98

5.4.2	Simulations and performance measure	99
5.4.3	Results: performance of RLOS and benchmarks	100
5.4.4	Trading around news events	105
5.4.5	Economic significance	107
5.4.6	Robustness checks	108
5.5	Conclusions	109
6	Concluding remarks	110
	References	111
	Glossary	117
	Notation	120
A	Appendix for Chapter 3	121
A.1	Proofs and further results	121
A.1.1	System of ODEs satisfied by ansatz (3.18)	121
A.1.2	Proof of Theorem 1.	122
A.2	Further robustness checks	127
A.2.1	Infinite-horizon value function	127
A.2.2	Alternative models for volatility	129
A.2.3	Performance of the optimal strategy for alternative models of volatility	129
A.2.4	Parameter estimates and sensitivity analysis	130
A.2.4.1	Sensitivity to parameter estimates	131
B	Appendix for Chapter 4	133
B.1	Proof of Lemma 1	133
B.2	Existence and uniqueness of forward and backward parts of the FBSDE	135
B.3	Second Gâteaux derivative	137
B.4	Further results	137
C	Appendix for Chapter 5	140
C.1	Random latency	140
C.1.1	Dynamic programming principle	140
C.1.2	PDE system viscosity characterization	143
C.1.3	Viscosity properties of the value function v_0	144
C.2	Deterministic latency	147
C.2.1	Dynamic programming principle	148
C.2.1.1	PDE system viscosity characterization	148
C.3	Zero latency	149

Chapter 1

Introduction

With the advent of computerized trading, the speed at which agents operate in electronic markets is down to milliseconds or even microseconds. Speed to process information and make decisions is essential to the success of trading strategies and being faster than one's peers provides a valuable competitive edge. Traders require time to execute strategies and their instructions take time to arrive and to be processed by the exchange. This time delay, known as latency, affects the efficacy of trading strategies because, in the meantime, the exchange processes other instructions that update the LOB with information not known by the trader at the time the trader executed the strategy.

In that which follows, we construct three optimal trading policies that aim (i) to quantify the cost of latency, (ii) to work out how latency affects the marksmanship of liquidity takers, and (iii) to determine how best to liquidate a position when there are random execution delays.

There are many papers in the mathematical finance literature that develop optimal trading policies. The seminal paper of [Merton \(1971\)](#) uses stochastic control techniques to find an optimal consumption rate and investment weights that maximize a given performance criterion. A second example is that of [Black and Scholes \(1973\)](#), in which a dynamic trading strategy (the delta-hedging strategy) replicates the payoff of a European option. A third example is [Almgren and Chriss \(2000\)](#), where the authors find the optimal trading strategy required to liquidate a large block of shares minimizing execution costs.

In this thesis we build on stochastic control theory and its applications to solve a number of mathematical problems that relate to latency and uncertainty in executions. Along with other methods we make use of (a) continuous-time stochastic control with ambiguity aversion, (b) Gâteaux derivatives, (c) forward-backward stochastic differential equations (FBSDEs) to characterize optimal strategies, and (d) impulse control theory. In [Chapters 4 and 5](#), we contribute to the literature of stochastic analysis and applications (i) by presenting a new type of FBSDE and proving its existence and uniqueness, and (ii) by introducing a new type of impulse control problem where the times between the decisions to impulse the system, and the impulses taking place, are random.

The material is organized as follows. In [Chapter 2](#) we introduce the concepts we use throughout the thesis. In particular, in [Section 2.2](#) we define the type of order we use throughout this

investigation, that is, marketable limit orders (MLOs). MLOs are liquidity taking orders that specify a price limit and are for immediate execution only; however, if the price limit of the MLO precludes it from being filled, the exchange rejects the trade.¹ We show the importance that these type of orders have in the market, and how relevant they are for liquidity takers. This is the first time in the algorithmic trading literature where MLOs are shown to play such a key role in addressing the latency problem that liquidity takers face in financial markets.

Then we present the work of the following three papers [Cartea and Sánchez-Betancourt \(2021b\)](#), [Cartea et al. \(2019a\)](#), and [Cartea and Sánchez-Betancourt \(2021a\)](#). More specifically:

Chapter 3 constructs a model to compute the price that liquidity takers would be willing to pay to reduce their latency in the marketplace. To this end, we derive a latency-optimal strategy that specifies the limit price of liquidity taking orders to increase the chances of filling orders if, due to latency, prices or quantities in the LOB have worsened. The latency-optimal strategy balances the tradeoff between the costs of walking the LOB and targeting a desired percentage of filled orders over a period of time. We employ the cost of improving fills with the latency-optimal strategy to compute the shadow price of latency. Finally, we use a proprietary data set of foreign exchange (FX) to compute the maximum price that an FX trader would be willing to pay for co-location and hardware to reduce their latency in the marketplace.

Chapter 4 develops a latency-optimal trading strategy that improves the marksmanship of liquidity takers. We model the interaction between the LOB and MLOs 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 this model, the liquidity taker balances the tradeoff between the costs of missing trades and the costs of walking the book. We employ techniques of variational analysis to obtain the optimal price limit of each MLO the agent sends. The price limit of an MLO is characterized as part of the solution to an FBSDE. As far as we know, this random-measure driven FBSDE has not been covered in the literature to date. This may be on account of (i) the particular non-standard representation of the FBSDE, (ii) the fact that the forward part is adapted and the backward part needs to be predictable, and (iii) the discontinuity of the driver in the forward part of the FBSDE. We prove the existence and uniqueness of the solution to the FBSDE using standard tools, and as with the treatment of other FBSDEs, we require a short time horizon. We numerically solve the FBSDE to illustrate the performance of the latency-optimal strategies.

Chapter 5 shows how traders use MLOs to liquidate a position over a trading window when there is latency in the marketplace. We frame our model as an impulse control problem with stochastic latency where the trader controls the times and the price limits of the MLOs sent to the exchange. We show that impatient liquidity takers submit MLOs that may walk the

¹MLOs are similar to standard market orders, but differ in that MLOs have a limit price whereas market orders do not. Thus, a buy MLO never pays more than the limit price, and a sell MLO never receives less than the limit price.

book (capped by the limit price) to increase the probability of filling the trades. On the other hand, patient liquidity takers use speculative MLOs that are only filled if there has been an advantageous move in prices over the latency period. Patient traders who are fast do not use their speed to hit the quotes they observe, nor to finish the execution programme early; they use speed to complete the execution programme with as many speculative MLOs as possible. We use foreign exchange data to implement the random-latency-optimal strategy and to compare it with four benchmarks: execution with deterministic latency, execution with zero latency, time-weighted average price, and execution of the entire order at the best quote in the LOB at the beginning of the trading window. We find that the performances of the random-latency and the deterministic-latency strategies are similar. For patient traders, these two strategies outperform the other three benchmarks by an amount that is greater than the transaction costs paid by liquidity takers in foreign exchange markets. Around news announcements, the value of the outperformance is between two and ten times the value of the transaction costs. The superiority of the latency-optimal strategies is due to both the speculative MLOs that are filled and the price protection of the MLOs.

We present concluding remarks in Chapter 6. Each of the chapters introduces acronyms separately and the reader is reminded of the main definitions. For the reader's convenience, we include a [glossary](#) with acronyms and technical terms, and a mathematical [notation](#) page after the references. Finally, we collect various proofs and robustness checks in the appendices.

Chapter 2

Background material

In this section we introduce concepts at the core of any order-driven market. Examples of this include the limit order book (LOB), market orders (MOs), limit orders (LOs), and marketable limit orders (MLOs).

The interested reader can consult the following manuscripts for a detailed introduction to the subject; these include [Cartea et al. \(2015\)](#), [Guéant \(2016\)](#), [Lehalle and Laruelle \(2018\)](#), and more recently [Bouchaud et al. \(2018\)](#).¹

An exchange is a marketplace where securities, commodities, derivatives and other financial instruments are traded. The core function of an exchange is to ensure fair and orderly trading and the efficient dissemination of price information for any securities trading on that exchange. Exchanges give companies, governments, and other groups a platform from which to sell securities to the investing public.

2.1 Limit order books

Exchanges accommodate a compendium of market participants, each of which acts in their own best interest, subject to certain marketplace rules. These market participants have, in the simplest description, two ways to intervene in the market for a security:² (i) they post buy LOs (*resp.* sell LOs) when they have a price and quantity in mind that they want to buy (*resp.* sell), and they are willing to wait until matched against a counterpart with opposite interest, or (ii) they send MOs when they have a quantity they want to buy (*resp.* sell) no matter the price they pay (*resp.* receive). We define the LOB of a given security or instrument to be the collection of limit orders waiting to be matched. We refer to the different layers of liquidity in the ask side (*resp.* the bid side), as all the different price levels at which liquidity providers have buy LOs (*resp.* sell LOs) waiting to be matched. Similarly, we say that an order sent by a trader ‘walks the book’, when the size of the order is such that the trader consumes more than one layer of liquidity to fill the trade.

¹See [Gould et al. \(2013\)](#) for a detailed exposition of LOBs.

²In the next subsection, we introduce other order types that play a crucial role in this thesis. For now, in the interest of a clear exposition, we only mention MOs and LOs.

Most financial contracts, financial instruments, and financial securities are traded in LOBs. For example, equity, spot FX, futures, and options are often traded in venues with a central LOB.³ All models that we discuss are applicable to any venue with a central LOB. In this thesis, all data experiments, motivations, and illustrations are for spot FX trading. We employ data provided by LMAX – an electronic multilateral trading facility in the FX market (www.lmax.com).

2.2 Order types

Electronic exchanges have a set of orders that market participants can use. The range of order types is disclosed to all participants who then specify which type they wish to use with every order they send to the exchange. There are order types that are common in some markets and uncommon in others. For example, pegged orders are common in equity markets but they do not exist in spot FX markets. In what follows, we describe the main order types that market participants can use. The order types that we list and explain can be used in equity markets (e.g., Nasdaq and LSE), spot FX markets (e.g., LMAX and EBS), and futures markets (e.g., CME).

There are two tags associated with any given order sent to the exchange. These are the type of quote and the time-in-force. We describe the type of quote first, following ‘Guide to trading system’ on the website of the London Stock Exchange.⁴

Type of quote	Definition
Limit order	An anonymous priced order that is fully displayed when persistent in an order book and may execute at prices equal to or better than its limit price.
Market order	Un-priced, and therefore not price forming, but has price priority over all priced orders. Market Orders cannot persist on the order book during regular trading.
Stop Limit Orders	A Limit Order that will remain unelected (will not be entered into order book) until the stop price is reached. Once elected, a Stop Limit Order will be treated as a regular Limit Order.
Stop Orders	A Market Order that will remain unelected (will not be entered into order book) until the stop price is reached. Once elected, it will be treated similar to a regular Market Order.

Table 2.1: Definition of the type of quote.

There are other types of quotes such as ‘hidden limit orders’, ‘midprice pegged orders’, ‘iceberg orders’, and so on, that we do not mention because they are not used to trade in many exchanges.

Next, we describe the second tag of any given order: the “time-in-force” of the order.

³Spot FX is mostly traded over-the-counter, nonetheless, there are a large number of exchanges with a central LOB trading spot FX; see Figure 3.3 in [Donnelly \(2019\)](#) for an overview of the market structure in FX.

⁴See www.londonstockexchange.com/resources.

Time-in-force	Definition
Good for day (DAY)	Expires at the end of the trading day on which it was entered or when cancelled.
Good until time (GTT)	Expires at the specified time/date or when cancelled.
Immediate or Cancel (IoC)	Executed on entry, with any remaining unexecuted volume expired.
Fill or Kill (FoK)	Executed in full on entry or immediately expired.

Table 2.2: Definition of time-in-force.

Order types are created with one tag from Table 2.1 and one from Table 2.2. There are some combinations that are not allowed, and when the second tag (time-in-force) is missing, there are default values for it.

		Time-in-force			
		DAY	GTT	IoC	FoK
Order type	Limit order	✓	✓	✓	✓
	Market order	×	×	✓	✓
	Stop limit order	✓	✓	×	×
	Stop order	✓	✓	×	×

Table 2.3: Possible combinations of order type and time-in-force. We use ‘✓’ for a combination that can be used, and ‘×’ for one that can not be used or that is redundant.

Table 2.2 shows the combinations that are allowed with a ‘✓’ and those that are not with an ‘×’. If the time-in-force tag is missing, the default value for a limit order is DAY, the default value for a market order is IoC, the default value for a stop limit order is DAY, and the default value for a stop order is DAY – this can vary from exchange to exchange.

From this point forward we refer to the combinations (i) LO and IoC, and (ii) LO and FoK, as marketable limit orders (MLOs). Similarly, we refer to the combinations (i) MO and FoK, and (ii) MO and IoC, as market orders (MOs), and we refer to the rest of the combinations as “other order types”. This is summarized in the following table.

		Time-in-force			
		DAY	GTT	IoC	FoK
Order type	Limit order	Other	Other	MLOs	MLOs
	Market order	×	×	MOs	MOs
	Stop limit order	Other	Other	×	×
	Stop order	Other	Other	×	×

2.3 Definition of latency

Latency is the sum of the time delays that occur in the course of the communications between a trader and a market. Such delays happen between the time an exchange streams market data to a trader, the time at which the trader processes the information and decides to trade, and the time at which the exchange receives and processes the order from the trader.

Next, we discuss the one-to-one correspondence between the clock in the exchange and the trader's clock. Figure 2.1 depicts how latency affects the flow of information between the exchange and the trader. We map the time of the events in the exchange to the time of the events for a trader who experiences latency in the marketplace. At time $e_0 = t_0$ the exchange sends data to the trader, the trader receives the data at time $t_1 = e_1$, where $t_1 - t_0 = e_1 - e_0$ is the observational latency. From $t_1 = e_1$ to $t_2 = e_2$, the trader processes market information and decides to trade. At time $t_2 = e_2$ the trader sends an instruction to the exchange and the exchange receives the instruction at $e_3 = t_3$. Between $e_3 = t_3$ and $e_4 = t_4$ the exchange processes the trade attempt and at time $e_4 = t_4$ the exchange notifies the agent of the outcome of the trade attempt. The trader receives the notification at t_5 .




Figure 2.1: Diagram of the various layers of latency.

The latency e_0 to e_4 in the exchange (t_0 to t_4 for the trader) maps to the latency t_1 to t_5 in the trader's clock (e_1 to e_5 in the exchange). In this thesis, the models we develop are in the timeline of the trader (blue arrow), where the relevant delays are those from t_1 to t_5 , which, due to the one-to-map correspondence, accounts for the latency period of each trade attempt, including the observational delay. For the convenience of the reader, we repeat some of the material in this subsection in Chapter 5. In particular, we include an expanded version of Figure 2.1 to draw attention to the specific modelling assumptions of that chapter.

2.4 Incentives to use MLOs

In this section we discuss the incentives that liquidity takers have to send MLOs instead of MOs. Latency is at the core of these incentives. Traders need to protect their orders from adverse price movements over the latency period. If latency were to be zero, MLOs would have no reason to exist, and traders would only use MOs. On the other hand, when latency is greater than zero, MLOs offer price protection, preserving the immediacy of MOs. Given that there are no traders with zero latency, this suggests that MLOs are not redundant, but the competitive edge of the MLOs over the MOs is yet to be understood.

To gain some insight into the advantage of using MLOs over MOs we present Toy model 1 and Toy model 2. Then we present Toy model 3, which, by use of an *ad hoc* quadratic penalty, achieves a behaviour similar to that which we observe in Toy models 1 and 2. The *ad hoc* quadratic penalty in Toy model 3 is similar to penalising the slack of the MLO once it is processed by the exchange; if the slack is less than zero then the trader was not generous enough and misses the trade; alternatively, if slack is greater than zero, the trader was more generous than what was necessary, and their assessment of the execution price was not as accurate as it could have been. Thus, the trader penalises having a slack different than zero.⁵ We remark that the other terms in the performance criterion reward high sell prices more than low sell prices.

In all models there are three points in time that are of particular interest. At time $t = 0$ the trader decides the price limit l of the sell MLO to be sent to the exchange – note that a sell MO is the same as a sell MLO with price limit $l = 0$. The trade attempt is for one unit of an asset. Time Δ is when the trade transacts, and time T is when the trader performs a valuation of the position. The fundamental price is $S_t = S_0 + \sigma W_t$, where $\sigma > 0$ and $(W_t)_{t \geq 0}$ a standard Brownian motion.⁶ The initial inventory is Q_0 , and the trader controls the price limit of the sell MLO. In all models, the performance criterion has three components: (i) the cash received from the trade if the order is filled, (ii) the mark-to-market value of the outstanding position, and (iii) a penalty for walking the book at T . The penalty parameter for walking the book is $\alpha \geq 0$.⁷ Note that this penalty gives MOs an advantage over MLOs.

Toy model 1. The investor has a view on the value of the asset at time T , which is equal to $\bar{S} \in \mathbb{R}^+$. See [Ho et al. \(1985\)](#) for a similar construction. The investor chooses the price limit $l \in \mathbb{R}$ to attach to the sell MLO to maximize the money value of the position. This is given by the following performance criterion:

$$V_1(l) = \mathbb{E} \left[S_\Delta \mathbf{1}_{\{l \leq S_\Delta\}} + (Q_0 - \mathbf{1}_{\{l \leq S_\Delta\}}) (\bar{S} - \alpha (Q_0 - \mathbf{1}_{\{l \leq S_\Delta\}})) \right].$$

Here, the quantity $\mathbf{1}_{\{l \leq S_\Delta\}}$ is one if the order is filled, and zero if the order is not filled. If the order is filled, the cash increases by S_Δ .

Toy model 2. The execution price is subject to an independent random shock over the uncertainty window from $(0, \Delta)$. This can be due to other traders consuming the liquidity that the trader was aiming to consume, spoofers providing phantom liquidity, or, more generally, it can be thought of as microstructure noise.⁸ Particularly, we assume that the execution price is

⁵The slack of a sell MLO is the limit price minus the execution price.

⁶Note that for the parameters we use below, namely $S_0 = 1.1$, and $\sigma = 2 \times 10^{-5}$, the probability that $(S_t)_{t \in [0,1]}$ goes negative can be computed explicitly using the reflection principle of the Brownian motion. The probability is equal to $1 - 2\Phi^c(S_0/(\sigma\sqrt{T}))$, where Φ is the cumulative distribution function of a standard normal random variable and $\Phi^c = 1 - \Phi$. Using the values of the parameters just mentioned, one finds that the probability that the price goes negative is zero to the first 32 digits of precision offered by Matlab.

⁷This penalty models the well-studied market friction that the more quantity one wishes to sell, the lower the average price one receives from the sell. A trade ‘walks the book’ when it consumes liquidity from more than one price level.

⁸The truth is that the execution price is uncertain in any firm-liquidity venue. This is an immediate consequence of traders not having zero latency.

$S_\Delta + P$ where $P \sim \mathcal{N}(\mu, \bar{\sigma}^2)$ is independent of $(W_t)_{t \geq 0}$. We consider the following performance criterion:

$$V_2(l) = \mathbb{E} \left[(S_\Delta + P) \mathbf{1}_{\{l \leq S_\Delta + P\}} + (Q_0 - \mathbf{1}_{\{l \leq S_\Delta + P\}}) (S_T - \alpha (Q_0 - \mathbf{1}_{\{l \leq S_\Delta + P\}})) \right].$$

Toy model 3. The trader penalizes the use of more than the necessary price limit to fill the trade and incorporates a penalty $\zeta > 0$ to the squared deviation of the execution price S_Δ and the price limit used $l \in \mathbb{R}$. Explicitly we have:

$$V_3(l) = \mathbb{E} \left[S_\Delta \mathbf{1}_{\{l \leq S_\Delta\}} + (Q_0 - \mathbf{1}_{\{l \leq S_\Delta\}}) (S_T - \alpha (Q_0 - \mathbf{1}_{\{l \leq S_\Delta\}})) - \zeta (l - S_\Delta)^2 \right].$$




Figure 2.2: Plots for the functions $V_1, V_2, V_3 : \mathbb{R}^+ \rightarrow \mathbb{R}$. We use the following parameters: $a = 2 \times 10^{-5}$, $S_0 = 1.1$, $Q_0 = 2$, $\sigma = \bar{\sigma} = 2 \times 10^{-5}$, $\Delta = 0.5$, $T = 1$, $\bar{S} = 1.1 + 1 \times 10^{-5}$, $\mu = -3 \times 10^{-5}$, and $\zeta = 10^2$.

Figure 2.2 shows the functions V_1 , V_2 , and V_3 for a range of values of the limit price $l \in \mathbb{R}$. We observe that MOs are suboptimal in the three toy models.⁹ Toy models 1 and 2 show incentives for traders to use MLOs over MOs: namely, when (i) traders have differing beliefs on the price, or when (ii) traders seek price protection against random shocks over a period of uncertain execution. Another feature of the three toy models is that the optimal price to sell is below the current price $S_0 = 1.1$. In Chapter 5, when we study optimal liquidation with delay, we observe that for a portion of the trading horizon traders find that it is optimal to send sell MLOs above the current price.

We observe that Toy model 3 can replicate both the sub-optimality of MOs over MLOs, and the optimality of the maximum being to the left of S_0 . Both features are achieved by means of the *ad hoc* quadratic penalty $\zeta (l - S_\Delta)^2$.

Chapter 3 touches on stylized features of Toy model 3, namely, a quadratic penalty on what we call ‘discretion’, which is a proxy for the quantity $l - S_\Delta$. Chapter 4 and Chapter 5 leverage on the core incentives described through Toy model 1 and Toy model 2. In particular, Chapter 5 shows empirical evidence of what we define as ‘flickers’, which play a similar role to that of the shock P in Toy model 2.

⁹Recall that a sell MO is a sell MLO with limit price equal to zero.

2.5 Previous work on latency

A number of authors have looked at the problem of latency in electronic markets. [Moallemi and Sağlam \(2013\)](#) look at the cost of latency in high-frequency trading. The authors employ market data to quantify the costs of latency in equity markets and show that the latency costs are of the same order of magnitude as other trading costs, such as commissions and exchange fees. Their data are publicly available from the NYSE and do not contain trader identification. Thus, the authors do not analyze latency for individual market participants. We return to the work of [Moallemi and Sağlam \(2013\)](#) in the chapters that follow.

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 to study the effect of latency on the efficacy of the execution strategy. The authors backtest their strategy on publicly available data for US Treasury bonds and show that the advantage of observing the LOB to compute volume imbalance dissipates quickly as latency increases. [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.

On latency arbitrage, [Cohen and Szpruch \(2012\)](#) develop a model where two investors have different speeds to interact with the market of a given security; they derive the fast trader's optimal strategy when there is distributional knowledge of the slow trader's actions. They show that taxation can lead to the elimination of profits from front-running strategies, possibly leading to an increase in market efficiency.

The work of [Hautsch et al. \(2018\)](#) provides a general theory on latency-induced limits to arbitrage. They show that having stochastic latency imposes limits to arbitrage and that they are increasing with respect to the expected latency (among other factors). Their work shares similarities with Chapter 5 where we study stochastic delays.

In market making, [Gao and Wang \(2020\)](#) show how an agent provides liquidity to the LOB of large-tick stocks when latency is deterministic and fixed during the trading horizon. In their model, the market maker may at each discrete time step (i) cancel past orders, (ii) update the limit prices of current orders, or (iii) do nothing. The authors find that as latency increases, the profits from making markets decrease because adverse selection costs increase.

Recent literature on high-frequency trading and algorithmic trading discusses various characteristics of trading and how traders use speed to obtain informational advantages; see [Lehalle and Laruelle \(2018\)](#). 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 [Bayraktar and Ludkovski \(2010\)](#).

In the material that follows, we devise latency-optimal trading strategies. We do not study risk preferences over potential price improvements –although it is certainly an interesting topic for future research. We believe that the models we develop may lead to social benefits, since they offer a way to stop the arms race to being faster in the marketplace. Finally, our assumptions

are non-restrictive. For example, in all of the models that we present, we account for the cases where latency is random or deterministic.

Chapter 3

The shadow price of latency

In foreign exchange (FX) markets, the supply of spot currency pairs displayed in the LOB might undergo thousands of updates over very short periods of time, see [Gould et al. \(2016\)](#). When the best bid and best ask prices in the LOB are changing very rapidly, liquidity takers face a moving target problem as a consequence of their latency in the marketplace. Liquidity takers send MLOs that aim at a price and quantity they have observed in the LOB, but by the time their order is processed by the exchange, prices or quantities may have worsened, so the order cannot be filled; alternatively, prices or quantities may have improved, so the order is filled at better prices.

In this chapter, we show how to compute the price that a liquidity taker is willing to pay to reduce their latency in the marketplace. We summarize the steps we take. One, we develop a latency-optimal strategy to improve the chances that liquidity taking orders are filled. Two, we compute the costs borne by the strategy. These costs consist of the slippage received to ensure that the order is filled, which is negative if prices or quantities worsen, or is positive if prices or quantities improve. Three, we build a function that links latency, fill ratios, and costs borne by the latency-optimal strategy. With this function, we compute the shadow price of latency that a trader would be willing to pay for co-location and hardware to reduce their latency in the marketplace.

Specifically, in our approach we show how the FX trader chooses the price limit of the MLO in an optimal way to improve fill ratios over a period of time (days, weeks, months), while keeping orders exposed to receiving a price improvement. Traders can buffer the adverse effects of missing a trade by specifying a price limit in their market orders to increase the probability of filling the order when it is processed by the exchange. This price limit consists of the best price seen by the trader in the LOB, plus a discretionary buffer that determines the number of ticks the order can walk the LOB. This buffer does not preclude the order from being filled at better prices if the LOB is updated with more favourable prices or quantities.¹

Increasing fill ratios is costly. Everything else being equal, the chances of filling an MLO increase if the order can walk the LOB. Thus, there is a tradeoff between ensuring high fill ratios and the execution costs borne by the trading strategy. In our model, the problem solved by the

¹The model we develop can be used in continuous time, regardless of whether there are pending orders or not.

trader balances this tradeoff by optimizing the price limit specified in the MLO, while targeting a fill ratio over a trading horizon. The trader’s optimal strategy specifies the discretion for each transaction depending on the proportion of orders that have been filled, how far the strategy is from the target fill ratio, the cost of walking the LOB, and the volatility of the exchange rate.

We employ a proprietary data set of FX trades to analyze the performance of the optimal strategy developed here. The data are provided by LMAX. We use transaction data for two FX traders, trader 1 (T1) and trader 2 (T2), to compare the fill ratios they have achieved to those attainable with the optimal strategy derived in this chapter. Both traders employ MLOs with every trade attempt. T1 employs MLOs of FoK type, and T2 uses MLOs of IoC type. The data spans a set of dates from December 2016 to March 2017. During this period both traders filled between approximately 80% and 90% of their liquidity taking orders in the currency pair US dollar and Japanese yen, which we refer to as the USD/JPY pair. Latency causes T1 and T2 to miss trades because by the time the exchange processes their orders, the best bid and best ask prices in the LOB are updated. The effect of latency on trade fills is exacerbated during times of heightened volatility in the exchange rate. When volatility is arranged in quartiles, we find that 36.5% (*resp.* 40.0%) of T1’s (*resp.* T2’s) unfilled trades occur in the top quartile of volatility.

Clearly, due to latency, trades are not filled because the market moves away from the prices and quantities that T1 and T2 are attempting to achieve. If T1 and T2 were to return to the market to complete unfilled trades, the prices they would receive are highly likely worse than those at which they attempted in the original trade. We compare the costs of the latency-optimal strategy to the mark-to-market costs of orders that walk the LOB until filled. We assume these orders are sent 20 ms or 100 ms after the exchange receives the instruction to execute the orders that could not be filled. We find that the average mark-to-market cost for T1’s missed trades is 2.33 and 2.88 ticks and for T2 is 3.18 and 3.06 ticks, 20 ms and 100 ms later, respectively (in the currency pair USD/JPY one tick is 10^{-3} JPY).^{2,3}

We employ the optimal strategy developed here to show the tradeoff between increasing fill ratios and the costs incurred by the strategy. For example, for a particular choice of model parameters, we show that T1 and T2 can increase the percentage of filled trades, during the period 5 December 2016 to 31 March 2017, to 99% for both traders. The increase in the fill ratios is due to the discretion included in the liquidity taking orders to walk the book, which come at a cost. In this example, the average cost incurred by T1 to fill missed trades is 1.76 ticks and for T2 is 1.24 ticks. On the other hand, the mark-to-market average cost of filling the missed trades (which were filled by the optimal strategy) with market orders 20 ms and 100 ms later is 2.01 and 2.58 ticks respectively for T1, and 2.82 and 2.75 ticks respectively for T2.

²The average mark-to-market cost is calculated by taking the difference between the cost of filling all the volume missed in the order and the cost that the trader initially attempted to pay to fill the order, and dividing by the missed volume of the order. This quantity is multiplied by 10^3 to express it in ticks.

³From February 2018, the tick size for the pair USD/JPY is 10^{-4} JPY in LMAX Exchange Tokyo.

The performance of the optimal strategy is more remarkable during times of heightened volatility of the exchange rate. In the top quartile of volatility, the average cost of filling missed trades is 1.88 ticks when T1 employs the optimal strategy, while the mark-to-market average cost of filling these missed trades with market orders 100 ms later is 3.04 ticks. Similarly, for T2 the average cost of filling missed trades with the optimal strategy is 1.86 ticks, while the mark-to-market average cost of filling the missed trades with a market order 100 ms later is 3.31 ticks.

Finally, we use the proprietary data to map various levels of latency to the corresponding percentage of filled orders. We use this mapping to calculate the shadow price of latency that T1 would be willing to pay to reduce latency in the marketplace. We show that T1 would be better off employing the latency-optimal strategy developed here, instead of investing in hardware and co-location services to reduce latency.⁴ The latency-optimal strategy is superior because it not only achieves the same fill ratios as those obtained with better hardware and co-location, but it scoops price improvements that stem from orders arriving with latency at the exchange. Note that if latency is reduced with hardware and co-location, T1's liquidity taking strategy would be less affected by the moving target problem, so fill ratios are high, but would benefit little from changes in the LOB that, due to latency, provide price improvements.

Our work is the first to employ proprietary data to analyze the effect of latency on the fill rates of liquidity taking strategies in FX markets and to compute the shadow price of latency for particular traders. Latency is specific to each market participant, and the effects of latency depend on the type of strategy employed by each trader. Access to proprietary data of FX transactions allows us to quantify the effects of latency on the efficacy of trading strategies and to develop and test latency-optimal strategies to compute the shadow price of latency for types of market participants.

The remainder of the chapter is organized as follows. In Section 3.1 we define latency and propose various ways to measure the fill ratios of MLOs. In Section 3.2 we describe the proprietary data set that we employ and we present empirical evidence of the relationship between fill ratios and the volatility of the exchange rate. In Section 3.3 we present the model of fill ratios and solve the trader's dynamic optimization problem, where we also assume that the trader makes her model of fill ratios robust to misspecification. In Section 3.4 we show how we estimate model parameters, and in Section 3.5 we show the performance of the trading strategy. Section 3.6 discusses the tradeoff between latency and fill ratios and computes the shadow price of latency for T1 and Section 4.5 concludes.

⁴To conclude this, it is necessary to know the cost of investing in hardware and co-location.

3.1 Latency and fill ratios

3.1.1 Latency

Latency is usually defined as the time “delay between sending a message to the market and it being received and processed by the exchange. Sometimes the time it takes for the exchange to acknowledge receipt is also accounted for”, see [Cartea et al. \(2015\)](#). This concept is useful if traders are only interested in measuring how long it takes a message to reach the exchange once it has been sent. However, this definition is ‘narrow’ because it only takes into account hardware capacity and co-location to measure latency and does not account for other factors that affect how quickly the trader can process and react to market information before sending an instruction to the exchange.

In this chapter we employ a broader definition of latency. To the time delay considered in the narrow definition of latency we add the time it takes the trader to execute the trading strategy. Thus, latency consists of four layers of time delay associated to these events: (i) The exchange streams LOB data to the trader. (ii) The trader receives the data from the exchange and processes this and other relevant information to make a decision. (iii) Once a decision is made, the trader sends a message to the exchange. (iv) The exchange receives and processes the order. The decision of the trader could be to trade using limit orders or MLOs, or to amend or cancel a limit order already resting in the LOB, see [Moallemi and Sağlam \(2013\)](#) and [Stoikov and Waeber \(2016\)](#) who employ a similar definition of latency.

As we indicated in Chapter 2.2, MLOs consist of two types of orders: Immediate-or-Cancel (IoC) or Fill-or-Kill (FoK). IoC and FoK are orders in which the trader specifies the maximum (minimum) price they are willing to pay (receive) when buying (selling) currency pairs. IoC is an order to buy or sell currency pairs which must be executed immediately, obeying the price limit, and any portion of the volume of the order that cannot be filled at the desired price limit is cancelled. FoK is an order to buy or sell currency pairs which must be executed immediately in full or cancelled.

We provide an example of the steps in the life of a FoK order to illustrate the effect of latency on the fill of the trade.

Assume the exchange streams a quote to a trader at time t_0 . The trader receives the quote at time $t_1 > t_0$ and at time $t_2 > t_1$ the trader sends a buy FoK order for 100 lots of the currency pair USD/JPY. The limit price of the FoK order is the best ask streamed by the exchange at time t_0 . Finally, the exchange receives the order at time $t_3 > t_2$, and processes the order at time $t_4 > t_3$, thus our measure of latency is $t_4 - t_0$. The order will be filled in full if at time t_4 the best ask in the book is equal to or lower than (i.e., price improvement) the best ask at time t_0 and there are at least 100 lots posted at that price, otherwise the order is not filled. In either case, the exchange sends a message to the trader notifying the outcome of the trade.

The effect of latency on the outcome of strategies depends on a number of factors that affect the supply and demand of liquidity in the FX market, including: magnitude of latency, type of order sent by trader, volatility of the mid-exchange rate, and trend in the mid-exchange rate.

On the supply side of liquidity, the number of updates in the LOB will determine if the trade will be filled at the price, or a better price, than that specified in the FoK order, or if the order is not filled because prices have worsened. Changes in the best quotes of the LOB are difficult to predict and depend on many factors such as news announcements and other idiosyncratic needs of the liquidity providers who send, update, and amend their limit orders in the LOB.

On the demand side of liquidity, the effect of latency will depend on the type of strategy employed by the trader. For example, for informed trading strategies, i.e., buy before rates increase or sell before rates decrease, latency will increase the chances of missing trades and decrease the chances of obtaining price improvements.

3.1.2 Fill ratios and post-trade analysis

Fill ratios of orders sent by a trader to an exchange can be computed in a number of ways. A trader who only sends FoK orders will receive fills of all volume or nothing, as opposed to IoC orders, which can be partially filled.

A measure of fill ratio is important because it summarizes how effective an agent's strategy is in completing trades. Here, we define three measures of fill ratio and show how to extend these measures to include relevant post-trade information such as (i) orders that were filled at better prices (i.e., price improvement), and (ii) orders that were not filled, but could have been filled at worse prices. Below we provide more details about the post-trade information we include in the measures of fill ratio. Next, we describe the variables and notation we employ.

We denote by n the number of MLOs employed to compute the trader's fill ratio. Let $\tau_1, \tau_2, \dots, \tau_n$ be the times at which the trader sends liquidity taking orders to the exchange and denote by $\ell_1, \ell_2, \dots, \ell_n$ the times at which the orders are processed by the exchange, $\ell_i > \tau_i$.

MLOs specify an average limit price per unit, denoted by L_i , which is the maximum (minimum) average price the trader is willing to pay (receive) per unit of bought (sold) of the currency pair.⁵ When the exchange processes the trader's message, the order is filled if there is enough liquidity posted in the LOB at prices and volumes that obey the average price limit L_i . An MLO buy (sell) with limit price $L_i = \infty$ ($L_i = 0$) will walk all levels of the LOB until the order is filled in full, in which case we denote by B_i the average price per unit of currency pair obtained by the order that is filled in full.

The quantity $L_i - B_i$ (which can be computed only after the exchange receives and processes the order) is important because it gives post-trade information about different outcomes such

⁵In practice, exchanges and over-the-counter facilities have their own rules regarding price limits. Some employ marginal prices as price limit (that is, the price limit applies to the last unit bought or sold), and others employ average prices of the order as we assume here. Our model works in both cases.

as: (i) by how much did prices move to cause the liquidity taking order to miss its target, which we refer to as *potential slippage*, or (ii) by how much did the execution price improve, which we refer to as *price improvement*. For example, when the trader sends a buy order with an average price limit of L_i per unit of currency pair, there is a price improvement if $L_i - B_i > 0$. On the other hand, if the average price required to fill the order is higher than average limit price the trader is willing to pay per unit of currency pair, then the potential slippage is $L_i - B_i < 0$ – note that in this second case the trader does not incur the potential slippage because the order is not filled in full due to the limit price.

Moreover, the variable $I_i \in \{-1, 1\}$ is used to indicate the direction of the liquidity taking order, where 1 denotes a buy order and -1 denotes a sell order. Finally, V_i denotes the volume of currency pair lots the trader attempts to fill, and V_i^f denotes the filled volume, clearly $V_i^f \leq V_i$.

We define three ways to compute a fill measure of n consecutive trades. These measures are based on: (1) number of trades, (2) average volume of trades, and (3) volume of trades.

1.

$$F_n = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{I_i L_i \geq I_i B_i\}}, \quad (3.1)$$

2.

$$F_n^{AV} = \frac{1}{n} \sum_{i=1}^n \frac{V_i^f}{V_i}, \quad (3.2)$$

3.

$$F_n^V = \frac{\sum_{i=1}^n V_i^f}{\sum_{i=1}^n V_i}, \quad (3.3)$$

where $\mathbf{1}$ is the indicator function and the superscript V in F_n^V stands for volume and the superscript AV in F_n^{AV} stands for average volume.

One can enhance the three measures above by incorporating more information on the LOB, so strategies that monitor the fill ratio can also trade off price improvement and slippage. Specifically, one can add to the measures above the component

$$F(n) = \zeta \frac{1}{n} \sum_{i=1}^n I_i (L_i - B_i), \quad (3.4)$$

where ζ is a non-negative parameter. Note that $F(n) \in \mathbb{R}$. The units of ζ are such that the measure $F(n)$ is unitless. Specifically, the units of ζ are $1/c$, where c is the unit of currency B in the pair A/B . For example, for the USD/JPY pair, c denotes ¥. For simplicity, in our analysis below we employ the fill measures in (3.1) and (3.2), although we aim to describe the general case when one adds $F(n)$ to measures (3.1) and (3.2).

3.2 Empirical analysis

3.2.1 Data

We employ a set of proprietary data of transactions for T1 and T2 during the period 2 December 2016 to 31 March 2017 in the currency pair USD/JPY. Both traders are liquidity takers, i.e., they do not send limit orders to provide liquidity in the LOB, during the period we study. We have the following information: side of order (buy or sell), time-stamp of when the exchange processed the order, volume, currency pair, type of order (FoK, IoC), limit price L_i , and whether the order was filled and at what price(s). Finally, we have the LOB for the pair USD/JPY at millisecond intervals. A summary of the main variables is given in Table 3.1.

	Number of attempts N	Successful filled trades (incl. partial)	Successful filled trades (only full)	Avg. vol. attempted $\times 10^5$ USD	Avg. filled vol. $\times 10^5$ USD	Order type	F_N	F_N^{AV}	F_N^V
T1	116,912	106,878	106,878	4.5	4.1	FoK	90.78%	90.78%	90.54%
T2	68,129	61,440	59,303	1.2	1.0	IoC	87.04%	88.62%	79.00%

Table 3.1: Proprietary data and measures of fill ratios for period 2 December 2016 to 31 March 2017. Here, ‘vol.’ stands for volume.

From 2 December 2016 to 31 March 2017 the best prices and quantities were updated over 63 million times, i.e., a mean of 746,871 of daily updates with standard deviation of 218,298. During this period, both traders send orders that aim at the best bid or best ask price that was streamed to them by the exchange. T1 sent FoK orders with zero discretion and T2 sent IoC orders with zero discretion. For a buy MLO we define ‘discretion’ as the difference between the price limit and the best offer seen by the trader when she devised the strategy. Similarly, for a sell MLO we define discretion as the difference between the best bid and the price limit of the MLO.

3.2.2 Fill ratios and volatility

The liquidity provided in the LOB is constantly being updated as a result of news, changes in the willingness to trade of liquidity providers, and the arrival of liquidity taking orders, see for example Almgren (2012) for a discussion of trading with stochastic liquidity. Most of the updates of limit orders (cancellations, amendments, and arrival of new limit orders) take place in the best quotes and ticks closest to the prevailing mid-exchange rate. Traders with non-zero latency face a moving target problem when they attempt to hit bids and lift offers in the LOB. This problem is exacerbated during periods of high volatility of the exchange rate.

The top figure in the left-hand panel of Figure 3.1 shows the evolution of the fill ratio of T1’s trade attempts, using measure (3.1) with $n = 50$, during 26 January 2017. The first observation in the figure is computed when the trader has attempted the 50th trade and the fill ratio is updated every time there is a new attempt.

The bottom figure in the left-hand panel shows the volatility of the mid-exchange rate over the same period. Throughout this chapter volatility is computed as the standard deviation of the log-returns of the micro-exchange rate (sampled every 500 milliseconds) over a 10-minute rolling window.⁶ The micro-exchange rate at time t is denoted by m_t and given by

$$m_t = \frac{P_t^a Q_t^b + P_t^b Q_t^a}{Q_t^a + Q_t^b}, \quad (3.5)$$

where Q_t^a and Q_t^b are the quantities available at the best ask price (denoted by P_t^a) and the best bid price (denoted by P_t^b), respectively.

From Figure 3.1, left-hand panel, we observe the negative correlation between volatility and the fill ratios achieved by T1's strategy, which always aims at the best bid or best offer price she observed before sending the MLO. Clearly, volatility is high when the micro-exchange rate undergoes many changes and updates – this explains why the number of unfilled MLOs (with zero discretion to walk the book) is higher when volatility is high.

The right-hand panel of Figure 3.1 shows a scatter plot of fill ratio and volatility. The solid line in the figure is obtained by performing an Ordinary Least Squares (OLS) regression of fill ratio on volatility.




Figure 3.1: Left-hand top panel: fill ratio using measure (3.1) with $n = 50$. Left-hand bottom panel: volatility of micro-exchange rate. Right-hand panel: Scatter plot and OLS regression of volatility and fill ratio. Data are from 26 January 2017.

3.3 Model

In this section we describe the model of fill ratios and the dynamic optimization problem solved by the trader. We develop the model of fills in three steps. First, we present a model for the dynamics of volatility of the micro-exchange rate. Second, we propose a model of fill ratios that includes the effect of volatility of the micro-exchange rate on fills. Third, we specify how MLOs with discretion to walk the LOB affect the dynamics of fill ratios.

⁶We chose the number of observations n to calculate the measure of fill ratio and the time-window to calculate the volatility of the mid-exchange rate as follows. For each day of the sample, we compute the correlation between the measure of fill ratio and the volatility of the mid-exchange rate for various values of n and various values of the time window. When $n = 50$ and the volatility window is ten minutes, the average of the daily correlations is smallest (i.e., is closest to -1).

Volatility of micro-exchange rate

We denote the volatility of the micro-exchange rate, see (3.5), by $v = (v_t)_{\{0 \leq t \leq T\}}$ and assume it satisfies the stochastic differential equation (SDE)

$$dv_t = \kappa (\bar{v} - v_t) dt + \sigma_v v_t dW_t^v, \quad (3.6)$$

where $\kappa > 0$ is the exponential speed at which volatility reverts to its long-term level, denoted by $\bar{v} > 0$, $\sigma_v > 0$ is a dispersion parameter (i.e., volatility of volatility), and $W^v = (W_t^v)_{\{0 \leq t \leq T\}}$ is a standard Brownian motion.

In our analysis we consider other models for volatility of the micro-exchange rate. In A.2.3, we show results when volatility follows the Ornstein-Uhlenbeck process

$$dv_t^{\text{OU}} = \kappa^{\text{OU}} (\bar{v}^{\text{OU}} - v_t^{\text{OU}}) dt + \sigma_v^{\text{OU}} dW_t^v, \quad (3.7)$$

and

$$dv_t^{\text{H}} = \kappa^{\text{H}} (\bar{v}^{\text{H}} - v_t^{\text{H}}) dt + \sigma_v^{\text{H}} \sqrt{v_t^{\text{H}}} dW_t^v, \quad (3.8)$$

see Heston (1993). We assume that the parameters κ^{H} , \bar{v}^{H} , σ_v^{H} satisfy the Feller condition, so that the process $(v_t^{\text{H}})_{\{0 \leq t \leq T\}}$ is strictly positive. However, the value of the volatility process in model (3.7) is not necessarily positive. In A.2 we prove that the performance of the latency-optimal strategy is robust to the choice of the volatility model for the micro-exchange rate.

Fill ratio: orders with no discretion to walk the LOB

The fill ratio of MLOs (buy and sell) with no discretion to walk the LOB is denoted by $F = (F_t)_{\{0 \leq t \leq T\}}$ and satisfies the SDE

$$dF_t = (\lambda (\bar{F} - F_t) - b(v_t - \bar{v})) dt + \sigma_F dW_t^F. \quad (3.9)$$

Here, the parameter b determines the effect of the volatility of the micro-exchange rate on the fill ratio of the trader's strategy. The parameter $\lambda > 0$ is the exponential speed at which the fill rate reverts to the level $\bar{F} \in \mathbb{R}$, the dispersion parameter σ_F is a non-negative constant, and $W^F = (W_t^F)_{\{0 \leq t \leq T\}}$ is a standard Brownian motion uncorrelated to W^v .⁷ Recall that we propose three ways to compute fill ratios, see (3.1), (3.2), (3.3), so here F refers to one of those choices. We note that equations (3.1), (3.2), and (3.3) give values between zero and one, while (3.9) asserts that $F_t \in \mathbb{R}$. As stated before, we can add (3.4) to any of the three ways to compute

⁷Here we assume that the Brownian motions W^F and W^v are independent because the impact of volatility v_t on the dynamics of the fill ratio is included in the drift of the fill ratio dynamics in (3.9). However, note that when the Brownian motions W^F and W^v are correlated, one can derive the latency-optimal strategy in closed-form when the volatility of the mid-exchange rate is as in (3.6) or (3.7) – but not for the specification of volatility in (3.8). To show this, observe that when W^F and W^v are correlated with parameter ρ , the PDE in (3.17) would include an extra term, which for model (3.6) is $\sigma_F \sigma_v \rho v V_{fv}$, for model (3.7) is $\sigma_F \sigma_v \rho V_{fv}$, and for model (3.8) is $\sigma_F \sigma_v \rho \sqrt{v} V_{fv}$. The ansatz we proposed in (3.18) works for the volatility models (3.6) and (3.7). Specifically, for model (3.6), the term $\sigma_F \sigma_v \rho v h_5(t)$ collects in the ODE for h_3 which can be solved analytically. For model (3.7), the term $\sigma_F \sigma_v \rho h_5(t)$ collects in the ODE for h_0 which can be solved analytically. Finally, for model (3.8), the ansatz in (3.18) is not suitable and we cannot obtain a closed-form solution.

fill ratios, in which case the fill ratios are not bound to be in $[0, 1]$ anymore. Nonetheless, even if we do not include the component in (3.4) as part of the measure for fill ratios, the model still proves useful. In particular, with the parameters we estimate in Section 3.4, we observe that the probability of $(F_t^{\delta^*})_{t \in [0, T]}$ being above one or below zero is small.⁸

The parameters λ , \bar{F} , σ^F , and the function $h(v)$ are specific to each trader in the FX market. Latency and type of liquidity taking strategy are the key characteristics that make the fill ratio dynamics specific to each trader. For example, the dynamics of the fill ratio for a trader with zero latency are as in (3.9) with parameters $\sigma_F = 0$, $\lambda > 0$, $\bar{F} = 1$, and $b = 0$.

The term $\sigma^F dW^F$ represents the combined effect of two sources of uncertainty in the dynamics of the fill ratio, both of which depend on latency. One, the volatility of the micro-exchange rate will affect fill ratios via the term $-b(v_t - \bar{v})$ in the drift of (3.9), and we observe noise around this effect. Two, between the time the trader is streamed LOB data and the time the exchange receives the trader's MLO, the LOB will have processed instructions from other market participants, hence the fill of the order is uncertain.

Figure 3.2 shows F_{50} (top panels) and the volatility of the micro-exchange rate (bottom panels) using data from 4 January 2017 (left panels) and 5 January 2017 (right panels). Visual inspection lends strong support to the mean reverting feature we propose. The dotted lines are estimated values for the mean reversion levels \bar{F} (see equation (3.9)) and \bar{v} (see equation (3.6)).




Figure 3.2: Top panel: fill ratio using measure (3.1) with $n = 50$. Bottom panel: volatility of micro-exchange rate. Left-hand panels: Data are from 4 January 2017. Right-hand panels: Data are from 5 January 2017.

Fill ratio: orders with discretion to walk the LOB

Liquidity takers acknowledge that however fast their computers are, or efficient their software is, or whether their hardware is co-located, the time to reach the market and execute orders is not zero.

One alternative to compensate for not reaching the market in time is to specify a maximum (minimum) price the trader is willing to pay (receive) when sending an MLO. As discussed

⁸For example, with the parameters we estimate in Section 3.4, only less than 0.5% of simulated paths are outside $[0, 1]$.

above, the price limit consists of the best offer (*resp.* bid) price observed by the trader plus (*resp.* minus) the discretion δ . This discretion is the additional slippage she is willing to incur to fill the trade – the discretion to walk the LOB does not preclude the strategy from benefiting from price improvements. For example, if the trader sees the bid (ask) exchange rate at 1.000, then an immediate-execution limit sell (buy) order with discretion $\delta = 0.001$ will be filled by the exchange with limit buy (sell) orders resting in the LOB down (up) to a minimum (maximum) bid (offer) price of 0.999 (1.001).

When the trader includes positive discretion in her MLOs she will exert upward pressure on the fill ratio of her strategy because the order can walk the LOB. Thus, we assume that the trader's fill ratio, denoted by the controlled process $F^\delta = (F_t^\delta)_{\{0 \leq t \leq T\}}$, follows the SDE

$$dF_t^\delta = \left(\lambda (\bar{F} - F_t^\delta) + a \delta_t - b(v_t - \bar{v}) \right) dt + \sigma_F dW_t^F, \quad (3.10)$$

with a and b non-negative constants; when $a = 0$ we obtain (3.9). Note that according to (3.10) the permanent effect of δ_t on F_t^δ (when $a > 0$) is counteracted by the mean-reverting component $\lambda (\bar{F} - F_t^\delta)$. This is a desirable feature of the model for fill ratios, because when one exerts an upward pressure on the fill ratio, this effect will last for as long as we consider these trades in the construction of the fill ratio process; see (3.1), (3.2), (3.3).

As usual, we work on a completed filtered probability space $(\Omega, \mathbb{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$, with \mathcal{F}_t the natural filtration generated by the 2-dimensional Brownian motion $\mathbf{W} = (W^F, W^v)$, and we refer to \mathbb{P} as the trader's reference measure.

3.3.1 Model uncertainty

The trader acknowledges that the model of fill ratios (3.10) may be misspecified, so she considers alternative dynamics of the fill ratio to make the model robust to misspecification, see [Cartea et al. \(2017\)](#).⁹ This ambiguity about model choice, or model uncertainty, affects the optimal strategy, i.e., affects the discretion of the orders she sends to the exchange. We incorporate the trader's ambiguity about model choice in two steps. First, we characterize alternative measures that describe the fill ratio dynamics considered by the trader, and then we determine how the trader chooses the reference measure \mathbb{P} or one of the alternative measures.

The trader considers candidate measures $\mathbb{Q}(x^F)$ that are equivalent to \mathbb{P} and characterized by the Radon-Nikodym derivative

$$\frac{d\mathbb{Q}(x^F)}{d\mathbb{P}} \Big|_t = \exp \left\{ -\frac{1}{2} \int_0^t (x_u^F)^2 du - \int_0^t x_u^F dW_u^F \right\}, \quad (3.11)$$

where $(x_t^F)_{\{0 \leq t \leq T\}}$ is an \mathcal{F}_t -adapted process such that $\left(\frac{d\mathbb{Q}(x^F)}{d\mathbb{P}} \Big|_t \right)_{0 \leq t \leq T}$ is a martingale. Thus, we denote by \mathfrak{Q} the class of alternative measures

$$\mathfrak{Q} = \left\{ \mathbb{Q}(x^F) \mid x^F \text{ is } \mathcal{F} \text{ - adapted and } \left(\frac{d\mathbb{Q}(x^F)}{d\mathbb{P}} \Big|_t \right)_{0 \leq t \leq T} \text{ is a martingale} \right\}.$$

⁹Classical references for model uncertainty are in [Hansen and Sargent \(2001\)](#) and [Hansen and Sargent \(2011\)](#).

Next, the trader penalizes deviations from the reference measure using the relative entropy from t to T , which is given by

$$\mathcal{H}_{t,T}(\mathbb{Q}|\mathbb{P}) = \frac{1}{\varphi} \log \left(\frac{d\mathbb{Q}/d\mathbb{P}|_T}{d\mathbb{Q}/d\mathbb{P}|_t} \right). \quad (3.12)$$

Here the parameter φ is a non-negative constant that represents the trader's degree of ambiguity aversion. If the trader is confident about the reference measure \mathbb{P} , then the ambiguity aversion parameter φ is small and any deviation from the reference model of fill dynamics is very costly. In the extreme $\varphi \rightarrow 0$, the trader is very confident about the reference measure, so she chooses \mathbb{P} because the penalty that results from rejecting the reference measure is too high. On the other hand, if the trader is very ambiguous about the reference model, considering alternative models results in a very small penalty. In the extreme $\varphi \rightarrow \infty$, deviations from the reference model are costless, so the trader considers the worst case scenario.¹⁰

3.3.2 Value function

The trader's performance criterion under measure \mathbb{Q} is given by

$$\mathbb{E}^{\mathbb{Q}} \left[-\phi_{\delta} \int_0^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_0^T (F_s^{\delta} - \hat{F})^2 ds + \mathcal{H}_{0,T}(\mathbb{Q}|\mathbb{P}) \right],$$

and the trader's value function is

$$V(t, f, v) = \sup_{\delta \in \mathcal{A}_{t,T}} \inf_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{t,f,v}^{\mathbb{Q}} \left[-\phi_{\delta} \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^{\delta} - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q}|\mathbb{P}) \right], \quad (3.13)$$

where \hat{F} denotes the target fill ratio and $\hat{\delta}$ is a constant that denotes a fixed discretion to target. Here, the penalty parameters ϕ_{δ} and ϕ_F are non-negative constants. As the value of the parameter ϕ_{δ} increases, the trader submits orders with discretion closer to the target $\hat{\delta}$.¹¹ Similarly, as the value of the parameter ϕ_F increases, the discretion to walk the book increases to ensure that the fill ratio hits the target \hat{F} . The last term on the right-hand side of the value function is the penalty in (3.12). Straightforward choices for the values of the parameters \hat{F} and $\hat{\delta}$ are one and zero, respectively. In practice, under these choices for the values of the parameters \hat{F} and $\hat{\delta}$, one is always below the target fill ratio, in which case we only penalise being away from one. The performance criterion we use, taken together with the dynamics for the fill ratios and volatility, allows us to transform the problem to a standard linear-quadratic stochastic control problem.¹²

¹⁰There are other choices we could consider to account for uncertainty. We choose to work with this entropy penalty for tractability and to obtain closed-form solutions.

¹¹In LMAX Exchange there are many traders who employ a fixed discretion (for all orders) to walk the book. For example, there are traders who for long periods of time, send all orders with a discretion of five ticks to walk the book. This is a simple strategy that decreases, due to latency, the chances of missing trades. In our setup, this fixed discretion for all trades is given by the parameter $\hat{\delta}$ in the performance criterion.

¹²Deriving a performance criterion from first principles rapidly becomes intractable.

Recall that the control variable δ denotes the discretion of the order to walk the LOB, which is measured from the best quote observed by the trader when receiving the quotes from the exchange.

Finally, recall that \mathcal{Q} represents all the measures equivalent to the reference measure \mathbb{P} , and the set of admissible strategies is

$$\mathcal{A}_{0,T} = \left\{ \delta = (\delta_t)_{\{0 \leq t \leq T\}} \mid \delta_t = \mu \left(t, F_t^\delta, v_t \right), \text{ for } \mu \in \mathcal{M}_{0,T} \right\}, \quad (3.14)$$

where

$$\mathcal{M}_{0,T} = \left\{ \mu : [0, T] \times \mathbb{R} \times \mathbb{R}^+ \mapsto \mathbb{R} \mid \exists K \in \mathbb{R} \text{ s.t. } \forall \mathbf{y}, \mathbf{z} \in \mathbb{R} \times \mathbb{R}^+ \text{ and } t \in [0, T] \right. \\ \left. |\mu(t, \mathbf{y}) - \mu(t, \mathbf{z})| \leq K |\mathbf{y} - \mathbf{z}| \text{ and } |\mu(t, \mathbf{y})|^2 \leq K^2 (1 + |\mathbf{y}|^2) \right\}. \quad (3.15)$$

We require the set of admissible strategies to satisfy the condition appearing in (3.14), so that the change of measure specified by (3.11) is valid.

By standard results, the Hamilton-Jacobi-Bellman-Isaacs (HJBI) equation associated to problem (3.13) is given by

$$V_t + \sup_{\delta} \inf_x \left((\lambda (\bar{F} - f) + a \delta - b(v - \bar{v}) - \sigma_F x) V_f + \frac{1}{2} \sigma_F^2 V_{ff} + \frac{1}{2} x^2 \right. \\ \left. + \kappa (\bar{v} - v) V_v + \frac{1}{2} v^2 \sigma_v^2 V_{vv} - \phi_\delta (\delta - \hat{\delta})^2 - \phi_F (f - \hat{F})^2 \right) = 0, \quad (3.16)$$

with terminal condition $V(T, f, v) = 0$. Here, subscripts in the function V represent partial derivatives, for example, $V_t = \partial V / \partial t$.

Proposition 1. *The supremum and infimum in HJBI (3.16) are achieved (in feedback form) at*

$$\delta^* = \hat{\delta} + \frac{a V_f}{2 \phi_\delta}, \quad \text{and} \quad x^* = \varphi \sigma_F V_f,$$

respectively.

Proof. Apply first order conditions to the supremum and infimum terms in (3.16), and check second order conditions to verify that these are the maximizer and minimizer, respectively. ■

Substituting δ^* and x^* in (3.16), the HJBI becomes

$$V_t + (\lambda (\bar{F} - f) - b(v - \bar{v})) V_f + \frac{1}{2} \sigma_F^2 V_{ff} + \frac{a^2 V_f^2}{4 \phi_\delta} - \frac{1}{2} \varphi V_f^2 \sigma_F^2 + a \hat{\delta} V_f \\ + \kappa (\bar{v} - v) V_v + \frac{1}{2} v^2 \sigma_v^2 V_{vv} - \phi_F (f - \hat{F})^2 = 0, \quad (3.17) \\ V(T, f, v) = 0.$$

By inspection of the PDE (3.17) and its terminal condition, we propose the ansatz

$$V(t, f, v) = h_0(t) + h_1(t) f + h_2(t) f^2 + h_3(t) v + h_4(t) v^2 + h_5(t) f v, \quad (3.18)$$

which we substitute in equation (3.17) and obtain a coupled system of ordinary differential equations (ODEs). The system of ODEs can be solved explicitly to obtain closed-form solutions for the deterministic functions $h_0, h_1, h_2, h_3, h_4, h_5$, see A.1.1.

Theorem 1. (Verification) Define the time-dependent deterministic functions

$g_1, g_2, g_3 : [0, T] \mapsto \mathbb{R}$ as

$$g_1(t) = h_1(t) + 2\hat{F}h_2(t) + \bar{v}h_5(t), \quad g_2(t) = -2h_2(t), \quad g_3(t) = h_5(t).$$

and let the ambiguity aversion parameter obey the bound

$$\varphi < a^2 / 2\sigma_F^2\phi_\delta,^{13}$$

which guarantees that we find a solution for h_2 . Then, the optimal discretion to walk the LOB is

$$\delta_t^* = \hat{\delta} + \frac{a}{2\phi_\delta} \left(g_1(t) + g_2(t) \left(\hat{F} - F_t^{\delta^*} \right) + g_3(t) (v_t - \bar{v}) \right), \quad (3.19)$$

and

$$x_t^* = \varphi\sigma_F \left(h_1(t) + 2h_2(t)F_t^\delta + h_5(t)v_t \right). \quad (3.20)$$

The controls in (3.19) and (3.20) are admissible and equation (3.18) is the trader's value function in (3.13).

Proof. For the proof see A.1.2. ■

We explain the intuition of the terms appearing in the optimal discretion (3.19). The first term on the right-hand side of (3.19) is the target discretion $\hat{\delta}$. For example, if the trader imposes an arbitrarily large penalty on deviations from this discretion target, i.e., penalty parameter $\phi_\delta \rightarrow \infty$, then the optimal discretion $\delta_t^* \rightarrow \hat{\delta}$.

The second term, $g_1(t)$, is the time-dependent baseline level of discretion the trader posts when her fill ratio is close to the target \hat{F} and volatility is around its long-term level \bar{v} .

The third term, $g_2(t) \left(\hat{F} - F_t^\delta \right)$, adjusts the strategy if the fill ratio is not on target. The function $g_2(t)$ is positive for $t < T$, so when the fill ratio is below (above) its target \hat{F} the discretion to walk the LOB is increased (decreased).

The last term of the optimal strategy shows how the optimal discretion depends on the volatility of the micro-exchange rate. The function $g_3(t)$ is positive for $t < T$, so when volatility is higher (lower) than its long-term level, the strategy increases (decreases) the discretion of the MLOs.¹⁴

Throughout this chapter we assume that the volatility of the micro-exchange rate follows (3.6). In A.2.2 we derive the optimal strategy when volatility follows (3.7) and (3.8). Furthermore, as a robustness check, in Section 3.5 we analyze the performance of the latency-optimal strategy for these two alternative volatility models.

¹³This bound is to fix the solution of the Riccati ODE satisfied by h_2 . See A.1.1 for the other cases.

¹⁴In Section 3.5, when we implement the strategy, we note that (3.19) does not go negative for either T1 or T2, in any of the scenarios we study.

3.3.2.1 Infinite-horizon latency-optimal strategy

We also derive the latency-optimal strategy when $T \rightarrow \infty$.¹⁵ In A.2.1 we present the infinite-horizon value function of the trader, where the optimal control is

$$\delta_t^* = \hat{\delta} + \frac{a}{2\phi_\delta} \left(c_1 + c_2 \left(\hat{F} - F_t^{\delta^*} \right) + c_3 (v_t - \bar{v}) \right), \quad (3.21)$$

with $c_1 = k_1 + 2\hat{F}k_2$, $c_2 = -2k_2$, and $c_3 = k_5$, and the constants k_1, k_2, k_5 are

$$k_2 = \begin{cases} \frac{2\lambda + \beta - \sqrt{(2\lambda + \beta)^2 + 8\alpha_1\phi_F}}{4\alpha_1} & \text{if } \alpha_1 > 0, \\ -\frac{\phi_F}{\beta + 2\lambda} & \text{if } \alpha_1 = 0, \\ \frac{2\lambda + \beta + \sqrt{(2\lambda + \beta)^2 + 8\alpha_1\phi_F}}{4\alpha_1} & \text{if } \alpha_1 < 0, \end{cases}$$

$$k_5 = \frac{2bk_2}{2k_2\alpha_1 - \lambda - \kappa - \beta}, \quad k_1 = \frac{2\phi_F\hat{F} + (2b\bar{v} + 2\lambda\bar{F} + 2a\hat{\delta})k_2 + \kappa\bar{v}k_5}{\lambda + \beta - 2k_2\alpha_1},$$

where $\alpha_1 = \frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2$, $\beta > 0$ is a discount factor (see (A.20)), and

$$\varphi \leq \frac{1}{2\sigma_F^2} \left(\frac{a^2}{\phi_\delta} + \frac{(2\lambda + \beta)^2}{4\phi_F} \right).$$

When the above bound holds, the solutions to the quadratic equation satisfied by k_2 are real; the other cases remain to be studied. It is straightforward to check that c_2 and c_3 are positive, and k_2 is continuous in α . At the end of Section 3.5 we discuss the performance of the infinite-horizon strategy.

3.4 Estimation of model parameters

In this section we show how to estimate the parameters of the models of volatility and fill ratio.

Volatility parameters. We discretize the continuous-time volatility model (3.6) and write

$$\tilde{v}_i = \tilde{v}_{i-1} + \kappa (\bar{v} - \tilde{v}_{i-1}) \Delta + \sigma_v \tilde{v}_{i-1} \sqrt{\Delta} \epsilon_i, \quad (3.22)$$

where $i \in \{1, 2, \dots, N-1\}$, $0 = t_1 < t_2 < \dots < t_N = T$, $T = 1$, $t_i = iT/N$, ϵ_i i.i.d. random variables, and $\Delta = T/N$.

Assume $\{\tilde{v}_i\}_{i \in \{1, 2, \dots, N-1\}}$ are positive and divide (3.22) by \tilde{v}_{i-1} , so that

$$\frac{\tilde{v}_i}{\tilde{v}_{i-1}} - 1 = -\kappa \Delta + \kappa \bar{v} \Delta \frac{1}{\tilde{v}_{i-1}} + \sigma_v \sqrt{\Delta} \epsilon_i,$$

which we write as the OLS problem

$$\mathbf{y} = \beta_0 \mathbf{1} + \beta_1 \mathbf{s} + \boldsymbol{\epsilon}. \quad (3.23)$$

Here \mathbf{y} is a $(N-1)$ -dimensional vector containing the dependent variables with entries $y_i = \tilde{v}_i/\tilde{v}_{i-1} - 1$, $\mathbf{1}$ is an $(N-1)$ -dimensional vector of ones, β_0 and β_1 are scalar parameters, \mathbf{s}

¹⁵The main purpose of studying this case is to prove that our results are robust to the time horizon T .

is an $(N - 1)$ -dimensional vector containing the regressors $s_i = 1/\tilde{v}_{i-1}$, and ϵ is the vector of residuals with i.i.d. entries.

To obtain parameter estimates for β_0 and β_1 we employ the volatility of the micro-exchange rate at 500 millisecond intervals (see subsection 3.2.2) and use a standard OLS routine to compute the parameter estimates $\hat{\beta}_0$ and $\hat{\beta}_1$. Now, because $\beta_0 = -\kappa \Delta$ and $\beta_1 = \kappa \bar{v} \Delta$, we obtain the estimates

$$\hat{\kappa} = -\hat{\beta}_0/\Delta, \quad \hat{v} = -\hat{\beta}_1/\hat{\beta}_0, \quad \text{and} \quad \hat{\sigma}_v = \sigma_\epsilon/\sqrt{\Delta}, \quad (3.24)$$

where σ_ϵ is the standard deviation of the residuals in (3.23).

Fill ratio parameters. We discretize the continuous-time fill ratio model (3.10) and write

$$\tilde{F}_i - \tilde{F}_{i-1} = \lambda \bar{F} \Delta + a \delta_{i-1} \Delta - \lambda \Delta \tilde{F}_{i-1} - b (\tilde{v}_{i-1} - \bar{v}) \Delta + \sigma_F \sqrt{\Delta} \eta_i, \quad (3.25)$$

where, as above, $i \in \{1, 2, \dots, N - 1\}$, $0 = t_1 < t_2 < \dots < t_N = T$, $T = 1$, $t_i = iT/N$, η_i i.i.d. random variables, and $\Delta = T/N$.

We write (3.25) as the OLS problem

$$\mathbf{z} = \gamma_0 \mathbf{1} + \gamma_1 \mathbf{u} + \gamma_2 \mathbf{w} + \gamma_3 \mathbf{r} + \boldsymbol{\eta}, \quad (3.26)$$

where the vector \mathbf{z} contains the dependent variables, γ_0 , γ_1 , γ_2 , and γ_3 are scalar parameters, the vectors \mathbf{u} , \mathbf{w} , \mathbf{r} are regressors, and $\boldsymbol{\eta}$ is the vector of residuals with i.i.d. entries.

To estimate the impact parameter a in (3.10) we compute the effect of the discretion δ on fill ratios for the set of values $\delta_t = 0, 1, 2, 3, 4, 5$, as follows. Throughout the day we assume that all FoK trades sent by T1 included a fixed discretion of $\delta = 0$ to walk the LOB, and then repeat the analysis assuming $\delta = 1$, and $\delta = 2$, and so on. Thus, we have six scenarios with fixed discretion for the whole trading day. We employ LOB data to determine if the orders would have been filled and create six $(N - 1)$ -dimensional vectors $\mathbf{z}_0, \mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_4, \mathbf{z}_5$, where each vector corresponds to each scenario with fixed discretion. Next, we pool all the data so that the dependent variable on the left-hand side of the regression and the regressors on the right-hand side are $6 \times (N - 1)$ -dimensional vectors. We run a standard OLS routine to compute the parameter estimates $\hat{\gamma}_0, \hat{\gamma}_1, \hat{\gamma}_2$, and $\hat{\gamma}_3$. Then, the parameters of the fill ratio dynamics are

$$\hat{\lambda} = -\hat{\gamma}_1/\Delta, \quad \hat{\bar{F}} = -\hat{\gamma}_0/\hat{\gamma}_1, \quad \hat{b} = -\hat{\gamma}_2/\Delta, \quad \hat{a} = \hat{\gamma}_3/\Delta, \quad \hat{\sigma}_F = \sigma_\eta/\sqrt{\Delta}, \quad (3.27)$$

where σ_η is the standard deviation of the errors in (3.26). To estimate model parameters for T2 we proceed in the same way and use the definition of fill ratio (3.2) because T2 sends IoC orders to the exchange.

Table 3.2 shows parameter estimates for T1 for the trading day starting at 22:05 JST of 3 January 2017 and ending at 22:00 JST on 4 January 2017, which for simplicity we refer to as the trading day of 4 January 2017.

Parameter	\hat{v}	$\hat{\kappa}$	$\hat{\sigma}_v$	$\hat{\lambda}$	\hat{F}	$\hat{\sigma}_F$	\hat{a}	\hat{b}
Estimate	0.0073	12.3887	1.3894	58.7206	0.9331	0.1944	1.0027	527.44

Table 3.2: Parameter estimates for T1’s model, 4 January 2017. All parameters are statistically significant.

Daily estimation of the parameters $(\hat{\kappa}, \hat{\sigma}_v, \hat{v})$, $(\hat{F}, \hat{b}, \hat{\sigma}_F, \hat{\lambda})$ and (\hat{a}, \hat{b}) takes on average 1.5 seconds, 0.5 seconds, and 0.3 seconds, respectively. The estimation is carried out offline with data from the previous day; thus, estimation times do not pose any challenges.

3.5 Performance of strategy

In this section we show the performance of the optimal discretion strategy developed above. In subsection 3.5.1 we employ two days of trading data for T1 to discuss features of the strategy. In subsection 3.5.2 we employ data during the period December 2016 to March 2017 to study the performance of the strategy for T1 and T2 and in subsection 3.5.3 we discuss the results of the infinite-horizon case and present robustness checks. Finally, in Section 3.6 we use trade data to infer the latency of T1 and calculate the shadow price of latency that T1 would be willing to pay to reduce her latency in the marketplace.

3.5.1 Stylized features of latency-optimal strategy

We showcase features of the strategy for T1. We employ parameter estimates for 4 January 2017 to compute the optimal discretion for the trading day 5 January 2017. During 5 January 2017, T1 sent 3,043 FoK orders (1,601 buys and 1,442 sells) to the exchange. All orders were sent with a discretion of $\delta = 0$, and due to latency, T1 missed 279 trades, thus the fill ratio for the whole day was $F_{3403} = 90.83\%$.

Had T1 employed the strategy developed here, the optimal discretion in the MLOs would have been as depicted in the top panel of Figure 3.3 and would have attained a fill ratio of $F_{3043}^{\delta^*} = 99.17\%$, i.e., miss 25 out of the 3,043 trade attempts.¹⁶ Compared to the fill ratio of the suboptimal strategy $\delta = 0$, which we refer to as the naive strategy, the latency-optimal strategy would have filled 254 more trades at an average cost of 1.74 ticks – recall that one tick in the USD/JPY pair is 10^{-3} JPY and that the naive and latency-optimal strategies receive the same price improvements. The average cost is calculated as the cost of filling the 254 trades, in JPY, divided by the total volume of currency pairs of the 254 trades, and then multiplied by 10^3 to express it in ticks. The bottom panel of the same figure shows the volatility of the micro-exchange rate. In the next subsection we analyze in more detail the costs incurred by the optimal strategy to fill missed trades.

¹⁶Here, the fill measure $F_n^{\delta^*}$ denotes the measure F_n controlled by the optimal discretion δ^* . Recall that we have the liquidity available in the LOB at all times; thus, we can compute the fill ratio of strategy δ^* .




Figure 3.3: Parameters are $\hat{F} = 100\%$, $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $T = 5$ days, and other model parameters are in Table 3.2.

The left-hand panel of Figure 3.4 shows a heatmap of the fill ratio $F_{3043}^{\delta*}$ (i.e., 5 January 2017) for a range of parameters ϕ_F and ϕ_δ . We observe that when the penalty parameter ϕ_δ is low, the strategy obtains the highest fill ratios because the trader is more willing to bear the costs of walking the LOB, i.e., incur more slippage. In the heatmap we also show, with a star, the pair $(\phi_\delta, \phi_F) = (2.5 \times 10^{-5}, 5)$, which are the parameters we use in the analysis of subsection 3.5.2.

The right-hand panel of the figure shows a heatmap of the average cost of the fill ratios attainable by the optimal strategy. The cost reported is the slippage of the order relative to T1's naive strategy.

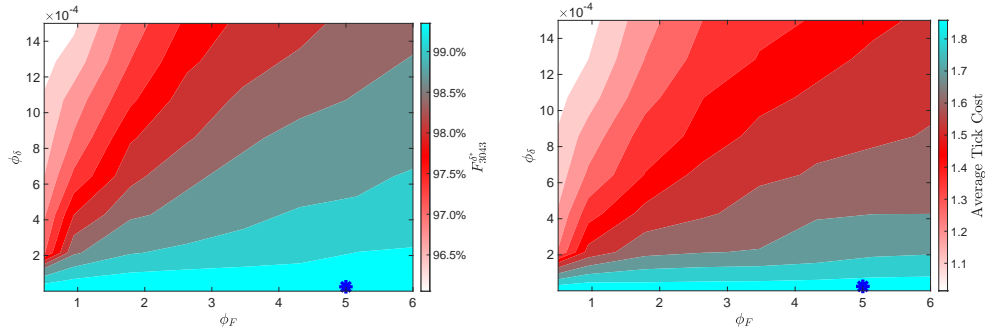


Figure 3.4: Fill ratio $F_{3043}^{\delta*}$ for 5 January 2017. Parameters are $\varphi = 10^{-3}$, $\hat{\delta} = 0$, $\hat{F} = 100\%$, $T = 5$ days, and other model parameters are in Table 3.2.

Finally, Figure 3.5 shows the fill ratio $F_{3043}^{\delta*}$ for the range $\varphi \in [0, a^2/\phi_\delta \sigma_F^2)$ of the ambiguity aversion parameter when the trader employs the optimal discretion strategy. Recall that the higher φ is, the less confident the trader is about the dynamics of F^δ . When $\varphi \rightarrow a^2/\phi_\delta \sigma_F^2$ the trader has very little confidence in her model of fill ratios, so the optimal strategy is to send orders with higher discretion to the exchange – the ambiguity averse trader is more willing to bear slippage costs to ensure high fill ratios. In A.1.1 we discuss the optimal control when $\varphi = a^2/2\phi_\delta \sigma_F^2$ and $a^2/2\phi_\delta \sigma_F^2 < \varphi$.

In the interest of space we only report results for T1 because those for T2 are similar. The computation time to implement the strategy is under 0.03 milliseconds per trade, which includes computing the volatility of the mid-exchange rate and the fill ratio.




Figure 3.5: Fill ratio for optimal strategy. Parameters are $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\hat{F} = 100\%$, $T = 5$ days, and other model parameters are in Table 3.2.

3.5.2 Performance of strategy for traders T1 and T2

We employ trade data for T1 and T2 between 5 December 2016 and 31 March 2017 to analyze the performance of the latency-optimal strategy. For each trading day we use the parameter estimates from the previous day and volatility of the exchange rate to compute the optimal discretion (3.19) in Proposition 1; recall that the volatility is calculated over a ten-minute window. We use LOB information to determine if the orders sent with the optimal discretion δ^* would have been filled, and if so, at what price(s).

In Table 3.1 we reported a number of 116,912 trade attempts for T1 and 68,129 trade attempts for T2. However, in the study that follows, we do not use all of these attempts. For each day we employ the first 10 minutes of the trading day (22:05 to 22:15) to compute the first observation of volatility of the micro-exchange rate. The first fill ratio observation (F_{50} for T1 and F_{50}^{AV} for T2) is computed the first time there are 50 trades after 22:15. Throughout the period we study we do not use 4,115 trades of T1 and 4,089 of T2 because each trading day we need to compute the first observations of volatility and fill ratio. Furthermore, to estimate the parameters of the model for the first day we employ 2,339 trades for T1 and 772 for T2. Thus, the number of trade attempts we employ to analyze the performance of the strategy is reduced to 110,458 for T1 and 63,268 for T2.

3.5.2.1 Performance of latency-optimal strategy for T1

Tables 3.3 and 3.4 provide information of the costs that T1 would have incurred had she employed the latency-optimal strategy and shows the mark-to-market costs of going back to the exchange to complete missed trades using market orders that walk the LOB until the order is filled in full.

During the period 5 December to 31 March 2017, T1 attempted 110,458 trades, all with discretion $\delta = 0$, and, due to latency, missed 9,483 trades, and received 3,938,559 JPY in price improvements. Had T1 employed market orders with infinite discretion, the cost of filling the 9,483 trades would have been 8,914,046 JPY, see row (a) in Table 3.3.¹⁷ The last two columns of the same row show the mark-to-market cost of market orders that walk the LOB 20 ms and

¹⁷A market order with infinite discretion walks the LOB until it is filled in full, that is, there is no price limit.

100 ms later to fill the 9,483 trades missed by the naive strategy – time is measured from the time-stamp recorded by the exchange when T1’s order was received. For example, if T1 were to return to the market (20 ms and 100 ms later) to retry the missed trade using a market order that can walk the LOB until filled in full, then the costs to fill the 9,483 missed trades would have been 9,841,755 JPY and 12,158,945 JPY, respectively.

In rows (b) to (e) we show the performance of the latency-optimal strategy for discretion targets $\hat{\delta} \in \{0, 1, 2, 3\}$, respectively. For example, row (b) shows that a strategy with target discretion $\hat{\delta} = 0$ would have filled 8,318 more trades than the naive strategy. The cost of filling the 8,318 trades is 6,528,965 JPY and the mark-to-market costs to fill these trades, with market orders 20 ms and 100 ms later, are 7,459,198 JPY and 9,601,222 JPY, respectively.

Table 3.4 presents the average costs of filling missed trades in ticks (10^{-3} JPY). Column 3 in row (a) shows that the average mark-to-market cost of market orders with no price limit to fill the 9,483 trades missed by the naive strategy is 2.11 ticks (i.e., $2.11 = 8,914,046 / 4,222 \times 10^3$). Row (b) shows that the average cost of filling 8,318 trades missed by the naive strategy, which were filled by the optimal strategy with discretion target $\hat{\delta} = 0$, would have been 1.76 ticks, and the mark-to-market cost to fill the 8,318 trades with market orders that walk the book until filled, submitted 20 ms or 100 ms later, is 2.01 or 2.58 ticks.

The last two columns in Table 3.4 show the ‘excess cost’ (in percentage terms relative to optimal strategy’s attainable cost) of returning to the market, 20 ms and 100 ms later, to complete unfilled trades with market orders that walk the book until filled in full. For example, the costs of filling missed trades with market orders 20 ms (100 ms) later are between 13% and 14% (44% and 47%) more than the cost of filling them with the optimal strategy.

The costs of returning to the market to fill missed trades employing market orders with infinite discretion is, on average, greater than the costs of completing the trades, on the first attempt, employing the latency-optimal strategy. Thus, not only does the optimal strategy attain a very high fill ratio, it does it at a cost which is lower than the mark-to-market value of completing the trades on a second attempt. These savings justify the extra cost of optimally walking the LOB in the first trade attempt to achieve higher fill ratios than the fill ratios achieved by the naive strategy. Note that we are not imputing other costs that the trader may incur if a trade is not filled. For example, some trades could be one leg of a trade, which if not completed, will affect other trades or strategies, see [Cartea et al. \(2019b\)](#).

		Missed by naive, filled by alternative strategies				
		Extra fills	Extra volume $\times 10^6$	Cost in JPY	Cost 20 ms later in JPY	Cost 100 ms later in JPY
(a)	$\delta_t = \infty \forall t$	9,483	4,222	8,914,046	9,841,755	12,158,945
	δ^* Target				$\delta_t = \infty \forall t$	$\delta_t = \infty \forall t$
(b)	$\hat{\delta} = 0$	8,318	3,718	6,528,965	7,459,198	9,601,222
(c)	$\hat{\delta} = 1$	8,603	3,839	6,807,416	7,770,984	9,959,191
(d)	$\hat{\delta} = 2$	8,758	3,906	6,999,804	7,950,543	10,185,240
(e)	$\hat{\delta} = 3$	8,938	3,982	7,232,545	8,181,285	10,447,450

Table 3.3: Performance of optimal strategy for T1. The total number of trade attempts is 110,458. Cost-related quantities are in JPY. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY. Computation times are under 10 seconds for each day.

		Missed by naive, filled by alternative strategies					
		Extra fills	Avg. cost in ticks	Avg. cost 20 ms later in ticks	Avg. cost 100 ms later in ticks	Cost increase 20 ms later %	Cost increase 100 ms later %
(a)	$\delta_t = \infty \forall t$	9,483	2.11	2.33	2.88	10%	36%
	δ^* Target					$\delta_t = \infty \forall t$	$\delta_t = \infty \forall t$
(b)	$\hat{\delta} = 0$	8,318	1.76	2.01	2.58	14%	47%
(c)	$\hat{\delta} = 1$	8,603	1.77	2.02	2.59	14%	46%
(d)	$\hat{\delta} = 2$	8,758	1.79	2.04	2.61	14%	46%
(e)	$\hat{\delta} = 3$	8,938	1.82	2.05	2.62	13%	44%

Table 3.4: Performance of optimal strategy for T1. The total number of trade attempts is 110,458. Cost-related quantities are weighted by volume and expressed in ticks, 10^{-3} JPY. The metrics in column 6 (*resp.* 7) are the proportional increase of the values in columns 5 and 6 (*resp.* 5 and 7) in Table 3.3. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY. Computation times are under 10 seconds for each day.

Table 3.5 shows the results of the optimal strategy arranged in volatility quartiles when the discretion target is $\hat{\delta} = 0$. As discussed above, when the micro-exchange rate is more volatile, the chances of missing trades are greater than when volatility is low because the LOB is undergoing many more updates. The optimal strategy developed here counterbalances the adverse effects of high volatility on fill ratios, so that in periods of high volatility, everything else being equal, the trader sends orders with more discretion to walk the book. The table shows that as volatility increases, the strategy maintains a stable level of fills, while T1's naive strategy misses more trades as volatility increases.

The last three columns in Table 3.5 show the cost of filling trades missed by T1's naive strategy that are filled by the optimal strategy with $\hat{\delta} = 0$. The first of the last three columns shows the cost incurred by the optimal strategy to execute missed trades for each volatility

quartile. The penultimate and last columns of the table show the cost of returning to the market to fill the missed trades with market orders that walk the LOB until they are filled. As discussed earlier, the cost incurred by the latency-optimal strategy to fill trades missed by the naive strategy are lower than the costs of returning to the market 20 ms and 100 ms later. The gap between the costs of returning to the market 20 ms and 100 ms later and the costs incurred by the optimal strategy widens as volatility increases.

Volatility Quartiles %	Trades sent to Exchange	Fills % $\delta_t = 0, \forall t$		Fills % $\delta_t^*, \hat{\delta} = 0$		Missed by naive, filled by δ^* with $\hat{\delta} = 0$		
						Avg. cost in ticks	Avg. cost 20 ms later in ticks	Avg. cost 100 ms later in ticks
0–25	27,615	26,151	94.7%	27,274	98.8%	1.61	1.77	2.02
25–50	27,614	25,588	92.7%	27,303	98.9%	1.73	1.92	2.20
50–75	27,615	25,086	90.8%	27,339	99.0%	1.67	1.91	2.49
75–100	27,614	24,150	87.5%	27,377	99.1%	1.88	2.19	3.04

Table 3.5: Performance of strategy in volatility quartiles for T1. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Ticks are 10^{-3} JPY. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

3.5.2.2 Performance of optimal strategy: T2

We repeat the analysis of the performance of the strategy using T2’s trades. Tables 3.6 and 3.7 are presented in the same format as Tables 3.3 and 3.4 respectively.

During the period 5 December 2016 to 31 March 2017, T2 attempted 63,268 IoC trades, all with discretion $\delta = 0$, and, due to latency, missed 7,038 trades, including partial misses. The cost of achieving better fill ratios using the optimal strategy (for various discretion targets $\hat{\delta}$) is justified by the cost of returning to the exchange to fill missed trades 20 ms or 100 ms later with market orders with infinite discretion to walk the book. For example, had T2 employed the optimal strategy with discretion target $\hat{\delta} = 0$, only 755 trades would have been missed, as opposed to the 7,038 misses by the naive strategy. The average cost of filling the 6,283 trades missed by the naive strategy is 1.24 ticks, as opposed to 2.82 and 2.75 ticks, which are the mark-to-market costs of returning to the market 20 ms and 100 ms later.

		Missed by naive, filled by alternative strategies					
		Extra Fills (incl. partial)	Extra Fills (only full)	Extra Volume $\times 10^6$ in JPY	Cost in JPY	Cost 20 ms later in JPY	Cost 100 ms later in JPY
(a)	$\delta_t = \infty, \forall t$	7,038	7,038	1,455	2,294,999	4,628,993	4,455,516
	δ^* Target					$\delta_t = \infty, \forall t$	$\delta_t = \infty, \forall t$
(b)	$\hat{\delta} = 0$	6,283	6,238	1,320	1,634,888	3,727,171	3,631,491
(c)	$\hat{\delta} = 1$	6,446	6,411	1,340	1,665,638	3,742,734	3,649,011
(d)	$\hat{\delta} = 2$	6,548	6,517	1,354	1,667,995	3,801,813	3,690,360
(e)	$\hat{\delta} = 3$	6,648	6,627	1,373	1,733,090	3,903,934	3,795,341

Table 3.6: Performance of optimal strategy for T2. The total number of trade attempts is 63,268. Cost-related quantities in ticks (10^{-3} JPY). Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies					
		Extra Fills (incl. partial)	Avg. cost in ticks	Avg. cost 20 ms later in ticks	Avg. cost 100 ms later in ticks	Growth 20 ms later %	Growth 100 ms later %
(a)	$\delta_t = \infty, \forall t$	7,038	1.58	3.18	3.06	102%	94%
δ^* Target						$\delta_t = \infty, \forall t$	$\delta_t = \infty, \forall t$
(b)	$\hat{\delta} = 0$	6,283	1.24	2.82	2.75	128%	122%
(c)	$\hat{\delta} = 1$	6,446	1.24	2.79	2.72	125%	119%
(d)	$\hat{\delta} = 2$	6,548	1.23	2.81	2.73	128%	121%
(e)	$\hat{\delta} = 3$	6,648	1.26	2.84	2.76	125%	119%

Table 3.7: Performance of optimal strategy for T2. The total number of trade attempts is 63,268. Cost-related quantities are weighted by volume and expressed in ticks (10^{-3} JPY). The metrics in column 6 (*resp.* 7) are the growth of the values in columns 5 and 6 (*resp.* 5 and 7) in Table 3.3. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

Table 3.8 shows the results of the optimal strategy arranged in volatility quartiles when the discretion target is $\hat{\delta} = 0$. The interpretation is the same as that of Table 3.5 for T1. The cost incurred by the latency-optimal strategy to fill the trades missed by the naive strategy is lower than the costs of returning to the market 20 ms and 100 ms later. The gap between the costs of returning to the market 20 ms and 100 ms later and the costs incurred by the optimal strategy widen as volatility increases.

Volatility Quartiles %	Trades sent to exchange	Fills %		Fills %		Missed by naive, filled by δ^* with $\hat{\delta} = 0$		
		$\delta_t = 0, \forall t$ (incl. partial)		$\delta_t^*, \hat{\delta} = 0$ (incl. partial)		Avg. cost in ticks	Avg. cost 20 ms later in ticks	Avg. cost 100 ms later in ticks
0%–25%	15,817	14,950	94.5%	15,658	99.0%	0.94	1.90	1.87
25%–50%	15,817	14,545	92.0%	15,631	98.8%	1.12	2.72	2.68
50%–75%	15,817	14,226	89.9%	15,629	98.8%	1.25	2.75	2.92
75%–100%	15,817	13,329	84.3%	15,610	98.7%	1.86	3.51	3.31

Table 3.8: Performance of strategy in volatility quartiles for T2. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

3.5.3 Infinite-horizon strategy and robustness checks

Infinite-horizon. Tables A.1 and A.2, which are the analogues of Tables 3.3 and 3.4, in A.2.1 show the results of the latency-optimal strategy for the infinite-horizon case discussed in Subsection 3.3.2.1. The results are broadly the same as those obtained for T1’s strategy discussed above, where we assume $T = 5$ days – in the interest of space we do not report the results for T2.

Robustness to volatility parameter estimates. In our set-up the investor makes her model of fill ratios robust to model misspecification. We also check the robustness of the performance of the latency-optimal to: (i) parameter uncertainty of the volatility model (3.6), and (ii) different models of the volatility of the micro-exchange rate.

To check robustness to parameter uncertainty in (3.6), we look at eight scenarios that result from shocking the parameter estimates $\hat{\kappa}$, \hat{v} , $\hat{\sigma}_v$. The shocked parameters are: $\hat{\kappa}^u = 1.2 \hat{\kappa}$, $\hat{\kappa}^d = 0.8 \hat{\kappa}$, $\hat{v}^u = 1.2 \hat{v}$, $\hat{v}^d = 1.2 \hat{v}$, $\hat{\sigma}_v^u = 1.2 \hat{\sigma}_v$, and $\hat{\sigma}_v^d = 0.8 \hat{\sigma}_v$, where $\hat{\kappa}$, \hat{v} , $\hat{\sigma}_v$ are the estimates we obtain with the methodology outlined in Section 3.4. Next, we compute the optimal strategy δ^* with the eight triplets of parameters resulting from all combinations of $(\hat{\kappa}, \hat{v}, \hat{\sigma}_v) \in \{\hat{\kappa}^d, \hat{\kappa}^u\} \times \{\hat{v}^d, \hat{v}^u\} \times \{\hat{\sigma}_v^d, \hat{\sigma}_v^u\}$. In A.2.4.1 we report the results for $\hat{\delta} = 0$. We find that the fill ratios lie in the range 98.5% to 99% and the average tick cost is approximately 1.76 ticks in all cases.

Finally, we assume that the volatility of the micro-exchange rate follows (3.7) and (3.8) and compute the performance of the latency-optimal strategy, see A.2.3. We find that the results are broadly the same as those discussed above when volatility follows (3.6).

3.6 The cost of latency

One approach to mitigate the effects of latency on the efficacy of trading strategies is to invest on hardware and co-location services to reduce the time delays associated to the different components of the life cycle of a trade, see definition of latency in Section 3.1. As latency shortens, traders that aim at the best prices they observe in the LOB will be on target more often. This

improvement in the marksmanship of the trader increases the fill ratio of the strategy, but also reduces the chances of receiving price improvements.

Another approach is to devise latency-optimal strategies that lessen the adverse effects of time delays, e.g., failing to complete trades. Although these strategies do not reduce time delays in the life cycle of trade, they are a substitute for cutting down latency when the objective is to increase the percentage of MLOs that are filled.

Neither approach is free. The first approach requires the trader to pay fixed costs to operate in the marketplace with lower latency, which increases the efficacy of the MLOs but reduces opportunities of obtaining price improvements. In the second approach, costs accrue to the strategy when MLOs walk the LOB, while benefits accrue because the latency-optimal approach scoops price improvements when the market, due to latency, moves in the trader's interest.

In this section we compute the shadow price that T1 would be willing to pay to reduce her latency in the marketplace. We proceed as follows. i) Consider hypothetical traders who employ the same naive strategy (i.e., zero discretion to walk LOB) as that of T1, but each trader acts in the market with a different fixed latency – for simplicity we assume latency is not stochastic. ii) Compute the fill ratio achieved by each hypothetical trader's naive strategy. iii) Map levels of latency to fill ratios. iv) Find the hypothetical trader whose fill ratio is the same as that achieved by T1's latency-optimal strategy developed above, and compute the price improvements of the hypothetical trader's strategy. v) Compute the difference of the latency-optimal costs of walking the LOB and its price improvements. Thus, the difference between the costs in (iv) and (v) is the shadow price of latency that T1 would be willing to pay to employ a naive strategy (no discretion to walk LOB), which obtains the same fill ratio as that of the latency-optimal strategy.

3.6.1 Latency and fill ratios

Recall that the attributes we have for each order include: type of trade, volume of the trade, price limit, and direction (buy or sell) of the trade. We also have the time-stamp of when the order is processed by the exchange (this time is denoted by t_3 in our discussion of latency in Section 3.1 above). However, we do not know the time it took the trader to process information and make a decision before instructing the exchange, or the time when the trader sent the orders to the exchange – all of which affect the latency of the trader in the marketplace.

To map latency to fill ratios we devise fifty hypothetical traders who are clones of T1. The only difference between T1 and the hypothetical traders is that each clone has a different latency in the market. We denote each hypothetical trader by T_j with $j = 0, 1, 2, \dots, 50$, and assume that the latency of T_j is j ms, i.e. $t_3 - t_0 = j$ ms. Note that although impossible for a trader to operate in the marketplace with a latency of 0 ms, we include it in our list of hypothetical traders for comparison purposes.

In our approach we assume that the clones employ the same strategy as T1 regardless of their latency. However, this may not be the case because if T1's latency was lower than her

actual latency, her trading strategy could be different from the one she employed during the period we analyze. For example, T1 may not be attempting potentially profitable trades because she acknowledges that her latency is too long to successfully complete trades that require more speed.

We use T_{10} to illustrate how we map a latency of 10 ms to a measure of fill ratio – the fill ratio of the other hypothetical traders is calculated in the same way. The fill ratio is computed as the percentage of filled orders in one trading day – this coincides with definition F_n , provided in (3.1), where n is the number of trades in the day. We assume that T_{10} replaces T1 in the market. Thus, T_{10} sends orders that have the same direction, volume, type (i.e., FoK), and discretion to walk the LOB, as those of T1, and assume the orders from T_{10} reach the LOB an instant before T1’s trade was filled (for which we have a time-stamp), so the hypothetical trader and T1 face the same market conditions. Therefore, orders sent by T_{10} and T1 only differ in that T_{10} attempts to fill the trade at the best quote streamed by the exchange 10 ms before the exchange receives the order.

For example, to compute the fill ratio that T_{10} would have obtained on 5 January 2017, we employ the time-stamps (when the exchange processes the order) of the 3,043 orders sent by T1 on that day and compute the fill ratio that T_{10} would have achieved. We repeat this approach for each hypothetical trader on the same day and depict the results in the left-hand panel of Figure 3.6. As expected, there is a (monotonic) decreasing relationship between latency and fill ratio. In one extreme, when latency is zero, T_0 would have obtained a fill ratio of 100%. At the other extreme T_{50} ’s fill ratio is 89%.

In the figure, the square marker shows the point where latency of 27 ms is mapped to 90%, which is the fill ratio obtained by T1’s naive strategy. The diamond marker in the upper-left corner of the figure shows the latency of a hypothetical trader who achieved a fill ratio of 99%. Note that this fill ratio corresponds (approximately) to the fill ratio T1 would have obtained ($F_{3043}^{\delta^*} = 98.73\%$) had she employed the latency-optimal strategy with targets $\hat{\delta} = 0$ and $\hat{F} = 100\%$. Thus, we say that, on 5 January 2017, T1’s implied latency is 27 ms and the implied latency of the optimal strategy is 3 ms.

We repeat the analysis for each day between 5 December 2016 and 31 March 2017. The right panel of Figure 3.6 shows the implied latencies of T1’s naive strategy (solid line) and that of the latency-optimal strategy with target $\hat{\delta} = 0$ (dotted line). The variations in ‘implied’ latency are due to a number of factors, including the volatility of the best quotes which, as discussed above, has a considerable effect on fill ratios.¹⁸ For ease of presentation the y -axis is capped at 50 ms – there were two days when T1’s implied latency was over 50 ms.

¹⁸In late February 2017 there is one day where the implied latency is above 10ms. On this day, there was a combination of factors that led to this spike, mainly, a considerable number of trade attempts around news announcements.

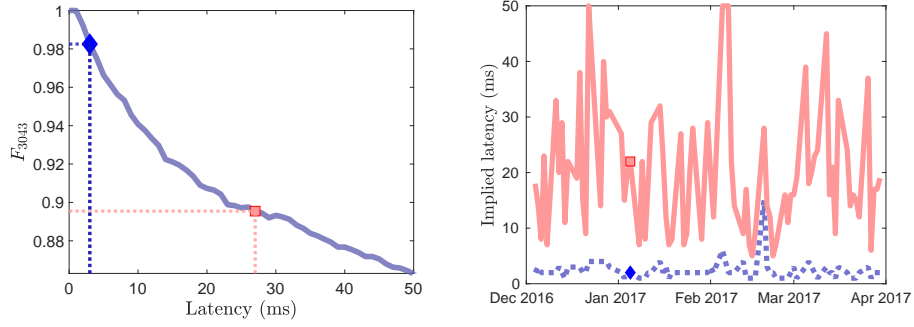


Figure 3.6: Left: Map of latency to fill ratio employing T1's trade time-stamps and LOB information on 5 January 2017. Square marker is T1's implied latency. Diamond marker is the implied latency of the optimal strategy. Right: implied latency for T1 and implied latency for latency-optimal strategy employing T1's trade time-stamps and LOB information during the period 5 December 2016 to 31 March 2017.

Table 3.9 summarizes the results. The top panel of the table shows the mean, standard deviation, and median of the latency implied by the naive strategy and by the latency-optimal strategies with various discretion targets $\hat{\delta}$. The bottom panel provides summary statistics of the daily fill ratio obtained by the strategies. Here, the fill ratio F_N refers to the percentage of orders that were filled each day, and the table shows the mean, standard deviation, and median of the daily percentage of fills. For the latency-optimal strategy the mean fill is between 98.9% and 99.5% and for T1's naive strategy the mean fill is 91.0%. We observe that the standard deviation of the fill ratio of the optimal strategy is also considerably lower than that of the naive strategy. Therefore, the optimal strategy achieves much higher fill ratios than the naive strategy, and produces more stable results because the optimal strategy depends on the volatility of the micro-exchange rate.

	Naive strategy	Optimal Strategy δ^* Target				$\delta_t = \infty$ $\forall t$
		$\hat{\delta} = 0$	$\hat{\delta} = 1$	$\hat{\delta} = 2$	$\hat{\delta} = 3$	
Latency						
mean	20.02	2.34	2.01	1.70	1.53	0
std	11.31	1.72	1.57	0.89	0.80	0
median	18.00	2.00	2.00	1.00	1.00	0
F_N						
mean	91.9%	98.9%	99.2%	99.3%	99.5%	100%
std	2.9%	0.4%	0.4%	0.4%	0.3%	0%
median	92.2%	98.9%	99.2%	99.4%	99.5%	100%

Table 3.9: Implied latency for T1. Period is 5 December 2016 to 31 March 2017, currency pair USD/JPY.

3.6.2 Shadow price of latency

The shadow price of latency is the price that T1 would be willing to pay to reduce her latency, so a naive and the latency-optimal strategies obtain the same fill ratio.¹⁹ To compute this price we need: Latency-optimal cost of walking the LOB and price improvements, (ii) price improvements received by the hypothetical trader with latency equal to that implied by the latency-optimal strategy.

For example, on 5 January 2017, the latency-optimal strategy incurs 92,771 JPY in costs from walking the LOB and receives 69,305 JPY in price improvements. On the other hand, the naive strategy of the hypothetical trader with latency 3 ms, receives 7,592 JPY in price improvements. Thus, on 5 January 2017, T1 would be willing to pay 31,058 JPY to reduce her latency from 22 ms to 3 ms. We repeat the analysis for all trading days between 5 December 2016 and 31 March 2017 and compute T1's shadow price of latency for each day. Table 3.10 contains summary statistics of the results.

The table shows that T1 would be willing to pay approximately $3 \cdot 10^6$ JPY (around 30,000 USD) to reduce her latency from 20.02 ms to 2.34 ms. The upper bound of approximately $5 \cdot 10^6$ JPY (around 50,000 USD) for the shadow price of latency is given by the last column and last row of the table, where we show the costs of always employing market orders. These shadow prices are, according to market participants, lower than co-location, hardware, and other costs that traders would have to incur to reduce latency in the life cycle of a trade.

	Naive strategy	Optimal Strategy δ^* Target				$\delta_t = \infty$ $\forall t$
		$\hat{\delta} = 0$	$\hat{\delta} = 1$	$\hat{\delta} = 2$	$\hat{\delta} = 3$	
Implied Latency (ms)	20.02	2.34	2.01	1.70	1.53	0
Walk LOB Cost	0	6,528,965	6,807,416	6,999,804	7,232,545	8,914,046
Price Improvement	3,938,559	3,938,559	3,938,559	3,938,559	3,938,559	3,938,559
Price Improvement with Reduced Latency	3,938,559	384,035	291,070	179,875	135,744	0
Shadow Price to Reduce Latency	0	2,974,441	3,159,927	3,241,120	3,429,730	4,975,487

Table 3.10: Shadow price to reduce latency reported in JPY. Period is 5 December 2016 to 31 March 2017, currency pair USD/JPY.

3.7 Conclusions

In this chapter, we showed how to compute the price that FX traders are willing to pay to reduce their latency in the marketplace. To achieve this we developed a latency-optimal strategy to map the costs of the strategy to a reduction in latency.

¹⁹Shadow prices are a well-known concept in economics. A shadow price is the estimated price of a good/service for which there is no market. See Kallsen and Muhle-Karbe (2010) for an example of shadow price processes being used in the context of the Merton problem.

In particular, we employed a set of proprietary data of FX transactions to show how latency affects the percentage of liquidity taking orders filled by a trader’s strategy. We analyzed the trades of two liquidity takers over a four-month period and showed that approximately 10% of their orders were not filled due to latency.

To calculate the shadow price of latency of traders we derived a strategy that specifies the limit price of the market orders sent to the exchange. The price limit for each liquidity taking order balances the tradeoff between the costs from walking the LOB and how far is the fill ratio from the trader’s desired target. The trader’s optimal strategy specifies the price limit for each transaction depending on the proportion of orders that have been filled, how far is the strategy from the target fill ratio, the cost of walking the LOB, and the volatility of the exchange rate. For example, everything else being equal, the optimal price limit to walk the LOB is higher (*resp.* lower) when the volatility of the exchange rate is above (*resp.* below) its long-term level.

We showed the performance of the latency-optimal strategy using trade data for two FX traders. We computed the costs incurred by the latency-optimal strategy to increase fill ratios and showed that the costs of the strategy are lower than the costs of returning to the market 20 ms or 100 ms later to fill a missed trade. Moreover, we showed that the latency-optimal strategy obtains similar fill ratios for different levels of volatility. This is in contrast to the performance of the naive strategy where the majority of the missed trades occur when the volatility of the exchange rate is high.

Finally, we employed the proprietary data set to build a function that maps latency of traders to fill ratios. We used this mapping to compute the shadow price of latency that traders are willing to pay to reduce latency in the marketplace. We computed the shadow price a particular trader is willing to pay to reduce her latency by investing in hardware, co-location, and other services. Our results showed how to compute the threshold at which traders would be better off employing the latency-optimal strategy developed here, instead of investing in hardware and co-location services to reduce latency.

Chapter 4

Tradeoff between costs and misses

Having 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. As defined before, 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 and processing the instruction from the trader. Thus, due to latency, there is no guarantee that liquidity providers can place a LO in 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. An 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. An 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 an 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 chapter, 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 agent sends buy/sell MLOs at random times to partly camouflage her order flow, and before the order reaches the exchange, the LOB undergoes quantity and price updates. 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

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.

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 an 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.

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, uniqueness and existence of the resulting random-measure driven FBSDE is not covered in the extant literature, and the particular form itself appears to be new.² The tools we use to prove existence and uniqueness are standard and similar to the results in the literature, we have a short time horizon restriction.

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.³ 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 \(2013\)](#), henceforth M-S, the objective is to quantify the cost of latency on transaction costs, while our work develops a latency-optimal trading strategy that improves the marksmanship of liquidity takers. Our work differs from that in M-S in four important aspects. First, in M-S, the trader employs a *limit order* (which may be cancelled and re-posted) that is pegged to the midprice of the asset to execute one

²We believe this to be case because of (i) the particular non-standard representation of the FBSDE, (ii) the fact that the forward part is adapted and the backward part needs to be predictable, and (iii) a discontinuous driver in the forward part of the FBSDE.

³Our framework is general; thus, under mild restrictions of the compensator of the MPP, these results hold. In particular, our results hold (i) for any bounded and predictable intensity process at which trades arrive, and (ii) for when the dynamics of the shocks exhibit stochastic volatility.

trade – if the limit order is not filled by the end of a fixed time horizon, the trader sends a market order. In our model, the trader is frequently trading and always employs liquidity taking orders, i.e., we do not focus on one trade alone. For each order, the trader specifies a **price limit** because of their latency in the marketplace. Thus, in the M-S model the trader controls the depth of one limit order, while in our model the trader controls by how much each liquidity taking order can walk the limit order book.

Second, in M-S, there is no dynamic optimisation for a trader who must complete a sequence of orders over a trading horizon, while our model is dynamically optimal. Third, the M-S model assumes that latency is fixed and known by the trader. In our model we assume that each trader experiences different levels of latency and that latency is random, which is indeed the case in electronic markets.

Finally, M-S focuses on equity order-driven markets and their model relies on the availability of limit orders pegged to the midprice. Pegged orders are common in equity exchanges, but are not available in some exchanges of other asset classes. For example, in many currency exchanges such as EBS, Reuters, and LMAX, there are no pegged orders to the mid-exchange rate. We note that our model works in all asset classes in which market participants trade in an order book.

The remainder of the chapter proceeds as follows. Section 4.1 proposes the agent’s performance criterion and characterizes the latency-optimal strategy as the solution to a FBSDE. Section 4.2 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 4.3 shows that the candidate control we find is the global optimum and Section 4.4 discusses the performance of the strategy for various scenarios. We conclude in Section 4.5 and collect some proofs in the Appendix.

4.1 Optimal discretion to walk the book

4.1.1 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.⁴ 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

⁴See Figure 3.6 in Chapter 3 for empirical evidence of random implied latencies; see also Hautsch et al. (2018).

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.

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 an 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 an 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 3.3 to 3.8 in Chapter 3, where we 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 chapter are available in exchanges that trade in a limit order book.⁵ Each exchange 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.⁶ 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 chapter is immediate execution.

⁵For example, MLOs are available in Cboe, LSE, LMAX Exchange, CME, and Nasdaq. Note that FoK, IoC orders have immediate execution in the time-in-force qualifier of the order.

⁶In this thesis we work in a framework with a central LOB with firm quotes; see Phillips et al. (2017) and Cartea et al. (2019b) for a comparison between last-look and firm liquidity venues.

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.

4.1.1.1 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.⁷

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 paid 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 4.1 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

⁷This raises the questions: To what extent is this strategy optimal? Are traders optimising the value of the constant discretion? Assuming traders prefer to miss few trades; is sending MLOs with a constant discretion optimal? Should the value of the discretion in the MLOs depend on time-of-day, activity, and volatility?

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.

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.⁸




Figure 4.1: 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.

4.1.2 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.

⁸In Chapter 3, we 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.

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 security, the mark Z represents a shock to the LOB commensurate with the volume of the MLO.⁹

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

$$\mathbf{p}(\mathrm{d}z, \mathrm{d}t) = \sum_{n \geq 1} \mathfrak{D}_{(Z_n, T_n)}(\mathrm{d}z, \mathrm{d}t),$$

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

$$\mathbb{E} \left[(\mathbf{p}(\mathbb{R}, [0, T]))^2 \right] < \infty \quad \text{and} \quad \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} |z| \mathbf{p}(\mathrm{d}z, \mathrm{d}t) \right] < \infty. \quad (4.1)$$

We denote by $\tilde{\mathbf{p}}$ the predictable compensator of the random measure \mathbf{p} , which admits the decomposition

$$\tilde{\mathbf{p}}(\mathrm{d}z, \mathrm{d}t) = \Phi_t(\mathrm{d}z) \mathrm{d}A_t, \quad (4.2)$$

see Theorem 22.22 in [Kallenberg \(1997\)](#) for existence of the predictable compensator and Theorem 13.3.7 in [Cohen and Elliott \(2015\)](#) for the decomposition property. Here, the compensator has the property that for $\mathbf{q} := \mathbf{p} - \tilde{\mathbf{p}}$, the stochastic integral $(H \star \mathbf{q})_t = \int_0^t \int_{\mathbb{R}} H_s \mathbf{q}(\mathrm{d}z, \mathrm{d}s)$ is a martingale for any predictable and sufficiently integrable process H . In (4.2), 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 $N_t := \mathbf{p}(\mathbb{R} \times [0, t])$.

We make the following three assumptions throughout the analysis.

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

⁹Note 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.

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

Assumption 3. *The function*

$$r_t(x) = \int_{-\infty}^x z \Phi_t(dz)$$

is Lipschitz in x , uniformly on $[0, T] \times \Omega$, with Lipschitz constant \tilde{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 consume the liquidity 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 consume the liquidity at 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, 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.¹⁰ 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.

4.1.3 Tradeoff: cost of walking the LOB and cost of missed trades

4.1.3.1 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.¹¹ For buy orders, the extra 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, extra 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 extra cost of the strategy is the cash paid (not received) to walk

¹⁰By ‘price per unit’ we mean the average price that the trade attempt would receive/pay if filled in full against the liquidity in the bid/ask side of the LOB.

¹¹We leave the joint optimisation of the discretion and volume of the orders for later investigations. For now, volume is taken to be fixed, which for simplicity we take to be one unit.

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

$$C_t^\delta = \int_0^t \int_{\mathbb{R}} z \hat{G}(\delta_s - z) \mathbf{p}(dz, ds). \quad (4.3)$$

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.

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.¹²

4.1.3.2 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

$$D_t^\delta = \int_0^t \int_{\mathbb{R}} G(\delta_s - z) \mathbf{p}(dz, ds), \quad (4.4)$$

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:

$$\alpha D_t^\delta + \gamma \left(D_t^\delta\right)^2, \quad (4.5)$$

where $\alpha \geq 0$ and $\gamma \geq 0$ are cost parameters that are specific to each trader.¹³

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.

¹²Trading with a fixed volume does not preclude orders from walking the book. For example, in the next chapter (Chapter 5) we observe that in LMAX London, a frequent amount of liquidity available in the best quotes of the EUR/USD currency pair is €500,000. As a consequence, a trade for €600,000 is likely to walk the LOB to be filled in full. Thus, a trader that sends trade attempts for €600,000 with each order will (in most cases) walk the LOB to fill the orders in full.

¹³Recall that agents assume that missing a trade is costly. This function penalises deviations from zero of the number of misses D_t^δ .

4.1.4 Performance criterion

The agent's performance criterion is

$$J(\delta) = \mathbb{E} \left[C_T^\delta + \alpha D_T^\delta + \gamma \left(D_T^\delta \right)^2 \right], \quad (4.6)$$

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

$$\mathcal{A} := \left\{ \delta = (\delta_t)_{t \in \mathfrak{T}} \mid \delta \text{ is } \mathcal{F}\text{-predictable \& } \mathbb{E} \left[\int_0^T (\delta_t)^2 dt \right] < \infty \right\}, \quad (4.7)$$

equipped with the norm

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

Throughout the chapter, 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 (4.6), 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 (4.1) 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.

4.1.5 Variational analysis approach

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

$$J(\delta) = J^C(\delta) + \alpha J^{LP}(\delta) + \gamma J^{QP}(\delta), \quad (4.8)$$

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

Proposition 1. *The following equation holds*

$$J^C(\delta) = \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} z \hat{G}(\delta_t - z) \mathbf{p}(dz, ds) \right] = \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} z \hat{G}(\delta_t - z) \Phi_t(dz) dA_t \right]. \quad (4.9)$$

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*

$$J^{LP}(\delta) = \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} G(\delta_t - z) \Phi_t(dz) dA_t \right], \quad \text{and} \quad (4.10)$$

$$J^{QP}(\delta) = \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \left(2 D_{t-}^{\delta} G(\delta_t - z) + G(\delta_t - z) \right) \Phi_t(dz) dA_t \right]. \quad (4.11)$$

Proof. Equation (4.10) follows from the predictability of the integrand. Next, we show (4.11). The number of missed trades D_t^{δ} satisfy the SDE

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

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

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

Then,

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

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

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

■

4.1.5.1 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

$$\langle \mathcal{D} J(\delta), w \rangle = \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} [J(\delta + \epsilon w) - J(\delta)], \quad (4.12)$$

when the limit exists.

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

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

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

Proof. See B.1. ■

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

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

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

$$\delta_t^* = 2\gamma \mathbb{E}_{t-} [D_T^{\delta^*}] + \gamma + \alpha, \quad (4.13)$$

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

Proof. By Lemma 1 and the performance criterion (4.8), the Gâteaux derivative of J vanishes at

$$\begin{aligned} \delta_t^* &= 2\gamma \mathbb{E}_{t-} \left[\int_t^T \int_{\mathbb{R}} G(\delta_s^* - z') \tilde{\mathbf{p}}(dz', ds) \right] + 2\gamma \left(D_{t-}^{\delta^*} + \frac{1}{2} \right) + \alpha \\ &= 2\gamma \mathbb{E}_{t-} \left[\int_t^T \int_{\mathbb{R}} G(\delta_s^* - z') \mathbf{p}(dz', ds) \right] + 2\gamma \left(D_{t-}^{\delta^*} + \frac{1}{2} \right) + \alpha \\ &= 2\gamma \mathbb{E}_{t-} \left[D_T^{\delta^*} - D_{t-}^{\delta^*} \right] + 2\gamma \left(D_{t-}^{\delta^*} + \frac{1}{2} \right) + \alpha \\ &= 2\gamma \mathbb{E}_{t-} [D_T^{\delta^*}] + \gamma + \alpha. \end{aligned} \quad (4.14)$$

Now we show that if the Gâteaux derivative at δ vanishes in every direction w , the control δ satisfies (4.14). 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

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

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 $\omega \in \mathfrak{D}$. ■

If the value of the quadratic cost parameter γ is zero, the candidate optimal control in (4.13) has the simple closed-form expression

$$\delta_t^* = \alpha, \quad (4.15)$$

which is independent of the number of missed trades.

In our framework, the trading pattern of the trader we analysed in subsection 4.1.1.1 would correspond to a performance criterion as that in (4.6) 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 (4.13), which affects the optimal discretion, i.e., determines the left-hand side of (4.13), which in turn affects the expected number of missed trades that appears on the right-hand side of (4.13), and so on. In Section 4.4 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 (4.15). This corresponds to the behaviour of traders described in the empirical analysis of subsection 4.1.1.1. 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 .

4.2 Existence and uniqueness of the FBSDE

To the best of our knowledge, the FBSDE

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

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.¹⁴ Finally, the function $G(x)$ –which is one if $x < 0$ or zero otherwise, is not Lipschitz and it is not continuous.

¹⁴Note that the conditional expectation is with respect to \mathcal{F}_{t-} .

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 chapter. The mathematical tools we use are different to those in their paper. 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, and (ii) the four-step scheme by [Ma et al. \(1994\)](#) that relies on partial differential equations and stochastic optimal control. Here, we employ the first approach. Similar to the literature on the existence of FBSDEs using fixed point arguments, we have restrictions on the time horizon T .¹⁵ 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 chapter.¹⁶ 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 within the mathematical finance literature, see [Kohlmann and Tang \(2002\)](#), [Bank et al. \(2017\)](#), [Belak et al. \(2018\)](#), [Casgrain and Jaimungal \(2018, 2020\)](#), and [Herdegen et al. \(2021\)](#).¹⁷

To streamline the results in this section, we define the spaces we use to prove existence and uniqueness of the solution to the FBSDE (4.13).

Define

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

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

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

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](#)

¹⁵The applicability of the FBSDE we study for large values of T remains to be studied.

¹⁶In B.2 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.

¹⁷The work of [Herdegen et al. \(2021\)](#) differs from the others in that they study non-linear FBSDEs under certain conditions of the non-linearity.

(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 $(C^\infty, \|\cdot\|_\infty)$ 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, see Chapter XI-7 in [Banach \(1955\)](#).

Next, we transform the FBSDE in (4.13), 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

$$\begin{aligned} \delta_t &= 2\gamma \mathbb{E}_{t-} [D_T^\delta] + \gamma + \alpha, \\ D_t^\delta &= \int_0^t \int_{\mathbb{R}} G(\delta_s - z) \mathbf{p}(dz, ds), \quad D_0^\delta = 0, \end{aligned} \tag{4.16}$$

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

$$\begin{aligned} \hat{\delta}_t &= 2\gamma \mathbb{E}_t [D_T^\delta] + \gamma + \alpha, \\ D_t^\delta &= \int_0^t \int_{\mathbb{R}} G(\hat{\delta}_{s-} - z) \mathbf{p}(dz, ds), \quad D_0^\delta = 0, \end{aligned} \tag{4.17}$$

exists and is unique, with $\hat{\delta}$ and D^δ \mathcal{F} -adapted. Finally, by a further change of variables, $\tilde{\delta}_t = \hat{\delta}_t - 2\gamma D_t^\delta$, a solution to (4.16) exists and is unique if a solution to

$$\begin{aligned} \tilde{\delta}_t &= 2\gamma \mathbb{E}_t \left[\int_t^T \int_{\mathbb{R}} G(\tilde{\delta}_{s-} + 2\gamma D_{s-}^\delta - z) \mathbf{p}(dz, ds) \right] + \gamma + \alpha, \\ D_t^\delta &= \int_0^t \int_{\mathbb{R}} G(\tilde{\delta}_{s-} + 2\gamma D_{s-}^\delta - z) \mathbf{p}(dz, ds), \quad D_0^\delta = 0, \end{aligned} \tag{4.18}$$

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

To analyse solutions to the FBSDE (4.18), we study the fixed points of the functional

$$\Upsilon(U, V)_t = \begin{pmatrix} H(U, V)_t \\ I(U, V)_t \end{pmatrix} = \begin{pmatrix} 2\gamma \mathbb{E}_t \left[\int_t^T \int_{\mathbb{R}} G(U_{s-} + 2\gamma V_{s-} - z) \mathbf{p}(dz, ds) \right] + \gamma + \alpha \\ \int_0^t \int_{\mathbb{R}} G(U_{s-} + 2\gamma V_{s-} - z) \mathbf{p}(dz, ds) \end{pmatrix}, \tag{4.19}$$

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 [B.2](#).

The next theorem shows the existence and uniqueness of the solution to the FBSDE (4.18).

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

$$k T \bar{\lambda} (\max\{1, 2\gamma\})^2 < 1, \quad (4.20)$$

there exists a unique solution to the FBSDE

$$\tilde{\delta}_t = 2\gamma \mathbb{E}_t \left[\int_t^T \int_{\mathbb{R}} G(\tilde{\delta}_{s-} + 2\gamma D_{s-}^{\tilde{\delta}} - z) \mathbf{p}(dz, ds) \right] + \gamma + \alpha, \quad (4.21a)$$

$$D_t^{\tilde{\delta}} = \int_0^t \int_{\mathbb{R}} G(\tilde{\delta}_{s-} + 2\gamma D_{s-}^{\tilde{\delta}} - z) \mathbf{p}(dz, ds), \quad D_0^{\tilde{\delta}} = 0. \quad (4.21b)$$

Proof. Consider the functional $\Upsilon : \mathcal{C}^1 \times \mathcal{C}^1 \rightarrow \mathcal{C}^1 \times \mathcal{C}^1$ defined in (4.19). 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

$$\begin{aligned} \|\Upsilon(U, V) - \Upsilon(X, Y)\|_{\mathcal{C}^1 \times \mathcal{C}^1} &= \mathbb{E} \left[\int_0^T |H(U, V)_t - H(X, Y)_t| dt \right] \\ &\quad + \mathbb{E} \left[\int_0^T |I(U, V)_t - I(X, Y)_t| dt \right]. \end{aligned} \quad (4.22)$$

The first term on the right-hand side of (4.22) 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 (4.22) satisfies the bound

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

where the last step follows from the Lipschitz property of Φ and the bound of $(\lambda_t)_{t \geq 0}$. 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 $\Upsilon(U^*, V^*) = (U^*, V^*)$. ■

At the core of the previous proof is that the cumulative distribution function Φ_t is uniformly Lipschitz in $[0, T] \times \Omega$, and that the intensity process $(\lambda_t)_{t \geq 0}$ is bounded. As stated before, the parameter constraint in (4.20) restricts the time horizon of the problem, which also arises in similar studies such as that of Antonelli (1993). Note that the inequality in (4.20) involves the upper bound $\bar{\lambda}$ (see Assumption 2) 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 (4.13) 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 (4.1). Therefore the control δ^* that satisfies (4.13) 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 (4.16), then $\|\delta^* - \delta^*\|_1 = 0$.

4.3 Optimality

In this section we show that the discretion δ^* in (4.13) is the global minimizer of the agent's performance criterion $J(\delta)$. First, Lemma 2 shows that it is suboptimal to send MLOs with negative discretion. Then, after two auxiliary lemmas, Theorem 5 demonstrates that if the infimum is attained in \mathcal{A} , it coincides with δ^* . Finally, we introduce a collection of coercive functionals to prove that δ^* is indeed the global minimizer of the performance criterion.

The next lemma shows that it is not optimal to send MLOs with negative discretion; i.e., it is suboptimal to send an MLO that bets on a price improvement relative to the best quote observed by the agent.¹⁸ An agent who sends an MLO with negative discretion is worse off than one who sends an MLO with zero discretion because: (i) an MLO with negative discretion precludes the agent from scooping price improvements that an MLO with zero discretion would receive, and (ii) a trade attempt with negative discretion has more chances of being rejected by the exchange than a trade attempt with zero discretion.

¹⁸We remark that this result is for a trader that directly penalises misses; see (4.6). In Chapter 5 we show that if we do not penalise misses, it can be optimal to send negative discretion.

Lemma 2. Let $\mathcal{A}^- = \{U \in \mathcal{A} \mid \exists(t, \omega) \in [0, T] \times \Omega \text{ such that } U_t(\omega) < 0\}$, and $\mathcal{A}^+ = \mathcal{A} \setminus \mathcal{A}^-$, then, for each $U \in \mathcal{A}^-$, $\exists V \in \mathcal{A}^+$ such that $J(V) \leq J(U)$.

Proof. Let $U \in \mathcal{A}^-$ and define $V = \max(0, U) \in \mathcal{A}^+$. As $U_t(\omega) < 0$ while $U_t(\omega) \leq V_t(\omega) \leq 0$ for all $(t, \omega) \in [0, T] \times \Omega$, we have $C_t^U \geq C_t^V$ and $D_t^U \geq D_t^V$ for all $t \in [0, T]$, which implies that $J(U) \geq J(V)$. ■

Theorem 4. The control δ^* in (4.13) 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 derivative¹⁹ at $\delta^* \in \mathcal{A}$ in the directions $\nu, \nu \in \mathcal{A}$ is

$$\langle \mathcal{D}^2 J(\delta^*), \nu, \nu \rangle = \mathbb{E} \left[\int_0^T \nu_t^2 \phi_t(\delta_t^*) dA_s \right] + \mathbb{E} \left[\left(\int_0^T \nu_t \phi_t(\delta_t^*) dA_s \right)^2 \right]. \quad (4.23a)$$

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

The next lemma shows that the functional J is Lipschitz continuous.

Lemma 3. Assumptions 1, 2, and 3 hold. Then, the functional $J : \mathcal{A} \rightarrow \mathbb{R}$ is Lipschitz continuous, i.e., there is $\eta > 0$ such that for all $w, \delta \in \mathcal{A}$, we have that $|J(\delta) - J(w)| \leq \eta \|\delta - w\|$.

Proof. Let $w, \delta \in \mathcal{A}$ and take

$$\eta = \sqrt{T} \bar{\lambda} (\tilde{k} + \alpha k + \gamma k) + 2\gamma k \bar{\lambda} \sqrt{T} \left(2\bar{\lambda} T + \sqrt{2} \sqrt{\mathbb{E}[N_T^2] + (\bar{\lambda} T)^2} \right) > 0. \quad (4.24)$$

Recall that \tilde{k} is the Lipschitz constant of $r_t(x) := \int_{-\infty}^x z \phi_t(dz)$ from Assumption 3. Write $J = J^C + \alpha J^{LP} + \gamma J^{QP}$ and observe that

$$|J(\delta) - J(w)| \leq |J^C(\delta) - J^C(w)| + \alpha |J^{LP}(\delta) - J^{LP}(w)| + \gamma |J^{QP}(\delta) - J^{QP}(w)|.$$

Next, we bound each term on the right-hand side of the inequality. First,

$$\begin{aligned} |J^C(\delta) - J^C(w)| &= \left| \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} z \left(\hat{G}(\delta_t - z) - \hat{G}(w_t - z) \right) \Phi_t(dz) \lambda_t dt \right] \right| \\ &\stackrel{(\text{Assump. 3})}{\leq} \tilde{k} \mathbb{E} \left[\int_0^T |\delta_t - w_t| dA_t \right] \\ &\stackrel{(\text{Assump. 1})}{\leq} \tilde{k} \bar{\lambda} \mathbb{E} \left[\int_0^T |\delta_t - w_t| dt \right]. \end{aligned}$$

Second,

$$\begin{aligned} |J^{LP}(\delta) - J^{LP}(w)| &\leq \mathbb{E} \left[\int_0^T |\Phi(\delta_t) - \Phi(w_t)| dA_t \right] \\ &\stackrel{(\text{Assump. 1 \& 2})}{\leq} \bar{\lambda} k \mathbb{E} \left[\int_0^T |\delta_t - w_t| dt \right]. \end{aligned}$$

¹⁹See B.3 for details of the second Gâteaux derivative.

Finally,

$$|J^{QP}(\delta) - J^{QP}(w)| \leq \mathfrak{A}^{\delta,w} + k \mathbb{E} \left[\int_0^T |\delta_t - w_t| dA_t \right] \leq \mathfrak{A}^{\delta,w} + k \bar{\lambda} \mathbb{E} \left[\int_0^T |\delta_t - w_t| dt \right], \quad (4.25)$$

where

$$\mathfrak{A}^{\delta,w} := \left| \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} \left(2G(\delta_t - z) D_{t-}^{\delta} - 2G(w_t - z) D_{t-}^w \right) \Phi_t(dz) dA_t \right] \right|.$$

The first inequality follows from Assumption 2 and the second from Assumption 3.

Next, we bound the first term on the right-hand side of the inequality in (4.25):

$$\begin{aligned} \mathfrak{A}^{\delta,w} &\leq \left| \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} 2G(\delta_t - z) \left(D_{t-}^{\delta} - D_{t-}^w \right) \Phi_t(dz) dA_t \right] \right| \\ &\quad + \left| \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} 2D_{t-}^w (G(\delta_t - z) - G(w_t - z)) \Phi_t(dz) dA_t \right] \right| \\ &\stackrel{(G(z) \leq 1 \text{ \& Assump. 1\&2})}{\leq} 2\bar{\lambda} \mathbb{E} \left[\int_0^T |D_{t-}^{\delta} - D_{t-}^w| dt \right] + 2k\bar{\lambda} \mathbb{E} \left[\int_0^T |\delta_t - w_t| D_{t-}^w dt \right] \\ &\leq 2\bar{\lambda} \mathbb{E} \left[\int_0^T \left| \int_0^{t-} \int_{\mathbb{R}} G(\delta_s - z) - G(w_s - z) \mathfrak{p}(dz, ds) \right| dt \right] \\ &\quad + 2k\bar{\lambda} \mathbb{E} \left[N_T \int_0^T |\delta_t - w_t| dt \right] \\ &\leq 2\bar{\lambda} \int_0^T \mathbb{E} \left[\int_0^{t-} \int_{\mathbb{R}} |G(\delta_s - z) - G(w_s - z)| \phi_s(dz) dA_s \right] dt \\ &\quad + 2k\bar{\lambda} \left(\sqrt{\text{Var}(N_T) \text{Var} \left(\int_0^T |\delta_t - w_t| dt \right)} + \mathbb{E}[N_T] \mathbb{E} \left[\int_0^T |\delta_t - w_t| dt \right] \right) \\ &\leq 2\bar{\lambda} k \int_0^T \mathbb{E} \left[\int_0^{t-} |\delta_s - w_s| \lambda_s ds \right] dt + 2\bar{\lambda}^2 k T \mathbb{E} \left[\int_0^T |\delta_t - w_t| dt \right] \\ &\quad + 2k\bar{\lambda} \sqrt{2T} \sqrt{\mathbb{E}[N_T^2] + (\bar{\lambda} T)^2} \|\delta - w\| \\ &\leq 4\bar{\lambda}^2 k T \mathbb{E} \left[\int_0^T |\delta_t - w_t| dt \right] + 2k\bar{\lambda} \sqrt{2T} \sqrt{\mathbb{E}[N_T^2] + (\bar{\lambda} T)^2} \|\delta - w\|. \end{aligned}$$

Hence,

$$|J^{QP}(\delta) - J^{QP}(w)| \leq (k\bar{\lambda} + 4\bar{\lambda}^2 k T) \mathbb{E} \left[\int_0^T |\delta_t - w_t| dt \right] + 2k\bar{\lambda} \sqrt{2T} \sqrt{\mathbb{E}[N_T^2] + (\bar{\lambda} T)^2} \|\delta - w\|,$$

and by the Cauchy–Schwarz inequality we have that

$$\mathbb{E} \left[\int_0^T |\delta_t - w_t| dt \right] \leq \sqrt{T} \|\delta - w\|, \quad (4.26)$$

and we conclude that

$$\begin{aligned} |J(\delta) - J(w)| &\leq |J^C(\delta) - J^C(w)| + \alpha |J^{LP}(\delta) - J^{LP}(w)| + \gamma |J^{QP}(\delta) - J^{QP}(w)| \\ &\leq \left(\sqrt{T} \bar{\lambda} (\tilde{k} + \alpha k + \gamma k) + 2\gamma k \bar{\lambda} \sqrt{T} \left(2\bar{\lambda} T + \sqrt{2} \sqrt{\mathbb{E}[N_T^2] + (\bar{\lambda} T)^2} \right) \right) \|\delta - w\| \\ &= \eta \|\delta - w\|, \end{aligned}$$

where the last equality follows from the choice of η in (4.24), which completes the proof. ■

Before proving the main result of this section, which shows that our candidate control is the global minimum of the performance criterion $J(\delta)$, we state the following well-known lemma.

Lemma 4. *If the functional $J(\delta)$ has a global minimum $\hat{\delta} \in \mathcal{A}$, then*

$$\langle \mathcal{D} J(\hat{\delta}), w \rangle \geq 0, \quad \forall w \in \mathcal{A}. \quad (4.27)$$

Theorem 5. Uniqueness of global optimality. *If J has a global minimum at $\hat{\delta} \in \mathcal{A}$, then $\hat{\delta} = \delta^*$ a.e. in $\mathfrak{T} \times \Omega$, with δ^* as in (4.13).*

Proof. The proof is by contradiction. Suppose the global minimum $\hat{\delta} \in \mathcal{A}$, but it is not true that $\hat{\delta} = \delta^*$ a.e. in $\mathfrak{T} \times \Omega$, with δ^* solving (4.13), i.e., there exists $(\mathbb{T}, \mathfrak{D}) \in \mathcal{B}(\mathfrak{T}) \times \mathcal{F}_T$ with $\mathbb{L}(\mathbb{T}) \mathbb{P}(\mathfrak{D}) > 0$ such that $\hat{\delta} \neq \delta^*$ on $\mathbb{T} \times \mathfrak{D}$. First, by Lemma 4

$$\langle \mathcal{D} J(\hat{\delta}), w \rangle \geq 0, \quad \forall w \in \mathcal{A}, \quad (4.28)$$

and because $\hat{\delta} \neq \delta^*$ on $\mathbb{T} \times \mathfrak{D}$, there exists $\hat{w} \in \mathcal{A}$ such that $\langle \mathcal{D} J(\hat{\delta}), \hat{w} \rangle > 0$. Now, take $\tilde{w} = -\hat{w} \in \mathcal{A}$, then

$$\begin{aligned} \langle \mathcal{D} J(\hat{\delta}), \tilde{w} \rangle &= \mathbb{E} \left[\int_0^T \tilde{w}_t \phi_t(\hat{\delta}_t) \left(\hat{\delta}_t - 2\gamma \mathbb{E}_{t-} \left[D_T^{\hat{\delta}} \right] - \alpha - \gamma \right) dA_t \right] \\ &= -\mathbb{E} \left[\int_0^T \hat{w}_t \phi_t(\hat{\delta}_t) \left(\hat{\delta}_t - 2\gamma \mathbb{E}_{t-} \left[D_T^{\hat{\delta}} \right] - \alpha - \gamma \right) dA_t \right] \\ &= -\langle \mathcal{D} J(\hat{\delta}), \hat{w} \rangle \\ &< 0, \end{aligned}$$

which contradicts Lemma 4. Therefore, if there is a global minimum at $\hat{\delta} \in \mathcal{A}$, then $\hat{\delta} = \delta^*$ a.e. in $\mathfrak{T} \times \Omega$. ■

Thus, we have shown: (i) δ^* is a local minimum for J , (ii) the functional $J : \mathcal{A} \rightarrow \mathbb{R}$ is Lipschitz continuous, (iii) if the infimum for $J : \mathcal{A} \rightarrow \mathbb{R}$ is attained in \mathcal{A} , then the infimum is attained at $\delta^* \in \mathcal{A}$. Next, we show that if J is weakly lower semicontinuous, then δ^* is the global minimizer of J over \mathcal{A} . To this end, we introduce a collection of functionals that are coercive and weakly lower semicontinuous in the reflexive Banach space $(\mathcal{A}, \|\cdot\|)$, which guarantees the existence of the infimum in the space \mathcal{A} , we use this to show that J attains the infimum in \mathcal{A} . For completeness we show how to proceed if J is weakly lower semicontinuous (which we assume from this point forward). If J is not weakly lower semicontinuous then the result of local optimality is the best we can show.

Consider the functional $J^\psi : \mathcal{A} \rightarrow \mathbb{R}$ defined as

$$J^\psi(\delta) = \mathbb{E} \left[C_T^\delta + \alpha D_T^\delta + \gamma \left(D_T^\delta \right)^2 \right] + \frac{1}{2} \psi \|\delta\|^2,$$

where $\psi > 0$ is a constant; this new functional can be regarded as a Tikhonov regularisation of the original functional in (4.6).²⁰ It follows that for any $\delta \in \mathcal{A}$, $J^\psi(\delta)$ is continuous in ψ ; in particular we have that

$$\lim_{\psi \rightarrow 0} J^\psi(\delta) = J(\delta). \quad (4.29)$$

Moreover, for every $\psi > 0$, J^ψ is a coercive functional in the reflexive Banach space $(\mathcal{A}, \|\cdot\|)$, i.e.,

$$\lim_{\|\delta\| \rightarrow \infty} J^\psi(\delta) = \infty.$$

Given that for every $\psi > 0$, J^ψ is a coercive and a weakly lower semicontinuous functional, it follows that the infimum is attained in \mathcal{A} (see Theorem 9.6.6 in [Suhubi \(2003\)](#)). Take the Gâteaux derivative of J^ψ and observe that it vanishes in all directions if

$$\delta_t^{*\psi} = \frac{\lambda_t \phi_t(\delta_t^{*\psi})}{\psi + \lambda_t \phi_t(\delta_t^{*\psi})} \left(2\gamma \mathbb{E}_{t-} \left[D_T^{\delta^{*\psi}} \right] + \gamma + \alpha \right), \quad \forall t \in [0, T]. \quad (4.30)$$

For every $\psi > 0$, the function

$$x \rightarrow \frac{\lambda_t \phi_t(x)}{\psi + \lambda_t \phi_t(x)}$$

is Lipschitz because assumptions 1, and 2 hold, and using the same arguments in the proof of Lemma 3, the FBSDE (4.30) has a unique solution. Furthermore, Theorem 5 applies to J^ψ and $\delta^{*\psi}$, and for each $\psi > 0$, the global minimizer exists (by Theorem 9.6.6 in [Suhubi \(2003\)](#)). Hence, for every $\psi > 0$, $\delta^{*,\psi}$ is the global minimizer of J^ψ .

Theorem 6. *Let $\{\psi_n\}_{n \geq 1}$ be a sub-sequence of $\{1/n\}_{n \geq 1}$ from Lemma 8. Then*

$$\lim_{n \rightarrow \infty} J^{\psi_n}(\delta^{*\psi_n}) = J(\delta^*). \quad (4.31)$$

Proof. The proof follows immediately from the inequality

$$\left| J^{\psi_n}(\delta^{*\psi_n}) - J(\delta^*) \right| \leq \left| J^{\psi_n}(\delta^{*\psi_n}) - J^{\psi_n}(\delta^*) \right| + \left| J^{\psi_n}(\delta^*) - J(\delta^*) \right|,$$

from $\left| J^{\psi_n}(\delta^{*\psi_n}) - J^{\psi_n}(\delta^*) \right| \rightarrow 0$ as $n \rightarrow \infty$ (as a result of Lemma 8), and from $\left| J^{\psi_n}(\delta^*) - J(\delta^*) \right| \rightarrow 0$ as $n \rightarrow \infty$ from (4.29). ■

We use the above result to prove that δ^* is indeed a global minimizer of J .

Theorem 7. Global minimizer. *For all $\nu \in \mathcal{A}$ we have that $J(\delta^*) \leq J(\nu)$.*

Proof. We prove the theorem by contradiction. Assume there is $\nu^* \in \mathcal{A}$ such that $\epsilon = J(\delta^*) - J(\nu^*) > 0$. Then, by (4.29) and (4.31) there exists $\bar{n} > 0$ such that for $\bar{\psi} = 1/\bar{n}$

$$\left| J(\nu^*) - J^{\bar{\psi}}(\nu^*) \right| < \frac{\epsilon}{3} \quad \text{and} \quad \left| J(\delta^*) - J^{\bar{\psi}}(\delta^{*\bar{\psi}}) \right| < \frac{\epsilon}{3},$$

which implies that

$$J^{\bar{\psi}}(\nu^*) < J^{\bar{\psi}}(\delta^{*\bar{\psi}}),$$

and contradicts that $\delta^{*\bar{\psi}}$ is the infimum of $J^{\bar{\psi}}$. Thus, δ^* is the global minimizer of J . ■

²⁰The goal here is to have a coercive functional in a reflexive Banach space.

In general, adding a regularisation does not necessarily help to find optimality. This is because as the regularising term goes to zero, the solution typically misbehaves. However, in our framework, it turned out to be helpful to study the regularising sequence of optimal solutions because they help to identify the global optimality of the original problem.

4.4 Performance of strategy

The expectation that appears in (4.13) is conditional on the information \mathcal{F}_{t-} . Here, we study a slight variation of the FBSDE in (4.13) 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 (4.13) 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 $\check{\delta}$:

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

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

Assumption 4. *The stochastic intensity $(\lambda_t)_{t \in \mathfrak{T}}$ has the Markov property.*

By Assumption 4, 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

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

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

Theorem 8. *Let $\check{\delta}_t = h(t, D_t^{\check{\delta}}, \lambda_t)$. When assumption 4 and (4.1) hold, and h is once differentiable in the first argument and twice differentiable in the third argument, the function h satisfies the PIDE*

$$0 = \partial_t h(t, D, \lambda) + \mathcal{L}_t^\lambda h(t, D, \lambda) + \left(\int_{h(t, D, \lambda)}^\infty \lambda \phi_t(z) dz \right) (h(t, D+1, \lambda) - h(t, D, \lambda)), \quad (4.32)$$

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.²¹ ■

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 $\delta_t^* = h(t, D_{t-}^{\delta^*}, \lambda_{t-})$.

4.4.1 Poisson arrival of trades

We solve the PIDE in (4.32) 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 independent and identically distributed normal with parameters μ and σ^2 , $n = 1, 2, \dots$, and the trading horizon is $T = 1$ day.²²

We employ a subset of the data in subsection 4.1.1.1 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 assume that $Z_n \sim \mathbf{N}(0.1, 1)$. In Section 4.4.1.1 we explore the results for other values of the cost parameters α and γ .

Figure 4.2 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 4.3 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

²¹Note that the martingale property of $(\check{\delta}_t)_{t \in \mathfrak{T}}$ follows from its definition and a straightforward calculation to prove the integrability condition – the number of misses under any strategy can always be bounded by the number of trade attempts that we know to be integrable.

²²We employ a homogeneous Poisson process for simplicity. Self-exciting intensities and arrivals that reflect the diurnal patterns observed in FX markets remain to be studied.

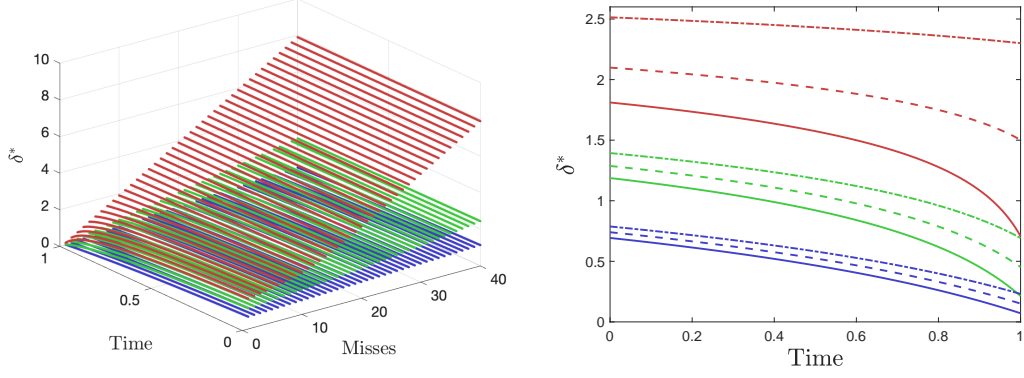


Figure 4.2: Left panel: Optimal strategy δ^* as a function of time and the number of missed trades for $\gamma = 0.01$ (bottom surface), $\gamma = 0.05$ (middle surface), and $\gamma = 0.15$ (top surface). The remaining parameters are: $\lambda = 100$, $\alpha = 0$, and $Z_n \sim \mathbf{N}(0.1, 1)$ 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 $D = 12$.

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.

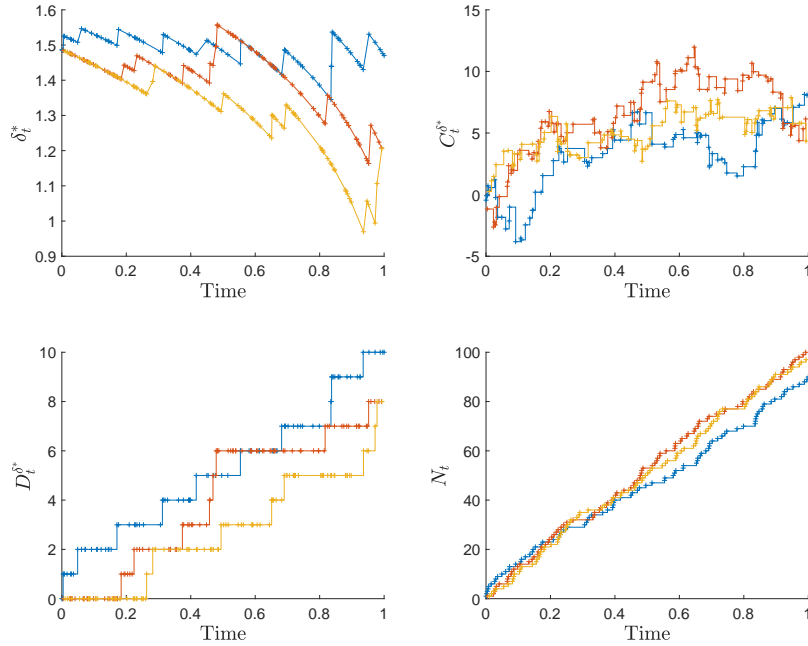


Figure 4.3: Sample paths for the optimal discretion δ^* (top left panel), number of missed trades D^{δ^*} (lower left panel), cost of strategy C^{δ^*} (top right panel), and number of trade attempts N (lower right panel) for three simulations of the MPP. Parameters: $\alpha = 0$, $\gamma = 0.1$, $\lambda = 100$, $T = 1$.

Figure 4.4 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 4.4: 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 Chapter 3. 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.

4.4.1.1 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.

Before discussing the results of the various simulation exercises, we prove a lemma. The lemma shows that the expected number of misses of a cost-neutral *fixed*-discretion strategy is lower than the expected number of misses of a cost-neutral *variable*-discretion strategy.

Lemma 5. *Let $\hat{\delta} \in \mathcal{A}$ be a positive cost-neutral fixed-discretion strategy, that is, $\mathbb{E}[C_T^{\hat{\delta}}] = 0$ and $\forall(t, \omega) \in \mathfrak{T} \times \Omega$ we have that $\hat{\delta}_t(\omega) = \hat{\alpha} \in \mathbb{R}^+$. Then $\forall \delta \in \mathcal{A}$ such that $\mathbb{E}[C_T^\delta] = 0$ we have that $\mathbb{E}[D_T^{\hat{\delta}}] \leq \mathbb{E}[D_T^\delta]$.*

Proof. Let $\hat{\delta}$ be a positive cost-neutral fixed-discretion strategy with $\hat{\delta} = \hat{\alpha} \in \mathbb{R}^+$. Suppose that there $\exists \tilde{\delta} \in \mathcal{A}$ with $\mathbb{E}[C_T^{\tilde{\delta}}] = 0$ and $\mathbb{E}[D_T^{\tilde{\delta}}] > \mathbb{E}[D_T^{\hat{\delta}}]$. Then, $J^C(\tilde{\delta}) = J^C(\hat{\delta}) = 0$ and $J^{LP}(\tilde{\delta}) < J^{LP}(\hat{\delta})$. Fix $\alpha = \hat{\alpha}$ and $\gamma = 0$, then $J(\tilde{\delta}) < J(\hat{\delta})$, which is a contradiction, because the global minimizer is $\hat{\delta}$ – see (4.15). ■

Lemma 5 pinpoints one of the insights of the strategy we develop in this chapter. For traders who seek cost-neutral strategies and also wish to temper their exposure to the number of missed trades over a trading horizon, it is optimal to employ cost-neutral strategies with fixed discretion. This could explain the strategies of the foreign exchange traders discussed in subsection 4.1.1.1.

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

The top panels in Figure 4.5 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 (4.13) 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 4.5: 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 (4.15). 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 $\mathbb{E}[C_T^{\delta^*}]$ when $\mathbb{P}(D_T^{\delta^*} < 0.1 N_T) \geq 0.99$. 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.

4.4.2 Pinned arrival rates

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

$$\lambda_t^* = \frac{M - N_{t-}}{T - t + \epsilon}, \quad (4.33)$$

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

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

with

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

Figure 4.6 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 4.2.




Figure 4.6: 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 (4.33) with $\epsilon = 0.1$. Figures 4.7 and 4.8 report the results, which have a similar interpretation to that of Figures 4.3 and 4.4, respectively.

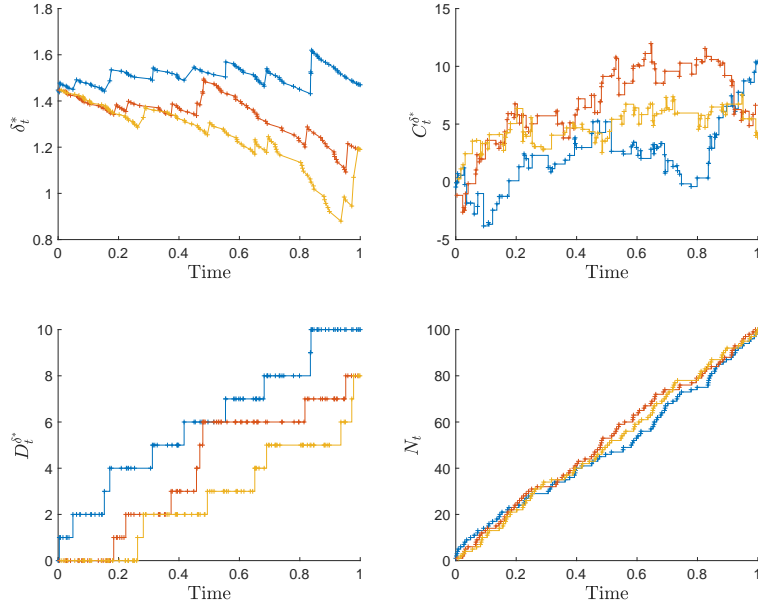


Figure 4.7: Sample paths for the optimal discretion δ^* (top left panel), number of missed trades D^{δ^*} (lower left panel), cost of strategy C^{δ^*} (top right panel), and number of trade attempts N (lower right panel) for three simulations of the MPP. Parameters: $\alpha = 0$, $\gamma = 0.1$, $\epsilon = 0.1$, $M = 100$, $T = 1$.




Figure 4.8: 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 $D_T^{\delta^*}/N_T$.

4.5 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 chapter 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 chapter. 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. \(2019b\)](#). 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.

Chapter 5

Optimal trading with delay

As we have seen so far, in electronic trading, agents are exposed to the risks that stem from latency in the marketplace. Latency is the aggregate of the time lags associated with the various stages of a trade. These stages, which occur in succession and are separated in time by random delays, include: the exchange sends quotes to an agent, the agent receives the quotes, the agent processes information and sends an instruction to the exchange, the exchange receives and handles the instruction, and finally, the outcome is notified to the agent. During the latency period, it is likely that the exchange will process other instructions that modify the limit order book (LOB) and possibly affect the outcome of the agent's instructions. In particular, liquidity takers face the risk that the prices they target are not available by the time the exchange processes their order because the best quotes they observed were stale and therefore updated during the latency period.¹ If the market moves against the trader's interest, the order will not be filled or will be filled at worse prices – the outcome depends on the type of liquidity taking order sent by the trader.

At high frequencies (milliseconds, microseconds), the best quotes in the LOB tend to flicker. Flickers are unpredictable short-lived deviations of a few ticks from the best quotes, and are the result of very frequent occurrences of rapid sequences of post-and-cancel activity in the LOB, and the less frequent arrival of aggressive orders that consume liquidity that is immediately replenished. Thus, over the latency period, the main source of risk faced by liquidity takers is from flickers in the book, and, to a lesser extent, the risk that stems from unexpected changes in the fundamental best quotes by the time their orders are processed. If latency were zero, traders would not face these risks because they would always take liquidity at the prices and quantities they observe in the LOB; however, all market participants face latency and latency is random. Indeed, the latency faced by each market participant consists of an idiosyncratic and a common component. A trader's software, hardware, and co-location are pivotal ingredients that regulate their speed in the marketplace, while the speed of all traders is affected by the exchange's capability to process messages and by the quality of the communication network.

¹The risks faced by liquidity providers are different. In particular, their limit orders may be adversely selected by a taker.

In this chapter, we show how a trader can execute a large position in a financial instrument when the trader faces random latency in the marketplace. The dynamics of the best bid price in the LOB consist of the ‘fundamental’ best quote of the instrument and the flickers (similarly for the best ask price). Innovations in the fundamental best quote are driven by a stochastic process, and the size and arrival times of flickers in the quotes are represented by a marked point process. The trader employs marketable limit orders (MLOs) to execute the position over a time window; MLOs are liquidity taking orders that specify a price limit and are for immediate execution only. Filled MLOs have permanent price impact: a filled buy (sell) MLO exerts upward (downward) pressure on the fundamental value of the instrument. However, if the price limit of the MLO precludes it from being filled, the exchange cancels the order and there is no price impact because other market participants cannot observe rejected trades.

In our model, the trader controls the price limit of the MLO and controls when to submit it to the exchange, both of which largely depend on the trader’s degree of urgency to execute the position.² An impatient liquidity taker will send sell (buy) MLOs with price limits that are below (above) the best bid (ask) price they observed in the LOB. This strategy increases the probability that the MLO is filled when processed by the exchange and caps how far the MLO is allowed to walk the LOB when there are adverse price changes over the latency period. On the other hand, we show that patient liquidity takers use MLOs predominantly to send speculative trades. Speculative MLOs seek a price improvement relative to the price observed by the trader when sending the order to the exchange. Patient liquidity takers do not use their speed to finalise their execution programme ahead of time, nor to hit the best quotes they observe in the LOB, they use speed mainly to send as many speculative MLOs as possible during the trading horizon – fast traders have many opportunities to retry rejected orders before reaching the end of the trading window.

We employ proprietary foreign exchange data from LMAX Exchange (henceforth LMAX) between September 2019 and February 2020 for ten currency pairs to study the order types that liquidity takers use in the foreign exchange market and to analyse the effect of latency on the efficacy (i.e., hit and miss) of liquidity taking orders.³ In all pairs, MLOs trade more volume than any other type of liquidity taking order; in particular, more than market orders (MOs), which are for immediate execution and walk the LOB until filled in full. Though MLOs offer protection against adverse price moves, traders concede that the order may be cancelled by the exchange because clearing prices would breach the price limit instructed in the MLO. For example, in the EUR/USD pair, 41.0% of the MLOs were filled in full, 1.9% were partially

²Our model differs from the well-known Almgren-Chriss framework in a number of modeling assumptions, namely, (i) we employ MLOs as opposed to MOs (i.e., we can miss trade attempts), (ii) stochastic and deterministic latency can both be accommodated in the framework, and (iii) we introduce the concept of ‘flickers’ and the key role of price protection over the latency period. Our models are similar in that (i) we both search for the optimal strategy to liquidate a position under the presence of transaction costs and price impact, and (ii) the liquidity taker does not employ LOs to execute the position. In an Almgren-Chriss framework with execution delays, the trader would not choose the optimal times to send the liquidity taking orders.

³See www.lmax.com.

filled, and 57.1% were rejected.⁴ The MLOs that were partially or fully filled represent 56.9% of the volume traded in the pair, while MOs represent 13.5% of the total volume. The limit rates of rejected MLOs were, on average, three ticks away from the best quote in the LOB when they were processed by the exchange.

We use the LMAX data to implement and to benchmark the performance of the trader's random-latency-optimal execution strategy against four strategies: (i) send MOs over a trading window and assume latency is zero, (ii) send MLOs over a trading window and assume that latency is deterministic, (iii) employ a discrete time-weighted average price (TWAP) strategy that sends MOs at equally spaced time intervals over the execution window, (iv) send one MO to execute the entire order at the beginning of the trading window, where, unrealistically, we assume that there is enough liquidity at the best bid in the LOB to fill the entire order.

Estimates of model parameters for the EUR/USD currency pair are obtained with data between 1 August 2019 and 31 August 2019, and trading strategies are implemented with data between 1 September 2019 and 29 February 2020. We show that the performance of the random-latency-optimal strategy and that of the deterministic-latency-optimal strategy are statistically the same.⁵ When the trader is patient, both strategies outperform the other benchmarks by a cash amount that is greater than the transaction costs paid by liquidity takers in foreign exchange markets. And around news events, the value of the outperformance increases to between two and ten times the value of the transaction costs. The source of the better performance of the latency-optimal strategies stem from the speculative MLOs that are filled during the trading horizon, and the price protection of the MLOs against adverse price moves. In the EUR/USD pair, the number and the value of filled speculative MLOs increase when market activity increases because the probability and the size of observing positive flickers in the best bid rate increase during heightened market activity – this explains the considerable outperformance of the latency-optimal strategies around news events. Finally, we show that the value of the outperformance decreases as the degree of impatience of the trader increases because the strategy sends fewer speculative MLOs – very impatient traders do not send speculative MLOs, all their MLOs are for price protection.

As far as we are aware, this is the first work to appear in the literature that shows how to optimally execute a position in an asset when the trader faces random latency. Closest to our work, are those by Øksendal and Sulem (2008) and Bruder and Pham (2009), where the authors investigate general impulse control problems with deterministic delay, while ours assumes stochastic delay. Øksendal and Sulem (2008) study an infinite horizon control problem

⁴Note that these percentages are describing all market participants. While brokers tend to care about not missing trade attempts, there are many other market participants that send speculative MLOs that aim to find hidden liquidity; later in the chapter we see that these practices constitute the majority of the misses we observe here. Thus, these rejections should not be thought of as trader indifference to misses.

⁵As stated above, latency is random; thus, assuming a deterministic latency is a simplification. However, our results suggest that as long as the trader knows the expected latency, the performance of both the random-latency-optimal strategies and the deterministic-latency-optimal strategies is similar. The difference between random-latency-optimal strategies and deterministic-latency-optimal strategies needs to be studied further in other financial markets/instruments.

with deterministic delay and arbitrary number of pending orders, while [Bruder and Pham \(2009\)](#) investigate a finite horizon stochastic control problem with any finite number of pending orders –these two works differ in the assumptions on the impulse operator and general dynamics of the system. Here, our notation and framework is closer to that in [Bruder and Pham \(2009\)](#). The key mathematical novelty that distinguishes our work is that delays are stochastic. From a modelling perspective, this requires an explicit definition of the sigma algebra of the trader and, compared to the literature, our proofs require a careful study of the new terms that arise as a consequence of stochastic delays.

Although our model focuses on an optimal execution problem with stochastic latency, our framework is general and can be applied to control problems in which the outcomes of the controlled actions are observed with stochastic delay. In financial applications, our model may be implemented in all asset classes that trade in electronic LOBs with liquidity taking orders that consider price limits and immediate time-in-force of the liquidity taking orders.⁶ In contrast, all models of optimal execution in the extant literature assume that the trader operates in the marketplace with zero latency and employ only MOs to take liquidity.⁷

In the next section, we employ high-frequency foreign exchange data to show that MLOs protect orders from adverse price moves and also receive price improvements when prices move in the trader’s interest over the latency period. Section [5.2](#) introduces the trader’s optimal execution problem with random latency and characterises the random-latency-optimal strategy as the solution of a Hamilton-Jacobi-Bellman quasi-variational inequality (HJBQVI). Section [5.3](#) develops two benchmarks where the investor faces deterministic latency or faces no latency. Section [5.4](#) compares the performance of the trader’s random-latency-optimal strategy with that of the benchmarks. Section [5.5](#) concludes and we collect some proofs in the Appendices.

5.1 Order types and data

5.1.1 Orders and revealed preferences

In order-driven electronic markets, the basic building blocks to trade are orders that provide liquidity and orders that take liquidity. When a trader sends an order to the exchange, the order contains two specific instructions: quote type and time-in-force.⁸ The quote type specifies the main feature of the order: price limit (i.e., LO), no price limit (i.e., MO), a trigger (i.e., stop orders), among others. 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

⁶Major equity exchanges and foreign exchange venues that offer MLOs include LMAX, LSE, Chi-X, Cboe, Nasdaq, CME, and NYSE.

⁷For example, see [Almgren and Chriss \(2000\)](#), [Almgren \(2003\)](#), [Bayraktar and Ludkovski \(2010\)](#), [Alfonsi et al. \(2010\)](#), [Guéant et al. \(2012\)](#), [Guilbaud and Pham \(2013\)](#), [Cartea et al. \(2014\)](#), [Cartea and Jaimungal \(2015\)](#), [Guéant \(2015\)](#), [Cartea and Jaimungal \(2016\)](#), [Barger and Lorig \(2019\)](#), and the monographs [Cartea et al. \(2015\)](#) and [Guéant \(2016\)](#).

⁸See Section [2.2](#).

the rest of the trading day.⁹ Two prominent liquidity taking order types that have an immediate execution time-in-force are immediate-or-cancel (IoC) and fill-or-kill (FoK). IoC is an order to buy or to sell assets that must be executed immediately, in full or in part, while obeying the order's price limit, and any portion of the volume of the order that cannot be filled at the desired price limit is cancelled. FoK is an order to buy or to sell assets that must be executed immediately in full or it is cancelled.

Combinations of the various features constitute the type of order that market participants send to the exchange. Here, we focus on the types of liquidity taking order that are designed to protect traders against the frictions that stem from latency. MLOs are liquidity taking orders that have a price limit and the time-in-force is IoC or FoK.¹⁰ MLOs protect liquidity takers from adverse price movements that occur between the time the trader makes a decision to trade (on a possibly stale quote) and the time the exchange matches the order with a limit order resting in the book (when possible). If latency were zero for all liquidity takers, MLOs would be redundant because traders would take liquidity at the price and quantity, they observe in the book; i.e., trades would not need price protection.¹¹

There are, however, no traders with zero latency in the marketplace. Indeed, each market participant is exposed to their own random latency, which depends on two components: (i) Idiosyncratic component: hardware, software, and co-location of each trader are key determinants of the baseline latency that they experience in the market. (ii) Common component: all traders are exposed to the random delays in the matching engine when it becomes overcrowded (the capacity of exchanges to process messages is limited), and are exposed to the variability in the quality and speed of the communication channels between traders and the exchange. Therefore, the time it takes traders to receive and 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 (and of each trade) in the marketplace; see Chapter 3 for empirical evidence in 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. Their 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 an MO 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 an MLO. Thus, in exchange for price protection against adverse price movements, the trader concedes that the trade might not get filled. On the other hand, if prices move in a favourable direction, the trader receives a

⁹Venues such as LSE allow long time-in-force. For example, LSE has the option of sending orders that can be active for up to 89 days.

¹⁰An MO is an MLO without price limit; here the MLOs we consider have a finite price limit and are for immediate execution.

¹¹An advantage of trading with MLOs, as opposed to trading with LOs (without IoC/FoK time-in-force), is that attempts to trade using MLOs are not revealed to other market participants unless a trade happens, whereas all LOs that rest in the LOB are disclosed to other market participants.

price improvement; i.e., the MLO is executed at a better price than the trader’s decision price. Indeed, MLOs are also used by traders to complete a trade only if the price improves over the latency period. In the next section we provide summary statistics of the use of MLOs in foreign exchange markets.

In asset classes such as equity, fixed income, foreign exchange, and commodities, MLOs are available in exchanges that trade in a limit order book.¹² Each exchange has its own rules, but all acknowledge that liquidity takers seek to protect their trades from the risks that arise from latency. Thus, exchanges provide market participants with order types where they can specify the price limit and time-in-force of their liquidity taking orders.

5.1.2 Data

We present descriptive statistics for ten currency pairs in the foreign exchange spot market at LMAX. The data are stamped at a microsecond frequency and the range is from 1 September 2019 to 29 February 2020 – in Section 5.4 we use data from August 2019 to estimate model parameters. For each currency pair we have: the aggregate volume of the limit sell orders posted at the best ask rate; the aggregate volume of the limit buy orders posted at the best bid rate; liquidity taking orders sent to the LOB, including the type of order and trader identification; rate limits of MLOs; full fills, partial fills, and rejected liquidity taking trades. Finally, our data set does not contain the liquidity posted at the LOB beyond the best bid and best ask rates, but does contain the average rate paid or received for all liquidity taking orders, including those that walked the LOB beyond the best quotes.

Table 5.1 shows the percentage of the total number of trades that are MLOs, MOs, and others, and also shows the percentages by traded volume. Recall that the MLOs we consider are IoC and FoK with an immediate time-in-force. In practice, the exchange processes the order as soon as it arrives and matches it (if possible). This processing time usually takes under 80 microseconds and it is followed by a message from the exchange to the trader to notify the outcome.

The category ‘others’ includes the remaining order types that take liquidity, for example, marketable: (i) limit orders that are good-until-cancel and good-for-day, (ii) stop orders good-for-day and good-until-cancel, and (iii) dark limit orders good-for-day and good-until-cancel.

¹²E.g., see order types in <https://markets.cboe.com/>, <https://www.londonstockexchange.com/>, www.lmax.com/, www.cmegroup.com, www.nasdaq.com. Note that FoK, IoC are the time-in-force qualifier of MLOs.

	Number of trades			Volume traded		
	MLOs	MOs	Other	MLOs	MOs	Other
EUR/USD	36.7%	32.1%	31.2%	56.9%	13.5%	29.6%
USD/JPY	41.4%	6.9%	51.7%	55.8%	3.1%	41.1%
EUR/GBP	38.3%	8.0%	53.7%	73.0%	5.2%	21.8%
GBP/USD	53.6%	21.9%	24.5%	72.8%	7.1%	20.2%
USD/MXN	44.7%	24.6%	30.7%	52.0%	26.3%	21.7%
USD/CAD	62.2%	10.9%	26.9%	71.2%	2.6%	26.1%
GBP/JPY	49.1%	19.4%	31.5%	61.2%	8.5%	30.3%
EUR/JPY	33.0%	7.3%	59.8%	65.0%	4.3%	30.8%
AUD/USD	66.7%	10.6%	22.7%	71.0%	4.4%	24.6%
AUD/JPY	52.8%	11.7%	35.5%	59.9%	6.1%	34.0%

Table 5.1: Percentage of liquidity taking by order type. Period: 1 September 2019 to 29 February 2020. For each currency pair, we use bold to highlight the highest percentage by number of trades and by volume traded.

By volume, MLOs are employed more than any other type of liquidity taking order in all currency pairs. In highly liquid currency pairs such as EUR/USD, USD/JPY, EUR/GBP, GBP/USD, the volume traded with MLOs is between 55% and 73%, which is between four and fourteen times higher than the volume traded with MOs.

The information in Table 5.1 shows that traders in LMAX have a strong preference for MLOs over any other type of order to consume liquidity. As discussed above, these revealed preferences show that traders balance the tradeoff between the costs of missing a trade as a result of latency in the marketplace and both the costs of price protection and the benefits of speculative trades. A trader who must complete a trade will choose an MO. A trader who can afford not to complete a trade, if the price is ‘not right’, will choose an MLO.

5.1.3 Flickering of best rates in the LOB

Figure 5.1 shows snippets of the evolution of best quotes and liquidity taking activity for the currency pair EUR/USD on 3 September 2019 at around 2.00pm British summer time – the data are timestamped to the microsecond. Solid circles denote filled sell MLOs, squares denote filled sell MOs, and empty circles denote sell MLOs that were not filled because the limit sell rate in the order was higher than the best bid rate in the LOB when the exchange processed the order.

The left and centre panels of the figure show the best bid rate and sell liquidity taking orders between 2.00.00pm and 2.01.00pm (i.e., one minute) and between 2.00.08pm and 2.00.14pm (i.e., six seconds), respectively. It is clear that the best bid rate goes through unpredictable flickers, which are extremely short-lived deviations of a significant number of ticks from the ‘fundamental’ or ‘true’ value of the best bid rate. These flickers are the result of: cancellations and arrivals of limit orders at the best bid rates, and liquidity that is consumed by aggressive orders and immediately replenished by the arrival of limit orders.

The panel on the right-hand side shows a one-second snippet, between 2.00.12pm and 2.00.13pm, of the best bid rate and the best ask rate and the liquidity taking orders, all of

which were buy MOs. In most cases, flickers on either side of the LOB cause a short-lived widening of the quoted spread of the LOB.

From visual inspection, flickers in the best bid (best ask) rate are more often negative (positive) than positive (negative). This is a prevalent feature in the dynamics of the best quotes in the LOB. We return to this in Section 5.4 when we estimate model parameters and discuss the distribution of the flickers in the performance of the execution strategies.




Figure 5.1: Best bid rate (blue line), best ask rate (red line), missed MLOs (empty circles), filled MLOs (full circles), filled MOs (squares). Left panel: 2:00pm to 2:01pm (60 seconds). Centre panel: 2:00.08pm to 2:00.14pm (6 seconds). Right panel: 2:00.12pm to 2:00.13pm (1 second). Microsecond data of EUR/USD pair, 3 September 2019.

5.1.4 Distribution of flickers: hit and miss

We proceed to discuss the limit rates employed by traders who send MLOs to buy and to sell the EUR/USD currency pair.¹³ We employ the same data as that in Table 5.1 – the results for the remaining currency pairs are similar. To analyse both the hit and miss performance of the MLOs and the distribution of the flickers in the best quotes, we compute the slippage-price-improvement (SPI) measure

$$\text{SPI}_i = (L_i - M_i) I_i. \quad (5.1)$$

Here, L_i denotes the exchange rate limit of the order, i.e., the maximum (minimum) rate that the trader is willing to pay (receive) per unit that she wishes to buy (sell); M_i is the exchange rate that the order would pay (receive) per unit bought (sold) if the order is filled in full. The indicator I_i determines the direction of the order: I_i takes the value $+1$ when the trader sends a buy MLO and takes the value -1 when the trader sends a sell MLO, and $i = 1, \dots, n$, where n is the total number of MLOs. Thus, when SPI_i is positive, the order was filled (relative to the limit rate L_i) with a slack of SPI_i ticks; which we refer to as price improvement. Similarly, when SPI_i is negative, the order missed (fully or in part) by SPI_i ticks (relative to the limit rate L_i); which we refer to as slippage.

For each MLO, the quantity in (5.1) could understate or overstate slippage and price improvement relative to the trader's decision rate; note that we compute SPI relative to the limit rate instructed in the order. Our data set does not contain the time-stamps of when the trader

¹³We choose to work with EUR/USD because it is the most traded currency pair in LMAX. The results and insights we obtain are the same for the other currency pairs.

submitted the MLO to the exchange, nor does it contain the rate observed by the trader when the trader decided to trade (i.e., the decision rate). We have the time-stamp of when the order is processed by the exchange. Thus, we do not know the best bid and best ask rates and quantities posted in the LOB when the trader decided to send the order.

The trader may find it optimal to send a buy MLO with a limit rate that is below the observed best ask rate or to send a sell MLO with a limit rate that is above the observed best bid rate. In other words, traders may send speculative MLOs that are filled only if there is a price improvement relative to the quotes they observe when deciding to trade – in Section 5.3 below we return to this point and show that it is optimal to send MLOs that target a price improvement.




Figure 5.2: Histograms of SPI between -25 ticks and 25 ticks in log-scale for the EUR/USD currency pair, 1 September 2019 to 29 February 2020. One tick in EUR/USD is 10^{-5} USD. Left panel: all MLOs. Right panel: MLOs with limit rate equal to trader's decision rate.

Figure 5.2 shows histograms in log-scale of SPI for the MLOs in the EUR/USD currency pair, where the x -axes in both panels are truncated to lie between -25 and 25 ticks. To gain insights into the limit rates of the MLOs and into the distribution of the flickers, the left panel of the figure shows the histogram of SPI with all the MLOs sent to the exchange, and the right panel shows a histogram of SPI with the MLOs that aimed at the best quote, i.e., the trader's decision rate and the limit rate she instructed in the MLO are the same. We consider that the limit rate of an MLO is equal to the decision rate of the trader if at any point within the previous 150 ms of the processing time of the MLO there was a best quote in the LOB equal to the limit rate of the MLO.¹⁴ The histograms are for both IoCs and FoKs. We do not show a histogram for IoCs and one for FoKs because their shapes are similar to those of the histograms shown here.

In both panels, the histograms are skewed to the left. The skewness may be due to informed traders in the market and due to the distribution of the flickers. Informed traders send buy (sell) MLOs in anticipation of an increase (decrease) in the exchange rate. Thus, all else being equal, their trading strategies skew the distribution of SPI to the left because it is more likely that a buy (sell) MLO misses the trade when the exchange rate drifts up (down). Informed MLOs are included in the left panel. In addition, the skew in the histogram also results from

¹⁴We obtain similar results when we assume that the look-back window is 100 ms.

asymmetry in the distribution of the size of the flickers – this is better represented in the right panel where the limit rate of the MLO is equal to the trader’s decision rate; see also right-hand panel of Figure 5.1, where most flickers on the best bid (ask) rate are negative (positive). Below we show that flickers in the best bid rate are negatively skewed and flickers in the best ask rate are positively skewed.

The histograms provide a consolidated view of the distribution of SPI for all traders. However, the distribution of SPI faced by each individual trader depends on their liquidity taking strategy and, more importantly, on their latency. We return to this point below when we implement execution strategies and show how a trader uses LOB data to determine the distribution of SPI they face.

5.1.5 Liquidity: make and take

In this subsection, we report various statistics of the volume posted at the best quotes in the LOB and the volume of the filled liquidity taking orders of the ten currency pairs we study. The key message is that liquidity taking orders hardly ever consume all the liquidity available at the best quotes in the LOB.

Table 5.2 reports statistics of the liquidity taking and the liquidity provision activity during the period 1 September 2019 to 29 February 2020 between 9.00am and 4.00pm British summer time. Columns two to four describe the liquidity taking activity. The mean, median, and standard deviation are in 10,000 units of the base currency. For example, the mean value of the traded quantity in the pair EUR/USD is 20.94, which, in the base currency EUR, is a mean of 209,400 EUR.

The last four columns of the table summarise statistics of the spread and of the best quotes in the LOB: the time-weighted spread for each currency pair, the time-weighted quantity, the median quantity available at the best quotes, and the mode quantity at the best quotes.

Pair	Volume of liquidity taking trades			Liquidity provision			
	Mean	Std	Median	TWS	TWQ ask/bid	Median Q ask/bid	Mode Q ask/bid
EUR/USD	20.94	26.64	4.00	3.94×10^{-5}	160 / 159	100 / 100	50 / 50
USD/JPY	26.81	23.78	27.00	4.20×10^{-3}	110 / 109	100 / 100	50 / 50
EUR/GBP	12.93	21.50	0.60	9.30×10^{-5}	94 / 94	80 / 82	50 / 50
GBP/USD	20.65	21.97	14.40	9.83×10^{-5}	91 / 91	75 / 66	50 / 50
USD/MXN	37.09	23.55	50.00	4.67×10^{-3}	80 / 80	50 / 50	50 / 50
USD/CAD	31.30	27.42	30.00	8.34×10^{-5}	101 / 104	100 / 100	50 / 50
GBP/JPY	20.03	16.35	30.00	1.55×10^{-2}	71 / 71	50 / 50	50 / 50
EUR/JPY	14.44	22.33	0.60	8.77×10^{-3}	100 / 100	100 / 100	50 / 50
AUD/USD	32.77	28.23	32.10	5.65×10^{-5}	113 / 114	100 / 100	50 / 50
AUD/JPY	31.08	23.27	40.00	8.62×10^{-3}	88 / 90	94 / 99	50 / 50

Table 5.2: TWS: time-weighted spread. TWQ: time-weighted quantity available at best quotes. Period: 1 September 2019 to 29 February 2020 between 9.00am and 4.00pm Each unit of quantity is a lot of 10,000 units of the base currency.

Table 5.3 shows the percentage of occurrences when the traded quantity is greater than: the time-weighted quantity at the best quotes; the median quantity at the best quotes; and the

mode of the quantity at the best quotes. Note that the information in columns 1, 3, 5 is the same as that in the last three columns of Table 5.2.

Currency pair	TWQ ask/bid	TQ > TWQ	Median Q ask/bid	TQ > MQ	Mode Q ask/bid	TQ > MoQ
EUR/USD	160 / 159	0.1%	100 / 100	0.2%	50 / 50	4.6%
USD/JPY	110 / 109	0.1%	100 / 100	0.1%	50 / 50	2.6%
EUR/GBP	94 / 94	1.2%	80 / 82	1.4%	50 / 50	1.9%
GBP/USD	91 / 91	1.2%	75 / 66	1.5%	50 / 50	2.3%
USD/MXN	80 / 80	4.6%	50 / 50	5.9%	50 / 50	5.9%
USD/CAD	101 / 104	0.2%	100 / 100	0.2%	50 / 50	7.6%
GBP/JPY	71 / 71	0.7%	50 / 50	0.9%	50 / 50	0.9%
EUR/JPY	100 / 100	1.8%	100 / 100	0.2%	50 / 50	2.8%
AUD/USD	113 / 114	0.1%	100 / 100	0.1%	50 / 50	8.8%
AUD/JPY	88 / 90	1.5%	94 / 99	1.4%	50 / 50	0.1%

Table 5.3: TWQ: time-weighted quantity available at best quotes. TQ: traded quantity. Q ask/bid: quantity posted at best ask rate and best bid rate. Period: 1 September 2019 to 29 February 2020 between 9.00am and 4.00pm. Each unit of quantity is for 10,000 units of the base currency.

We observe that in all currency pairs, the proportion of trades that cannot be filled with the liquidity available at the best quotes is between 0.1% and 8.8%. In the model that follows, we assume that the volume of the MLOs sent by the trader to the exchange is such that the orders will not walk the LOB.

5.2 Optimal execution with random latency

We present a model of optimal execution with random latency. We focus on the execution of a large order that is divided in child orders that are sent to the market over a finite trading horizon. However, we remark that our model is useful to decide the optimal strategy to execute as little as one order over a trading horizon.

Our discussion is framed within foreign exchange markets; however, our model is general and applicable in all asset classes in which instruments are traded in an visible electronic LOB. Also, our framework may be applied in general stochastic control problems in which the outcomes of the agent's actions are known at a random future date. The remainder of this subsection proceeds as follows. Subsection 5.2.1 discusses the correspondence between the clock of the exchange and the clock of the trader. Subsection 5.2.2 presents a model of the best bid rate and in subsection 5.2.3 we discuss the modelling framework to determine fills or misses of trade attempts. Subsection 5.2.4 characterizes the set of admissible strategies of the trader, and subsection 5.2.5 presents the system dynamics of the control problem faced by the trader. Finally, we define the value function of the trader and discuss stylised features of the random-latency-optimal strategy.

5.2.1 Trader's clock

As discussed above, traders observe data that could be stale because it takes time for the information sent by the exchange to reach the trader. In this subsection, we discuss the one-to-one correspondence between the clock in the exchange and the trader's clock. Due to this correspondence, it suffices to model all processes and derive the random-latency-optimal strategy in the trader's clock.

Figure 5.3 depicts how latency affects the flow of information between the exchange and the trader; the left (right) panel assumes that latency is deterministic (stochastic). In both panels of the figure, we map the time of the events in the exchange to the time of the events for a trader who experiences latency in the marketplace. At time $e_0 = t_0$ the exchange sends data to the trader, the trader receives the data at time $t_1 = e_1$, where $t_1 - t_0 = e_1 - e_0$ is the observational latency. From $t_1 = e_1$ to $t_2 = e_2$, the trader processes market information and decides to trade. At time $t_2 = e_2$ the trader sends an instruction to the exchange and the exchange receives the instruction at $e_3 = t_3$. Between $e_3 = t_3$ and $e_4 = t_4$ the exchange processes the trade attempt and at time $e_4 = t_4$ the exchange notifies the agent the outcome of the trade attempt. The trader receives the notification at t_5 .




Figure 5.3: Diagram of the various layers of latency. Left panel: deterministic latency. Right panel: random latency.

In both cases depicted in the figure, the latency e_0 to e_4 in the exchange maps to the latency t_1 to t_5 in the trader's clock. Thus, our model is developed in the timeline of the trader (blue arrow), where the relevant delays are those from t_1 to t_5 , which, due to the one-to-map correspondence, accounts for the latency period of each trade attempt, including the observational delay.

5.2.2 Exchange rate dynamics and execution of MLOs

The left-hand side of Figure 5.4 shows a simulated path of the best bid rate, which we depict as: the fundamental best bid rate (continuous line) plus short-lived flickers (dash line). In the figure, t_0 denotes the time when the trader sends a sell MLO to the exchange, and \tilde{t}_0 denotes the time when the trader receives the notification of the outcome of the MLO. For clarity, we say that flickers happen at 'processing times' and the trader learns about it at 'notification times' – see Figure 5.3, where the flicker that potentially affects the trade outcome happens at time e_4 , but it is learned by the trader at the notification time t_5 ; note that the pair (t_0, \tilde{t}_0) corresponds to the pair (t_1, t_5) in Figure 5.3.

As discussed above, flickers are short-lived deviations, so here we assume that only the flicker that occurs at the processing time \tilde{t}_0 is relevant when the exchange processes the trader's sell MLO. Therefore, the trader models the best bid rate $\hat{S} = (\hat{S}_t)_{t \geq 0}$ as

$$\hat{S}_t = S_t + F_t, \quad (5.2)$$

where $S = (S_t)_{t \geq 0}$ denotes the fundamental best bid rate, and $F = (F_t)_{t \geq 0}$ denotes the flickers that affect the bid rate only at processing times of MLOs.¹⁵ Model (5.2) does not include flickers that occur between processing times of the MLOs sent by the trader – it models paths that are relevant for the trader; see e.g., right-hand side of Figure 5.4. In the sequel, we refer to \hat{S}_t as the observed best bid rate and in Section 5.4 we show how the trader employs the data of the LOB to obtain the fundamental best bid rate S_t .¹⁶




Figure 5.4: Left panel: best bid rate in LOB, which consists of fundamental best bid and all the flickers. Right panel: best bid rate observed by the trader $\hat{S}_t = S_t + F_t$.

Formally, the notification time of the order and the size of the flicker are modelled by the background marked point process (MPP) $\mathcal{N} = (T_n, Z_n)_{n \geq 1}$ with random measure $\mathbf{p}(dz, dt)$. The background MPP dictates the stopping time when the outcome of each order is notified to the trader and dictates the flicker that affects the fundamental bid rate when each order is processed.¹⁷ That is, the trader sends a sell MLO to the exchange at time t and the exchange notifies the outcome of the order at time \tilde{t} , where \tilde{t} is the next stopping time (after t) when a mark (i.e., flicker) of the MPP arrives. More precisely, let $n \geq 1$ such that $T_{n-1} \leq t < T_n$, so the notification time \tilde{t} is T_n and the flicker $F_{\tilde{t}}$ is Z_n . Note that the value of the flicker $F_{\tilde{t}}$ at the notification time is independent of the trader's information at time t .

¹⁵Equation (5.2) is key to include the flickers in our model. This approach differs from the Almgren-Chriss framework in that we do not model the best bid plus a temporary price impact. Recall that unlike the Almgren-Chriss framework, we do not optimise over the quantity of each trade attempt; in our model the quantity is fixed and is for an amount that can be filled with the liquidity available at the best quotes. The Almgren-Chriss model helps to find the optimal quantity to trade at each time step to minimise executions costs, while our framework helps to find the optimal limit rates to attach to the MLOs in a liquidation programme.

¹⁶The model for the best ask rate is similar. Below we focus on a liquidation problem, so only the best bid rate is relevant.

¹⁷In the next sections we see that the background MPP helps to keep the framework tractable. In particular, having exponential waiting times simplifies the state space because of the memoryless property of exponential random variables.

The background process \mathcal{N} is non-explosive and the interarrival times $(T_i - T_{i-1})_{i \geq 1}$ are independent and identically distributed exponential random variables with parameter $\lambda > 0$ and $T_0 = 0$. The flickers $(Z_n)_{n \geq 1}$ that affect the fundamental best bid rate at processing times take values in \mathbb{R} and are i.i.d. with law:

$$\nu(dz) = p_0 \delta_{\{0\}}(dz) + \underbrace{p_+ \eta_+ e^{-\eta_+ z} \mathbf{1}_{\{z > 0\}} dz}_{\text{price improvement}} + \underbrace{p_- \eta_- e^{\eta_- z} \mathbf{1}_{\{z < 0\}} dz}_{\text{slippage}}, \quad (5.3)$$

where $\delta_{\{0\}}(dz)$ is the Dirac measure centered at zero, $\eta_+, \eta_- \in (0, \infty)$, and $p_0, p_+, p_- \in \mathbb{R}^+$ with $p_0 + p_+ + p_- = 1$.¹⁸

Let $T \in (0, \infty)$. It is straightforward to see that

$$\mathbb{E} \left[(\mathbf{p}(\mathbb{R}, [0, T]))^2 \right] < \infty \quad \text{and} \quad \mathbb{E} \left[\int_0^T \int_{\mathbb{R}} |z| \mathbf{p}(dz, dt) \right] < \infty. \quad (5.4)$$

The law in (5.3) is a modelling choice; the framework we develop here is valid for any law $\nu(dz)$ that satisfies the two conditions in (5.4).

Finally, one could specify other models for the flickers of the best bid rates (similarly for the best ask rates). For example, one could assume that all flickers that affect the best bid quotes in the LOB follow a fast-mean-reverting to zero Ornstein–Uhlenbeck process. Another alternative is to include a further MPP in (5.2) to model ‘all’ the flickers in the best quote, in addition to those that affect the MLO when it is processed by the exchange. However, although these alternatives include desirable properties of the observed best bid price, they would increase the complexity of the problem without adding further insights; these alternatives are very likely to produce the same qualitative results we find below when we implement the latency-optimal strategies.

5.2.3 Outcome of trade attempts: miss or fill

The information flow that the trader observes is encoded in the filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$, defined by

$$\mathcal{F}_t = \sigma(W_s, \mathbf{p}(A, [0, s]), A \in \mathcal{B}(\mathbb{R}), s \leq t).$$

Let τ be an \mathbb{F} -stopping time, and set $N_t = \mathbf{p}(\mathbb{R}, (0, t])$ for $t \in (0, T]$ and $N_0 = 0$, where N_t denotes the total number of marks of the background process \mathcal{N} up to time t . We define the notification time $\tilde{\tau}$ of a trade attempt by

$$\tilde{\tau} = \inf \{t > \tau \mid \Delta N_t > 0\},$$

where $\Delta N_t = N_t - N_{t-}$. Thus, $\tilde{\tau}$ denotes the time when the exchange notifies the outcome of the trade attempt sent at time $\tau < \tilde{\tau}$. Henceforth, we put a tilde $\tilde{\cdot}$ on stopping times to denote that they are notification times.

¹⁸In Figure 5.4 we assume that S_t is a standard Brownian motion and the flickers arrive at exponential times with law

$$\nu(dz) = 0.76 \delta_{\{0\}}(dz) + 0.08 \times 2 e^{-2z} \mathbf{1}_{\{z > 0\}} dz + 0.16 \times 2 e^{2z} \mathbf{1}_{\{z < 0\}} dz.$$

The following lemma shows that: (i) the notification times are stopping times, (ii) the times between trade attempts and notification times are i.i.d. with exponential distribution.

Lemma 6. *Let $(\tau_i)_{i \in \mathbb{N}}$ be a sequence of \mathbb{F} -stopping times where the index i denotes the i th trade attempt, and let $(\tilde{\tau}_i)_{i \in \mathbb{N}}$ be the sequence of notification times of the outcome of each trade attempt. Let $(\tau_i)_{i \in \mathbb{N}}$ and $(\tilde{\tau}_i)_{i \in \mathbb{N}}$ satisfy $\tilde{\tau}_i \leq \tau_{i+1}$ for all $i \in \mathbb{N}$, then: (i) $(\tilde{\tau}_i)_{i \in \mathbb{N}}$ are \mathbb{F} -stopping times, and (ii) $(\tilde{\tau}_i - \tau_i)_{i \in \mathbb{N}}$ constitute a collection of independent and identically distributed random variables that are exponentially distributed with parameter $\lambda > 0$.*

Proof. (i) Let $j \in \mathbb{N}$ and $t \geq 0$. Then

$$\{\tilde{\tau}_j \leq t\} = \{\tau_j \leq t\} \cap \{N_t > N_{t \wedge \tau_j}\} \in \mathcal{F}_t$$

because $\{\tau_j \leq t\} \in \mathcal{F}_t$ and $\{N_t > N_{t \wedge \tau_j}\} \in \mathcal{F}_t$.

(ii) For $i \in \mathbb{N}$ and $t \geq 0$ we have

$$\begin{aligned} \mathbb{P}(\tilde{\tau}_i - \tau_i \geq t) &= \mathbb{P}(\inf\{s > 0 \mid \Delta N_{s+\tau_i} > 0\} \geq t) \\ &= \mathbb{P}(\inf\{s > 0 \mid \Delta N_s > 0\} \geq t) = e^{-\lambda t} \end{aligned}$$

because $(N_{\tau_i+s} - N_{\tau_i})_{s \geq 0}$ is a Poisson process with parameter λ . Finally, the collection $(\tilde{\tau}_i - \tau_i)_{i \in \mathbb{N}}$ has the independence property because $[\tau_i, \tilde{\tau}_i)_{i \in \mathbb{N}}$ are over non-overlapping intervals. ■

Next, we define the auxiliary process $F = (F_t)_{t \geq 0}$, which we employ to determine the outcome of the trade attempts. The value of F_t is the most recent mark up to time t . Thus, the process F satisfies

$$dF_t = \int_{\mathbb{R}} (z - F_{t-}) \mathbf{p}(dz, dt),$$

with $F_0 = 0$, and recall that $\mathbf{p}(dz, dt)$ is the random measure of the background process \mathcal{N} . When the trader sends a trade at time τ , and the notification time is $\tilde{\tau}$, then the value of the flicker $F_{\tilde{\tau}}$ is $F_{\tilde{\tau}}$.

At the notification time $\tilde{\tau}$, if the sell MLO is an FoK with limit rate \mathfrak{l} , it will be filled in full if $\mathfrak{l} \leq S_{\tilde{\tau}-} + F_{\tilde{\tau}}$; otherwise the MLO misses the trade and it is cancelled by the exchange. To formally define a miss and a fill of a trade attempt, we proceed as follows. Let $\tilde{H}(x) = \mathbf{1}_{x \leq 0}$ be a step function. The outcome of the sell MLO is determined by the step function $\tilde{H}(\mathfrak{l} - S_{\tilde{\tau}-} - F_{\tilde{\tau}})$, which takes the value of one when the trader receives the notification of a fill or takes the value of zero when the notification is of a miss.¹⁹ Next, approximate \tilde{H} with a \mathcal{C}^∞ function f^ε , where the parameter ε controls the convergence of f^ε to \tilde{H} as $\varepsilon \searrow 0$. For example, here we employ

$$f^\varepsilon(x) = \frac{e^{-\frac{2x}{\varepsilon}}}{1 + e^{-\frac{2x}{\varepsilon}}}, \quad (5.5)$$

¹⁹The MPP helps to define trade executions in a dynamically consistent manner. We remark that when the value of the flicker is positive and the MLO is filled, the liquidity providers that were offering liquidity at the best quote are likely to be adversely selected.

instead of \tilde{H} to determine the fill or miss outcome of a trade attempt. We work with an approximation to the step function \tilde{H} to preserve the continuity of the impulse operator in the trader's optimal execution problem.²⁰ To simplify notation, we drop the superscript ε and refer to the approximation in (5.5) as f .²¹

5.2.4 Admissible strategies

The trader controls when to send MLOs and controls the limit rate of the MLOs; this is summarised in the execution strategy $\alpha = (\tau_i, \mathfrak{l}_i)_{i \geq 1}$. The limit rates $(\mathfrak{l}_i)_{i \geq 1}$ of the MLOs are \mathcal{F}_{τ_i} -measurable and take values in a non-empty compact set $L \subset \mathbb{R}$. We assume that the trader does not attempt a trade until the outcome of the previous MLO is known. Thus, we require that the \mathbb{F} -stopping times in the sequence $(\tau_i)_{i \geq 1}$ satisfy $\tau_{i+1} \geq \tilde{\tau}_i > \tau_i$ for $i \in \mathbb{N}$. Specifically, the trader's set of admissible strategies is

$$\mathcal{A} = \left\{ \alpha = (\tau_i, \mathfrak{l}_i)_{i \geq 1} \mid \text{For each } i \geq 1, \tau_i \text{ is an } \mathbb{F}\text{-stopping time,} \right. \\ \left. \tau_{i+1} \geq \tilde{\tau}_i > \tau_i, \quad \mathfrak{l}_i \in L, \text{ and } \mathfrak{l}_i \text{ is } \mathcal{F}_{\tau_i}\text{-measurable} \right\}.$$

If at the terminal time T there is an outstanding MLO in the exchange, the order is cancelled. The trader keeps track of pending orders in the exchange with the process $k_t(\alpha)$, which returns the value one if the trader is waiting to be notified the outcome of a trade attempt or returns the value zero if there is no pending order waiting to be processed by the exchange. Thus,

$$k_t(\alpha) = \text{card} \{i \in \mathbb{N} : \tau_i \leq t, \tilde{\tau}_i > t\} \in \{0, 1\},$$

which is adapted to \mathbb{F} because for a fixed $t \in [0, T]$ and $i \in \mathbb{N}$, we have that τ_i and $\tilde{\tau}_i$ are \mathbb{F} -stopping times. Thus $\{\tau_i \leq t\} \cap \{\tilde{\tau}_i > t\} \in \mathcal{F}_t$.

5.2.5 System dynamics

In general, the lifetime of the impact that liquidity taking orders have on the price of a traded instrument determines if the impact is considered permanent or temporary; see e.g., [Bertsimas and Lo \(1998\)](#), [Obizhaeva and Wang \(2013\)](#), [Cartea et al. \(2015\)](#), and [Guéant \(2016\)](#). Permanent impact refers to the long-lived price effect that an aggressive order (i.e., consumes liquidity) has on the fundamental value of the traded instrument. On the other hand, temporary price impact refers to short-lived deviations from the value of the traded instrument. These deviations are caused by the aggressive orders that consume various levels of the LOB, after which new limit

²⁰The equations we derive in this chapter are yet to be studied when $\varepsilon \rightarrow 0$. We remark that the continuity of f^ε is used in the viscosity solutions proofs we collect in the appendix – see the supersolution property in Theorem 13.

²¹Below, we focus on MLOs of FoK type, so we assume a small value of the parameter ε . The intuition is as follows: (i) when $\varepsilon \approx 0$, the function f^ε is, for practical purposes, either zero or one, which perfectly describes MLOs of FoK type, where one has orders filled in full or rejections only; (ii) when $\varepsilon \gg 0$, then $f^\varepsilon(1 - S_{\tilde{\tau}-} - F_{\tilde{\tau}})$ takes values in $(0, 1)$ with non-negligible probability, allowing one to account for partial fills. The results we develop work for any choice $\varepsilon > 0$.

orders arrive and replenish the LOB back to the price levels before the arrival of the aggressive order.

In our framework, MLOs have permanent price impact when the order is filled. Note that if the MLO is not filled, the trader is the only market participant who knows that the trade attempt was not successful; thus, missed MLOs do not have price impact. If market participants had access to the information on missed trades, they would learn about buy and sell pressure in the market and would adjust their liquidity provision and liquidity taking strategies.

All trade attempts are of size one unit, which we refer to as child orders. The size of the parent order is $\mathfrak{M} > 0$ child orders, which is the initial inventory the trader seeks to liquidate. If the limit rate of the sell MLO to execute a child order is less than or equal to the best bid rate, we assume that there is enough liquidity to fill the MLO without walking the book, in which case the trader is indifferent between sending an IoC or an FoK to the exchange. Thus, we assume that trades do not have temporary price impact; it is straightforward to include this price impact in the price dynamics of our model – for example, we can include it in the flicker. In Table 5.3 we see that the majority of liquidity taking orders do not walk the LOB. Traders specify quantities that can be filled with the volumes displayed at the best bid and best ask rates in the market.²²

For an execution strategy $\alpha = (\tau_i, l_i)_{i \geq 1} \in \mathcal{A}$, the trader monitors the system

$$\mathbf{X}_t^\alpha = (S_t^\alpha, Q_t^\alpha, C_t^\alpha),$$

where $(S_t^\alpha)_{t \geq 0}$ is the fundamental best bid rate process, $(Q_t^\alpha)_{t \geq 0}$ is the inventory of the agent, and $(C_t^\alpha)_{t \geq 0}$ is the cash process. The dynamics of the fundamental best bid rate process are

$$S_t^\alpha = S_0 + \int_0^t b(S_u^\alpha) du + \int_0^t \sigma(S_u^\alpha) dW_u - \kappa \sum_{\tilde{\tau}_i \leq t} f(l_i - S_{\tilde{\tau}_i}^\alpha - F_{\tilde{\tau}_i}), \quad (5.6)$$

and the cash and the inventory process satisfy, respectively,

$$C_t^\alpha = \sum_{\tilde{\tau}_i \leq t} f(l_i - S_{\tilde{\tau}_i}^\alpha - F_{\tilde{\tau}_i}) (S_{\tilde{\tau}_i}^\alpha + F_{\tilde{\tau}_i}) \quad \text{and} \quad Q_t^\alpha = \mathfrak{M} - \sum_{\tilde{\tau}_i \leq t} f(l_i - S_{\tilde{\tau}_i}^\alpha - F_{\tilde{\tau}_i}).$$

In the fundamental bid rate in (5.6), the functions $b, \sigma : \mathbb{R} \rightarrow \mathbb{R}$ are Lipschitz continuous. The function b is the drift of the fundamental bid price and σ is the volatility of the innovations. The last term on the right-hand side represents the permanent price impact that filled MLOs have on the fundamental best bid rate, where $\kappa \geq 0$ is the permanent impact parameter of the bid rate.²³ Recall that f is the \mathcal{C}^∞ approximation of the step function that flags when a trade

²²This assumption is motivated by Table 5.3. There we see that the probability that a liquidity taking order consumes liquidity beyond the best quotes is around 0.1%. Accounting for quantity in the optimisation is still an open problem in our framework.

²³In our framework we may assume that the parameter κ is stochastic or that it is a function of the best bid rate dynamics before an MLO is sent to the exchange. Either choice would increase the complexity of the model without necessarily providing additional insights into the latency-optimal strategies we develop. The fundamental bid rate dynamics can be taken to be driven by random measures; we do not take this approach for tractability but the advantages of having point processes as opposed to diffusive dynamics is yet to be understood in the current framework.

is filled or missed and recall that all MLOs sent by the trader are of size one. We assume that the affected fundamental best bid $(S_t^\alpha)_{t \geq 0}$ rate is bounded by the unaffected fundamental best bid rate $(S_t)_{t \geq 0}$ as follows, for any execution strategy $\alpha \in \mathcal{A}$ we have that $\sup_{0 \leq t \leq T} |S_t^\alpha| \leq \sup_{0 \leq t \leq T} |S_t| + \kappa N_T$.²⁴

The trader intervenes the system at time τ_i (i.e., sends a sell MLO), and, due to latency, the outcome is known at the notification time $\tilde{\tau}_i > \tau_i$ when the system evolves from $\mathbf{X}_{\tilde{\tau}_i}^\alpha$ to $\mathbf{X}_{\tilde{\tau}_i}^\alpha = \Gamma(\mathbf{X}_{\tilde{\tau}_i-}, \mathbf{F}_{\tilde{\tau}_i}, \mathbf{l}_i)$. Here, the function $\Gamma : \mathbb{R}^3 \times \mathbb{R} \times L \rightarrow \mathbb{R}^3$ is the impulse operator

$$\Gamma(s, q, c, f, l) = (s - \kappa f(l - s - f), q - f(l - s - f), c + (s + f)f(l - s - f)), \quad (5.7)$$

which describes how the system (i.e., best bid rate, inventory, cash) changes at notification times. If the MLO sell order is filled: the first argument on the right-hand side of (5.7) shows that the fundamental bid rate decreases by one tick; the second argument shows that the inventory decreases by one unit; and the last argument shows that the amount of cash increases by the fundamental bid rate plus the flicker. If the sell MLO is not filled, the fundamental bid rate is not affected, and the cash and inventory positions do not change.

The operator Γ is a continuous function and it satisfies

$$\sup_{(\mathbf{y}, l) \in \mathbb{R}^4 \times L} \frac{\|\Gamma(\mathbf{y}, l)\|}{1 + \|\mathbf{y}\|} < \infty, \quad (5.8)$$

where $\|\cdot\|$ is the Euclidean norm.

The initial state of the system is $\mathbf{X}_0 = (S_0, \mathfrak{M}, 0)$. Here, S_0 is the initial value of the fundamental bid rate, recall that \mathfrak{M} is the number of child orders (lots of equal size) that the trader wishes to liquidate, and the initial value of the cash account is zero. The controlled system \mathbf{X}_t^α is the solution to the SDE

$$\mathbf{X}_t^\alpha = \mathbf{X}_0 + \int_0^t \mathbf{b}(\mathbf{X}_u^\alpha) du + \int_0^t \boldsymbol{\sigma}(\mathbf{X}_u^\alpha) dW_u + \sum_{\tilde{\tau}_i \leq t} \left(\Gamma(\mathbf{X}_{\tilde{\tau}_i-}^\alpha, \mathbf{F}_{\tilde{\tau}_i}, \mathbf{l}_i) - \mathbf{X}_{\tilde{\tau}_i-}^\alpha \right), \quad (5.9)$$

where \mathbf{b} , and $\boldsymbol{\sigma} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ are given by $\mathbf{b}(x_1, x_2, x_3) = (b(x_1), 0, 0)$ and $\boldsymbol{\sigma}(x_1, x_2, x_3) = (\sigma(x_1), 0, 0)$, respectively, and both functions are Lipschitz continuous. We now fix a finite horizon $T < \infty$ and an execution strategy $\alpha \in \mathcal{A}$. It follows that

$$\mathbb{E} \left[\sup_{0 \leq s \leq T} \|\mathbf{X}_s^\alpha\|^2 \right] < \infty; \quad (5.10)$$

a result we employ below to show that the value function of the trader is well-defined.

²⁴This assumption is non-restrictive. If the affected fundamental best bid rate is assumed to be always positive, then, $S_t^\alpha \leq S_t$ for $t \in [0, T]$ implies that $\sup_{0 \leq t \leq T} |S_t^\alpha| \leq \sup_{0 \leq t \leq T} |S_t| \leq \sup_{0 \leq t \leq T} |S_t| + \kappa N_T$. Alternatively, if b, σ are constant functions, then the assumption is also satisfied. In the numerical experiments below, we set b, σ to be constant functions.

5.2.6 Liquidity taking with stochastic delay

For all $t \in [0, T)$, the agent has either one or no pending order in the exchange; thus, we define two sets of admissible strategies: (i) admissible strategies with one pending order, and (ii) admissible strategies with no pending order.²⁵ If at time $t \in [0, T)$ there is one pending order with price limit $l \in L$, then the set of admissible strategies is

$$\mathcal{A}_{t,l} = \{\alpha = (\tau_i, l_i)_{i \geq 1} \in \mathcal{A} : \tau_1 = t, l_1 = l\},$$

and if there is no pending order, the set of admissible strategies is

$$\mathcal{A}_t = \{\alpha = (\tau_i, l_i)_{i \geq 1} \in \mathcal{A} : \tau_1 \geq t\}.$$

For any $(t, \mathbf{x}, \mathbf{f}) \in [0, T] \times \mathbb{R} \times (-\infty, \mathfrak{M}] \times \mathbb{R} \times \mathbb{R}$ with $\mathbf{x} = (s, q, c)$, and for a pending order with limit price $l \in L$, and $\alpha \in \mathcal{A}_{t,l}$, we denote by $\mathbf{X}^{t,\mathbf{x},\mathbf{f},l,\alpha}$ the solution to (5.9) for $t \leq s \leq T$, with initial data $\mathbf{X}_t = \mathbf{x}$, and $\mathbf{F}_t = \mathbf{f}$. One can drop the dependence on \mathbf{f} in $\mathbf{X}^{t,\mathbf{x},l,\alpha}$ because for any $\mathbf{f}_1, \mathbf{f}_2 \in \mathbb{R}$, we have that $\mathbf{F}_t^{t,\mathbf{f}_1} = \mathbf{F}_t^{t,\mathbf{f}_2}$. That is, write

$$\mathbf{X}_s^{t,\mathbf{x},l,\alpha} = \mathbf{x} + \int_t^s \mathbf{b}(\mathbf{X}_u^{t,\mathbf{x},l,\alpha}) du + \int_t^s \boldsymbol{\sigma}(\mathbf{X}_u^{t,\mathbf{x},l,\alpha}) dW_u + \sum_{t < \tilde{\tau}_i \leq s} \left(\Gamma(\mathbf{X}_{\tilde{\tau}_i^-}^{t,\mathbf{x},l,\alpha}, \mathbf{F}_{\tilde{\tau}_i}^{t,0}, l_i) - \mathbf{X}_{\tilde{\tau}_i^-}^{t,\mathbf{x},l,\alpha} \right).$$

Similarly, when there is no pending order at time t , we denote by $\mathbf{X}^{t,\mathbf{x},\alpha}$ the solution to (5.9) for $t \leq s \leq T$, with initial data $\mathbf{X}_t = \mathbf{x}$, and for any $(t, \mathbf{x}) \in [0, T] \times \mathbb{R} \times (-\infty, \mathfrak{M}] \times \mathbb{R}$.

Fix $(t, \mathbf{x}) \in [0, T] \times \mathbb{R} \times (-\infty, \mathfrak{M}] \times \mathbb{R}$, $l \in L$, $\alpha_1 \in \mathcal{A}_{t,l}$, and $\alpha_0 \in \mathcal{A}_t$. It follows from equation (5.8), Gronwall's lemma, the Burkholder–Davis–Gundy inequality, and equation (5.4), that there exists a constant $C > 0$ such that

$$\mathbb{E} \left[\sup_{t \leq s \leq T} \left\| \mathbf{X}_s^{t,\mathbf{x},l,\alpha_1} \right\|^2 \right] < C \left(1 + \|\mathbf{x}\|^2 \right) \quad \text{and} \quad \mathbb{E} \left[\sup_{t \leq s \leq T} \left\| \mathbf{X}_s^{t,\mathbf{x},\alpha_0} \right\|^2 \right] < C \left(1 + \|\mathbf{x}\|^2 \right).$$

The trader evaluates the performance of the execution strategy α with the function

$$\Pi(\alpha) = g(\mathbf{X}_T^\alpha) + \int_0^T h(\mathbf{X}_s^\alpha) ds + \sum_{\tilde{\tau}_i \leq T} \mathfrak{C}(\mathbf{X}_{\tilde{\tau}_i^-}, \mathbf{F}_{\tilde{\tau}_i}, l_i), \quad (5.11)$$

where $g : \mathbb{R}^3 \rightarrow \mathbb{R}$, $h : \mathbb{R}^3 \rightarrow \mathbb{R}$, $\mathfrak{C} : \mathbb{R}^4 \rightarrow \mathbb{R}$ are

$$g(s, q, c) = c + q(s + \zeta - a(q-1))(1-\mathfrak{c}), \quad h(s, q, c) = -\phi q^2, \quad \mathfrak{C}(s, q, c, \mathbf{f}, l) = -\mathfrak{c}(s+\mathbf{f})f(l-s-\mathbf{f}). \quad (5.12)$$

The function g consists of three terms. These are (i) c , the cash accumulated by the strategy, (ii) $q(s + \zeta)$, the mark-to-market value of the inventory (net of fees), which includes the expected value of the flicker $\zeta = \int_{\mathbb{R}} z \nu(dz)$, and (iii) $a q(q-1)$, the costs of walking the LOB (net of fees) when terminal inventory is greater than one child order, where $a \geq 0$ is the terminal inventory penalty parameter.

²⁵This is a significant simplification of the strategy space. Otherwise, we would need to provide an execution framework that allows a finite number of pending orders.

The function h represents the urgency of the agent to complete the execution programme. Everything else being equal, the higher the value of the running inventory penalty parameter $\phi \geq 0$, the quicker the trader will liquidate inventory at the beginning of the trading horizon. Each trader decides the urgency of their trade or set of trades. For example, the value of the urgency parameter is arbitrarily high for a trader who requires immediate execution. On the other hand, a patient liquidity taker who has time to retry missed trades will set the value of ϕ close to or at zero. We return to this point in subsection 5.2.7, where we discuss stylised facts of the strategy for various values of ϕ , and in Section 5.4 when we discuss the performance of the latency-optimal strategies.

The function \mathfrak{C} represents transaction costs per unit of inventory traded. These costs are based on the notional traded, which is $(s + f) f(l - s - f)$ per one unit of inventory and $\mathfrak{c} \in (0, 1)$ is the transaction cost parameter.

The controlled processes $(C_t^\alpha)_{0 \leq t \leq T}$, $(S_t^\alpha)_{0 \leq t \leq T}$, and $(Q_t^\alpha)_{0 \leq t \leq T}$ obey the bounds, which do not depend on the strategy α or $t \in [0, T]$,

$$|C_t^\alpha| \leq \sum_{i=1}^{N_T} |Z_i| + N_T \sup_{0 \leq u \leq T} S_u + \kappa N_T^2, \quad |S_t^\alpha| \leq \sup_{0 \leq u \leq T} |S_u| + \kappa N_T, \quad |Q_t^\alpha| \leq \mathfrak{M} + N_T. \quad (5.13)$$

Furthermore, we see that

$$\sup_{(\mathbf{x}, \mathbf{f}, \mathbf{l}) \in \mathbb{R}^3 \times \mathbb{R} \times L} \frac{|g(\mathbf{x})| + |h(\mathbf{x})| + |\mathfrak{C}(\mathbf{x}, \mathbf{f}, \mathbf{l})|}{1 + |\mathbf{f}| + \|\mathbf{x}\|^2} < \infty, \quad (5.14)$$

which ensures that the value functions that we present below is well defined, i.e., the value functions are finite as a consequence of the inequality above and of the expressions in (5.10) and (5.13).

When there is one pending order in the exchange, the trader's performance criterion and value function are, respectively,

$$J_1(t, \mathbf{x}, l, \alpha) = \mathbb{E} \left[g(\mathbf{X}_T^{t, \mathbf{x}, l, \alpha}) + \int_t^T h(\mathbf{X}_s^{t, \mathbf{x}, l, \alpha}) ds + \sum_{\tilde{\tau}_i \leq T} \mathfrak{C}(\mathbf{X}_{\tilde{\tau}_i}^{t, \mathbf{x}, l, \alpha}, \mathbf{F}_{\tilde{\tau}_i}^{t, 0}, \mathbf{l}_i) \right],$$

and

$$v_1(t, \mathbf{x}, l) = \sup_{\alpha \in \mathcal{A}_{t, l}} J_1(t, \mathbf{x}, l, \alpha), \quad (5.15)$$

respectively, for $(t, \mathbf{x}) \in [0, T] \times \mathbb{R} \times (-\infty, \mathfrak{M}] \times \mathbb{R}$, $l \in L$, $\alpha \in \mathcal{A}_{t, l}$.

Similarly, when the trader does not have an order pending in the exchange, we have

$$J_0(t, \mathbf{x}, \alpha) = \mathbb{E} \left[g(\mathbf{X}_T^{t, \mathbf{x}, \alpha}) + \int_t^T h(\mathbf{X}_s^{t, \mathbf{x}, \alpha}) ds + \sum_{\tilde{\tau}_i \leq T} \mathfrak{C}(\mathbf{X}_{\tilde{\tau}_i}^{t, \mathbf{x}, \alpha}, \mathbf{F}_{\tilde{\tau}_i}^{t, 0}, \mathbf{l}_i) \right],$$

for $(t, \mathbf{x}) \in [0, T] \times \mathbb{R} \times (-\infty, \mathfrak{M}] \times \mathbb{R}$, $\alpha \in \mathcal{A}_t$, with the corresponding value function

$$v_0(t, \mathbf{x}) = \sup_{\alpha \in \mathcal{A}_t} J_0(t, \mathbf{x}, \alpha). \quad (5.16)$$

In C.1.1 we show that the trader's value functions satisfy the dynamic programming principle,²⁶ in C.1.2 we derive the HJBQVI satisfied by the value functions, and in C.1.3 we study the viscosity properties of the value functions.²⁷

5.2.7 Latency-optimal strategy: price protection and speculation

Below, in Section 5.4, we discuss in detail the performance of the random-latency-optimal strategy against various benchmarks. We implement the random-latency-optimal strategy using a finite-difference scheme.²⁸

The success of the latency-optimal strategy relies on the use of MLOs to limit exposure to adverse price moves and the use of speculative MLOs, which are orders that target better rates than those observed by the trader in the LOB. Speculative MLOs are the key driver in the superior performance of the strategy and a prominent feature of strategies that employ MLOs when the trader can retry rejected orders over a trading window.

Figure 5.5 shows the limit rates of MLOs for various levels of inventory when: the target is to exchange $\mathfrak{M} = 10$ lots of currency pairs over a trading horizon of $T = 6$ seconds; expected latency is 30 ms; the fundamental best bid rate follows an arithmetic Brownian motion with volatility $\sigma = 2 \times 10^{-4}$, and throughout the trading window the rate is fixed at $S_t = 1.1$; the permanent rate impact parameter is $\kappa = 0$; the parameters of the distribution of the flickers that affect the MLO are: $p^+ = 0.05$, $p^- = 0.08$, $1/\eta_+ = 1.87 \times 10^{-5}$, $1/\eta_- = 2.66 \times 10^{-5}$; the terminal liquidation penalty parameter is $a = 1.80 \times 10^{-5}$, and the urgency parameter is $\phi \in \{0, 1 \times 10^{-5}, 1 \times 10^{-4}\}$. In the left, centre, and right panels, the value of the running inventory penalty parameter is $\phi = 0$, $\phi = 1 \times 10^{-5}$, and $\phi = 1 \times 10^{-4}$, respectively. Recall that ϕ is interpreted as an urgency parameter: as the value of ϕ increases, the trader becomes more impatient, so the rate of execution increases at the beginning of the trading window.

In the left panel, when there is no urgency to liquidate early, all sell MLOs are speculative during approximately the first three quarters of the trading window. Only near expiry, does the strategy send MLOs that may walk the LOB to increase the probability of a fill and to complete the liquidation programme on time. On the other hand, when there is urgency to liquidate early, the right panel shows that the limit rates of the sell MLOs are below the fundamental best bid rate when inventory is greater than one lot, so the orders may walk the LOB (if necessary) to increase the probability of a fill.

In the three panels, when the trader is near the end of the trading horizon, the strategy sends sell MLOs that are willing to receive rates below the fundamental best bid rate to increase the probability of a fill. The limit sell rates are not zero, in which case the MLO would be equivalent

²⁶The part of the proof that deals with measurable selection arguments is omitted. See Bouchard and Touzi (2011) for the derivation of the dynamic programming equation in the sense of viscosity solutions.

²⁷Future work includes the proof of higher regularity of the value functions.

²⁸We remark that if one were to extend the current framework to accommodate more than one pending order, then we would need to follow an algorithm like than in Section 5 of Bruder and Pham (2009). The convergence of these algorithms in our current framework remains to be studied. For penalty schemes of viscosity solutions to HJBQVIs (with enough regularity, e.g., Lipschitz continuity or semiconvexity) see Reisinger and Zhang (2020).

to a sell MO, because at time T the strategy liquidates terminal inventory with an MO that pays the expected value of the flicker $q\zeta$ and pays the penalty $qa(q-1)$, which provides a lower bound for the limit rate of the sell MLO; see the discussion below (5.12).




Figure 5.5: Expected latency is 30 ms. LR denotes the limit rate of the MLO. The black dash line is at $S_t = 1.1$ and the terminal inventory penalty parameter $a = 1.80 \times 10^{-5}$. Left panel $\phi = 0$, centre panel $\phi = 1 \times 10^{-5}$, right panel $\phi = 1 \times 10^{-4}$.

5.3 Benchmarks

In this section, we develop two strategies to benchmark the performance of the random-latency-optimal trading strategy derived above. Throughout, most of the notation is the same as that employed in the previous sections. In Sections 5.3.1 and 5.3.2, we derive optimal execution strategies when latency is zero in the marketplace, and when latency is greater than zero and deterministic, respectively. In both cases, the trader solves an impulse control problem: in the first model the strategy determines the optimal times to send MOs to the exchange; and in the second model the strategy determines the timing and the limit rate of each MLO. We remark that most approaches in the execution literature assume that agents send MOs at a continuous rate. Although agents cannot continuously trade in the market, this assumption is convenient because in some cases one finds strategies in closed-form, see Cartea et al. (2015) and Guéant (2016). One notable exception is the paper Cartea and Jaimungal (2015), where the authors solve an impulse control problem in which the trader employs both MOs and limit orders to execute a large position in an financial instrument when there is zero latency in the marketplace.

5.3.1 Optimal execution with zero latency

Assume there is no latency in the marketplace.²⁹ The trader controls the times at which she sends sell MOs to the exchange. The size of all MOs is one and the filled trades have permanent price impact as in the model of Section 5.2.

Let $\tau = (\tau_i)_{i \geq 1}$ denote an increasing sequence of stopping times when the trader sends sell MOs to the exchange. Observe that if latency is zero, a sell MO of size one at time τ_i is the

²⁹This is the case that most of the literature in algorithmic trading has studied. It is well-known that if we employ continuous trading, many of the optimal strategies are of the form TWAP plus adjustments; see Cartea et al. (2015). Here, we use impulses with zero latency to have a framework that is similar those presented above. In any case, we use TWAP as one of the benchmarks in the study that follows.

same as a sell MLO for a unit size with a limit price $S_{\tau_i}^{\tau-}$ because the trader observes and acts with no delay in the marketplace and the exchange processes the order with no delay. The dynamics of the best bid price process, denoted by $S^{\tau} = (S_t^{\tau})_{t \geq 0}$, follow

$$S_t^{\tau} = S_0 + \int_0^t b(S_u^{\tau}) du + \int_0^t \sigma(S_u^{\tau}) dW_u - \kappa \sum_{\tau_i \leq t} 1,$$

where the functions $b, \sigma : \mathbb{R} \rightarrow \mathbb{R}$ are Lipschitz continuous and κ is the permanent bid rate impact parameter. The cash process $C^{\tau} = (C_t^{\tau})_{t \geq 0}$ and the inventory process $Q^{\tau} = (Q_t^{\tau})_{t \geq 0}$ are given by

$$C_t^{\tau} = \sum_{\tau_i \leq t} S_{(\tau_i)-}^{\tau} \quad \text{and} \quad Q_t^{\tau} = \mathfrak{M} - \sum_{\tau_i \leq t} 1.$$

For each execution strategy τ over the trading window $[0, T]$, the trader computes

$$\Pi(\tau) = g(\mathbf{X}_T^{\tau}) + \int_0^T h(\mathbf{X}_s^{\tau}) ds + \sum_{\tau_i \leq T} \mathfrak{C}(\mathbf{X}_{\tau_i}^{\tau-}),$$

where $g, h : \mathbb{R}^3 \rightarrow \mathbb{R}$ are as in (5.12), and the function $\mathfrak{C}(s, q, c) = -c s$ represents a transaction cost every time a trade of unit size is executed.

The trader solves the control problem

$$\sup_{\tau \in \mathcal{A}} \mathbb{E}[\Pi(\tau)], \quad \mathcal{A}_t = \{\tau = (\tau_i)_{i \geq 1} \in \mathcal{A} : \tau_1 \geq t\}.$$

For any $(t, \mathbf{x}) \in [0, T] \times \mathbb{R}^+ \times (-\infty, \mathfrak{M}] \times \mathbb{R}^+$, and $\tau \in \mathcal{A}_t$, the system $\mathbf{X}^{t, \mathbf{x}, \tau}$ is given by

$$\mathbf{X}_s^{t, \mathbf{x}, \tau} = \mathbf{x} + \int_t^s \mathbf{b}(\mathbf{X}_u^{t, \mathbf{x}, \tau}) du + \int_t^s \boldsymbol{\sigma}(\mathbf{X}_u^{t, \mathbf{x}, \tau}) dW_u + \sum_{t < \tau_i \leq s} \left(\Gamma(\mathbf{X}_{(\tau_i)-}^{t, \mathbf{x}, \tau}) - \mathbf{X}_{(\tau_i)-}^{t, \mathbf{x}, \tau} \right), \quad (5.17)$$

where $\Gamma(s, q, c) = (s - \kappa, q - 1, c + s)$.

The trader's performance criterion and value function are

$$J(t, \mathbf{x}, \tau) = \mathbb{E} \left[g(\mathbf{X}_T^{t, \mathbf{x}, \tau}) + \int_t^T h(\mathbf{X}_s^{t, \mathbf{x}, \tau}) ds + \sum_{\tau_i \leq T} \mathfrak{C}(\mathbf{X}_{\tau_i}^{t, \mathbf{x}, \tau-}) \right] \quad \text{and} \quad v(t, \mathbf{x}) = \sup_{\tau \in \mathcal{A}_t} J(t, \mathbf{x}, \tau).$$

5.3.2 Optimal execution with deterministic latency

The trader employs MLOs and operates in the market with a fixed and known latency. The execution strategy has at most one pending order at any one time – we can easily extend this framework so the agent can have up to $m \in \mathbb{N}$ pending orders. Our approach is based on the work of Bruder and Pham (2009), who develop a general framework for impulse control with deterministic delay.

Let $\Delta > 0$ be the fixed and known delay of the trader in the marketplace and let $(Z_n)_{n \in \mathbb{N}}$ be a collection of i.i.d. random variables with law as in (5.3), such that for $i \in \mathbb{N}$, Z_i is \mathcal{F}_t -measurable for $t \geq i\Delta$. As in Section 5.2, we use the sequence of marks $(Z_n)_{n \in \mathbb{N}}$ to model the flickers that affect the trader's MLO when the exchange processes the order.

The agent's execution strategy is $\beta = (\tau_i, \mathfrak{l}_i)_{i \geq 1}$, where (τ_i) is an increasing sequence of stopping times when the agent sends MLOs with limit price \mathfrak{l}_i .

We require that $\tau_{i+1} - \tau_i \geq \Delta$ a.s., for all values of i , because the trader can have at most one pending order in the exchange. The set of admissible strategies is

$$\mathcal{A} = \{\beta = (\tau_i, \mathfrak{l}_i)_{i \geq 1} : \forall i \geq 1 \quad \tau_{i+1} - \tau_i \geq \Delta \text{ a.s. and } \mathfrak{l}_i \text{ is } \mathcal{F}_{\tau_i}\text{-measurable}\} . \quad (5.18)$$

Interventions in the system $\mathbf{X}_t^\beta = (S_t^\beta, Q_t^\beta, C_t^\beta)$ are at time τ_i , but processed by the exchange and notified to the trader at $\tau_i + \Delta$, so the system evolves from $\mathbf{X}_{(\tau_i + \Delta)^-}^\beta$ to $\mathbf{X}_{(\tau_i + \Delta)}^\beta = \Gamma(\mathbf{X}_{(\tau_i + \Delta)^-}^\beta, Z_{\lfloor (\tau_i + \Delta)/\Delta \rfloor}, \mathfrak{l}_i)$, where Γ is as in (5.7), and $\lfloor \cdot \rfloor$ is the floor function. The value of $j = \lfloor (t + \Delta)/\Delta \rfloor$ is the index of the random variable Z_j , which is measurable at time $j\Delta \in (t, t + \Delta]$.

Let $\beta = (\tau_i, \mathfrak{l}_i)_{i \geq 1} \in \mathcal{A}$. The fundamental best bid price process is

$$S_t^\beta = S_0 + \int_0^t b(S_u^\beta) du + \int_0^t \sigma(S_u^\beta) dW_u - \kappa \sum_{\tau_i + \Delta \leq t} f\left(\mathfrak{l}_i - S_{(\tau_i + \Delta)^-}^\beta - Z_{\lfloor (\tau_i + \Delta)/\Delta \rfloor}\right) ,$$

where f is as in (5.5). The cash process satisfies

$$C_t^\beta = \sum_{\tau_i + \Delta \leq t} f\left(\mathfrak{l}_i - S_{(\tau_i + \Delta)^-}^\beta - Z_{\lfloor (\tau_i + \Delta)/\Delta \rfloor}\right) \left(S_{(\tau_i + \Delta)^-}^\beta + Z_{\lfloor (\tau_i + \Delta)/\Delta \rfloor}\right) ,$$

and the inventory is

$$Q_t^\beta = \mathfrak{M} - \sum_{\tau_i + \Delta \leq t} f\left(\mathfrak{l}_i - S_{(\tau_i + \Delta)^-}^\beta - Z_{\lfloor (\tau_i + \Delta)/\Delta \rfloor}\right) .$$

For $\mathbf{X}_0 = (S_0, \mathfrak{M}, 0)$, the controlled system \mathbf{X}_t^β is the solution to the SDE

$$\mathbf{X}_t^\beta = \mathbf{X}_0 + \int_0^t \mathbf{b}(\mathbf{X}_u^\beta) du + \int_0^t \boldsymbol{\sigma}(\mathbf{X}_u^\beta) dW_u + \sum_{\tau_i + \Delta \leq t} \left(\Gamma(\mathbf{X}_{(\tau_i + \Delta)^-}^\beta, Z_{\lfloor (\tau_i + \Delta)/\Delta \rfloor}, \mathfrak{l}_i) - \mathbf{X}_{(\tau_i + \Delta)^-}^\beta \right) , \quad (5.19)$$

where $\mathbf{b}, \boldsymbol{\sigma} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ are given by $\mathbf{b}(x_1, x_2, x_3) = (b(x_1), 0, 0)$ and $\boldsymbol{\sigma}(x_1, x_2, x_3) = (\sigma(x_1), 0, 0)$, and b, σ are Lipschitz continuous. Finally, $\Gamma : \mathbb{R}^4 \times L \rightarrow \mathbb{R}^3$ is as in (5.7).

Now, fix a finite horizon $T < \infty$ and assume that $T - \Delta \geq 0$ to avoid trivialities. One may check that

$$\mathbb{E} \left[\sup_{s \leq T} \left\| \mathbf{X}_s^\beta \right\|^2 \right] < \infty ,$$

and as above, the agent maximises

$$\Pi(\beta) = g(\mathbf{X}_T^\beta) + \int_0^T h(\mathbf{X}_s^\beta) ds + \sum_{\tau_i + \Delta \leq T} \mathfrak{C}\left(\mathbf{X}_{(\tau_i + \Delta)^-}^\beta, Z_{\lfloor (\tau_i + \Delta)/\Delta \rfloor}, \mathfrak{l}_i\right) ,$$

where $g, h : \mathbb{R}^3 \rightarrow \mathbb{R}$ and $\mathfrak{C} : \mathbb{R}^4 \times L \rightarrow \mathbb{R}$ is as in (5.12). Thus, the agent solves the control problem

$$\sup_{\beta \in \mathcal{A}} \mathbb{E} [\Pi(\beta)] .$$

The system in (5.19) is non-Markovian.³⁰ Given an impulse, the future of \mathbf{X}^β is influenced by a possible pending order attempted between $t - \Delta$ and t . For any $t \in [0, T]$, $k = 0, 1$, we let $P_t(0) = \emptyset$ and

$$P_t(1) = \{p = (t_1, l_1) \in [0, T - \Delta] \times L : t - \Delta < t_1 \leq t\}$$

denote whether there is a pending order in the exchange.

For any $p = (t_1, l_1) \in P_t(1)$, $t \in [0, T]$, the set of admissible strategies at time t with pending order p is

$$\mathcal{A}_{t,p} = \{\beta = (\tau_i, l_i)_{i \geq 1} \in \mathcal{A} : (\tau_1, l_1) = (t_1, l_1) \text{ and } \tau_2 \geq t_1 + \Delta\}.$$

Similarly, if there is no pending order in the exchange,

$$\mathcal{A}_{t,\emptyset} = \mathcal{A}_t = \{\beta = (\tau_i, l_i)_{i \geq 1} \in \mathcal{A} : \tau_1 \geq t\}.$$

For any $(t, \mathbf{x}) \in [0, T] \times \mathbb{R}^3$, $p \in P_t(k)$ for $k = 0, 1$, and $\beta \in \mathcal{A}_{t,p}$, we denote by $\mathbf{X}^{t,\mathbf{x},p,\beta}$ the solution to

$$\begin{aligned} \mathbf{X}_s^{t,\mathbf{x},p,\beta} = \mathbf{x} &+ \int_t^s \mathbf{b}(\mathbf{X}_u^{t,\mathbf{x},p,\beta}) \, du + \int_t^s \boldsymbol{\sigma}(\mathbf{X}_u^{t,\mathbf{x},p,\beta}) \, dW_u \\ &+ \sum_{t < \tau_i + \Delta \leq s} \left(\Gamma(\mathbf{X}_{(\tau_i + \Delta)^-}^{t,\mathbf{x},p,\beta}, Z_{\lfloor (\tau_i + \Delta)/\Delta \rfloor}, l_i) - \mathbf{X}_{(\tau_i + \Delta)^-}^{t,\mathbf{x},p,\beta} \right). \end{aligned} \quad (5.20)$$

Next, consider the performance criterion

$$J_k(t, \mathbf{x}, p, \beta) = \mathbb{E} \left[g(\mathbf{X}_T^{t,\mathbf{x},p,\beta}) + \int_t^T h(\mathbf{X}_s^{t,\mathbf{x},p,\beta}) \, ds + \sum_{\tau_i + \Delta \leq T} \mathfrak{c} \left(\mathbf{X}_{(\tau_i + \Delta)^-}^{t,\mathbf{x},p,\beta}, Z_{\lfloor (\tau_i + \Delta)/\Delta \rfloor}, l_i \right) \right],$$

for $(t, \mathbf{x}) \in [0, T] \times \mathbb{R}^3$, $p \in P_t(k)$, $k = 0, 1$, $\beta \in \mathcal{A}_{t,p}$, and the corresponding value functions

$$v_k(t, \mathbf{x}, p) = \sup_{\beta \in \mathcal{A}_{t,p}} J_k(t, \mathbf{x}, p, \beta), \quad k = 0, 1 \quad \text{and } (t, \mathbf{x}, p) \in \mathcal{D}_k,$$

where

$$\mathcal{D}_k = \{(t, \mathbf{x}, p) : (t, \mathbf{x}) \in [0, T] \times \mathbb{R}^3, p \in P_t(k)\}.$$

Recall that for $k = 0$, $P_t(0) = \emptyset$, therefore we write $v_0(t, \mathbf{x}) = v_0(t, \mathbf{x}, \emptyset)$, $\mathcal{D}_0 = [0, T] \times \mathbb{R}^3$.

Finally, the proof of the dynamic programming principle and the PDE system viscosity characterisation of the value functions are in C.2.

5.4 Performance of execution strategy

We employ ultra-high frequency data for the currency pair EUR/USD to compare the performance of the random-latency-optimal strategy (RLOS) developed in Section 5.2 with four benchmarks (of which the first two were derived above): (i) deterministic-latency optimal strategy (DLOS), (ii) zero latency-optimal strategy (ZLOS), (iii) TWAP, and (iv) ENOW.³¹ TWAP

³⁰In the analysis that follows we increase the state space to produce a Markovian system.

³¹The strategies RLOS, DLOS, and ZLOS are computed using finite differences. Convergence results for these algorithms remain to be studied.

sends MOs of equal size and at equally spaced time intervals over the trading window. ENOW is a hypothetical benchmark that assumes there is enough liquidity at the best bid rate in the LOB to execute all the inventory with one MO, with zero latency, at the beginning of the trading window.

Recall that the fundamental bid rate S_t is an input to RLOS, DLOS, ZLOS. However, market participants observe the best bid rate posted on the LOB of the exchange; a rate that consists of the fundamental bid rate S_t and short-lived deviations. Here, we employ LOB observations of the best bid rates and of the quoted spreads in the LOB of the exchange to estimate the fundamental bid rate S_t . Above, the right-hand panel of Figure 5.1 depicted the evolution of the best rates in the LOB of the EUR/USD currency pair, where we saw that most flickers cause short-lived widening of the quoted spread: most flickers in the best bid rate are positive and most flickers in the best ask rate are negative. Thus, here we assume that at time t , the fundamental best bid rate S_t is the bid rate in the LOB for $u \leq t$ when the spread in the LOB was less than v ticks, i.e., $u = \operatorname{argmax} \{s \leq t : D_t < v\}$, where $(D_t)_{t \geq 0}$ is the quoted spread. Throughout, we assume that $v = 11$ ticks, and in subsection 5.4.6 we show that our results are robust to various choices of the value of v .

The trader's execution horizon is $T = 6$ seconds and the inventory to liquidate is $\mathfrak{M} = 10$ lots, where each lot is €500,000, so the objective is to exchange €5,000,000 into USD.³² The execution value of ENOW in a frictionless market is the value of \mathfrak{M} lots exchanged at the best bid rate in the LOB at the start of the trading horizon. For TWAP and DLOS, the MO for each lot is processed with the liquidity posted at the best bid rate in the LOB of the exchange, and if necessary, the order will walk down until filled in full. Similarly, if the limit rate of the MLO allows it, RLOS and DLOS will also walk down the LOB if the liquidity in the best quotes is not enough – we do so, despite the strategies derived above assumed that orders of size one do not walk the LOB.

To compute RLOS, DLOS, ZLOS we assume that the fundamental best bid rate satisfies

$$S_t = S_0 + \sigma W_t, \quad (5.21)$$

where σ is the volatility parameter. However, we remark that when we implement the liquidation strategies, we do not simulate the fundamental best bid rate, and do not simulate the flickers that affect the MLO when it is processed by the exchange. Instead, we use the best bid rate in the LOB of LMAX.

We employ data from 1 August 2019 to 31 August 2019 to estimate the parameters of the model, and data between 1 September 2019 and 29 February 2020 to implement and to compute the performance of RLOS and the benchmark strategies.

³²Our results are robust to the quantity of each lot as long as it is between €1,000 and €500,000 because there is enough liquidity in the best quotes to fill such an order; see Table 5.2. The average daily traded quantity in the EUR/USD currency pair in LMAX London is €1.5 billion. Thus, €5,000,000 is about 0.3% of the average daily traded volume. It is plausible to observe such a quantity being traded over a six second window. Below we show that trading activity varies according to the time-of-day; see Figure 5.6.

Each day, we implement the liquidation strategy over three ten-minute windows. We use the data in the ten minutes starting at 9.00am, at 2.00pm, and at 7.00pm for each trading day between 1 September 2019 and 29 February 2020.³³ At 9.00am and 2.00pm the foreign exchange activity tends to be very high because by 9.00am many financial markets in London are open and very active, and many financial markets in the US open at around 2.00pm. Later in the afternoon, foreign exchange activity decreases, and by 7.00pm many markets world-wide are outside the hours of their main trading activity.

Figure 5.6 shows foreign exchange market activity in the EUR/USD currency pair from 1 September 2019 to 29 February 2020; the data are in 15-minute bins. The left-hand panel shows the number of updates in the best quotes (changes in rates or in volumes), and the right-hand panel shows the number of fully and partially filled liquidity taking trades. Clearly, periods of very high activity in the EUR/USD currency pair are around 9.00am and 2.00pm, while there is little activity at 7.00pm.




Figure 5.6: Left panel: number of updates in the best quotes (changes in rates or in volumes). Right panel: number of filled and partially filled trades. EUR/USD currency pair, 1 September 2019 to 29 February 2020.

5.4.1 Model parameters

We employ data from August 2019 to estimate the parameters of the model as follows:

(i) Volatility of fundamental best bid rate dynamics. We compute the quadratic variation of the fundamental best bid rate to estimate the volatility parameter σ in (5.21). The estimates over the three ten-minute windows are: (9.00am to 9.10am) $\hat{\sigma}_{9\text{am}} = 2.1 \times 10^{-4}$; (2.00pm to 2.10pm) $\hat{\sigma}_{2\text{pm}} = 2.9 \times 10^{-4}$; (7.00pm to 7.10pm) $\hat{\sigma}_{7\text{pm}} = 1.7 \times 10^{-4}$. And the estimate of volatility from 9.00am to 8.00pm is $\hat{\sigma} = 1.9 \times 10^{-4}$; we use this parameter estimate in subsection 5.4.4 when trading around news events.

(ii) Distribution of the size of the flicker that affects the MLOs at processing times. Recall that the law of the flickers is

$$\nu(dz) = p_0 \delta_{\{0\}}(dz) + \underbrace{p_+ \eta_+ e^{-\eta_+ z} \mathbf{1}_{\{z>0\}} dz}_{\text{price improvement}} + \underbrace{p_- \eta_- e^{-\eta_- z} \mathbf{1}_{\{z<0\}} dz}_{\text{slippage}} .$$

³³All times in this chapter are in BST.

Agents use their trade activity to estimate their expected latency in the marketplace; here, we assume that their estimate is one of the following: 10, 30, 60, 90 ms. Next, for a given expected latency, and for a value of v to obtain the fundamental best bid rate, agents estimate the parameters of the distribution of flickers as follows: employ LOB data to study the hit and miss of MLOs with limit rates equal to the fundamental best bid rate when latency is exponentially distributed; use ten minutes of LOB data (beginning at 9.00am, 2.00pm, and 7.00pm) of each trading day in August to simulate over one million MLOs; compute the SPIs of the MLOs; and finally, obtain maximum likelihood estimates of the parameters $p_0, p_+, p_-, \eta_+, \eta_-$ for each ten-minute trading horizon and for each expected latency.

(iii) Terminal penalty. We employ the liquidity taking trades that walk the LOB to compute the terminal penalty parameter a . For each order, we look at the excess traded quantity above the liquidity posted at the best quotes, and perform a linear regression where the explanatory variable is the excess traded quantity, bought or sold, and the dependent variable is the excess rate paid or received in the quote currency for walking the book. The slope of the regression for data in August 2019 is $\hat{a} = 1.80 \times 10^{-5}$ (which we use to compute the optimal strategies), and from September 2019 to February 2020 is $\hat{a} = 1.61 \times 10^{-5}$ (which we use to compute execution costs when orders walk the book).

We assume that the value of the permanent price impact parameter κ is zero because we use LOB data to run the simulations, so we cannot include the permanent impact that the trader's orders would have had on the market.³⁴ The central scenarios we study assume that the liquidity trader is patient, so we set the value of the urgency parameter ϕ to zero. Finally, the transaction cost parameter is $\mathfrak{c} = 3 \times 10^{-6}$; i.e., 3 USD per million USD traded, which are the fees paid by active traders in the foreign exchange market. Table 5.4 summarises the model parameters for RLOS, DLOS, ZLOS to trade between 9.00am and 9.10am when expected latency is 30 ms. In the interest of space, we do not report the parameters for the other scenarios we study.

$\hat{\sigma}_{9\text{am}}$	\hat{p}_+	\hat{p}_-	$\hat{\eta}_+, \hat{\eta}_-$	\hat{a}	Expected latency	ϕ	κ	\mathfrak{c}
2.1×10^{-4}	0.05	0.08	$\{1.87, 2.66\} \times 10^{-5}$	1.80×10^{-5}	30 ms	0	0	3×10^{-6}

Table 5.4: Model parameters for trading window 9.00am to 9.10am. Fundamental best bid rate is obtained with spread filter of $v = 11$ ticks.

5.4.2 Simulations and performance measure

There are 128 days of market data in the period from 1 September 2019 to 29 February 2020. For each day, we split the ten-minute window into 100 intervals of six seconds. For each interval of six seconds, we combine the one market run of the best bid rate posted on the LOB with

³⁴Recall that in the implementations the MLOs walk the book if necessary (while obeying the price limit), and the strategies incorporate the fees paid to the exchange. Thus, we account for all trading frictions except the friction arising from the permanent price impact.

100 simulated runs of the random latency that the trader would face in each trade attempt. Thus, the total number of runs for each ten-minute window (9.00am to 9.10am, 2.00pm to 2.10pm, 7.00pm to 7.10pm) is $128 \times 100 \times 100 = 1,280,000$. In each run, the strategy exchanges 5,000,000 EUR into USD, which is 10 times the volume that one normally observes posted at the best bid rate in the exchange; i.e., $\mathfrak{M} = 10$ lots of 500,000 EUR each.

For RLOS, DLOS, ZLOS we recall that the cost of liquidating terminal inventory with one MO at the terminal time T is as follows: if $Q_{T-} = 1$ lot, the order does not walk the LOB. However, if $Q_T \geq 2$, the first lot does not walk the book, and the remaining $Q_T - 1$ lots walk the LOB, so that the USD received from exchanging the final position in EUR is $Q_{T-} (S_T + \zeta - a(Q_{T-} - 1)) (1 - \mathfrak{c})$ – see (5.12).

The measure of performance is

$$\frac{\mathbb{E}[TC_T(\text{RLOS})] - \mathbb{E}[TC_T(\text{B})]}{V_0} \times 10^6, \quad (5.22)$$

where TC_T is the terminal cash in USD obtained by the strategy, V_0 is the value of the position to exchange at the beginning of the trading window, and B represents one of the four benchmark strategies. The numerator of (5.22) is in USD and the denominator is in EUR, so the units of the performance measure are $\$/\text{€M}$, i.e., USD per million of EUR exchanged. Finally, when the benchmark is ENOW, the performance measure is ‘implementation shortfall’, which is a widely used benchmark to assess implicit costs and the liquidity of markets, but it is very seldom achieved because there is not enough liquidity at the best quotes of the LOB to fill a large order; see Almgren (2010) and Cartea et al. (2015).³⁵

5.4.3 Results: performance of RLOS and benchmarks

In this subsection, we present the results of the execution strategy for the three ten-minute trading windows that start at 9.00am, 2.00pm, and 7.00pm for various levels of the trader’s expected latency in the marketplace. Later, in subsection 5.4.4, we examine the performance of the strategies around news announcements. The main finding we report is the superior performance of the latency-optimal strategies RLOS and DLOS. The cash value of the outperformance is approximately the same as the fees paid by traders in foreign exchange markets; the value of the outperformance considerably increases during news events. The outperformance stems from the speculative MLOs sent by RLOS and DLOS, which are more valuable when the market is more active. There are two key insights. First, the performance of MLOs depends on the distribution of the flickers. When activity in the markets increase, the value of MLOs increases because both the probability of receiving a price improvement and the size of the improvements increase – the value of MLOs is highest during news events. On the other hand, an increase in the size and the probability of slippage during times of heightened activity has little effect on the performance of latency-optimal strategies because the rate limits in the sell MLOs prevent

³⁵Generally, one computes implementation shortfall with the midprice. Here we use the best bid rate for a liquidation problem.

orders from walking down the LOB. This is in contrast with strategies that send MOs (e.g., TWAP and ZLOS) because these orders are exposed to adverse changes in the rates.

Second, traders use speed to complete the execution programme with as many MLOs as possible because they can retry rejected speculative MLOs before reaching the end of the trading horizon. Thus, the lower the latency of the trader, the more opportunities the trader will have to fill speculative MLOs, so the better the performance of latency-optimal strategies that use MLOs over strategies that employ MOs.

In Chapter 3 we estimate the range of the stochastic latency of an active trader in the foreign exchange market to be between 10 ms and 50 ms. Table 5.5 reports the results of the performance measure (5.22) over the window 9.00am to 9.10am for a patient trader with 10, 30, 60, 90 ms of expected latency. RLOS receives approximately between 5.14 and 2.65 USD more than TWAP for every million EUR exchanged into USD, and RLOS receives between 4.56 and 2.34 USD more than ZLOS per million EUR exchanged, and RLOS outperforms ENOW by between 1.42 and 3.91 USD per million EUR exchanged. Finally, the average USD amount that RLOS receives is statistically the same as that received by DLOS, so the table does not report the results for DLOS. Here, two means are statistically the same if the Student's t -test does not reject the hypothesis that the two means are the same with 95% confidence. ³⁶

	Performance 9.00am-9.10am			
	Expected latency in ms			
	10	30	60	90
ZLOS	4.56	3.40	2.74	2.34
TWAP	5.14	3.80	3.09	2.65
ENOW	3.91	2.57	1.86	1.42

Table 5.5: Values of performance measure (5.22) from 9.00am to 9.10am, $T = 6$ seconds, $\phi = 0$. The revenue from RLOS and the revenue of DLOS (not shown) are statistically the same.

The results in Table 5.5 show that the value of the outperformance of RLOS (and DLOS) decreases as the value of the expected latency increases – we provide two stylised facts of the strategies to explain the results in the table. First, as expected latency increases, it is more likely that at the end of the trading horizon all strategies will fall short of liquidating the target inventory. In the extreme case where expected latency is arbitrarily large, it is very likely that by the terminal time T most of the target inventory will remain in Euros; i.e., not exchanged into USD. Thus, at time T , the remaining lots of currency pairs will be liquidated with one MO, where the trader pays the costs of walking the book; the costs are the same for all strategies (except for ENOW).

³⁶All the outperformances we report below are statistically different from zero, for example, in Table 5.5, the outperformance of RLOS over TWAP when the expected latency is 30ms is \$3.80/€M. Here, we reject the null hypothesis that the two means are the same with a p -value of 2.15×10^{-5} .

Second, for a fixed trading window and for a fixed liquidation target, faster traders (i.e., lower expected latency) have, on average, more opportunities than slow traders to send MLOs to the exchange (recall that in our model the trader cannot send another trade if her previous order is pending in the exchange). Therefore, compared with slow traders, faster traders send more (and fill more) MLOs that seek a rate improvement, which explains the results we report in the table – we return to this point below. In an extreme case, a patient trader with nearly zero latency will use her superior speed advantage to send a large amount of MLOs that seek price improvements, and due to the flickers in the best bid rate in the LOB, the trader will exchange a large proportion of the initial inventory in Euros into USD with these speculative MLOs. A very fast and very patient trader will not employ her speed in the marketplace to liquidate the target inventory quickly and ahead of time, the trader will use her low latency to fill as many speculative trades as possible to complete the execution target.

Our results show that the slower is the trader, the less she benefits from latency-optimal strategies. However, on the other hand, an implication of our findings is that slow traders may profit from investing in co-location, software, and other services to increase their speed in the marketplace if the improvement in the performance of the execution strategies outweighs the costs to reduce their latency, which largely depends on the volumes traded. Below, in subsection 5.4.5, we discuss the economic significance of our results and provide summary statistics of traded volumes in the spot foreign exchange market to assess the potential gains that liquidity takers could receive if they were to employ latency-optimal strategies.

To gain further insights, we report summary statistics of the behaviour of the strategies from 9.00am to 9.10am when expected latency is 30 ms. Table 5.6 reports: the average number of trade attempts (there are 200 periods of 30 ms in an execution horizon of six seconds); the average number of filled orders; the proportion of runs for which the strategy sends an MO at time T to liquidate terminal inventory and pays extra costs from walking the LOB (i.e., average number of runs for which $Q_{T-} \geq 2$); the proportion of the number of attempts for which the limit rate (LR) of the MLO is lower than the fundamental best bid rate S_t (the trader is willing to walk the LOB to complete the trade); the average number of attempts for which $LR = S_t$, and the average number of attempts where $LR > S_t$, which is the proportion of speculative trades that target a price improvement.

Expected latency is 30 ms, start time 9.00am						
	MLO Attempts	MLO Fills	$Q_{T-} \geq 2$	$LR < S_t$	$LR = S_t$	$LR > S_t$
RLOS	195.07	9.95	0.00	1.00	5.71	188.35
DLOS	197.44	9.64	0.07	1.26	5.01	191.17
ZLOS	10.00	10.00	0.00	10.00	-	-
TWAP	10.00	10.00	0.00	10.00	-	-
ENOW	1.00	1.00	0.00	-	1.00	-

Table 5.6: Expected latency is 30 ms, $T = 6$ seconds, so expected number of attempts is 200, and $\phi = 0$. LR: limit rate of MLO. S_t : fundamental best bid rate when the trader decides to send the MLO.

On average, RLOS and DLOS employ a similar number of MLO attempts, of which RLOS fills 9.95 and DLOS fills 9.65 (these figures exclude the MOs sent, if necessary, at time T). Approximately 97% of the attempted MLOs by both strategies are speculative, of which 2.4% (2.3%) sent by RLOS (DLOS) were filled. As discussed above, filled speculative MLOs are the main source of the superior performance of latency-optimal strategies. Within the trading window, the limit rates of the speculative MLOs become less ambitious if the inventory is not on target and the terminal date approaches; see also subsection 5.2.7. However, the latency-optimal strategies cannot guarantee that at time T the inventory in EUR is drawn to $Q_{T-} < 2$. The table shows that in approximately 0.3% (6.9%) of the runs, RLOS (DLOS) arrives at the end of the trading horizon with two or more lots that are liquidated with one MO. Finally, the percentage of runs where RLOS performs better than: DLOS is 57%, ZLOS is 61%, TWAP is 67%, and ENOW is 55%.

As Table 5.6 reports, when the value of the urgency parameter ϕ is zero, RLOS and DLOS miss approximately 95% of attempts with MLOs – the majority of the MLOs seek a price improvement that does not materialise. To examine the role of the urgency parameter, we implement the execution strategies described above for a range of values of ϕ and report a summary of the findings for $\phi = 5 \times 10^{-5}$ and $\phi = 1 \times 10^{-3}$. These two choices represent impatient liquidity takers who have more urgency to liquidate the position than the patient trader with $\phi = 0$.

When $\phi = 1 \times 10^{-4}$, RLOS misses 85% of the MLOs and outperforms TWAP by 1.68 USD per million Euro exchanged. Similarly, when $\phi = 1 \times 10^{-3}$, RLOS misses 2% of the MLOs and outperforms TWAP by 0.79 USD per million Euro exchanged. The performance of latency-optimal strategies is similar to that of TWAP because the urgency with which RLOS must liquidate inventory precludes the strategy from sending speculative trades, and because the limit rates of the MLOs are set low to increase the probability of a fill at the expense of price protection. In Table 5.7, we observe that RLOS sends nearly all MLOs with a limit rate below the fundamental best bid rate to ensure a prompt liquidation and does not send speculative

MLOs – due to the trader’s impatience, in 99.97% of the simulations the strategy completes the liquidation early. These results are in stark contrast with those reported in Table 5.6, where the trader was patient and used the entire window to complete the execution programme with as many speculative MLOs as possible, and where the majority of orders sent are retries of rejected speculative MLOs.

When we amalgamate the MLOs of all traders in our data set between 1 September 2019 and 29 February 2020 in the EUR/USD currency pair, we find that 57.1% of MLOs sent by all traders were rejected by the exchange. On the other hand, when we look at the MLOs sent by each trader, we find that approximately 35% of the traders miss up to 25% of their MLOs, and approximately 6% of traders miss between 75% and 100% of their MLOs – these traders also send other types of orders to the exchange (e.g., MOs,); see Table 5.1. Finally, one cannot tell from the data which trades are part of an execution programme (i.e., child orders) or which trades are a single execution that is not part of a larger order. However, as discussed above, our model is designed to execute one stand-alone order or to execute a large parent order that is split into child orders.

Expected latency is 30 ms, start time 9.00am, $\phi = 10^{-3}$						
	MLO Attempts	MLO Fills	$Q_{T-} \geq 2$	LR < S_t	LR = S_t	LR > S_t
RLOS	10.20	10.00	0.00	10.20	-	-
DLOS	10.21	10.00	0.00	10.21	-	-
ZLOS	10.00	10.00	0.00	10.00	-	-
TWAP	10.00	10.00	0.00	10.00	-	-
ENOW	1.00	1.00	0.00	-	1.00	-

Table 5.7: Expected latency is 30 ms, $T = 6$ seconds, so expected number of attempts is 200, and $\phi = 10^{-3}$. LR: limit rate of MLO. S_t : fundamental best bid rate when the trader decides to send the MLO.

In the remainder of this subsection, we focus on a trader with 30 ms of expected latency. Table 5.8 reports the performance measure when the ten-minute execution windows start at 9.00am, 2.00pm, and 7.00pm. We find that the outperformance of RLOS over ZLOS, TWAP, ENOW is worse when the start time is 7.00pm, which is the less active period we study and one for which the probability of positive flickers in the best bid rate is lowest, so MLOs are less valuable.

	Performance measure (5.22)		
	9.00am to 9.10am	2.00pm to 2.10pm	7.00pm to 7.10pm
ZLOS	3.40	3.36	1.93
TWAP	3.80	3.51	2.25
ENOW	2.57	2.84	1.17

Table 5.8: Values of performance measure (5.22) from 9.00am to 9.10am (first column), from 2.00pm to 2.10pm (second column), and from 7.00pm to 7.10pm (third column), $T = 6$ seconds, and expected latency is 30 ms.

Next, we study the performance of RLOS when a the trader estimates 30 ms of expected latency in the market, but her latency is: 20 ms (underestimated speed in marketplace), 30 ms (correct speed estimate), 40 ms (overestimated speed in marketplace). The results, reported in Table 5.9, show that the value of the outperformance of the RLOS and DLOS over the other strategies is slightly higher (lower) when the agent underestimates (overestimates) expected latency. Recall that over the trading window, the lower (higher) is the value of the expected latency, the trader has more (fewer) attempts, on average, to execute her inventory.

	Trader assumes expected latency of 30 ms		
	True expected latency in marketplace is:		
	20 ms	30 ms	40 ms
ZLOS	3.73	3.39	3.14
TWAP	4.20	3.79	3.50
ENOW	2.97	2.55	2.27

Table 5.9: Values of performance measure (5.22) starting at 9.00am, $T = 6$ seconds, and $\phi = 0$. Expected latency in marketplace: 20 ms, 30 ms, 40 ms.

5.4.4 Trading around news events

Throughout the calendar year, many pieces of news are released according to a schedule. Market participants know the timing of the release, but do not know the content – they may have a view on how the news will affect exchange rates and trade accordingly. In general, trade activity and volatility of the exchange rates tend to increase around the time of the news release. Here, we employ the EUR/USD news event information provided by FXStreet between 1 September 2019 and 29 February 2020.³⁷ During this period, there are 117 news events marked as high impact for EUR and for USD and the future time of the release of each event is timestamped with minute precision.

³⁷<https://www.fxstreet.com/economic-calendar>.

We use the estimate $\hat{\sigma} = 1.9 \times 10^{-4}$, which is the volatility of the fundamental best bid between 9.00am to 8.00pm (see subsection 5.4.1), to compute the trading strategies around the time of the arrival of the scheduled news.³⁸ The execution horizon is $T = 6$ seconds and the agent trades during 1, 3, 5, 7, and 10 minutes around the news event. For example, the ISM manufacturing index was scheduled to be released at 2.00pm on 3 September 2019, so for the one-minute window, we employ market data from 1.59.30pm to 1.59.36pm for the first simulation, data from 1.59.36pm to 1.59.42pm for the second simulation, and so on, until the tenth simulation from 2.00.24pm to 2.00.30pm. For each six-second execution horizon we perform 100 simulations of the random latency that the trader faces. Thus, for the one-minute window, the news event study consists of $117 \times 10 \times 100 = 117,000$ simulations. Table 5.10 reports the results of the performance measure in (5.22).

	News events: Performance measure (5.22)				
	Expected latency is 30 ms				
	Minutes around news announcements:				
	1 min	3 mins	5 mins	7 mins	10 mins
ZLOS	36.31	15.16	10.62	8.66	7.21
TWAP	29.45	13.35	9.62	8.12	6.92
ENOW	23.86	10.23	7.28	6.06	5.15

Table 5.10: Performance measure around news events, $T = 6$ seconds, $\phi = 0$, and expected latency is 30 ms.

The value of the outperformance of RLOS over ZLOS and TWAP is approximately between two and ten times the value reported in Table 5.5. This significant improvement in the performance of RLOS follows from an increase in the number of speculative MLOs that are filled. Table 5.11 shows the ex-post parameters of the distribution of flickers faced by the trader with 30 ms of expected latency. It is clear that during news announcements, the value of speculative MLOs is higher because the probability of a price improvement (*resp.* slippage) and the size of the positive (*resp.* negative) flickers are greater than those of the distribution of flickers for the ten-minute trading windows that start at 9.00am, 2.00pm, and 7.00pm.

³⁸The results are robust to employing volatility estimates computed with the data around news events in August 2019.

	Parameters of the distribution of flickers (ex-post)				
	Insights from market data				
	\hat{p}_- slippage	\hat{p}_0 on target	\hat{p}_+ price improvement	$1/\hat{\eta}_-$ slippage	$1/\hat{\eta}_+$ price improvement
9am	0.057	0.901	0.042	2.68×10^{-5}	1.92×10^{-5}
2pm	0.064	0.889	0.047	2.66×10^{-5}	2.04×10^{-5}
7pm	0.054	0.923	0.023	2.81×10^{-5}	2.00×10^{-5}
News (1 min)	0.229	0.715	0.056	1.62×10^{-4}	2.52×10^{-5}
News (3 mins)	0.158	0.792	0.050	1.08×10^{-4}	2.25×10^{-5}
News (5 mins)	0.131	0.821	0.047	8.92×10^{-5}	2.18×10^{-5}
News (7 mins)	0.118	0.835	0.047	7.94×10^{-5}	2.14×10^{-5}
News (10 mins)	0.107	0.844	0.049	6.89×10^{-5}	2.12×10^{-5}

Table 5.11: Period: 1 September 2019 to 29 February 2020. Expected latency is 30ms.

5.4.5 Economic significance

In our study, when expected latency is 30 ms, latency-optimal strategies perform better than TWAP by \$3.80 per million of EUR exchanged during the trading window 9.00am to 9.10am; see Table 5.5. The value of the outperformance of the latency-optimal strategies are of the same order of magnitude as liquidity taking fees, which are between \$2 and \$5 per million of Euros exchanged. The most active traders in the exchanges tend to pay fees on the lower end of this range.³⁹ In addition, we showed that the value of the outperformance of RLOS and DLOS increases to \$6.92 per million of EUR exchanged when trading around news events for ten minutes (the outperformance is considerably better for shorter trading windows); see Table 5.10.

To further highlight the economic significance of the performance of RLOS and DLOS, Table 5.12 reports various statistics of the activity of traders in the currency pair EUR/USD during the period 1 September 2019 to 29 February 2020. For example, the trader who consumed the most liquidity in the six months of our sample, traded approximately €35 billion in the EUR/USD currency pair, and the median trader approximately €2 billion. Thus, if we assume that the median trader does not employ latency-optimal strategies, the trader could have saved up to 140,000 USD in six months had she employed RLOS or DLOS – we use as reference \$3.80/€M traded; see Table 5.5 when expected latency is 30 ms.

Savings that stem from RLOS and DLOS could be substantial for traders who are very active in the market across many currency pairs, and who are very likely to operate in more than one electronic exchange. Note, however, that these potential savings are an upper bound on the

³⁹Transaction fees are not the same for all market participants; fees are negotiated between the exchange and the trader on a bilateral basis.

savings a trader could make in one currency pair, and similarly, in other currency pairs. If all liquidity takers were to implement latency-optimal strategies, one expects liquidity providers to adjust their strategies. Savings that accrue to liquidity takers due to latency optimal strategies, is revenue lost by liquidity makers. Thus, liquidity makers will adjust their quotes to recover the drop in revenue.

	Values by trader, over six months	
	Traded quantity in €	How much is \$ 3.80/€M in \$?
Maximum	€35,740,466,000	\$ 135,813.77
75% quantile	€2,909,457,000	\$ 11,055.94
Median	€462,418,000	\$ 1,757.19
25% quantile	€67,130,000	\$ 255.09
Minimum	€2,000	\$ 0.01
Mean	€3,016,943,016	\$ 11,464.38
Standard deviation	€6,097,287,418	\$ 23,169.69

Table 5.12: Statistics of volumes traded by trader. Period: 1 September 2019 to 29 February 2020.

5.4.6 Robustness checks

In the results we showed above, the fundamental best bid rate was the current or most recent observed best bid rate when the spread was less than $v = 11$ ticks; recall that one tick is 10^{-5} USD. Table 5.13 shows that the outperformance of RLOS over the benchmarks is robust for $v \in \{7, 9, 11, 13, \infty\}$ ticks, where expected latency is 30 ms. Note that when $v = \infty$ ticks, the fundamental best bid rate is the observed bid rate, i.e., $\hat{S}_t = S_t$.

	Filter with v ticks to obtain fundamental price				
	7	9	11	13	∞
ZLOS	3.03	3.45	3.61	3.55	3.34
TWAP	3.29	3.71	3.87	3.81	3.60
ENOW	1.89	2.31	2.47	2.41	2.20

Table 5.13: Performance measure from 9.00am to 9.10am, $T = 6$ seconds, and expected latency 30 ms.

The value of the outperformance of RLOS and DLOS over other benchmarks peaks when $v = 11$ ticks. In addition, we note that for the outperformance of RLOS over TWAP, any two consecutive results are statistically the same. For example, the outperformance of RLOS over TWAP for $v = 9$ and $v = 11$ are statistically the same.

Finally, the performances of RLOS and DLOS are robust to misspecifications of the volatility parameter of the fundamental best rate. We ran simulations with volatility estimates: $\sigma = \{1, 2, 3\} \times 10^{-4}$ and expected latency 30 ms. The performance measure are as those reported above plus minus \$ 0.4/€M.

5.5 Conclusions

We have solved a general stochastic impulse control problem in which there is a stochastic delay between the action and its outcome. As an application, an investor executes a large position in a financial instrument when there is latency in the marketplace. We have derived the optimal strategies for stochastic and for deterministic latencies, and compared their performance with that of three benchmarks, including TWAP.

During normal trading hours, we have found that the latency-optimal strategies outperform the benchmarks by an amount similar to that of the fees paid by traders in the foreign exchange market. Around news events, the value of the outperformance increases to between two and ten times the value of the fees. The superior performance of the latency-optimal strategies stems from both the speculative MLOs sent by the trader and the rate protection provided by the MLOs. We have shown that fast and patient traders use their superior speed to execute the target inventory with as many speculative MLOs as possible; they do not use their speed advantage to finish the execution programme early, nor to hit the best quotes they observe in the LOB.

In financial applications, interesting research problems include a study of the effect that latency has on the financial performance of electronic trading strategies, e.g., market making, pairs trading and other statistical arbitrage strategies. In particular, extend the recent works of [Kalsi et al. \(2020\)](#) and [Cartea et al. \(2020\)](#) to account for latency. In these works, the authors use techniques from rough path theory to solve relevant algorithmic trading problems, where it is assumed that there is no delay between trade attempts and executions. A starting point is to compute the signature of the market with deterministic delay.

Finally, from a mathematical point of view, a challenging problem is to extend the current framework to include other distributions of the delay in the marketplace, and extend the model, so that there can be more than one pending order in the system when latency is stochastic.

Chapter 6

Concluding remarks

In this thesis we studied the effects of latency on liquidity takers. We presented three pieces of work that aim to (i) quantify the cost of latency, (ii) improve the marksmanship of liquidity taking orders, and (iii) incorporate latency into the optimal trading framework.

We employed marketable limit orders (MLOs) to show how to optimally account for latency. Recall that if latency is zero, then MLOs would be redundant. However, as shown in the empirical analysis in Chapter 5, in all currency pairs we study, MLOs are the type of orders with which most of the volume is traded. This thesis is, to the best of our knowledge, the first time in the literature where MLOs are used to offer solutions to the latency problem faced by liquidity takers.

There are various directions in which the works in this thesis can be extended. Arguably, the most relevant would be to leverage on tools and methods from model-based reinforcement learning (RL) to derive the results in Chapter 3, Chapter 4, and Chapter 5. Recently, the works of Leal et al. (2020), Guéant and Manziuk (2019), and Cartea et al. (2021) have shown the potential advantages that these techniques can bring over traditional stochastic control techniques. Certainly, many of the insights we draw from the models in this thesis could be replicated with RL approaches, others, such as (i) the insights derived from the closed-form solutions, or (ii) semi-analytic characterisations of the optimal strategies, will not be possible, but ultimately, it is likely that new insights can be drawn from studying the strategies ‘learned’ by these methods.

Another natural way to extend our work is to incorporate signals. The studies in Cartea and Jaimungal (2016) and Cartea et al. (2017) are examples of the improvements that trading strategies have when the filtration of the trader is enhanced with order flow and volume imbalance –see also Lehalle and Neuman (2019) and Forde et al. (2020) for recent advances on this front. It is relevant to compare the savings of the optimal use of MLOs against those that the incorporation of signals would offer. For example, in Chapter 5 we showed that if traders (i) acknowledge latency, and (ii) use MLOs optimally, the performance of trading strategies improve (over strategies that ignore latency and employ MOs) by a quantity that is in the same order of magnitude as transaction costs. Then, one would like to know by how much more can this outperformance increase if we incorporate signals.

Bibliography

- Alfonsi, A., Fruth, A., and Schied, A. (2010). Optimal execution strategies in limit order books with general shape functions. *Quantitative Finance*, 10(2):143–157.
- Almgren, R. (2003). Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance*, 10:1–18.
- Almgren, R. (2010). Execution costs. *Encyclopedia of quantitative finance*.
- Almgren, R. (2012). Optimal trading with stochastic liquidity and volatility. *SIAM Journal on Financial Mathematics*, 3(1):163–181.
- Almgren, R. and Chriss, N. (2000). Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–39.
- Antonelli, F. (1993). Backward-forward stochastic differential equations. *The Annals of Applied Probability*, 3(3):777–793.
- Banach, S. (1955). *Théorie des opérations linéaires*. Chelsea Publishing Company.
- Bandini, E. (2016). Probabilistic representation of HJB equations for optimal control of jump processes, BSDEs and related stochastic calculus. *PhD Thesis*.
- Bank, P., Soner, H. M., and Voß, M. (2017). Hedging with temporary price impact. *Mathematics and financial economics*, 11(2):215–239.
- Bardi, M. and Capuzzo-Dolcetta, I. (2008). *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*. Springer Science & Business Media.
- Barger, W. and Lorig, M. (2019). Optimal liquidation under stochastic price impact. *International Journal of Theoretical and Applied Finance*, 22(02):1850059.
- Bayraktar, E. and Ludkovski, M. (2010). Optimal trade execution in illiquid markets. *Mathematical Finance*, 21(4):681–701.
- Belak, C., Muhle-Karbe, J., and Ou, K. (2018). Liquidation in target zone models. *Market Microstructure and Liquidity*, 4(03n04):1950010.

- Bertsekas, D. P. and Shreve, S. (1978). *Stochastic optimal control: the discrete-time case*. Academic Press.
- Bertsimas, D. and Lo, A. W. (1998). Optimal control of execution costs. *Journal of Financial Markets*, 1(1):1–50.
- Black, F. and Scholes, M. S. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637–654.
- Bouchard, B., Fukasawa, M., Herdegen, M., and Muhle-Karbe, J. (2018). Equilibrium returns with transaction costs. *Finance and Stochastics*, 22(3):569–601.
- Bouchard, B. and Touzi, N. (2011). Weak dynamic programming principle for viscosity solutions. *SIAM Journal on Control and Optimization*, 49(3):948–962.
- Bouchaud, J.-P., Bonart, J., Donier, J., and Gould, M. (2018). *Trades, Quotes and Prices: Financial Markets Under the Microscope*. Cambridge University Press.
- Bruder, B. and Pham, H. (2009). Impulse control problem on finite horizon with execution delay. *Stochastic Processes and their Applications*, 119(5):1436 – 1469.
- Carmona, R. (2016). *Lectures on BSDEs, stochastic control, and stochastic differential games with financial applications*. SIAM.
- Cartea, Á., Donnelly, R., and Jaimungal, S. (2017). Algorithmic trading with model uncertainty. *SIAM Journal on Financial Mathematics*, 8(1):635–671.
- Cartea, Á. and Jaimungal, S. (2015). Optimal execution with limit and market orders. *Quantitative Finance*, 15(8):1279–1291.
- Cartea, Á. and Jaimungal, S. (2016). Incorporating order-flow into optimal execution. *Mathematics and Financial Economics*, 10(3):339–364.
- Cartea, Á., Jaimungal, S., and Penalva, J. (2015). *Algorithmic and High-Frequency Trading*. Cambridge University Press.
- Cartea, Á., Jaimungal, S., and Ricci, J. (2014). Buy low, sell high: A high frequency trading perspective. *SIAM Journal on Financial Mathematics*, 5(1):415–444.
- Cartea, Á., Jaimungal, S., and Sánchez-Betancourt, L. (2019a). Latency and Liquidity Risk. *Available at SSRN*.
- Cartea, Á., Jaimungal, S., and Sánchez-Betancourt, L. (2021). Deep reinforcement learning in algorithmic trading. In Lehalle, C.-A. and Capponni, A., editors, *Machine Learning in Financial Markets: A guide to contemporary practices (forthcoming)*. Cambridge University Press.

- Cartea, Á., Jaimungal, S., and Walton, J. (2019b). Foreign exchange markets with last look. *Mathematics and Financial Economics*, 13(1):1–30.
- Cartea, Á., Pérez Arribas, I., and Sánchez-Betancourt, L. (2020). Optimal execution of foreign securities: A double-execution problem with signatures and machine learning. *Available at SSRN*.
- Cartea, Á. and Sánchez-Betancourt, L. (2021a). Optimal execution with stochastic delay. *Available at SSRN*.
- Cartea, Á. and Sánchez-Betancourt, L. (2021b). 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*.
- Casgrain, P. and Jaimungal, S. (2020). Mean-field games with differing beliefs for algorithmic trading. *Mathematical Finance*, 30(3):995–1034.
- Cheridito, P., Filipović, D., and Kimmel, R. L. (2007). Market price of risk specifications for affine models: Theory and evidence. *Journal of Financial Economics*, 83(1):123 – 170.
- Cohen, S. N. and Elliott, R. J. (2015). *Stochastic calculus and applications*, volume 2. Springer.
- Cohen, S. N. and Szpruch, L. (2012). A limit order book model for latency arbitrage. *Mathematics and Financial Economics*, 6(3):211–227.
- 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.
- Crandall, M. G., Ishii, H., and Lions, P.-L. (1992). User’s guide to viscosity solutions of second order partial differential equations. *Bulletin of the American mathematical society*, 27(1):1–67.
- Donnelly, B. (2019). *The Art of Currency Trading: A Professional’s Guide to the Foreign Exchange Market*. John Wiley & Sons.

- Duffie, D. and Epstein, L. G. (1992). Stochastic differential utility. *Econometrica: Journal of the Econometric Society*, pages 353–394.
- Fleming, W. H. and Soner, H. M. (2006). *Controlled Markov processes and viscosity solutions*, volume 25. Springer Science & Business Media.
- Forde, M., Sánchez-Betancourt, L., and Smith, B. (2020). Optimal trade execution for gaussian signals with power-law resilience. *Preprint*.
- Gao, X. and Wang, Y. (2020). Optimal market making in the presence of latency. *Quantitative Finance*, pages 1–18.
- Gould, M. D., Porter, M. A., and Howison, S. D. (2016). The Long Memory of Order Flow in the Foreign Exchange Spot Market. *Market Microstructure and Liquidity*, 02(01):1650001.
- Gould, M. D., Porter, M. A., Williams, S., McDonald, M., Fenn, D. J., and Howison, S. D. (2013). Limit order books. *Quantitative Finance*, 13(11):1709–1742.
- Guéant, O. (2015). Optimal execution and block trade pricing: a general framework. *Applied Mathematical Finance*, 22(4):336–365.
- Guéant, O. (2016). *The financial mathematics of market liquidity: From optimal execution to market making*, volume 33. CRC Press.
- Guéant, O., Lehalle, C.-A., and Fernandez-Tapia, J. (2012). Optimal portfolio liquidation with limit orders. *SIAM Journal on Financial Mathematics*, 3(1):740–764.
- Guéant, O. and Manziuk, I. (2019). Deep reinforcement learning for market making in corporate bonds: beating the curse of dimensionality. *Applied Mathematical Finance*, 26(5):387–452.
- Guilbaud, F. and Pham, H. (2013). Optimal high-frequency trading with limit and market orders. *Quantitative Finance*, 13(1):79–94.
- Hansen, L. and Sargent, T. J. (2001). Robust control and model uncertainty. *American Economic Review*, 91(2):60–66.
- Hansen, L. P. and Sargent, T. J. (2011). *Robustness*. Princeton university press.
- Hautsch, N., Scheuch, C., and Voigt, S. (2018). Limits to arbitrage in markets with stochastic settlement latency. Technical report, CFS Working Paper Series.
- Herdegen, M., Muhle-Karbe, J., and Possamaï, D. (2021). Equilibrium asset pricing with transaction costs. *Finance and Stochastics*, 25(2):231–275.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6.

- Ho, T. S., Schwartz, R. A., and Whitcomb, D. K. (1985). The trading decision and market clearing under transaction price uncertainty. *The Journal of Finance*, 40(1):21–42.
- 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.
- Kallenberg, O. (1997). *Foundations of modern probability*. Springer Science & Business Media.
- Kallsen, J. and Muhle-Karbe, J. (2010). On using shadow prices in portfolio optimization with transaction costs. *Annals of Applied Probability*, 20(4):1341–1358.
- Kalsi, J., Lyons, T., and Perez Arribas, I. (2020). Optimal execution with rough path signatures. *SIAM Journal on Financial Mathematics*, 11(2):470–493.
- Kohlmann, M. and Tang, S. (2002). Global adapted solution of one-dimensional backward stochastic riccati equations, with application to the mean–variance hedging. *Stochastic Processes and their Applications*, 97(2):255–288.
- Leal, L., Laurière, M., and Lehalle, C.-A. (2020). Learning a functional control for high-frequency finance. *arXiv preprint*.
- Lehalle, C.-A. and Laruelle, S. (2018). *Market Microstructure in Practice*. World Scientific, 2 edition.
- 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.
- Lehalle, C.-A. and Neuman, E. (2019). Incorporating signals into optimal trading. *Finance and Stochastics*, 23(2):275–311.
- Ma, J., Protter, P., and Yong, J. (1994). Solving forward-backward stochastic differential equations explicitly—a four step scheme. *Probability theory and related fields*, 98(3):339–359.
- Ma, J. and Yong, J. (2007). *Forward-Backward Stochastic Differential Equations and their Applications*. Springer.
- Merton, R. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3:373–413.
- Moallemi, C. C. and Sağlam, M. (2013). The cost of latency in high-frequency trading. *Operations Research*, 61(5):1070–1086.

- Obizhaeva, A. A. and Wang, J. (2013). Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets*, 16(1):1–32.
- Øksendal, B. and Sulem, A. (2008). Optimal stochastic impulse control with delayed reaction. *Applied Mathematics and Optimization*, 58(2):243–255.
- 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.
- Pham, H. (2005). On some recent aspects of stochastic control and their applications. *Probability Surveys*, 2:506–549.
- Phillips, A., Stewart, A., and Adams, S. (2017). TCA and fair execution. The metrics that the FX industry must use. LMAX Exchange. Available at: <https://www.lmax.com/FX-TCA-Transaction-Cost-Analysis-white-paper.pdf>.
- 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.
- Reisinger, C. and Zhang, Y. (2020). Error estimates of penalty schemes for quasi-variational inequalities arising from impulse control problems. *SIAM Journal on Control and Optimization*, 58(1):243–276.
- 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.
- Şuhubi, E. S. (2003). *Functional Analysis*. Springer.
- 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.
- Touzi, N. (2012). *Optimal stochastic control, stochastic target problems, and backward SDE*, volume 29. Springer Science & Business Media.
- 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.

Glossary

Here we provide a brief glossary of common acronyms and terms used throughout the thesis.

Acronyms

a.e. : almost everywhere

a.s. : almost surely

BSDE : backward stochastic differential equation

FBSDE : forward-backward stochastic differential equation

FoK : fill-or-kill

FX : foreign exchange

DAY : good for day

DPP : dynamic programming principle

GTT : good until time

HJBI : Hamilton-Jacobi-Bellman-Isaacs

HJBQVI : Hamilton-Jacobi-Bellman quasi-variational inequality

IoC : immediate-or-cancel

LO : limit order

LOB : limit order book

LR : limit rate

LMAX : LMAX Exchange, London

LSE : London Stock Exchange

MLO : marketable limit order

MO : market order

MPP : marked point process

ms : millisecond

ODE : ordinary differential equation

PDE : partial differential equation

PIDE : partial integro-differential equation

SDE : stochastic differential equation

SPI : slippage-price-improvement

RL : reinforcement learning

Currencies

AUD : Australian dollar

CAD : Canadian dollar

EUR : Euro

GBP : British pound sterling

JPY : Japanese yen

MXN : Mexican peso

USD : United States of America dollar

Trading strategies

DLOS : deterministic-latency optimal strategy

RLOS : random-latency optimal strategy

TWAP : time-weighted average price

ZLOS : zero-latency optimal strategy

Technical terms

càdlàg : right continuous with left limits

càglàd : left continuous with right limits

discretion : we use the term ‘discretion’ to refer to: (i) the difference between the limit price and the best bid price for a sell MLO, or (ii) the difference between the best ask price and the limit price for a buy MLO

flicker : unpredictable short-lived deviations of a few ticks from the best quotes. Flickers are the result of very frequent occurrences of rapid sequences of post-and-cancel activity in the limit order book, and the less frequent arrival of aggressive orders that consume liquidity that is immediately replenished

FoK : the time-in-force ‘fill-or-kill’ denotes an execution upon entry (i.e., it does not rest in the limit order book), and executes either (i) all the volume of the order, or (ii) it gets cancelled entirely

IoC : the time-in-force ‘immediate-or-cancel’ denotes an execution upon entry (i.e., it does not rest in the limit order book), and executes as much volume as possible, the unfilled part of the order is cancelled

MLO : a marketable limit order is a limit order that has one of the following time-in-force qualifiers: (i) fill-or-kill, or (ii) immediate-or-cancel

LOB : at a given point in time, the limit order book is the collection of limit orders waiting to be matched

shadow price : an estimated price for an asset or service for when there is no active market

time-in-force : period of time (and other relevant factors) determining when an order is active in the market

walk the book: when a trade consumes liquidity from more than one price level

Notation

\int_A	Integral over a set A
$\mathbf{1}_B$	Indicator function over statement B
\mathcal{A}	A set of admissible strategies
$\langle \mathcal{D} J(\delta), w \rangle$	Gâteaux derivative of the functional J centered at δ in the direction of w
$\langle \mathcal{D}^2 J(\delta), w, v \rangle$	Second Gâteaux derivative of the functional J centered at δ in the directions of w and v
$\mathbb{E}_t [\cdot]$	Conditional expectation given the information up to time t
$\mathcal{L} \cdot$	Infinitesimal generator
$\mathfrak{p}(\mathrm{d}z, \mathrm{d}t)$	Random measure associated with a marked point process
$\tilde{\mathfrak{p}}(\mathrm{d}z, \mathrm{d}t)$	Compensator of the random measure $\mathfrak{p}(\mathrm{d}z, \mathrm{d}t)$
$\mathfrak{q}(\mathrm{d}z, \mathrm{d}t)$	Compensated random measure $\mathfrak{q} := \mathfrak{p} - \tilde{\mathfrak{p}}$
$\$/\text{€M}$	x USD per million EUR traded
$\$/\text{\$M}$	x USD per million USD traded
$\Psi^{(n)}(U)$	For a functional $\Psi : \mathcal{C} \rightarrow \mathcal{C}$, we let $\Psi^{(1)} = \Psi$, and $\Psi^{(n)}(U) = \Psi(\Psi^{(n-1)}(U))$
V_x, V_{xx}	V_x refers to $\partial V / \partial x$ and V_{xx} refers to $\partial^2 V / \partial x^2$ when $V : \mathbb{R}^n \rightarrow \mathbb{R}$
$\underline{w}(t, \mathbf{x}), \overline{w}(t, \mathbf{x})$	lower-semicontinuous envelope and upper-semicontinuous envelope of $w(t, \mathbf{x})$

Appendix A

Appendix for Chapter 3

A.1 Proofs and further results

A.1.1 System of ODEs satisfied by ansatz (3.18)

Here, we substitute the ansatz (3.18) in the PDE (3.17) to obtain a system of ODEs that we solve explicitly – we use these solutions in Theorem 1. Insert ansatz (3.18) into HJBI (3.17) and collect terms to obtain the following coupled system of six ODEs:

$$h'_0(t) + (b\bar{v} + a\hat{\delta} + \lambda\bar{F})h_1(t) + \sigma_F^2 h_2(t) + \kappa\bar{v}h_3(t) + \frac{1}{2} \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) (h_1(t))^2 - \phi_F \hat{F}^2 = 0, \quad (\text{A.1})$$

$$h'_1(t) + (2b\bar{v} + 2\lambda\bar{F} + 2a\hat{\delta})h_2(t) - \lambda h_1(t) + \kappa\bar{v}h_5(t) + 2 \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) h_1(t)h_2(t) + 2\phi_F \hat{F} = 0, \quad (\text{A.2})$$

$$h'_2(t) - 2\lambda h_2(t) + 2 \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) (h_2(t))^2 - \phi_F = 0, \quad (\text{A.3})$$

$$h'_3(t) - b h_1(t) + (b\bar{v} + a\hat{\delta} + \lambda\bar{F})h_5(t) + 2\kappa\bar{v}h_4(t) - \kappa h_3(t) + \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) h_1(t)h_5(t) = 0, \quad (\text{A.4})$$

$$h'_4(t) - b h_5(t) + (\sigma_v^2 - 2\kappa)h_4(t) + \frac{1}{2} \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) (h_5(t))^2 = 0, \quad (\text{A.5})$$

$$h'_5(t) - 2b h_2(t) - (\lambda + \kappa)h_5(t) + 2 \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) h_2(t)h_5(t) = 0, \quad (\text{A.6})$$

with terminal conditions

$$h_0(T) = 0, h_1(T) = 0, h_2(T) = 0, h_3(T) = 0, h_4(T) = 0, h_5(T) = 0.$$

Here the notation $h'(t)$ is short-hand for $dh(t)/dt$. Next we provide solutions for h_2 , h_5 , and h_1 , which are the functions required to write the optimal control in closed-form.

Explicit solution for h_2 .

The function h_2 satisfies the Riccati ODE in (A.3) with terminal condition $h_2(T) = 0$ and its solution depends on the sign of the coefficient $\varphi\sigma_F^2 - a^2/2\phi_\delta$. Therefore:

Case 1: $\varphi < a^2/2\sigma_F^2\phi_\delta$, then

$$h_2(t) = c_1 \tanh(c_2 t + c_3) + c_4,$$

with

$$\alpha_1 = \frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2, \quad c_1 = \frac{c_2}{2\alpha_1}, \quad c_2 = -\sqrt{2\alpha_1\phi_F + \lambda^2}, \quad c_3 = \text{atanh}\left(-\frac{\lambda}{c_2}\right) + T\sqrt{2\alpha_1\phi_F + \lambda^2}, \quad c_4 = \frac{\lambda}{2\alpha_1}.$$

The inequality $2\alpha_1\phi_F + \lambda^2 > 0$ follows from $\varphi < a^2/2\sigma_F^2\phi_\delta$, thus, c_2 is well defined.

Case 2: $\varphi = a^2/2\sigma_F^2\phi_\delta$, then

$$h_2(t) = \frac{\phi_F}{2\lambda} \left(e^{-2\lambda(T-t)} - 1 \right).$$

Case 3: $a^2/2\sigma_F^2\phi_\delta < \varphi$, then

$$h_2(t) = \tilde{c}_1 \tan(\tilde{c}_2 t + \tilde{c}_3) + \tilde{c}_4,$$

with

$$\tilde{\alpha}_1 = \varphi \sigma_F^2 - \frac{a^2}{2\phi_\delta}, \quad \tilde{c}_1 = \frac{\tilde{c}_2}{2\tilde{\alpha}_1}, \quad \tilde{c}_2 = -\sqrt{2\tilde{\alpha}_1\phi_F - \lambda^2}, \quad \tilde{c}_3 = -\operatorname{atan}\left(-\frac{\lambda}{\tilde{c}_2}\right) + T\sqrt{2\tilde{\alpha}_1\phi_F - \lambda^2}, \quad \tilde{c}_4 = \frac{\lambda}{2\tilde{\alpha}_1},$$

when $(\lambda^2/\phi_F + a^2/\phi_\delta)/2\phi_F^2 > \varphi$. We do not solve for h_2 when $(\lambda^2/\phi_F + a^2/\phi_\delta)/2\phi_F^2 \geq \varphi > a^2/2\sigma_F^2\phi_\delta$ – throughout Chapter 3 we employ the solution of the ODE discussed in Case 1.

We use the integrating-factor technique to write

$$h_5(t) = -e^{\int_t^T (2\alpha_1 h_2(u) - \kappa - \lambda) du} \int_t^T 2b h_2(s) e^{-\int_s^T (2\alpha_1 h_2(u) - \kappa - \lambda) du} ds,$$

and similarly we write

$$h_1(t) = e^{\int_t^T (2\alpha_1 h_2(u) - \lambda) du} \int_t^T (\kappa \bar{v} h_5(s) + e_1 h_2(s) + e_2) e^{-\int_s^T (2\alpha_1 h_2(u) - \lambda) du} ds,$$

where $e_1 = 2b\bar{v} + 2\lambda\bar{F} + 2a\hat{\delta}$ and $e_2 = 2\phi_F\hat{F}$.

From now on we focus on Case 1 to obtain closed-form formulae for h_5 and h_1 ; Cases 2 and 3 are similar.

Explicit solution for h_5 . Function h_5 can be explicitly written as

$$h_5(t) = -\frac{e^{-\kappa(T-t)}}{\cosh(tc_2 + c_3)} \times \left[\left(\frac{e^{d_2(T-t)+d_3} - e^{d_3}}{d_2} \right) b(c_1 + c_4) + \left(\frac{e^{d_1(T-t)-d_3} - e^{-d_3}}{d_1} \right) b(c_4 - c_1) \right],$$

where the constants $d_1 = \kappa + c_2$, $d_2 = \kappa - c_2$, and $d_3 = c_3 + Tc_2$.

Explicit solution for h_1 . Function h_1 can be solved explicitly to obtain

$$\begin{aligned} h_1(t) = & \frac{-\kappa \bar{v}}{\cosh(c_2 t + c_3)} \\ & \times \left[\frac{b(c_1 + c_4)e^{d_3}}{d_2} \left(\frac{e^{-\kappa(T-t)} - 1}{\kappa} - \frac{e^{-c_2(T-t)} - 1}{c_2} \right) \right. \\ & \quad \left. + \frac{b(c_4 - c_1)e^{-d_3}}{d_1} \left(\frac{e^{c_2(T-t)} - 1}{c_2} + \frac{e^{-\kappa(T-t)} - 1}{\kappa} \right) \right] \\ & + \frac{e_1 c_1}{c_2 \cosh(c_2 t + c_3)} (\cosh(c_2 T + c_3) - \cosh(c_2 t + c_3)) \\ & + \frac{e_1 c_4 + e_2}{c_2 \cosh(c_2 t + c_3)} (\sinh(c_2 T + c_3) - \sinh(c_2 t + c_3)). \end{aligned}$$

Finally, the functions h_0 and h_3 are obtained explicitly in a similar way using an integrating factor technique. Note that to obtain the optimal controls explicitly we only require h_1 , h_2 , h_5 .

A.1.2 Proof of Theorem 1.

Here we provide a class of controls that produces strong solutions for F^δ and v , see Lemma 7. We employ Lemma 7 to prove the admissibility of the control δ^* (see Corollary 1). Then we prove that the control x^* is also admissible and conclude this section with the verification of the value function.

Lemma 7. *Let $\delta_t = \mu(t, F_t, v_t)$ with $\mu \in \mathcal{M}_{0,T}$. Then, for all $t \in [0, T]$ there exist strong solutions to the SDEs of F^δ and v .*

Proof. Let $\delta_t = \mu(t, F_t, v_t)$, for $\mu \in \mathcal{M}_{0,T}$. Thus, there exists $K \in \mathbb{R}$ such that for all $t \in [0, T]$, $\mathbf{y}, \mathbf{z} \in \mathbb{R}^2$

$$|\mu(t, \mathbf{y}) - \mu(t, \mathbf{z})| \leq K |\mathbf{y} - \mathbf{z}|, \quad \text{and} \quad |\mu(t, \mathbf{y})|^2 \leq K^2 (1 + |\mathbf{y}|^2). \quad (\text{A.7})$$

Let $\mathbf{Z}_t^\delta = (F_t^\delta, v_t)^\top$, thus

$$d\mathbf{Z}_t^\delta = (A(t) + B(t) \mathbf{Z}_t^\delta + \boldsymbol{\mu}(t, \mathbf{Z}_t^\delta)) dt + C_t d\mathbf{W}_t,$$

with $A(t)$, $B(t)$, C_t , and $\boldsymbol{\mu}$ given by

$$A(t) = \begin{bmatrix} \lambda \bar{F} + b \bar{v} \\ \kappa \bar{v} \end{bmatrix}, \quad B(t) = \begin{bmatrix} -\lambda & -b \\ 0 & -\kappa \end{bmatrix}, \quad C_t = \begin{bmatrix} \sigma_F & 0 \\ 0 & \sigma_v v_t \end{bmatrix}, \quad \boldsymbol{\mu}(t, \mathbf{Z}_t^\delta) = \begin{bmatrix} a \mu(t, \mathbf{Z}_t^\delta) \\ 0 \end{bmatrix}.$$

Define the functions $c : [0, T] \times \mathbb{R}^2 \rightarrow \mathbb{R}^2$, and $\sigma : [0, T] \times \mathbb{R}^2 \rightarrow \mathbb{R}^{2 \times 2}$, which are given by

$$c(t, \mathbf{Z}) = (A(t) + B(t) \mathbf{Z} + \boldsymbol{\mu}(t, \mathbf{Z})), \quad \sigma(t, \mathbf{Z}) = \begin{bmatrix} \sigma_F & 0 \\ 0 & \sigma_v v_t \end{bmatrix}.$$

Now, let $t \in [0, T]$, $\mathbf{Z}_1, \mathbf{Z}_2 \in \mathbb{R}^2$. We first verify the Lipschitz condition on the function c ,

$$\begin{aligned} |c(t, \mathbf{Z}_1) - c(t, \mathbf{Z}_2)| &= |B(t) (\mathbf{Z}_1 - \mathbf{Z}_2) + \boldsymbol{\mu}(t, \mathbf{Z}_1) - \boldsymbol{\mu}(t, \mathbf{Z}_2)| \\ &\leq |B(t)| |\mathbf{Z}_1 - \mathbf{Z}_2| + |\boldsymbol{\mu}(t, \mathbf{Z}_1) - \boldsymbol{\mu}(t, \mathbf{Z}_2)| \\ &\leq \bar{B} |\mathbf{Z}_1 - \mathbf{Z}_2| + K |\mathbf{Z}_1 - \mathbf{Z}_2| = \hat{K} |\mathbf{Z}_1 - \mathbf{Z}_2|, \end{aligned}$$

for $\hat{K} = K + \bar{B}$, and $\bar{B} = \max_{t \in [0, T]} |B(t)|$, which is clearly attained. For the function σ , we have

$$|\sigma(t, \mathbf{Z}_1) - \sigma(t, \mathbf{Z}_2)| \leq \sigma_v |\mathbf{Z}_1 - \mathbf{Z}_2|.$$

We test the linear growth condition on the function c . We use triangle inequalities and simple bounds to obtain

$$\begin{aligned} |c(t, \mathbf{Z}_1)|^2 &= |A(t) + B(t) \mathbf{Z}_1 + \boldsymbol{\mu}(t, \mathbf{Z}_1)|^2 \leq (|A(t)| + |B(t) \mathbf{Z}_1| + |\boldsymbol{\mu}(t, \mathbf{Z}_1)|)^2 \\ &\leq 3 |A(t)|^2 + 3 |B(t)|^2 |\mathbf{Z}_1|^2 + 3 |\boldsymbol{\mu}(t, \mathbf{Z}_1)|^2 \\ &\leq 3 \bar{A}^2 + 3 \bar{B}^2 |\mathbf{Z}_1|^2 + 3 K^2 (1 + |\mathbf{Z}_1|^2) \\ &\leq \bar{D}^2 + \bar{D}^2 |\mathbf{Z}_1|^2 + 3 K^2 (1 + |\mathbf{Z}_1|^2) \\ &= (\bar{D}^2 + 3 K^2) (1 + |\mathbf{Z}_1|^2), \end{aligned}$$

for $\bar{A} = \max_{t \in [0, T]} |A(t)|$, which is clearly attained. Finally, we take $\bar{D}^2 = \max \{\bar{A}^2, \bar{B}^2\}$.

On the other hand, for σ , we have

$$\begin{aligned} |\sigma(t, \mathbf{Z}_1)|^2 &\leq 2 \sigma_F^2 + 2 \sigma_v^2 |\mathbf{Z}_1|^2 \\ &\leq \bar{E}^2 (1 + |\mathbf{Z}_1|^2), \end{aligned}$$

for $\bar{E} = \max \{\sigma_F^2, \sigma_v^2\}$. Lipschitz and linear growth conditions follow by choosing

$$\bar{K} = \max \left\{ \bar{E}, \sqrt{\bar{D}^2 + 3 K^2}, \sigma_v, \hat{K} \right\} < \infty.$$

Thus the Lipschitz and linear growth conditions are satisfied, therefore, for all $t \in [0, T]$, there exists a strong solution to the SDE satisfied by \mathbf{Z}^δ , and subsequently F^δ and v_t . ■

Corollary 1. *Let $\delta \in \mathcal{A}_{t,T}$. Then, for all $t \in [0, T]$ there exist strong solutions to the SDEs of F^δ and v .*

This shows that as long as the control δ is admissible, F^δ and v have strong solutions.

Proof. Let $\delta \in \mathcal{A}_{t,T}$, thus there exist $\mu \in \mathcal{M}_{0,T}$ such that $\delta_t = \mu(t, F_t, v_t)$. By Lemma 7, for all $t \in [0, T]$ there exist strong solutions to the SDEs of F^δ and v . ■

Proof. **The control δ^* is admissible.**

Recall that δ^* is given by

$$\delta_t^* = \hat{\delta} + \frac{a}{2\phi_\delta} \left(h_1(t) + 2h_2(t) F_t^{\delta^*} + h_5(t) v_t \right).$$

Define $f : [0, T] \times \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$ as

$$f(t, x, y) = \hat{\delta} + \frac{a}{2\phi_\delta} h_1(t) + \frac{a h_2(t)}{\phi_\delta} x + \frac{a h_5(t)}{2\phi_\delta} y,$$

so $\delta_t^* = f(t, F_t, v_t)$, and, since f is linear in x and y , and a continuous function of t , then $f \in \mathcal{M}_{0,T}$. The control δ^* is \mathcal{F} -adapted, because it is a continuous function of two \mathcal{F} -adapted processes. Then, we only need to prove that

$$\mathbb{E} \left[\int_0^T (\delta_s^*)^2 ds \right] < \infty.$$

Based on Young's inequality we find the following bound for the optimal control:

$$(\delta_t^*)^2 \leq 2 \left(\hat{\delta} + \frac{a}{2\phi_\delta} h_1(t) \right)^2 + \frac{2a^2}{\phi_\delta^2} (h_2(t))^2 (F_t^{\delta^*})^2 + \frac{a^2}{2\phi_\delta^2} (h_5(t))^2 (v_t)^2,$$

and because h_1 , h_2 and h_5 are continuous functions in the closed interval $[0, T]$, they are bounded. Let U be an upper bound for $|h_1|$, $|h_2|$, and $|h_5|$ on $[0, T]$.

As $\delta_t^* = f(t, F_t, v_t)$ for $f \in \mathcal{M}_{0,T}$, by Lemma 7, we have that F^{δ^*} and v have strong solutions for all $t \in [0, T]$, thus

$$\mathbb{E} \left[\sup_{0 \leq t \leq T} (F_t^{\delta^*})^2 \right] < \infty \quad \text{and} \quad \mathbb{E} \left[\sup_{0 \leq t \leq T} (v_t)^2 \right] < \infty. \quad (\text{A.8})$$

Finally, using the upper bound U and the inequalities in (A.8), we conclude that

$$\begin{aligned} \mathbb{E} \left[\int_0^T (\delta_s^*)^2 ds \right] &\leq \mathbb{E} \left[\int_0^T 2 \left(\hat{\delta} + \frac{a}{2\phi_\delta} h_1(t) \right)^2 + \frac{2a^2}{\phi_\delta^2} (h_2(t))^2 (F_t^{\delta^*})^2 + \frac{a^2}{2\phi_\delta^2} (h_5(t))^2 (v_t)^2 \right] \\ &\leq \mathbb{E} \left[2T \left(\hat{\delta} + \frac{a}{2\phi_\delta} U \right)^2 + \frac{2a^2}{\phi_\delta^2} U^2 \sup_{0 \leq t \leq T} (F_t^{\delta^*})^2 + T \frac{a^2}{2\phi_\delta^2} U^2 \sup_{0 \leq t \leq T} (v_t)^2 \right] \\ &< \infty. \end{aligned}$$

■

Proof. **The control x^* is admissible.**

This proof is similar to that of Theorem 1 in Cheridito et al. (2007). Let $\delta \in \mathcal{A}_{t,T}$, then $\exists \mu \in \mathcal{M}_{t,T}$, such that $\delta_t = \mu(t, F_t, v_t)$. Recall that the optimal control x^* is given by

$$\begin{aligned} x_t^* &= \varphi \sigma_F (h_1(t) + 2h_2(t) F_t^\delta + h_5(t) v_t) \\ &= \eta(t, F_t, v_t), \end{aligned}$$

which is adapted because it is a continuous function of adapted processes. Therefore, it remains to show that

$$Z_t = \exp \left\{ -\frac{1}{2} \int_0^t (x_u^*)^2 du - \int_0^t x_u^* dW_u^F \right\},$$

is a martingale. As the control x^* is a well-defined continuous process, Z is a well-defined local martingale with respect to \mathbb{P} , and subsequently, because Z is non-negative, it follows that it is a super-martingale, and for it to be a martingale, we only need to prove that

$$\mathbb{E}^\mathbb{P} [Z_T] = 1.$$

Consider the function $\mu^\mathbb{Q}$ given by

$$\mu^\mathbb{Q}(t, F_t, v_t) = \mu(t, F_t, v_t) - \frac{\sigma_F}{a} \eta(t, F_t, v_t),$$

and because the term $\eta(t, F_t, v_t)$ is continuous in time and linear in F and v , it belongs to $\mathcal{M}_{t,T}$, therefore $\mu^\mathbb{Q} \in \mathcal{M}_{t,T}$.

Given that $\mu, \mu^\mathbb{Q} \in \mathcal{M}_{t,T}$, by Lemma 7, we know there is a strong solution for the SDEs given by

$$dF_t = (\lambda(\bar{F} - F_t) + a\mu(t, F_t, v_t) - b(v_t - \bar{v})) dt + \sigma_F dW_t^F, \quad (\text{A.9})$$

$$d\tilde{F}_t = (\lambda(\bar{F} - \tilde{F}_t) + a\mu^\mathbb{Q}(t, \tilde{F}_t, v_t) - b(v_t - \bar{v})) dt + \sigma_F dW_t^F. \quad (\text{A.10})$$

Now define

$$\tilde{x}_t^* = \varphi \sigma_F (h_1(t) + 2h_2(t) \tilde{F}_t^\delta + h_5(t) v_t) \quad (\text{A.11})$$

$$= \eta(t, \tilde{F}_t, v_t). \quad (\text{A.12})$$

For $n \in \mathbb{N}$, consider the stopping times τ_n , and $\tilde{\tau}_n$

$$\tau_n = \inf \{t > 0 \mid x_t^* \geq n\} \wedge T \quad \text{and} \quad \tilde{\tau}_n = \inf \{t > 0 \mid \tilde{x}_t^* \geq n\} \wedge T, \quad (\text{A.13})$$

and because x^* and \tilde{x}^* are \mathbb{P} -a.s. finite, then

$$\lim_{n \rightarrow \infty} \mathbb{P}(\tau_n = T) = \lim_{n \rightarrow \infty} \mathbb{P}(\tilde{\tau}_n = T) = 1. \quad (\text{A.14})$$

Define for each $n \in \mathbb{N}$ the process

$$x_t^{(n)} = x_t^* \mathbf{1}_{\{t \leq \tau_n\}},$$

which satisfies the Novikov's condition

$$\mathbb{E}^\mathbb{P} \left[\exp \left\{ \frac{1}{2} \int_0^T (x_u^{(n)})^2 du \right\} \right] \leq \exp \left\{ \frac{n^2 T}{2} \right\} < \infty.$$

Then it follows that

$$\begin{aligned} \frac{d\mathbb{Q}(x^{(n)})}{d\mathbb{P}} \Big|_t &= \exp \left\{ -\frac{1}{2} \int_0^t (x_u^{(n)})^2 du - \int_0^t x_u^{(n)} dW_u^F \right\} \\ &= Z_t^{(n)}, \end{aligned}$$

defines an equivalent measure to \mathbb{P} , and by Girsanov's theorem, $W_t^n = W_t^F + \int_0^t x_s^{(n)} ds$ is a Brownian motion under the measure $\mathbb{Q}(x^{(n)})$, abbreviated as \mathbb{Q}^n . Finally,

$$\begin{aligned} \mathbb{E}^\mathbb{P} [Z_T] &= \mathbb{E}^\mathbb{P} \left[\lim_{n \rightarrow \infty} Z_T \mathbf{1}_{\{\tau_n = T\}} \right] \\ &= \mathbb{E}^\mathbb{P} \left[\lim_{n \rightarrow \infty} Z_T^n \mathbf{1}_{\{\tau_n = T\}} \right] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^\mathbb{P} [Z_T^n \mathbf{1}_{\{\tau_n = T\}}] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{Q}^n} [\mathbf{1}_{\{\tau_n = T\}}] \\ &= \lim_{n \rightarrow \infty} \mathbb{Q}^n(\tau_n = T), \end{aligned}$$

where the first two equalities hold because the control x^* is \mathbb{P} -a.s. finite, and definition (A.13). The third equality follows from the monotone convergence theorem. The fourth equality is the change of measure imposed by Z_T^n .

Finally, starting at F_0 , the stopped process $F_{t \wedge \tau_n}$ under \mathbb{Q}^n is given by

$$\begin{aligned} F_{t \wedge \tau_n} &= F_0 + \int_0^{t \wedge \tau_n} (\lambda(\bar{F} - F_s) + a\mu(s, F_s, v_s) - \sigma_F x_s^{(n)} - b(v_s - \bar{v})) ds + \int_0^{t \wedge \tau_n} \sigma_F dW_s^n \\ &= F_0 + \int_0^{t \wedge \tau_n} (\lambda(\bar{F} - F_s) + a\mu(s, F_s, v_s) - \sigma_F x_s^* - b(v_s - \bar{v})) ds + \int_0^{t \wedge \tau_n} \sigma_F dW_s^n \\ &= F_0 + \int_0^{t \wedge \tau_n} (\lambda(\bar{F} - F_s) + a\mu^\mathbb{Q}(s, F_s, v_s) - b(v_s - \bar{v})) ds + \int_0^{t \wedge \tau_n} \sigma_F dW_s^n, \end{aligned}$$

and similarly, the stopped process $\tilde{F}_{t \wedge \tilde{\tau}_n}$, under \mathbb{P} is given by

$$\tilde{F}_{t \wedge \tilde{\tau}_n} = F_0 + \int_0^{t \wedge \tilde{\tau}_n} \left(\lambda \left(\bar{F} - \tilde{F}_s \right) + a \mu^{\mathbb{Q}}(s, \tilde{F}_s, v_s) - b(v_s - \bar{v}) \right) ds + \int_0^{t \wedge \tilde{\tau}_n} \sigma_F dW_s^F.$$

Therefore $(\tilde{F}_{t \wedge \tilde{\tau}_n})_{0 \leq t \leq T}$, and $(F_{t \wedge \tau_n})_{0 \leq t \leq T}$ have the same distribution under \mathbb{Q}^n and \mathbb{P} respectively, thus

$$\lim_{n \rightarrow \infty} \mathbb{Q}^n(\tau_n = T) = \lim_{n \rightarrow \infty} \mathbb{P}(\tilde{\tau}_n = T) = 1,$$

and

$$\mathbb{E}^{\mathbb{P}}[Z_T] = \lim_{n \rightarrow \infty} \mathbb{Q}^n(\tau_n = T) = \lim_{n \rightarrow \infty} \mathbb{P}(\tilde{\tau}_n = T) = 1,$$

which concludes the proof. \blacksquare

Finally, we present the proof of Theorem 1, i.e., we show that (3.18) is the trader's value function given in (3.13).

Proof. **V is the value function of the trader's control problem in (3.13).**

Let $\mathbb{Q}^* = \mathbb{Q}^*(\delta)$ be the measure inferred by $x_t^*(\delta) = \varphi \sigma_F (h_1(t) + 2h_2(t) F_t^\delta + h_5(t) v_t)$, in other words,

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}} \Big|_T = \exp \left\{ -\frac{1}{2} \int_0^T \mathbf{x}_u^{*'} \mathbf{x}_u^* du - \int_0^T \mathbf{x}_u^* d\mathbf{W}_u \right\}, \quad (\text{A.15})$$

where $\mathbf{x}_t^* = (x_t^*, 0)$. Clearly $x_t^*(\delta)$ represents the infimum measure of (3.13) for an admissible δ .

The dynamics of v under \mathbb{Q}^* do not change, but the dynamics of F^{δ, \mathbb{Q}^*} do change. Specifically

$$dF_t^{\delta, \mathbb{Q}^*} = \left(\lambda \bar{F} - (\lambda + 2\varphi \sigma_F^2 h_2(t)) F_t^{\delta, \mathbb{Q}^*} + a \delta_t + b \bar{v} - (b + \varphi \sigma_F^2 h_5(t)) v_t - \varphi \sigma_F^2 h_1(t) \right) dt + \sigma_F dW_t^{F, \mathbb{Q}^*}.$$

As $V \in \mathcal{C}^{1,2,2}$, apply Itô's Lemma under \mathbb{Q}^* , and consider $\delta \in \mathcal{A}_{t, T}$, hence

$$V(T, F_T^\delta, v_T) = V(t, f, v) + \int_t^T \partial_s V + \mathcal{L}^{\mathbb{Q}^*} V ds + \int_t^T \partial_v V \sigma_v v_s dW_s^{v, \mathbb{Q}^*} + \int_t^T \partial_f V \sigma_F dW_s^{F, \mathbb{Q}^*}, \quad (\text{A.16})$$

where W_s^{F*} and W_s^{v*} are standard \mathbb{Q}^* -Brownian motions. The function V is a solution of (3.16), and δ is an arbitrary control, thus

$$\partial_s V + \mathcal{L}^{\mathbb{Q}^*} V \leq \phi_\delta \left(\delta - \hat{\delta} \right)^2 + \phi_F \left(f - \hat{F} \right)^2 - \frac{1}{2\varphi} (x^*)^2.$$

Now, take expectations on both sides of (A.16) to obtain

$$\begin{aligned} 0 &= \mathbb{E}_{t, f, v}^{\mathbb{Q}^*} [V(T, F_T^\delta, v_T)] , \\ &= V(t, f, v) + \mathbb{E}_{t, f, v}^{\mathbb{Q}^*} \left[\int_t^T \partial_s V + \mathcal{L}_s^{\mathbb{Q}^*} V ds \right] , \\ &\leq V(t, f, v) + \mathbb{E}_{t, f, v}^{\mathbb{Q}^*} \left[\phi_\delta \int_t^T \left(\delta_s - \hat{\delta} \right)^2 ds + \phi_F \int_t^T \left(F_s^\delta - \hat{F} \right)^2 ds - \frac{1}{2\varphi} \int_t^T (x_s^*)^2 ds \right] , \\ &= V(t, f, v) + \mathbb{E}_{t, f, v}^{\mathbb{Q}^*} \left[\phi_\delta \int_t^T \left(\delta_s - \hat{\delta} \right)^2 ds + \phi_F \int_t^T \left(F_s^\delta - \hat{F} \right)^2 ds - \frac{1}{2\varphi} \int_t^T (x_s^*)^2 ds + \frac{1}{\varphi} \int_t^T x_s^* dW_s^{\mathbb{Q}^*} \right] , \end{aligned}$$

therefore

$$\mathbb{E}_{t, f, v}^{\mathbb{Q}^*} \left[-\phi_\delta \int_t^T \left(\delta_s - \hat{\delta} \right)^2 ds - \phi_F \int_t^T \left(F_s^\delta - \hat{F} \right)^2 ds + \frac{1}{2\varphi} \int_t^T (x_s^*)^2 ds - \frac{1}{\varphi} \int_t^T x_s^* dW_s^{\mathbb{Q}^*} \right] \leq V(t, f, v),$$

or alternatively,

$$\begin{aligned} \mathbb{E}_{t,f,v}^{\mathbb{Q}^*} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds - \frac{1}{2\varphi} \int_t^T (x_s^*)^2 ds - \frac{1}{\varphi} \int_t^T x_s^* dW_s \right] &\leq V(t, f, v), \\ \mathbb{E}_{t,f,v}^{\mathbb{Q}^*} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q} | \mathbb{P}) \right] &\leq V(t, f, v), \end{aligned}$$

and because $V(t, f, v)$ is a bound for any $\delta \in \mathcal{A}_{t,T}$, we have that

$$\begin{aligned} \sup_{\delta \in \mathcal{A}_{t,T}} \mathbb{E}_{t,f,v}^{\mathbb{Q}^*} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q}^* | \mathbb{P}) \right] &\leq V(t, f, v) \\ \sup_{\delta \in \mathcal{A}_{t,T}} \inf_{\mathbb{Q} \in \mathfrak{Q}} \mathbb{E}_{t,f,v}^{\mathbb{Q}} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q} | \mathbb{P}) \right] &\leq V(t, f, v), \quad (\text{A.17}) \end{aligned}$$

which completes the first part of the proof.

Now, fix the strategy δ^* , then for any $\mathbb{Q} \in \mathfrak{Q}$, because V is a solution of (3.16), and because the inf and sup can be interchanged in (3.16) (x and δ are decoupled in (3.16)), we have the inequality

$$\partial_s V + \mathcal{L}^{\mathbb{Q}} V \geq \phi_\delta (\delta^* - \hat{\delta})^2 + \phi_F (f - \hat{F})^2 - \frac{1}{2\varphi} (x)^2, \quad (\text{A.18})$$

and, following the same steps as before together with (A.18), we have, under the measure \mathbb{Q} ,

$$\mathbb{E}_{t,f,v}^{\mathbb{Q}} \left[-\phi_\delta \int_t^T (\delta_s^* - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^{\delta^*} - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q} | \mathbb{P}) \right] \geq V(t, f, v),$$

and because this is true for all $\mathbb{Q} \in \mathfrak{Q}$, we write

$$\begin{aligned} \inf_{\mathbb{Q} \in \mathfrak{Q}} \mathbb{E}_{t,f,v}^{\mathbb{Q}} \left[-\phi_\delta \int_t^T (\delta_s^* - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^{\delta^*} - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q} | \mathbb{P}) \right] &\geq V(t, f, v) \\ \sup_{\delta \in \mathcal{A}_{t,T}} \inf_{\mathbb{Q} \in \mathfrak{Q}} \mathbb{E}_{t,f,v}^{\mathbb{Q}} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q} | \mathbb{P}) \right] &\geq V(t, f, v), \quad (\text{A.19}) \end{aligned}$$

which completes the second part of the proof.

Combine (A.17) and (A.19) to obtain the desired result.

■

A.2 Further robustness checks

A.2.1 Infinite-horizon value function

We solve the infinite-horizon problem for the trader, which is the analogue to (3.13) with $T \rightarrow \infty$ and here we include an exponential discount factor $\beta > 0$, so the problem is well defined. Thus, the trader's value function is

$$H(f, v) = \sup_{\delta \in \mathcal{A}_{0,\infty}} \inf_{x^F \in \mathcal{X}_{0,\infty}} \mathbb{E}_{f,v} \left[- \int_0^\infty e^{-\beta s} \left(\phi_\delta (\delta_s - \hat{\delta})^2 ds + \phi_F (F_s^{\delta, x^F} - \hat{F})^2 ds + \mathcal{H}^\beta(x^F) \right) \right], \quad (\text{A.20})$$

where we now define the penalty as

$$\mathcal{H}^\beta(x^F) = \frac{1}{\varphi} \log \left(\exp \left(-\frac{1}{2} \int_0^\infty e^{-\beta s} (x_s^F)^2 ds \right) \right).$$

Similar to Bouchard et al. (2018) we define the set of admissible strategies $\mathcal{A}_{0,\infty}$ and $\mathcal{X}_{0,\infty}$ to be

$$\mathcal{A}_{0,\infty} = \left\{ \delta = (\delta_t)_{\{0 \leq t < \infty\}} \mid x^F \text{ is } \mathcal{F} \text{ - adapted, } \mathbb{E} \left[\int_0^\infty e^{-\beta s} (\delta_s)^2 ds \right] < \infty \right\}, \quad (\text{A.21})$$

$$\mathcal{X}_{0,\infty} = \left\{ x^F = (x_t^F)_{\{0 \leq t < \infty\}} \mid x^F \text{ is } \mathcal{F} \text{ - adapted, } \mathbb{E} \left[\int_0^\infty e^{-\beta s} (x_s^F)^2 ds \right] < \infty \right\}, \quad (\text{A.22})$$

and $F^{\delta, x^F} = (F_t^{\delta, x^F})_{\{0 \leq t \leq T\}}$, follows the SDE

$$dF_t^{\delta, x^F} = \left(\lambda (\bar{F} - F_t^{\delta, x^F}) + a \delta_t - b(v_t - \bar{v}) - \sigma_F x_t^F \right) dt + \sigma_F dW_t^F. \quad (\text{A.23})$$

As above, the value function satisfies the HJBI

$$\begin{aligned} -\beta H + \sup_{\delta} \inf_x \left(\left(\lambda (\bar{F} - f) + a \delta - b(v - \bar{v}) - \sigma_F x \right) H_f + \frac{1}{2} \sigma_F^2 H_{ff} + \frac{1}{2} x^2 \right. \\ \left. + \kappa (\bar{v} - v) H_v + \frac{1}{2} v^2 \sigma_v^2 H_{vv} - \phi_{\delta} (\delta - \hat{\delta})^2 - \phi_F (f - \hat{F})^2 \right) = 0, \end{aligned} \quad (\text{A.24})$$

which, after substituting the feedback controls, reduces to

$$\begin{aligned} -\beta H + (\lambda (\bar{F} - f) - b(v - \bar{v})) H_f + \frac{1}{2} \sigma_F^2 H_{ff} + \frac{a^2 V_f^2}{4 \phi_{\delta}} - \frac{1}{2} \varphi H_f^2 \sigma_F^2 + a \hat{\delta} H_f \\ + \kappa (\bar{v} - v) H_v + \frac{1}{2} v^2 \sigma_v^2 H_{vv} - \phi_F (f - \hat{F})^2 = 0. \end{aligned}$$

We use a similar ansatz to that in (3.18)

$$H(f, v) = k_0 + k_1 f + k_2 f^2 + k_3 v + k_4 v^2 + k_5 f v, \quad (\text{A.25})$$

where $k_0, k_1, k_2, k_3, k_4, k_5$ are constants. Substitute the ansatz in the HJBI and collect terms to obtain the system of equations:

$$\begin{aligned} -\beta k_0 + (b \bar{v} + a \hat{\delta} + \lambda \bar{F}) k_1 + \sigma_F^2 k_2 + \kappa \bar{v} k_3 + \frac{1}{2} \left(\frac{a^2}{2 \phi_{\delta}} - \varphi \sigma_F^2 \right) (k_1)^2 - \phi_F \hat{F}^2 &= 0, \\ -\beta k_1 + (2 b \bar{v} + 2 \lambda \bar{F} + 2 a \hat{\delta}) k_2 - \lambda k_1 + \kappa \bar{v} k_5 + 2 \left(\frac{a^2}{2 \phi_{\delta}} - \varphi \sigma_F^2 \right) k_1 k_2 + 2 \phi_F \hat{F} &= 0, \\ -\beta k_2 - 2 \lambda k_2 + 2 \left(\frac{a^2}{2 \phi_{\delta}} - \varphi \sigma_F^2 \right) (k_2)^2 - \phi_F &= 0, \\ -\beta k_3 - b k_1 + (b \bar{v} + a \hat{\delta} + \lambda \bar{F}) k_5 + 2 \kappa \bar{v} k_4 - \kappa k_3 + \left(\frac{a^2}{2 \phi_{\delta}} - \varphi \sigma_F^2 \right) k_1 k_5 &= 0, \\ -\beta k_4 - b k_5 + (\sigma_v^2 - 2 \kappa) k_4 + \frac{1}{2} \left(\frac{a^2}{2 \phi_{\delta}} - \varphi \sigma_F^2 \right) (k_5)^2 &= 0, \\ -\beta k_5 - 2 b k_2 - (\lambda + \kappa) k_5 + 2 \left(\frac{a^2}{2 \phi_{\delta}} - \varphi \sigma_F^2 \right) k_2 k_5 &= 0. \end{aligned}$$

In Tables A.1 and A.2, which are the analogues of Tables 3.3 and 3.4, we compare the performance of the strategy in the finite- and infinite-horizon cases. The results are broadly the same – recall that in the finite-horizon case we assume that every day the latency-optimal strategy assumes $T = 5$.

Missed by naive, filled by alternative strategies						
		Extra Fills	Extra Volume $\times 10^6$	Cost in JPY	Cost 20ms later in JPY	Cost 100ms later in JPY
(a)	$\delta_t = \infty, \forall t$	9,483	4,222	8,914,046	9,841,755	12,158,945
	$\delta^*, \hat{\delta} = 0$				$\delta_t = \infty, \forall t$	$\delta_t = \infty, \forall t$
	Value function					
(b)	Finite-horizon	8,318	3,718	6,528,965	7,459,198	9,601,222
(c)	Infinite-horizon	8,317	3,718	6,527,965	7,458,198	9,600,222

Table A.1: Performance of optimal strategy for T1 for finite and infinite horizon cases. The total number of trade attempts is 110,458. Cost-related quantities are in JPY. Parameters: $\phi_{\delta} = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, $\beta = 0.01$, and $T = 5$ days in the finite-horizon case. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies					
		Extra Fills	Avg. cost in ticks	Avg. cost 20ms later in ticks	Avg. cost 100ms later in ticks	Cost increase 20ms later %	Cost increase 100ms later %
(a)	$\delta_t = \infty, \forall t$	9,483	2.11	2.33	2.88	10%	36%
	$\delta^*, \hat{\delta} = 0$					$\delta_t = \infty, \forall t$	$\delta_t = \infty, \forall t$
	Value function						
(b)	Finite-horizon	8,318	1.76	2.01	2.58	14%	47%
(c)	Infinite-horizon	8,317	1.76	2.01	2.58	14%	47%

Table A.2: Performance of optimal strategy for T1 for finite and infinite horizon cases. The total number of trade attempts is 110,458. Cost-related quantities are weighted by volume and expressed in ticks, 10^{-3} JPY. The metrics in column 6 (*resp.* 7) are the proportional increase of the values in columns 5 and 6 (*resp.* 5 and 7) in Table 3.3. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, $\beta = 0.01$, and $T = 5$ days in the finite-horizon case. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

A.2.2 Alternative models for volatility

If we model the volatility of the micro-exchange rate with the Ornstein-Uhlenbeck process (3.7), then the trader's value function satisfies the HJBI

$$V_t + \sup_{\delta} \inf_x \left((\lambda(\bar{F} - f) + a\delta - b(v - \bar{v}^{\text{OU}}) - \sigma_F x) V_f + \frac{1}{2} \sigma_F^2 V_{ff} + \frac{1}{2\varphi} x^2 + \kappa^{\text{OU}} (\bar{v}^{\text{OU}} - v) V_v + \frac{1}{2} v^2 (\sigma_v^{\text{OU}})^2 V_{vv} - \phi_\delta (\delta - \hat{\delta})^2 - \phi_F (f - \hat{F})^2 \right) = 0. \quad (\text{A.26})$$

This HJBI is similar to (3.16), where the three parameters κ , \bar{v} , σ_v in (3.16) become κ^{OU} , \bar{v}^{OU} , σ_v^{OU} , respectively, and the term $\frac{1}{2} v^2 \sigma_v^2 V_{vv}$ becomes $\frac{1}{2} (\sigma_v^{\text{OU}})^2 V_{vv}$.

Clearly, the optimal controls that result from (A.26) have the same functional form as those we derived above in Proposition 1. To solve (A.26) we employ ansatz (3.18) and derive a system of ODEs similar to the one shown in A.1.1.

When we employ Heston's volatility model (3.8), the trader's value function satisfies the HJBI

$$V_t + \sup_{\delta} \inf_x \left((\lambda(\bar{F} - f) + a\delta - b(v - \bar{v}^{\text{H}}) - \sigma_F x) V_f + \frac{1}{2} \sigma_F^2 V_{ff} + \frac{1}{2\varphi} x^2 + \kappa^{\text{H}} (\bar{v}^{\text{H}} - v) V_v + \frac{1}{2} v^2 (\sigma_v^{\text{H}})^2 V_{vv} - \phi_\delta (\delta - \hat{\delta})^2 - \phi_F (f - \hat{F})^2 \right) = 0. \quad (\text{A.27})$$

As before, this HJBI is similar to (3.16), where the three parameters κ , \bar{v} , σ_v are now κ^{H} , \bar{v}^{H} , σ_v^{H} , and the term $\frac{1}{2} v^2 \sigma_v^2 V_{vv}$, is now $\frac{1}{2} (\sigma_v^{\text{H}})^2 v V_{vv}$.

The optimal controls that result from (A.27) have the same functional form as those we derived above in Proposition 1. In A.2.3 we discuss the performance of the strategy for these two additional models of volatility.

A.2.3 Performance of the optimal strategy for alternative models of volatility

Table A.3 presents the performance of the optimal strategy for three volatility models of the micro-exchange rate. We observe small differences with respect to the baseline model in row (b). Finally, Table A.4 summarizes the results of Table A.3 using average figures.

		Missed by naive, filled by alternative strategies				
		Extra Fills	Extra Volume $\times 10^6$	Cost in JPY	Cost 20ms later in JPY	Cost 100ms later in JPY
(a)	$\delta_t = \infty, \forall t$	9,483	4,222	8,914,046	9,841,755	12,158,945
	$\delta^*, \hat{\delta} = 0$				$\delta_t = \infty, \forall t$	$\delta_t = \infty, \forall t$
	Volatility model					
(b)	Equation (3.6)	8,318	3,718	6,528,965	7,459,198	9,601,222
(c)	Equation (3.7)	8,434	3,792	6,711,593	7,663,781	9,868,183
(d)	Equation (3.8)	8,434	3,713	6,529,054	7,465,914	9,593,287

Table A.3: Performance of optimal strategy for T1 for three volatility models of the micro-exchange rate. The total number of trade attempts is 110,458. Cost-related quantities are in JPY. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies					
		Extra Fills	Avg. cost in ticks	Avg. cost 20ms later in ticks	Avg. cost 100ms later in ticks	Cost increase 20ms later %	Cost increase 100ms later %
(a)	$\delta_t = \infty, \forall t$	9,483	2.11	2.33	2.88	10%	36%
	$\delta^*, \hat{\delta} = 0$					$\delta_t = \infty, \forall t$	$\delta_t = \infty, \forall t$
	Parameters						
(b)	Equation (3.6)	8,318	1.76	2.01	2.58	14%	47%
(c)	Equation (3.7)	8,395	1.77	2.02	2.60	14%	47%
(d)	Equation (3.8)	8,307	1.76	2.01	2.58	14%	47%

Table A.4: Performance of optimal strategy for T1 under different models for the volatility of the micro-exchange rate. The total number of trade attempts is 110,458. Cost-related quantities are weighted by volume and expressed in ticks, 10^{-3} JPY. The metrics in column 6 (*resp.* 7) are the proportional increase of the values in columns 5 and 6 (*resp.* 5 and 7) in Table 3.3. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

A.2.4 Parameter estimates and sensitivity analysis

Recall that we estimate daily parameters of the model for volatility during the period from December 2016 to March 2017, i.e., 83 trading days. In total, we employ around 170,000 observations of the micro-exchange rate because we sample prices every 500 milliseconds, and around 2,800 observations for the fill ratio because we sample fills and misses every 30 seconds. Here, we use measure F_n for T1, and measure F_n^{AV} for T2. In our analysis, the OLS regressions are not restricted to obtain non-negative parameter estimates. In Tables A.5 and A.6 we report the number of occurrences when a parameter estimate is negative. In these few cases, to run the strategy, we replace the set of parameters with those of the previous day.

Volatility model	(3.6): $\hat{\kappa}, \hat{v}, \hat{\sigma}_v$	(3.7): $\hat{\kappa}^{OU}, \hat{v}^{OU}, \hat{\sigma}_v^{OU}$	(3.8): $\hat{\kappa}^H, \hat{v}^H, \hat{\sigma}_v^H$
Days with at least one estimated negative parameter	1	7	7

Table A.5: Number of days with at least one negative parameter estimate (out of 83 days).

Table A.6 shows the number of days in which the parameter estimates of the dynamics of fill ratio dynamics are negative.

	\hat{a}	\hat{b}	$\hat{\lambda}$	\hat{F}	$\hat{\sigma}_F$
Days with at least one estimated negative parameter	1	12	1	1	1

Table A.6: Number of days with at least one negative parameter estimate (out of 83 days)

A.2.4.1 Sensitivity to parameter estimates

In Tables A.7 and A.8 we show the performance of the optimal strategy δ^* when we stress the parameter estimates of the volatility model (3.6). For each day, we estimate the parameters $\hat{\kappa}$, \hat{v} , and $\hat{\sigma}_v$ as described in Section 3.4 and analyze eight scenarios, where we shock each parameter by multiplying it by 1.2 and 0.80, See subsection 3.5.3. Next, we compute the optimal strategy δ^* with the eight triplets of parameters resulting from all the different combinations of $(\hat{\kappa}, \hat{v}, \hat{\sigma}_v) \in \{\hat{\kappa}^d, \hat{\kappa}^u\} \times \{\hat{v}^d, \hat{v}^u\} \times \{\hat{\sigma}_v^d, \hat{\sigma}_v^u\}$.

Tables A.7 and A.8 are the analogue of Tables 3.3 and 3.4.

		Missed by naive, filled by alternative strategies				
		Extra Fills	Extra Volume $\times 10^6$	Cost in JPY	Cost 20ms later in JPY	Cost 100ms later in JPY
(a)	$\delta_t = \infty, \forall t$	9,483	4,222	8,914,046	9,841,755	12,158,945
$\delta^*, \hat{\delta} = 0$					$\delta_t = \infty, \forall t$	$\delta_t = \infty, \forall t$
(b)	Parameters $\hat{\kappa}, \hat{v}, \hat{\sigma}_v$	8,318	3,718	6,528,965	7,459,198	9,601,222
(b.1)	$\hat{\kappa}^d, \hat{v}^d, \hat{\sigma}_v^d$	8,434	3,765	6,633,578	7,579,730	9,741,145
(b.2)	$\hat{\kappa}^d, \hat{v}^d, \hat{\sigma}_v^u$	8,434	3,765	6,633,578	7,579,730	9,741,145
(b.3)	$\hat{\kappa}^d, \hat{v}^u, \hat{\sigma}_v^d$	8,225	3,688	6,507,051	7,421,755	9,541,736
(b.4)	$\hat{\kappa}^d, \hat{v}^u, \hat{\sigma}_v^u$	8,225	3,688	6,507,051	7,421,755	9,541,736
(b.5)	$\hat{\kappa}^u, \hat{v}^d, \hat{\sigma}_v^d$	8,041	3,757	6,621,327	7,563,658	9,722,004
(b.6)	$\hat{\kappa}^u, \hat{v}^d, \hat{\sigma}_v^u$	8,041	3,757	6,621,327	7,563,658	9,722,004
(b.7)	$\hat{\kappa}^u, \hat{v}^u, \hat{\sigma}_v^d$	7,839	3,683	6,494,442	7,408,846	9,527,502
(b.8)	$\hat{\kappa}^u, \hat{v}^u, \hat{\sigma}_v^u$	7,839	3,683	6,494,442	7,408,846	9,527,502

Table A.7: Sensitivity analysis (volatility parameters) of the performance of optimal strategy for T1. The total number of trade attempts is 110,458. Cost-related quantities are in JPY. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies					
		Extra Fills	Avg. cost in ticks	Avg. cost 20ms later in ticks	Avg. cost 100ms later in ticks	Cost increase 20ms later %	Cost increase 100ms later %
(a)	$\delta_t = \infty, \forall t$	9,483	2.11	2.33	2.88	10%	36%
$\delta^*, \hat{\delta} = 0$						$\delta_t = \infty, \forall t$	$\delta_t = \infty, \forall t$
	Parameters						
(b)	$\hat{\kappa}, \hat{v}, \hat{\sigma}_v$	8,318	1.76	2.01	2.58	14%	47%
(b.1)	$\hat{\kappa}^d, \hat{v}^d, \hat{\sigma}_v^d$	8,434	1.76	2.01	2.59	14%	47%
(b.2)	$\hat{\kappa}^d, \hat{v}^d, \hat{\sigma}_v^d$	8,434	1.76	2.01	2.59	14%	47%
(b.3)	$\hat{\kappa}^d, \hat{v}^u, \hat{\sigma}_v^d$	8,225	1.76	2.01	2.59	14%	47%
(b.4)	$\hat{\kappa}^d, \hat{v}^u, \hat{\sigma}_v^d$	8,225	1.76	2.01	2.59	14%	47%
(b.5)	$\hat{\kappa}^u, \hat{v}^d, \hat{\sigma}_v^d$	8,041	1.76	2.01	2.59	14%	47%
(b.6)	$\hat{\kappa}^u, \hat{v}^d, \hat{\sigma}_v^d$	8,041	1.76	2.01	2.59	14%	47%
(b.7)	$\hat{\kappa}^u, \hat{v}^u, \hat{\sigma}_v^d$	7,839	1.76	2.01	2.59	14%	47%
(b.8)	$\hat{\kappa}^u, \hat{v}^u, \hat{\sigma}_v^d$	7,839	1.76	2.01	2.59	14%	47%

Table A.8: Sensitivity analysis (of volatility parameter) of the performance of optimal strategy for T1. The total number of trade attempts is 110,458. Cost-related quantities are weighted by volume and expressed in ticks, 10^{-3} JPY. The metrics in column 6 (*resp.* 7) are the proportional increase of the values in columns 5 and 6 (*resp.* 5 and 7) in Table 3.3. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

Appendix B

Appendix for Chapter 4

B.1 Proof of Lemma 1

We prove the lemma in three parts. First we work out the Gâteaux derivative of the cost function. We use (4.9) 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

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

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

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

Subtract and add

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

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

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

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

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

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

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

Finally, the limit of QP3 is given by

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

which concludes the proof.

B.2 Existence and uniqueness of forward and backward parts of the FBSDE

Following [Antonelli \(1993\)](#) and for mathematical interest, we study the forward and backward parts of the FBSDE separately. Thus, the following proofs are not related to the existence and uniqueness result of the full FBSDE and can be regarded as separate problems.

The following theorem shows the existence and uniqueness of the solution to the backward part of the FBSDE (4.18).

Theorem 9. Fix $V \in \mathcal{C}^\infty$. Assumptions 1 and 2 hold. Then, the functional $\Psi : \mathcal{C}^\infty \rightarrow \mathcal{C}^\infty$ given by

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

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

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

which after n iterations becomes

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

Finally,

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

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 unique¹ 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 (4.18).

Theorem 10. Fix $U \in \mathcal{C}^\infty$. Assumptions 1 and 2 hold. Then, the functional $\Theta : \mathcal{C}^\infty \rightarrow \mathcal{C}^\infty$ given by

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

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

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

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:

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

¹Unique in the sense of indistinguishability.

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:

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

By Borel-Cantelli arguments, there is $\mathfrak{D} \subset \Omega$ such that for all $\omega \in \mathfrak{D}$ the functions $t \rightarrow \Theta_t^{(n)}(w)$ 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. ■

B.3 Second Gâteaux derivative

The first Gâteaux derivative of the functional J is given by

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

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

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

B.4 Further results

Lemma 8. *There is a sub-sequence $\{\psi_n\}_{n \geq 1}$ of $\{1/n\}_{n \geq 1}$ such that*

$$\lim_{n \rightarrow \infty} \|\delta^{*\psi_n} - \delta^*\| = 0.$$

Proof. Let $\psi > 0$, the processes $\delta^{*\psi}$ and δ^* are bounded below by zero and above by $B = (B_t)_{t \geq 0}$, $B_t = 2\gamma N_{t-} + 2\gamma \bar{\lambda} T + \gamma + \alpha$. By Assumption 4.1, we have $\|B\| < \infty$ and $\|B\|_1 < \infty$. Let $b = \|B\|$, then $\{\|\delta^{*\psi_n} - \delta^*\|\}_{n \geq 1}$ is a sequence of real numbers in $[0, 2b]$, hence there is a convergent sub-sequence $\{\psi_n^0\}_{n \geq 1}$ of $\{\frac{1}{n}\}_{n \geq 1}$. By Cauchy-Schwarz (see (4.26)), furthermore, the sequence of real numbers $\left\{ \left\| \delta^{*\psi_n^0} - \delta^* \right\|_1 \right\}_{n \geq 1}$ is in $[0, 2b\sqrt{T}]$, and hence there is a sub-sequence $\{\psi_n\}_{n \geq 1}$ of $\{\psi_n^0\}_{n \geq 1}$ such that both $\|\delta^{*\psi_n} - \delta^*\|_1$ and $\|\delta^{*\psi_n} - \delta^*\|$ converge.

As $\{\psi_n\}_{n \geq 1}$ is sub-sequence of $\{\frac{1}{n}\}_{n \geq 1}$, $\psi_n \rightarrow 0$ as $n \rightarrow \infty$. Next, by the dominated convergence theorem we have that

$$\begin{aligned}
& \lim_{n \rightarrow \infty} \|\delta^{*\psi_n} - \delta^*\|^2 \\
&= \mathbb{E} \left[\int_0^T \lim_{n \rightarrow \infty} \left| \frac{\lambda_t \phi_t(\delta_t^{*\psi_n})}{\psi_n + \lambda_t \phi_t(\delta_t^{*\psi_n})} \left(2\gamma \mathbb{E}_{t-} [D_T^{\delta^{*\psi_n}}] + \gamma + \alpha \right) - \left(2\gamma \mathbb{E}_{t-} [D_T^{\delta^*}] + \gamma + \alpha \right) \right|^2 dt \right] \\
&\leq 2 \mathbb{E} \left[\int_0^T \lim_{n \rightarrow \infty} \left\{ 4\gamma^2 \left| \frac{\lambda_t \phi_t(\delta_t^{*\psi_n})}{\psi_n + \lambda_t \phi_t(\delta_t^{*\psi_n})} \mathbb{E}_{t-} [D_T^{\delta^{*\psi_n}}] - \mathbb{E}_{t-} [D_T^{\delta^*}] \right|^2 \right. \right. \\
&\quad \left. \left. + \left| (\gamma + \alpha) \frac{\psi_n}{\psi_n + \lambda_t \phi_t(\delta_t^{*\psi_n})} \right|^2 \right\} dt \right] \\
&= 8\gamma^2 \mathbb{E} \left[\int_0^T \lim_{n \rightarrow \infty} \left| \frac{\lambda_t \phi_t(\delta_t^{*\psi_n})}{\psi_n + \lambda_t \phi_t(\delta_t^{*\psi_n})} \mathbb{E}_{t-} [D_T^{\delta^{*\psi_n}}] - \mathbb{E}_{t-} [D_T^{\delta^*}] \right|^2 dt \right] \\
&\leq 16\gamma^2 \mathbb{E} \left[\int_0^T \lim_{n \rightarrow \infty} \left\{ \left| \frac{\lambda_t \phi_t(\delta_t^{*\psi_n})}{\psi_n + \lambda_t \phi_t(\delta_t^{*\psi_n})} \mathbb{E}_{t-} [D_T^{\delta^{*\psi_n}} - D_T^{\delta^*}] \right|^2 + \left| \frac{\psi_n}{\psi_n + \lambda_t \phi_t(\delta_t^{*\psi_n})} \mathbb{E}_{t-} [D_T^{\delta^*}] \right|^2 \right\} dt \right] \\
&= 16\gamma^2 \mathbb{E} \left[\int_0^T \lim_{n \rightarrow \infty} \left| \frac{\lambda_t \phi_t(\delta_t^{*\psi_n})}{\psi_n + \lambda_t \phi_t(\delta_t^{*\psi_n})} \mathbb{E}_{t-} [D_T^{\delta^{*\psi_n}} - D_T^{\delta^*}] \right|^2 dt \right]
\end{aligned}$$

Next, writing the counting processes in terms of the random measures, we have

$$\begin{aligned}
& \lim_{n \rightarrow \infty} \|\delta^{*\psi_n} - \delta^*\|^2 \\
&\leq 16\gamma^2 \mathbb{E} \left[\int_0^T \lim_{n \rightarrow \infty} \left| \mathbb{E}_{t-} [D_T^{\delta^{*\psi_n}} - D_T^{\delta^*}] \right|^2 dt \right] \\
&\leq 16\gamma^2 \mathbb{E} \left[\int_0^T \lim_{n \rightarrow \infty} \int_0^T \int_{\mathbb{R}} |G(\delta_s^{*\psi_n} - z) - G(\delta_s^* - z)|^2 \mathbf{p}(dz, ds) dt \right] \\
&\leq 16\gamma^2 T \mathbb{E} \left[\lim_{n \rightarrow \infty} \int_0^T \int_{\mathbb{R}} |G(\delta_t^{*\psi_n} - z) - G(\delta_t^* - z)| \phi_t(dz) \lambda_t dt \right] \\
&\leq 16\gamma^2 T \bar{\lambda} \mathbb{E} \left[\lim_{n \rightarrow \infty} \int_0^T |\Phi_t(\delta_t^{*\psi_n}) - \Phi_t(\delta_t^*)| dt \right] \\
&\stackrel{(\text{Lipschitz cont.})}{\leq} 16\gamma^2 T \bar{\lambda} k \lim_{n \rightarrow \infty} \mathbb{E} \left[\int_0^T |\delta_t^{*\psi_n} - \delta_t^*| dt \right] \\
&= 16\gamma^2 T \bar{\lambda} k \lim_{n \rightarrow \infty} \|\delta^{*\psi_n} - \delta^*\|_1.
\end{aligned}$$

Following similar techniques, we further have

$$\begin{aligned}
& \lim_{n \rightarrow \infty} \|\delta^{*\psi_n} - \delta^*\|_1 \\
&= \mathbb{E} \left[\int_0^T \lim_{n \rightarrow \infty} \left| \frac{\lambda_t \phi_t(\delta_t^{*\psi_n})}{\psi_n + \lambda_t \phi_t(\delta_t^{*\psi_n})} \left(2\gamma \mathbb{E}_{t-} [D_T^{\delta^{*\psi_n}}] + \gamma + \alpha \right) - \left(2\gamma \mathbb{E}_{t-} [D_T^{\delta^*}] + \gamma + \alpha \right) \right| dt \right] \\
&\leq 2\gamma \mathbb{E} \left[\int_0^T \lim_{n \rightarrow \infty} \left| \frac{\lambda_t \phi_t(\delta_t^{*\psi_n})}{\psi_n + \lambda_t \phi_t(\delta_t^{*\psi_n})} \mathbb{E}_{t-} [D_T^{\delta^{*\psi_n}} - D_T^{\delta^*}] \right| dt \right] \\
&\leq 2\gamma \bar{\lambda} T k \lim_{n \rightarrow \infty} \mathbb{E} \left[\int_0^T |\delta_t^{*\psi_n} - \delta_t^*| dt \right] \\
&= 2\gamma \bar{\lambda} T k \lim_{n \rightarrow \infty} \|\delta^{*\psi_n} - \delta^*\|_1.
\end{aligned}$$

As consequence of (4.20), $2\gamma\bar{\lambda}Tk < 1$, hence,

$$\lim_{n \rightarrow \infty} \|\delta^{*\psi_n} - \delta^*\|_1 = 0,$$

and using this in the previous limit, we have

$$\lim_{n \rightarrow \infty} \|\delta^{*\psi_n} - \delta^*\|^2 \leq 16\gamma^2 T \bar{\lambda} k \lim_{n \rightarrow \infty} \|\delta^{*\psi_n} - \delta^*\|_1 = 0,$$

which concludes the proof. ■

Appendix C

Appendix for Chapter 5

C.1 Random latency

The set L_k describes the set of pending orders – if any. In the interest of a more compact notation, we let $L_0 = \emptyset$ and $L_1 = L$, and we drop the argument of a mathematical object if it is the empty set. For example, we write \mathcal{A}_t instead of $\mathcal{A}_{t,\emptyset}$, or we write $\mathbf{X}^{t,\mathbf{x},\alpha}$ instead of $\mathbf{X}^{t,\mathbf{x},\emptyset,\alpha}$, etc. Finally, when the index of a function is the empty set, the function returns the empty set.¹

For $k \in \{0, 1\}$, the domains of the value functions v_k are

$$\mathcal{D}_k = \{(t, \mathbf{x}, l) : (t, \mathbf{x}) \in [0, T] \times \mathbb{R} \times (-\infty, \mathfrak{M}] \times \mathbb{R} \text{ and } l \in L_k\}. \quad (\text{C.1})$$

C.1.1 Dynamic programming principle

Here, we prove the dynamic programming principle satisfied by the value functions $v_k(t, \mathbf{x}, l)$, $k \in \{0, 1\}$, $t \in [0, T]$, and $l \in L_k$. Before proceeding, we introduce the function $\iota(t, \alpha)$, which returns the index of the pending order (if any) at time t and strategy α . That is,

$$\iota(t, \alpha) = \begin{cases} i & \text{if there is a pending order, i.e., } i \text{ s.t. } \tau_i \leq t \text{ and } \tilde{\tau}_i > t, \\ \emptyset & \text{otherwise.} \end{cases}$$

Theorem 11. *Dynamic programming principle. The value functions in (5.15) and (5.16) satisfy the dynamic programming principle. That is, for $k \in \{0, 1\}$, $(t, \mathbf{x}, l) \in \mathcal{D}_k$ (see (C.1)), and $\theta \in \mathcal{T}_{t,T}$ – where $\mathcal{T}_{t,T}$ is the set of stopping times valued in $[t, T]$, and $v_{k(\theta, \alpha)}(\theta, \dots)$ Borel measurable, we have:*

$$v_k(t, \mathbf{x}, l) = \sup_{\alpha \in \mathcal{A}_{t,l}} \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t,\mathbf{x},l,\alpha}) ds + \sum_{\tilde{\tau}_i \leq \theta} \mathfrak{C}(\mathbf{X}_{\tilde{\tau}_i}^{t,\mathbf{x},l,\alpha}, \mathbf{F}_{\tilde{\tau}_i}^{t,0}, \mathfrak{l}_i) + v_{k(\theta, \alpha)}(\theta, \mathbf{X}_\theta^{t,\mathbf{x},l,\alpha}, \mathfrak{l}_{\iota(\theta, \alpha)}) \right],$$

where \mathfrak{l}_\emptyset is \emptyset . In other words, for $k \in \{0, 1\}$, $(t, \mathbf{x}, l) \in \mathcal{D}_k$:

(a) **DPP1.** For all $\alpha \in \mathcal{A}_{t,l}$, we have

$$v_k(t, \mathbf{x}, l) \leq \sup_{\alpha \in \mathcal{A}_{t,l}} \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t,\mathbf{x},l,\alpha}) ds + \sum_{\tilde{\tau}_i \leq \theta} \mathfrak{C}(\mathbf{X}_{\tilde{\tau}_i}^{t,\mathbf{x},l,\alpha}, \mathbf{F}_{\tilde{\tau}_i}^{t,0}, \mathfrak{l}_i) + v_{k(\theta, \alpha)}(\theta, \mathbf{X}_\theta^{t,\mathbf{x},l,\alpha}, \mathfrak{l}_{\iota(\theta, \alpha)}) \right]$$

for all stopping times θ valued in $[t, T]$.

(b) **DPP2.** For all $\epsilon > 0$, there exists $\alpha \in \mathcal{A}_{t,l}$, such that for all stopping times θ valued in $[t, T]$,

$$v_k(t, \mathbf{x}, l) \geq \sup_{\alpha \in \mathcal{A}_{t,l}} \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t,\mathbf{x},l,\alpha}) ds + \sum_{\tilde{\tau}_i \leq \theta} \mathfrak{C}(\mathbf{X}_{\tilde{\tau}_i}^{t,\mathbf{x},l,\alpha}, \mathbf{F}_{\tilde{\tau}_i}^{t,0}, \mathfrak{l}_i) + v_{k(\theta, \alpha)}(\theta, \mathbf{X}_\theta^{t,\mathbf{x},l,\alpha}, \mathfrak{l}_{\iota(\theta, \alpha)}) \right] - \epsilon.$$

¹Recall that in our construction we enhanced the state space to work with a Markovian system. As a consequence of the memoryless property of the exponential random variable, we do not need to keep track of the times at which a pending order was sent, but we do need to know the limit price of the order that was sent (if any).

Before proceeding to the proof, we remark that from the dynamics of the controlled system \mathbf{X}^α in (5.9), the independence of the marks, and the memoryless property of the exponential distribution of latency, we have the following properties:

1. Markov property of $(\mathbf{X}^\alpha, k(\cdot, \alpha), \mathbf{l}_{\iota(\cdot, \alpha)})$. For any $\alpha \in \mathcal{A}$, we have that

$$\mathbb{E} [\psi(\mathbf{X}_{\theta_2}^\alpha) | \mathcal{F}_{\theta_1}] = \mathbb{E} [\psi(\mathbf{X}_{\theta_2}^\alpha) | (\mathbf{X}_{\theta_1}^\alpha, k(\theta_1, \alpha), \mathbf{l}_{\iota(\theta_1, \alpha)})], \quad (\text{C.2})$$

for any bounded measurable function ψ , and stopping times $\theta_1 \leq \theta_2$ a.s..

2. Causality of the control. For any $\alpha = (\tau_i, \mathbf{l}_i)_{i \geq 1} \in \mathcal{A}$, and θ a stopping time,

$$\alpha^\theta \in \mathcal{A}_{\theta, \mathbf{l}_{\iota(\theta, \alpha)}} \quad \text{and} \quad \mathbf{l}_{\iota(\theta, \alpha)} \in L_{k(\theta, \alpha)}, \quad (\text{C.3})$$

where $\alpha^\theta = (\tau_{i+\iota(\theta, \alpha)}, \mathbf{l}_{i+\iota(\theta, \alpha)})_{i \geq 1}$.

3. Pathwise uniqueness of the state process:

$$\mathbf{X}^{t, \mathbf{x}, l, \alpha} = \mathbf{X}^\theta, \mathbf{X}_\theta^{t, \mathbf{x}, l, \alpha, \mathbf{l}_{\iota(\theta, \alpha)}}, \alpha^\theta \quad (\text{C.4})$$

on $[0, T]$, for any $(t, \mathbf{x}, l) \in \mathcal{D}_k$, $k = 0, 1$, $\alpha \in \mathcal{A}_{t, l}$, and $\theta \in \mathcal{T}_{t, T}$.

Proof. DPP1: Fix $(t, \mathbf{x}, l) \in \mathcal{D}_k$, $k \in \{0, 1\}$ and take an arbitrary control $\alpha \in \mathcal{A}_{t, l}$. By the law of iterated conditional expectations we have

$$\begin{aligned} J_k(t, \mathbf{x}, l, \alpha) &= \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t, \mathbf{x}, l, \alpha}) ds + \sum_{\bar{\tau}_i \leq \theta} \mathfrak{C}(\mathbf{X}_{\bar{\tau}_i}^{t, \mathbf{x}, l, \alpha}, \mathbf{F}_{\bar{\tau}_i}^{t, 0}, \mathbf{l}_i) \right. \\ &\quad \left. + \mathbb{E} \left[\int_\theta^T h(\mathbf{X}_s^{t, \mathbf{x}, l, \alpha}) ds + \sum_{\theta < \bar{\tau}_i \leq T} \mathfrak{C}(\mathbf{X}_{\bar{\tau}_i}^{t, \mathbf{x}, l, \alpha}, \mathbf{F}_{\bar{\tau}_i}^{t, 0}, \mathbf{l}_i) + g(\mathbf{X}_T^{t, \mathbf{x}, l, \alpha}) \middle| \mathcal{F}_\theta \right] \right], \end{aligned}$$

and because of the joint Markov property of $(\mathbf{X}^\alpha, k(\cdot, \alpha), \mathbf{l}_{\iota(\cdot, \alpha)})$, the causality of the control, and the pathwise uniqueness of $\mathbf{X}^{t, \mathbf{x}, l, \alpha}$ (see (C.2), (C.3), (C.4)), we have

$$\begin{aligned} J_k(t, \mathbf{x}, l, \alpha) &= \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t, \mathbf{x}, l, \alpha}) ds + \sum_{\bar{\tau}_i \leq \theta} \mathfrak{C}(\mathbf{X}_{\bar{\tau}_i}^{t, \mathbf{x}, l, \alpha}, \mathbf{F}_{\bar{\tau}_i}^{t, 0}, \mathbf{l}_i) + J_{k(\theta, \alpha)}(\theta, \mathbf{X}_\theta^{t, \mathbf{x}, l, \alpha}, \mathbf{l}_{\iota(\theta, \alpha)}, \alpha^\theta) \right] \\ &\leq \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t, \mathbf{x}, l, \alpha}) ds + \sum_{\bar{\tau}_i \leq \theta} \mathfrak{C}(\mathbf{X}_{\bar{\tau}_i}^{t, \mathbf{x}, l, \alpha}, \mathbf{F}_{\bar{\tau}_i}^{t, 0}, \mathbf{l}_i) + v_{k(\theta, \alpha)}(\theta, \mathbf{X}_\theta^{t, \mathbf{x}, l, \alpha}, \mathbf{l}_{\iota(\theta, \alpha)}) \right]. \end{aligned}$$

Take supremum over $\alpha \in \mathcal{A}_{t, l}$ on both sides of the inequality above to obtain the desired result.

DPP2: Fix $(t, \mathbf{x}, l) \in \mathcal{D}_k$, $k = 0, 1$, and take an arbitrary control $\alpha \in \mathcal{A}_{t, l}$. By the definition of v_k , for any $\epsilon > 0$ and $\omega \in \Omega$, there exists $\alpha_{\epsilon, \omega} \in \mathcal{A}_{\theta(\omega), \mathbf{l}_{\iota(\theta(\omega), \alpha(\omega))}}$, which is an ϵ -optimal control for $v_{k(\theta(\omega), \alpha(\omega))}$ at the point $(\theta(\omega), \mathbf{X}_{\theta(\omega)}^{t, \mathbf{x}, l, \alpha}, \mathbf{l}_{\iota(\theta(\omega), \alpha(\omega))})$. Next, by a measurable selection argument,² one can show that there is $\bar{\alpha}_\epsilon \in \mathcal{A}_{\theta, \mathbf{l}_{\iota(\theta, \alpha)}}$ such that $\bar{\alpha}_\epsilon(\omega) = \alpha_{\epsilon, \omega}(\omega)$ for almost all $\omega \in \Omega$, so that

$$v_{k(\theta, \alpha)}(\theta, \mathbf{X}_\theta^{t, \mathbf{x}, l, \alpha}, \mathbf{l}_{\iota(\theta, \alpha)}) - \epsilon \leq J_{k(\theta, \alpha)}(\theta, \mathbf{X}_\theta^{t, \mathbf{x}, l, \alpha}, \mathbf{l}_{\iota(\theta, \alpha)}, \bar{\alpha}_\epsilon). \quad (\text{C.5})$$

Next, concatenate the controls α and $\bar{\alpha}_\epsilon$; that is, take $\bar{\alpha}$ to be α for the impulses that are before (or include) time θ , and take $\bar{\alpha}_\epsilon$ for the impulses that are strictly after time θ . By construction $\bar{\alpha} \in \mathcal{A}_{t, l}$, so $\forall u \in [t, \theta]$ $\mathbf{X}_u^{t, \mathbf{x}, l, \alpha} = \mathbf{X}_u^{t, \mathbf{x}, l, \bar{\alpha}}$, $k(\theta, \alpha) = k(\theta, \bar{\alpha})$, $\mathbf{l}_{\iota(\theta, \alpha)} = \mathbf{l}_{\iota(\theta, \bar{\alpha})}$, and $\bar{\alpha}^\theta = \bar{\alpha}_\epsilon$. Therefore, following the same steps as in the proof of DPP1, we have

$$J_k(t, \mathbf{x}, l, \bar{\alpha}) = \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t, \mathbf{x}, l, \bar{\alpha}}) ds + \sum_{\bar{\tau}_i \leq \theta} \mathfrak{C}(\mathbf{X}_{\bar{\tau}_i}^{t, \mathbf{x}, l, \bar{\alpha}}, \mathbf{F}_{\bar{\tau}_i}^{t, 0}, \mathbf{l}_i) + J_{k(\theta, \bar{\alpha})}(\theta, \mathbf{X}_\theta^{t, \mathbf{x}, l, \bar{\alpha}}, \mathbf{l}_{\iota(\theta, \bar{\alpha})}, \bar{\alpha}_\epsilon) \right],$$

²See Chapter 7 in Bertsekas and Shreve (1978). As mentioned before, we do not present the proof of this step. One could also follow the weak dynamic programming principle for viscosity solutions by Bouchard and Touzi (2011).

which together with (C.5) yields

$$J_k(t, \mathbf{x}, l, \bar{\alpha}) \geq \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t, \mathbf{x}, l, \bar{\alpha}}) ds + \sum_{\tilde{\tau}_i \leq \theta} \mathfrak{C}(\mathbf{X}_{\tilde{\tau}_i}^{t, \mathbf{x}, l, \bar{\alpha}}, F_{\tilde{\tau}_i}^{t, 0}, \mathfrak{l}_i) + v_{k(\theta, \bar{\alpha})}(\theta, X_\theta^{t, \mathbf{x}, l, \bar{\alpha}}, \mathfrak{l}_{\ell(\theta, \bar{\alpha})}) \right] - \epsilon,$$

where $\bar{\alpha}$ can be replaced by α because $\bar{\alpha}$ is less than or equal to θ . Take supremum on both sides of the inequality above to conclude that

$$v_k(t, \mathbf{x}, l) \geq \sup_{\alpha \in \mathcal{A}_{t, l}} \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t, \mathbf{x}, l, \alpha}) ds + \sum_{\tilde{\tau}_i \leq \theta} \mathfrak{C}(\mathbf{X}_{\tilde{\tau}_i}^{t, \mathbf{x}, l, \alpha}, F_{\tilde{\tau}_i}^{t, 0}, \mathfrak{l}_i) + v_{k(\theta, \alpha)}(\theta, X_\theta^{t, \mathbf{x}, l, \alpha}, \mathfrak{l}_{\ell(\theta, \alpha)}) \right] - \epsilon,$$

which completes the proof because ϵ is arbitrary. ■

We characterize the value functions with the following two corollaries, which are stated without proof (they follow from Theorem 11). In what follows, Z is a random variable with law $\nu(dx)$ as in (5.3), independent of the Brownian motion W and of the exponentially distributed delays between trade attempts and notifications.

Corollary 2. For $k = 1$, $(t, \mathbf{x}, l) \in \mathcal{D}_1$, $p \in \mathbb{R}$ and $\theta \in \mathcal{T}_{t, T}$, we have

$$\begin{aligned} v_1(t, \mathbf{x}, l) &= \mathbb{E} \left[\int_t^{\theta \wedge \tilde{t}} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + \left(v_0(\tilde{t}, \Gamma(\mathbf{X}_{\tilde{t}-}^{t, \mathbf{x}}, F_{\tilde{t}}^{t, 0}, l)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}-}^{t, \mathbf{x}}, F_{\tilde{t}}^{t, 0}, l) \right) \mathbf{1}_{\tilde{t} \leq \theta} + v_1(\theta, \mathbf{X}_\theta^{t, \mathbf{x}}, l) \mathbf{1}_{\tilde{t} > \theta} \right] \\ &= \mathbb{E} \left[\int_t^{\theta \wedge \tilde{t}} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + \left(v_0(\tilde{t}, \Gamma(\mathbf{X}_{\tilde{t}-}^{t, \mathbf{x}}, Z, l)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}-}^{t, \mathbf{x}}, Z, l) \right) \mathbf{1}_{\tilde{t} \leq \theta} + v_1(\theta, \mathbf{X}_\theta^{t, \mathbf{x}}, l) \mathbf{1}_{\tilde{t} > \theta} \right]. \end{aligned}$$

In particular, when $\theta = T$ characterize v_1 in terms of v_0 :

$$v_1(t, \mathbf{x}, l) = \mathbb{E} \left[\int_t^{T \wedge \tilde{t}} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + \left(v_0(\tilde{t}, \Gamma(\mathbf{X}_{\tilde{t}-}^{t, \mathbf{x}}, Z, l)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}-}^{t, \mathbf{x}}, Z, l) \right) \mathbf{1}_{\tilde{t} \leq T} + v_0(T, \mathbf{X}_T^{t, \mathbf{x}}) \mathbf{1}_{\tilde{t} > T} \right], \quad (\text{C.6})$$

and $v_1(T, \mathbf{x}, l) = v_0(T, \mathbf{x}) = c + q(s + \zeta - a(q - 1))(1 - c)$, $\forall l \in L$.

Corollary 3. For $(t, \mathbf{x}) \in \mathcal{D}_0$, and $\theta \in \mathcal{T}_{t, T}$, we have the following relationship between v_0 and v_1 :

$$v_0(t, \mathbf{x}) = \sup_{(\tau, \mathfrak{l}) \in \mathcal{I}_t} \mathbb{E} \left[\int_t^{\theta \wedge \tau} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + v_0(\theta, \mathbf{X}_\theta^{t, \mathbf{x}}) \mathbf{1}_{\theta < \tau} + v_1(\tau, \mathbf{X}_\tau^{t, \mathbf{x}}, \mathfrak{l}) \mathbf{1}_{\tau \leq \theta} \right], \quad (\text{C.7})$$

where \mathcal{I}_t is the set of pairs (τ, \mathfrak{l}) , and where τ is a stopping time, $t \leq \tau < T$, or $\tau = \infty$ a.s., and \mathfrak{l} is an \mathcal{F}_τ -measurable random variable in L .

The following two corollaries are crucial to prove the subsolution property of the value function – see Theorem 13.

Corollary 4. For $(t, \mathbf{x}) \in \mathcal{D}_0$ and $\theta \in \mathcal{T}_{t, T}$, we have the following relationship between v_0 and v_1 :

$$v_0(t, \mathbf{x}) = \sup_{(\tau, \mathfrak{l}) \in \mathcal{I}_t} \mathbb{E} \left[\int_t^{\theta \wedge \tilde{\tau}} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + v_0(\theta, \mathbf{X}_\theta^{t, \mathbf{x}}) \mathbf{1}_{\theta < \tau} + v_1(\theta, \mathbf{X}_\theta^{t, \mathbf{x}}, \mathfrak{l}) \mathbf{1}_{\tau \leq \theta < \tilde{\tau}} + v_0(\tilde{\tau}, \Gamma(\mathbf{X}_{\tilde{\tau}-}^{t, \mathbf{x}}, Z, \mathfrak{l})) \mathbf{1}_{\tau < \tilde{\tau} < \theta} \right].$$

Corollary 5. For $(t, \mathbf{x}) \in \mathcal{D}_0$ and $\theta \in \mathcal{T}_{t, T}$, we have

$$\begin{aligned} v_0(t, \mathbf{x}) &= \sup_{(\tau, \mathfrak{l}) \in \mathcal{I}_t} \mathbb{E} \left[\left(\int_t^\theta h(\mathbf{X}_s^{t, \mathbf{x}}) ds + v_0(\theta, \mathbf{X}_\theta^{t, \mathbf{x}}) \right) \mathbf{1}_{\theta < \tau} \right. \\ &\quad \left. + \left(\int_t^{\theta \wedge \tilde{\tau}} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + \left(v_0(\tilde{\tau}, \Gamma(\mathbf{X}_{\tilde{\tau}-}^{t, \mathbf{x}}, Z, \mathfrak{l})) + \mathfrak{C}(\mathbf{X}_{\tilde{\tau}-}^{t, \mathbf{x}}, Z, \mathfrak{l}) \right) \mathbf{1}_{\tilde{\tau} \leq \theta} + v_1(\theta, \mathbf{X}_\theta^{t, \mathbf{x}}) \mathbf{1}_{\tilde{\tau} > \theta} \right) \mathbf{1}_{\tau \leq \theta} \right], \end{aligned}$$

as a result of property (C.4) and Corollaries 2, 3.

In particular, for $\theta = T$ we have

$$\begin{aligned} v_0(t, \mathbf{x}) &= \sup_{(\tau, \mathfrak{l}) \in \mathcal{I}_t} \mathbb{E} \left[\left(\int_t^T h(\mathbf{X}_s^{t, \mathbf{x}}) ds + v_0(T, \mathbf{X}_T^{t, \mathbf{x}}) \right) \mathbf{1}_{T < \tau} \right. \\ &\quad \left. + \left(\int_t^{T \wedge \tilde{\tau}} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + \left(v_0(\tilde{\tau}, \Gamma(\mathbf{X}_{\tilde{\tau}-}^{t, \mathbf{x}}, Z, \mathfrak{l})) + \mathfrak{C}(\mathbf{X}_{\tilde{\tau}-}^{t, \mathbf{x}}, Z, \mathfrak{l}) \right) \mathbf{1}_{\tilde{\tau} \leq T} + v_0(T, \mathbf{X}_T^{t, \mathbf{x}}) \mathbf{1}_{\tilde{\tau} > T} \right) \mathbf{1}_{\tau \leq T} \right]. \end{aligned}$$

C.1.2 PDE system viscosity characterization

In Corollaries 2 and 3 we send θ to t to obtain the following characterization of the value functions v_0 and v_1 :

Theorem 12. *Dynamic programming equation. The value functions satisfy:*

$$\min \left\{ -\frac{\partial v_0}{\partial t}(t, \mathbf{x}) - \mathcal{L}v_0(t, \mathbf{x}) - h(\mathbf{x}), v_0(t, \mathbf{x}) - \sup_{l \in L} v_1(t, \mathbf{x}, l) \right\} = 0, \quad \text{on } \mathcal{D}_0, \quad (\text{C.8})$$

and

$$\frac{\partial v_1}{\partial t}(t, \mathbf{x}, l) + \mathcal{L}v_1(t, \mathbf{x}, l) + h(\mathbf{x}) - \lambda v_1(t, \mathbf{x}, l) + \lambda \mathbb{E} \left[v_0(t, \Gamma(\mathbf{x}, Z, l)) + \mathfrak{C}(\mathbf{x}, Z, l) \right] = 0, \quad \text{on } \mathcal{D}_1, \quad (\text{C.9})$$

where \mathcal{L} is the infinitesimal generator

$$\mathcal{L} \cdot = \mu(s) \frac{\partial \cdot}{\partial s} + \frac{1}{2} \sigma^2(s) \frac{\partial^2 \cdot}{\partial s^2},$$

the random variable Z has law $\nu(dx)$ as in (5.3), and the terminal conditions are

$$v_1(T, \mathbf{x}, l) = v_0(T, \mathbf{x}) = c + q(s + \zeta - a(q - 1))(1 - \mathfrak{c}), \quad \forall l \in L. \quad (\text{C.10})$$

The intuition for the result in the theorem above is as follows. With no pending orders, the trader will submit an order with limit rate l^* when it is better to have one order pending in the exchange than none; this decision is determined by the QVI in (C.8). The crucial tradeoff is that once the agent sends an order, she must wait until it is processed by the exchange before sending another order; thus, a pending order precludes the trader from taking advantage of any opportunities that arise while a pending order is being processed during the period of latency. Finally, the trader's value function satisfies the dynamic programming equation in (C.9) when there is an order pending in the exchange.

A priori, in equation C.8, it is not clear that the candidate optimal control is well defined. If $v_1(t, \mathbf{x}, l)$ is continuous in l , then, given that L is compact, we can define the candidate optimal control as $l^* = \arg\max_l v_1(t, \mathbf{x}, l)$, whenever $v_0(t, \mathbf{x}) = v_1(t, \mathbf{x}, l^*)$. On the other hand, if $v_1(t, \mathbf{x}, l)$ is not continuous in l , we can take the set L to be a finite collection of points in \mathbb{R} . Recall that L is the set of possible limit rates, and limit rates in EUR/USD can take values with a precision of (at most) five decimals. Thus, we can define L to be all limit rates with five decimals of precision between zero and a fixed \bar{l} – in practice, any number that is twice the current level of the exchange rate suffices.³

We use equation (C.6) to make a key simplification to the system in (C.8) and (C.9). We plug equation (C.6) in equation (C.8) to obtain a HJBQVI that characterizes the value function v_0 . We remark that once we know v_0 we compute v_1 by means of (C.6). Equation (C.8) becomes

$$\min \left\{ -\frac{\partial v_0}{\partial t}(t, \mathbf{x}) - \mathcal{L}v_0(t, \mathbf{x}) - h(\mathbf{x}), v_0(t, \mathbf{x}) - \sup_{l \in L} \mathbb{E} \left[\int_t^{T \wedge \tilde{t}} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + (v_0(\tilde{t}, \Gamma(\mathbf{X}_{\tilde{t}-}^{t, \mathbf{x}}, Z, l)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}-}^{t, \mathbf{x}}, Z, l)) \mathbf{1}_{\tilde{t} \leq T} + v_0(T, \mathbf{X}_T^{t, \mathbf{x}}) \mathbf{1}_{\tilde{t} > T} \right] \right\} = 0, \quad (\text{C.11})$$

on \mathcal{D}_0 , with terminal condition $v_0(T, \mathbf{x}) = c + q(s + \zeta - a(q - 1))(1 - \mathfrak{c})$, which is equivalent to

$$\begin{aligned} \min \left\{ -\frac{\partial v_0}{\partial t}(t, \mathbf{x}) - \mathcal{L}v_0(t, \mathbf{x}) - h(\mathbf{x}), \right. \\ v_0(t, \mathbf{x}) + q \left\{ q \frac{\phi}{\lambda} \left(1 - e^{-\lambda(T-t)} \right) + a(q - 1) e^{-\lambda(T-t)} \right\} - q e^{-\lambda(T-t)} \mathbb{E} [S_T^{t, s}] - c e^{-\lambda(T-t)} \\ \left. - \sup_{l \in L} \mathbb{E} \left[(v_0(\tilde{t}, \Gamma(\mathbf{X}_{\tilde{t}-}^{t, \mathbf{x}}, Z, l)) - \mathfrak{c}(S_{\tilde{t}-}^{t, s} + Z)) f(l - S_{\tilde{t}-}^{t, s} - Z) \right] \mathbf{1}_{\tilde{t} \leq T} \right\} = 0. \end{aligned}$$

³Note that all financial instruments trade in a discretised grid according to the tick size.

We propose the ansatz $v_0(t, s, q, c) = c + h_0(t, s, q)$ to reduce the system (C.11) to

$$\begin{aligned} \min \bigg\{ & -\frac{\partial h_0}{\partial t}(t, s, q) - \mathcal{L}h_0(t, s, q) + \phi q^2, \\ & h_0(t, s, q) + q \left\{ q \frac{\phi}{\lambda} \left(1 - e^{-\lambda(T-t)} \right) + a(q-1) e^{-\lambda(T-t)} \right\} - q e^{-\lambda(T-t)} \mathbb{E} [S_T^{t,s}] \\ & - \sup_{l \in L} \mathbb{E} \left[\left((1-\mathfrak{c}) (S_{\tilde{t}^-}^{t,s} + Z) f(l - S_{\tilde{t}^-}^{t,s} - Z) + h_0(\tilde{t}, \tilde{\Gamma}(S_{\tilde{t}^-}^{t,s}, Z, q, l)) \right) \mathbf{1}_{\tilde{t} \leq T} \right] \bigg\} = 0, \end{aligned}$$

with terminal condition $h_0(T, s, q) = (q s + q \zeta - a q (q-1)) (1-\mathfrak{c})$, and

$$\tilde{\Gamma}(s, Z, q, l) = (s - \kappa f(l - s - Z), q - f(l - s - Z)). \quad (\text{C.12})$$

C.1.3 Viscosity properties of the value function v_0

The value functions v_0 and v_1 do not need to be smooth a priori. We focus on the viscosity properties of the HJBQVI satisfied by v_0 because we fully characterized v_1 in terms of v_0 (see Corollary 2). In Bruder and Pham (2009), the authors use backward and forward iterations on the domains and value functions because they have more than two coupled value functions. In our case, a more ‘standard’ approach suffices because of the reduction in (C.11). For classical references on the subject of viscosity solutions we refer the reader to Crandall et al. (1992), Fleming and Soner (2006), Bardi and Capuzzo-Dolcetta (2008), and Touzi (2012).

For a locally bounded function w on \mathcal{D}_0 , we denote by \underline{w} (resp. \bar{w}) the lower-semicontinuous (resp. upper-semicontinuous) envelope of w , given by

$$\begin{aligned} \underline{w}(t, \mathbf{x}) &= \liminf_{(t', \mathbf{x}') \rightarrow (t, \mathbf{x})} w(t', \mathbf{x}'), \\ \bar{w}(t, \mathbf{x}) &= \limsup_{(t', \mathbf{x}') \rightarrow (t, \mathbf{x})} w(t', \mathbf{x}'), \quad (t, \mathbf{x}) \in \mathcal{D}_0. \end{aligned}$$

For any locally bounded function u on \mathcal{D}_0 , define the locally bounded function $\mathcal{H}u$ on \mathcal{D}_0 by

$$\mathcal{H}u(t, \mathbf{x}) = \sup_{l \in L} \mathbb{E} \left[\int_t^{T \wedge \tilde{t}} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + (u(\tilde{t}, \Gamma(\mathbf{X}_{\tilde{t}^-}^{t, \mathbf{x}}, Z, l)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}^-}^{t, \mathbf{x}}, l)) \mathbf{1}_{\tilde{t} \leq T} + u(T, \mathbf{X}_T^{t, \mathbf{x}}) \mathbf{1}_{\tilde{t} > T} \right].$$

Lemma 9. *Let w be a locally bounded function on \mathcal{D}_0 . Then, $\mathcal{H}\bar{w}$ is upper-semicontinuous, and $\overline{\mathcal{H}w} \leq \mathcal{H}\bar{w}$.*

Proof. First, we prove that $\mathcal{H}\bar{w}$ is upper-semicontinuous. Let $(t, \mathbf{x}) \in \mathcal{D}_0$, and let $(t_n, \mathbf{x}_n)_{n \geq 1}$ be a sequence in \mathcal{D}_0 converging to (t, \mathbf{x}) as $n \rightarrow \infty$. There exists a sequence $(l_n)_{n \in \mathbb{N}}$ valued in L , such that

$$\begin{aligned} \mathcal{H}\bar{w}(t_n, \mathbf{x}_n) &= \mathbb{E} \left[\int_{t_n}^{T \wedge \tilde{t}_n} h(\mathbf{X}_s^{t_n, \mathbf{x}_n}) ds + \left(\bar{w}(\tilde{t}_n, \Gamma(\mathbf{X}_{\tilde{t}_n^-}^{t_n, \mathbf{x}_n}, Z, l_n)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}_n^-}^{t_n, \mathbf{x}_n}, Z, l_n) \right) \mathbf{1}_{\tilde{t}_n \leq T} + \bar{w}(T, \mathbf{X}_T^{t_n, \mathbf{x}_n}) \mathbf{1}_{\tilde{t}_n > T} \right] \end{aligned}$$

for $n \geq 1$ because \bar{w} is upper-semicontinuous and because L is compact.

By Bolzano-Weierstrass, the sequence $(l_n)_{n \in \mathbb{N}}$ converges (up to a subsequence) to an element $\check{l} \in L$, thus

$$\begin{aligned} \mathcal{H}\bar{w}(t, \mathbf{x}) &\geq \mathbb{E} \left[\int_t^{T \wedge \tilde{t}} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + \left(\bar{w}(\tilde{t}, \Gamma(\mathbf{X}_{\tilde{t}^-}^{t, \mathbf{x}}, Z, \check{l})) + \mathfrak{C}(\mathbf{X}_{\tilde{t}^-}^{t, \mathbf{x}}, Z, \check{l}) \right) \mathbf{1}_{\tilde{t} \leq T} + \bar{w}(T, \mathbf{X}_T^{t, \mathbf{x}}) \mathbf{1}_{\tilde{t} > T} \right] \\ &\geq \limsup_{n \rightarrow \infty} \mathbb{E} \left[\int_{t_n}^{T \wedge \tilde{t}_n} h(\mathbf{X}_s^{t_n, \mathbf{x}_n}) ds + \left(\bar{w}(\tilde{t}_n, \Gamma(\mathbf{X}_{\tilde{t}_n^-}^{t_n, \mathbf{x}_n}, Z, l_n)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}_n^-}^{t_n, \mathbf{x}_n}, Z, l_n) \right) \mathbf{1}_{\tilde{t}_n \leq T} + \bar{w}(T, \mathbf{X}_T^{t_n, \mathbf{x}_n}) \mathbf{1}_{\tilde{t}_n > T} \right] \\ &= \limsup_{n \rightarrow \infty} \mathcal{H}\bar{w}(t_n, \mathbf{x}_n), \end{aligned}$$

where the first inequality follows because $\check{l} \in L$. The second inequality follows because: \bar{w} is upper-semicontinuous, Γ is continuous, and Fatou's lemma. Finally, the last equality follows by definition. Hence, $\mathcal{H}\bar{w}$ is upper-semicontinuous as required.

Next, to prove that $\overline{\mathcal{H}w} \leq \mathcal{H}\bar{w}$, fix $(t, \mathbf{x}) \in \mathcal{D}_0$, and let $(t_n, \mathbf{x}_n)_{n \geq 1}$ be a sequence in \mathcal{D}_0 converging to (t, \mathbf{x}) as $n \rightarrow \infty$ such that $\mathcal{H}w(t_n, \mathbf{x}_n) \rightarrow \overline{\mathcal{H}w}(t, \mathbf{x})$. Then, by the properties of the sequence and the upper-semicontinuity of \bar{w} , we have that

$$\overline{\mathcal{H}w}(t, \mathbf{x}) = \lim_{n \rightarrow \infty} \mathcal{H}w(t_n, \mathbf{x}_n) \leq \limsup_{n \rightarrow \infty} \mathcal{H}\bar{w}(t_n, \mathbf{x}_n) \leq \mathcal{H}\bar{w}(t, \mathbf{x}),$$

where the last inequality follows from the upper-semicontinuity of $\mathcal{H}\bar{w}$. Thus, $\overline{\mathcal{H}w} \leq \mathcal{H}\bar{w}$. ■

Definition 1. A locally bounded function w on \mathcal{D}_0 is:

- (a) A viscosity supersolution of (C.11) on \mathcal{D}_0 , if for all $(t_0, \mathbf{x}_0, \psi) \in \mathcal{D}_0 \times \mathcal{C}^2(\mathcal{D}_0)$ such that $\underline{w} - \psi$ attains a minimum at (t_0, \mathbf{x}_0) , we have that

$$\min \left\{ -\frac{\partial \psi}{\partial t}(t_0, \mathbf{x}_0) - \mathcal{L}\psi(t_0, \mathbf{x}_0) - h(\mathbf{x}_0), \underline{w}(t_0, \mathbf{x}_0) - \mathbb{E} \left[\int_{t_0}^{T \wedge \tilde{t}_0} h(\mathbf{X}_s^{t_0, \mathbf{x}_0}) ds \right] \right. \\ \left. - \sup_{l \in L} \mathbb{E} \left[\left(\underline{w}(\tilde{t}_0, \Gamma(\mathbf{X}_{\tilde{t}_0}^{t_0, \mathbf{x}_0}, Z, l)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}_0}^{t_0, \mathbf{x}_0}, Z, l) \right) \mathbf{1}_{\tilde{t}_0 \leq T} + \underline{w}(T, \mathbf{X}_T^{t_0, \mathbf{x}_0}) \mathbf{1}_{\tilde{t}_0 > T} \right] \right\} \geq 0.$$

- (b) A viscosity subsolution of (C.11) on \mathcal{D}_0 , if for all $(t_0, \mathbf{x}_0, \psi) \in \mathcal{D}_0 \times \mathcal{C}^2(\mathcal{D}_0)$ such that $\bar{w} - \psi$ attains a maximum at (t_0, \mathbf{x}_0) , we have that

$$\min \left\{ -\frac{\partial \psi}{\partial t}(t_0, \mathbf{x}_0) - \mathcal{L}\psi(t_0, \mathbf{x}_0) - h(\mathbf{x}_0), \bar{w}(t_0, \mathbf{x}_0) - \mathbb{E} \left[\int_{t_0}^{T \wedge \tilde{t}_0} h(\mathbf{X}_s^{t_0, \mathbf{x}_0}) ds \right] \right. \\ \left. - \sup_{l \in L} \mathbb{E} \left[\left(\bar{w}(\tilde{t}_0, \Gamma(\mathbf{X}_{\tilde{t}_0}^{t_0, \mathbf{x}_0}, Z, l)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}_0}^{t_0, \mathbf{x}_0}, Z, l) \right) \mathbf{1}_{\tilde{t}_0 \leq T} + \bar{w}(T, \mathbf{X}_T^{t_0, \mathbf{x}_0}) \mathbf{1}_{\tilde{t}_0 > T} \right] \right\} \leq 0.$$

Finally, we say that w is a viscosity solution of (C.11) on \mathcal{D}_0 , if it is a viscosity supersolution and a viscosity subsolution of (C.11) on \mathcal{D}_0 .

Theorem 13. The value function v_0 is a viscosity solution of the HJBQVI in (C.11) on \mathcal{D}_0 , with terminal condition $v_0(T, \mathbf{x}) = c + q(s + \zeta - a(q - 1))(1 - \mathfrak{c})$.

Proof. Local boundedness: Let $(t, \mathbf{x}) \in \mathcal{D}_0$. Use Jensen's inequality and the bounds in (5.13) to write

$$|v_0(t, \mathbf{x})| \leq \mathbb{E} \left[\int_t^T \int_{\mathbb{R}} |z| \mathfrak{p}(dz, du) + N_T \left(\sup_{t \leq u \leq T} |S_u^{t, s}| + \kappa N_T \right) + (\mathfrak{M} + N_T) \left(\sup_{t \leq u \leq T} |S_u^{t, s}| + \kappa N_T + |\zeta| \right) \right. \\ \left. + a(\mathfrak{M} + N_T + 1)^2 + \phi T(\mathfrak{M} + N_T)^2 + \mathfrak{c} N_T \sup_{t \leq u \leq T} |S_u^{t, s}| + \mathfrak{c} \kappa N_T + \mathfrak{c} \int_t^T \int_{\mathbb{R}} |z| \mathfrak{p}(dz, du) \right].$$

From equation (5.4) we know that

$$r_1 := \mathbb{E} \left[(1 + \mathfrak{c}) \int_0^T \int_{\mathbb{R}} |z| \mathfrak{p}(dz, du) + \kappa N_T^2 + (\mathfrak{M} + N_T)(\kappa N_T + |\zeta|) \right. \\ \left. + a(\mathfrak{M} + N_T + 1)^2 + \phi T(\mathfrak{M} + N_T)^2 + \mathfrak{c} \kappa N_T \right] < \infty,$$

$$r_2 := \mathbb{E} [\mathfrak{M} + 2 N_T + \mathfrak{c} N_T] < \infty,$$

and r_1 and r_2 do not depend on (t, \mathbf{x}) . Therefore, we have that

$$|v_0(t, \mathbf{x})| \leq r_1 + r_2 \mathbb{E} \left[\sup_{t \leq u \leq T} |S_u^{t, s}| \right] \\ \leq r_1 + r_2 \left(r_3 + r_4 (T - t) e^{r_4 (T - t)} s^2 + r_4 s^2 \right), \\ \leq r_1 + r_2 \left(r_3 + r_4 (T - t) e^{r_4 (T - t)} \|\mathbf{x}\|^2 + r_4 \|\mathbf{x}\|^2 \right),$$

for some constants r_3 and r_4 that do not depend on (t, \mathbf{x}) . Thus, the value function v_0 is locally bounded. **Supersolution property:** Let $(t_0, \mathbf{x}_0, \psi) \in \mathcal{D}_0 \times \mathcal{C}^2(\mathcal{D}_0)$, such that $\underline{v}_0 - \psi$ attains a minimum at (t_0, \mathbf{x}_0) , with $\underline{v}_0(t_0, \mathbf{x}_0) = \psi(t_0, \mathbf{x}_0)$. By the definition of \underline{v}_0 , there is a sequence $(t_n, \mathbf{x}_n)_{n \geq 1} \in \mathcal{D}_0$ such that

$$v_0(t_n, \mathbf{x}_n) \rightarrow \underline{v}_0(t_0, \mathbf{x}_0) \quad \text{with} \quad (t_n, \mathbf{x}_n) \rightarrow (t_0, \mathbf{x}_0). \quad (\text{C.13})$$

Let $l_1 \in L$ and $(t, \mathbf{x}) \in \mathcal{D}_0$. Use equation (C.7) and Corollary 2 with an immediate impulse and price limit l_1 , (i.e., $\tau = t$, $\theta = T$, and $l = l_1$) to write

$$\begin{aligned} v_0(t, \mathbf{x}) &\geq v_1(t, \mathbf{x}, l_1) \\ &= \mathbb{E} \left[\int_t^{T \wedge \tilde{t}} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + (v_0(\tilde{t}, \Gamma(\mathbf{X}_{\tilde{t}}^{t, \mathbf{x}}, Z, l_1)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}}^{t, \mathbf{x}}, Z, l_1)) \mathbf{1}_{\tilde{t} \leq T} + v_0(T, \mathbf{X}_T^{t, \mathbf{x}}) \mathbf{1}_{\tilde{t} > T} \right]. \end{aligned}$$

Thus,

$$\begin{aligned} \liminf_{n \rightarrow \infty} v_0(t_n, \mathbf{x}_n) \\ \geq \liminf_{n \rightarrow \infty} \mathbb{E} \left[\int_{t_n}^{T \wedge \tilde{t}_n} h(\mathbf{X}_s^{t_n, \mathbf{x}_n}) ds + (v_0(\tilde{t}_n, \Gamma(\mathbf{X}_{\tilde{t}_n}^{t_n, \mathbf{x}_n}, Z, l_1)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}_n}^{t_n, \mathbf{x}_n}, Z, l_1)) \mathbf{1}_{\tilde{t}_n \leq T} + v_0(T, \mathbf{X}_T^{t_n, \mathbf{x}_n}) \mathbf{1}_{\tilde{t}_n > T} \right]. \end{aligned}$$

Next, use Fatou's lemma, the continuity of Γ , and the definition of \underline{v}_0 to obtain

$$\underline{v}_0(t_0, \mathbf{x}_0) \geq \mathbb{E} \left[\int_{t_0}^{T \wedge \tilde{t}_0} h(\mathbf{X}_s^{t_0, \mathbf{x}_0}) ds + (v_0(\tilde{t}_0, \Gamma(\mathbf{X}_{\tilde{t}_0}^{t_0, \mathbf{x}_0}, Z, l_1)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}_0}^{t_0, \mathbf{x}_0}, Z, l_1)) \mathbf{1}_{\tilde{t}_0 \leq T} + v_0(T, \mathbf{X}_T^{t_0, \mathbf{x}_0}) \mathbf{1}_{\tilde{t}_0 > T} \right].$$

Therefore,

$$\underline{v}_0(t_0, \mathbf{x}_0) \geq \sup_{l \in L} \mathbb{E} \left[\int_{t_0}^{T \wedge \tilde{t}_0} h(\mathbf{X}_s^{t_0, \mathbf{x}_0}) ds + (v_0(\tilde{t}_0, \Gamma(\mathbf{X}_{\tilde{t}_0}^{t_0, \mathbf{x}_0}, Z, l)) + \mathfrak{C}(\mathbf{X}_{\tilde{t}_0}^{t_0, \mathbf{x}_0}, Z, l)) \mathbf{1}_{\tilde{t}_0 \leq T} + v_0(T, \mathbf{X}_T^{t_0, \mathbf{x}_0}) \mathbf{1}_{\tilde{t}_0 > T} \right]$$

because the price limit l_1 is arbitrary.

Finally, to complete the viscosity supersolution property, use Corollary 3 with $\tau = \infty$, Ito's formula for $\psi \in \mathcal{C}^2(\mathcal{D}_0)$, the property that $\underline{v}_0 - \psi$ attains a minimum at (t_0, \mathbf{x}_0) with $\underline{v}_0(t_0, \mathbf{x}_0) = \psi(t_0, \mathbf{x}_0)$, and let $\theta \rightarrow t_0$ to obtain

$$-\frac{\partial \psi}{\partial t}(t_0, \mathbf{x}_0) - \mathcal{L}\psi(t_0, \mathbf{x}_0) - h(\mathbf{x}_0) \geq 0,$$

see Pham (2005).

We base the next part of the proof on Bruder and Pham (2009). **Subsolution property:** Let $(t_0, \mathbf{x}_0, \psi) \in \mathcal{D}_0 \times \mathcal{C}^2(\mathcal{D}_0)$, such that $\overline{v}_0 - \psi$ attains a maximum at (t_0, \mathbf{x}_0) , $\psi(t_0, \mathbf{x}_0) = \overline{v}_0(t_0, \mathbf{x}_0)$, and $\psi \geq \overline{v}_0$ on \mathcal{D}_0 .

If $\overline{v}_0(t_0, \mathbf{x}_0) \leq \mathcal{H} \overline{v}_0(t_0, \mathbf{x}_0)$, the desired inequality follows and the proof is complete. On the other hand, if $\overline{v}_0(t_0, \mathbf{x}_0) > \mathcal{H} \overline{v}_0(t_0, \mathbf{x}_0)$, we proceed by contradiction, i.e., we assume there is $\eta > 0$ such that

$$\eta = -\frac{\partial \psi}{\partial t}(t_0, \mathbf{x}_0) - \mathcal{L}\psi(t_0, \mathbf{x}_0) - h(\mathbf{x}_0) > 0.$$

By the definition of \overline{v}_0 and the continuity of ψ , there exists $\tilde{\epsilon} > 0$ and a sequence

$$(t_n, \mathbf{x}_n)_{n \geq 1} \subset ((t_0 - \tilde{\epsilon}, t_0 + \tilde{\epsilon}) \times B(\mathbf{x}_0, \tilde{\epsilon})) \cap \mathcal{D}_0,$$

such that

$$v_0(t_n, \mathbf{x}_n) \rightarrow \overline{v}_0(t_0, \mathbf{x}_0) \quad \text{with} \quad (t_n, \mathbf{x}_n) \rightarrow (t_0, \mathbf{x}_0), \quad (\text{C.14})$$

and

$$-\frac{\partial \psi}{\partial t} - \mathcal{L}\psi - h > \frac{\eta}{2} \quad \text{on} \quad ((t_0 - \tilde{\epsilon}, t_0 + \tilde{\epsilon}) \times B(\mathbf{x}_0, \tilde{\epsilon})) \cap \mathcal{D}_0, \quad (\text{C.15})$$

and

$$\overline{v}_0 - \psi < \frac{\eta}{4} \quad \text{on} \quad ((t_0 - \tilde{\epsilon}, t_0 + \tilde{\epsilon}) \times B(\mathbf{x}_0, \tilde{\epsilon})) \cap \mathcal{D}_0.$$

Here $B(\mathbf{x}_0, \tilde{\epsilon})$ is the open ball with centre in \mathbf{x}_0 and radius $\tilde{\epsilon}$ under the Euclidean metric. By continuity of ψ and (C.14), we also have that

$$\gamma_n := v_0(t_n, \mathbf{x}_n) - \psi(t_n, \mathbf{x}_n) \rightarrow 0 \text{ as } n \rightarrow \infty. \quad (\text{C.16})$$

By Corollary 4 we have that, for each $n \geq 1$, there is a control $(\tau_n, \mathbf{l}_n) \in \mathcal{I}_{t_n}$ such that

$$\begin{aligned} v_0(t_n, \mathbf{x}_n) - \frac{\eta}{4} \epsilon_n & \leq \mathbb{E} \left[\int_t^{\theta_n \wedge \tilde{\tau}_n} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + v_0(\theta_n, \mathbf{X}_{\theta_n}^{t, \mathbf{x}}) \mathbf{1}_{\theta_n < \tau_n} \right. \\ & \quad \left. + v_1(\theta_n, \mathbf{X}_{\theta_n}^{t, \mathbf{x}}, \mathbf{l}_n) \mathbf{1}_{\tau_n \leq \theta_n < \tilde{\tau}_n} + v_0(\tilde{\tau}_n, \Gamma(\mathbf{X}_{\tilde{\tau}_n}^{t, \mathbf{x}}, Z, \mathbf{l}_n)) \mathbf{1}_{\tau_n < \tilde{\tau}_n < \theta_n} \right]. \end{aligned} \quad (\text{C.17})$$

Choose $\theta_n = (t_n + \epsilon_n) \wedge \vartheta_n \wedge \tilde{t}_n$, with

$$\vartheta_n = \inf \{s \geq t_n : X_s^{t_n, \mathbf{x}_n} \notin B(\mathbf{x}_0, \tilde{\epsilon}/2), s \leq t_n + \tilde{\epsilon}/2, s \leq T\} \in \mathcal{T}_{t_n, T}$$

and $(\epsilon_n)_{n \in \mathbb{N}}$ a strictly positive sequence that satisfies

$$\epsilon_n \rightarrow 0, \quad \frac{\gamma_n}{\epsilon_n} \rightarrow 0, \quad \text{as } n \rightarrow \infty.$$

After simple inspection of (C.17), and that \tilde{t}_n is the first time there is a jump in N after t_n , write

$$v_0(t_n, \mathbf{x}_n) - \frac{\eta}{4} \epsilon_n \leq \mathbb{E} \left[\int_t^{\theta_n} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + v_0(\theta_n, \mathbf{X}_{\theta_n}^{t, \mathbf{x}}) \mathbf{1}_{\theta_n < \tau_n} + v_1(\theta_n, \mathbf{X}_{\theta_n}^{t, \mathbf{x}}, \mathbf{l}_n) \mathbf{1}_{\tau_n \leq \theta_n} \right]. \quad (\text{C.18})$$

Use Lemma 9 to write

$$\overline{\mathcal{H}v_0}(t_0, \mathbf{x}_0) \leq \mathcal{H}\overline{v_0}(t_0, \mathbf{x}_0) < \overline{v_0}(t_0, \mathbf{x}_0) = \psi(t_0, \mathbf{x}_0),$$

and given that $\overline{\mathcal{H}v_0}$ is upper semi-continuous, ψ is continuous, and $\overline{\mathcal{H}v_0}(t_0, \mathbf{x}_0) < \psi(t_0, \mathbf{x}_0)$, the inequality $\mathcal{H}v_0 \leq \psi$ holds in a neighbourhood of (t_0, \mathbf{x}_0) . Thus, for n sufficiently large and because $v_1(t_n, \mathbf{x}_n, \mathbf{l}_n) \leq \mathcal{H}v_0(t_n, \mathbf{x}_n)$, we write

$$v_1(\theta_n, \mathbf{X}_{\theta_n}^{t_n, \mathbf{x}_n}, \mathbf{l}_n) \mathbf{1}_{\tau_n \leq \theta_n} \leq \psi(\theta_n, \mathbf{X}_{\theta_n}^{t_n, \mathbf{x}_n}) \mathbf{1}_{\tau_n \leq \theta_n}. \quad (\text{C.19})$$

Equation (C.18), together with (C.16) and (C.19), yields

$$\psi(t_n, \mathbf{x}_n) + \gamma_n - \frac{\eta}{4} \epsilon_n \leq \mathbb{E} \left[\int_{t_n}^{\theta_n} h(\mathbf{X}_s^{t_n, \mathbf{x}_n}) ds + \psi(\theta_n, \mathbf{X}_{\theta_n}^{t_n, \mathbf{x}_n}) \right].$$

Use Ito's formula between t_n and θ_n , divide by ϵ_n , and use the bound in (C.15) to obtain

$$\frac{\gamma_n}{\epsilon_n} - \frac{\eta}{4} \leq \frac{1}{\epsilon_n} \mathbb{E} \left[\int_{t_n}^{\theta_n} (\partial_t \psi + \mathcal{L}\psi + h)(s, \mathbf{X}_s^{t_n, \mathbf{x}_n}) ds \right] \leq -\frac{\eta}{2} \mathbb{E} \left[\frac{\theta_n - t_n}{\epsilon_n} \right].$$

Let $n \rightarrow \infty$, so that

$$-\frac{\eta}{4} \leq -\frac{\eta}{2},$$

which is a contradiction. Thus, the subsolution property follows. ■

C.2 Deterministic latency

Here, we show the dynamic programming principle and a verification theorem when the trader executes orders with deterministic latency.⁴

⁴In general, a deterministic delay can be regarded as a stochastic delay where the stochasticity converges to a point mass. In our framework, the stochasticity is modelled with an exponential random variable with parameter $\lambda > 0$; thus, we cannot recover the deterministic latency results from the random latency case.

C.2.1 Dynamic programming principle

We state the dynamic programming principle that holds for the deterministic latency execution problem. Let $t \in [0, T]$, $\beta \in \mathcal{A}$, with $\beta = (\tau_i, \mathbf{l}_i)_{i \geq 1}$. Denote the number of executed orders by $\iota(t, \beta)$, the number of pending orders at time t by $k_t(\beta)$ (which is zero or one), and the pending order $p(t, \beta)$ (which is the empty set if there is no pending order):

$$\begin{aligned}\iota(t, \beta) &= \inf\{i \geq 1 : \tau_i > t - \Delta\} - 1 \in \mathbb{N} \cup \{\infty\}, \\ k_t(\beta) &= \text{card}\{i \geq 1 : t - \Delta < \tau_i \leq t\} \in \{0, 1\}, \\ p(t, \beta) &= (\tau_{i+\iota(t, \beta)}, \mathbf{l}_{i+\iota(t, \beta)})_{1 \leq i \leq k_t(\beta)} \in P_t(k_t(\beta)).\end{aligned}$$

Theorem 14. *Dynamic programming principle. The value functions satisfy the dynamic programming principle: for $k \in \{0, 1\}$, $(t, \mathbf{x}, p) \in \mathcal{D}_k$,*

$$v_k(t, \mathbf{x}, p) = \sup_{\beta \in \mathcal{A}_{t,p}} \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t, \mathbf{x}, p, \beta}) ds + \sum_{t < \tau_i + \Delta \leq \theta} \mathfrak{c} \left(\mathbf{X}_{(\tau_i + \Delta)^-}^{t, \mathbf{x}, p, \beta}, Z_{\lfloor (\tau_i + \Delta)/\Delta \rfloor}, \mathbf{l}_i \right) + v_{k(\theta, \alpha)}(\theta, \mathbf{X}_\theta^{t, \mathbf{x}, p, \beta}, p(\theta, \beta)) \right],$$

where θ is any stopping time valued in $[t, T]$, possibly depending on β .

The following Corollary is consequence of the dynamic programming principle. For all $t \in [0, T]$, we denote by \mathcal{I}_t the set of pairs (τ, \mathbf{l}) where τ is a stopping time, $t \leq \tau \leq T - \Delta$, or $\tau = \infty$ a.s. and \mathbf{l} is an \mathcal{F}_τ -measurable random variable in L .

Corollary 6. *Let $(t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d$. (1) For any stopping time θ valued in $[t, t_1 + \Delta]$ and for $p = (t_1, l_1) \in P_t(1)$, we have*

$$v_1(t, \mathbf{x}, p) = \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t, \mathbf{x}, 0}) ds + v_1(\theta, \mathbf{X}_\theta^{t, \mathbf{x}, 0}, p) \right].$$

(2) For any stopping time θ valued in $[t, t + \Delta]$:

$$v_0(t, \mathbf{x}, \emptyset) = \sup_{(\tau, \mathbf{l}) \in \mathcal{I}_t} \mathbb{E} \left[\int_t^\theta h(\mathbf{X}_s^{t, \mathbf{x}, 0}) ds + v_0(\theta, \mathbf{X}_\theta^{t, \mathbf{x}, 0}, \emptyset) \mathbf{1}_{\theta < \tau} + v_1(\theta, \mathbf{X}_\theta^{t, \mathbf{x}, 0}, (\tau, \mathbf{l})) \mathbf{1}_{\theta \geq \tau} \right].$$

The proof is a straightforward modification of Theorem 3.1 and Corollary 3.2 in [Bruder and Pham \(2009\)](#).

C.2.1.1 PDE system viscosity characterization

Define $\mathcal{D}_k^{1,2}$ for $k = 0, 1$ as follows:

$$\mathcal{D}_0^1 = (T - \Delta, T) \times \mathbb{R}^3, \quad \mathcal{D}_0^2 = [0, T - \Delta] \times \mathbb{R}^3, \quad \mathcal{D}_1^1 = \mathcal{D}_1, \quad \mathcal{D}_1^2 = \emptyset.$$

Theorem 15. *Dynamic programming equation. For $k \in \{0, 1\}$, $(t, \mathbf{x}, p) \in \mathcal{D}_k^{1,2}$, the value functions satisfy*

$$\min \left\{ -\frac{\partial v_k}{\partial t}(t, \mathbf{x}, p) - \mathcal{L}v_k(t, \mathbf{x}, p) - h(\mathbf{x}), v_k(t, \mathbf{x}, p) - \sup_{l \in L} v_{k+1}(t, \mathbf{x}, (t, l)) \right\} = 0, \quad \text{on } \mathcal{D}_0^2, k = 0 \quad (\text{C.20})$$

and

$$-\frac{\partial v_k}{\partial t}(t, \mathbf{x}, p) - \mathcal{L}v_k(t, \mathbf{x}, p) - h(\mathbf{x}) = 0 \quad \text{on } \mathcal{D}_k^1, k = 0, 1. \quad (\text{C.21})$$

From the system of equations (C.20) and (C.21), we have the following Feynman-Kac representation:

$$v_1(t, \mathbf{x}, (t_1, l_1)) = \mathbb{E} \left[\int_t^{t_1 + \Delta} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + \mathfrak{c} \left(\mathbf{X}_{(t_1 + \Delta)^-}^{t, \mathbf{x}}, Z, l_1 \right) + v_0 \left(t_1 + \Delta, \Gamma \left(\mathbf{X}_{(t_1 + \Delta)^-}^{t, \mathbf{x}}, Z, l_1 \right) \right) \right], \quad (\text{C.22})$$

for all $(t_1, l_1) \in [0, T - \Delta] \times L$, $t \in [t_1, t_1 + \Delta] \times \mathbb{R}^3$, and Z is an independent random variable with law as in (5.3). We substitute (C.22) in (C.20) to obtain the variational inequality satisfied by v_0 :

$$0 = \min \left\{ -\frac{\partial v_0}{\partial t} - \mathcal{L}v_0 - h, \right. \\ \left. v_0 - \sup_{l \in L} \mathbb{E} \left[\int_t^{t+\Delta} h(\mathbf{X}_s^{t, \mathbf{x}}) ds + \mathfrak{C}(\mathbf{X}_{(t+\Delta)-}^{t, \mathbf{x}}, Z, l) + v_0(t + \Delta, \Gamma(\mathbf{X}_{(t+\Delta)-}^{t, \mathbf{x}}, Z, l)) \right] \right\},$$

on $[0, T - \Delta] \times \mathbb{R}^3$, with

$$v_0(t, \mathbf{x}) = \mathbb{E} \left[\int_t^T h(\mathbf{X}_s^{t, \mathbf{x}}) ds + g(\mathbf{X}_T^{t, \mathbf{x}}) \right] \quad \text{on } (T - \Delta, T] \times \mathbb{R}^3.$$

Thus, we transformed the original problem into a non-delay impulse control problem for v_0 , and we related v_1 to v_0 . In our formulation, $h(\mathbf{x}) = -\phi q^2$, and

$$v_1(t, \mathbf{x}, (t_1, l_1)) = -\phi q^2(t_1 + \Delta - t) + \mathbb{E} \left[v_0(t_1 + \Delta, \Gamma(\mathbf{X}_{(t_1+\Delta)-}^{t, \mathbf{x}}, Z, l_1)) + \mathfrak{C}(\mathbf{X}_{(t_1+\Delta)-}^{t, \mathbf{x}}, Z, l_1) \right], \quad (\text{C.23})$$

for all $(t_1, l_1) \in [0, T - \Delta] \times L$, $t \in [t_1, t_1 + \Delta] \times \mathbb{R}^3$ because the inventory does not change during the latency period. Substitute (C.22) in (C.20) to obtain the variational inequality satisfied by v_0 :

$$0 = \min \left\{ -\frac{\partial v_0}{\partial t} - \mathcal{L}v_0 - h, \right. \\ \left. v_0 + \phi q^2 \Delta - \sup_{l \in L} \mathbb{E} \left[v_0(t + \Delta, \Gamma(\mathbf{X}_{(t+\Delta)-}^{t, \mathbf{x}}, Z, l)) - \mathfrak{C}(\mathbf{X}_{(t+\Delta)-}^{t, \mathbf{x}}, Z, l) \right] \right\}, \quad (\text{C.24})$$

on $[0, T - \Delta] \times \mathbb{R}^3$, with

$$v_0(t, \mathbf{x}) = -\phi q^2(T - t) + \mathbb{E} [g(\mathbf{X}_T^{t, \mathbf{x}})] \quad \text{on } (T - \Delta, T] \times \mathbb{R}^3. \quad (\text{C.25})$$

By inspection of (C.24) and (C.25), we make the following ansatz for v_0 :

$$v_0(t, \mathbf{x}) = v_0(t, s, q, c) = c + h_0(t, s, q);$$

thus

$$0 = \min \left\{ -\frac{\partial h_0}{\partial t} - \mathcal{L}h_0 + \phi q^2, \quad h_0 + \phi q^2 \Delta \right. \\ \left. - \sup_{l \in L} \mathbb{E} \left[h_0(t + \Delta, \tilde{\Gamma}(S_{(t+\Delta)-}^{t, s}, q, Z, l)) + (1 - \mathfrak{c})(S_{(t+\Delta)-}^{t, s} + Z) f(l - S_{(t+\Delta)-}^{t, s} - Z) \right] \right\}, \quad (\text{C.26})$$

on $[0, T - \Delta] \times \mathbb{R}^3$, $\tilde{\Gamma}$ is as in (C.12), and the terminal condition is

$$h_0(t, s, q) = -\phi q^2(T - t) + (q \mathbb{E}[S_T^{t, s}] + q\zeta - a q(q - 1))(1 - \mathfrak{c}) \quad \text{on } (T - \Delta, T] \times \mathbb{R}^3. \quad (\text{C.27})$$

C.3 Zero latency

Standard results in optimal stopping and impulse control theory yield the following theorem.

Theorem 16. *Dynamic programming equation. For all $(t, \mathbf{x}) \in \mathcal{D}$, the value function satisfies:*

$$\min \left\{ -\frac{\partial v}{\partial t}(t, \mathbf{x}) - \mathcal{L}v(t, \mathbf{x}) - h(\mathbf{x}), v(t, \mathbf{x}) - v(t, \Gamma(\mathbf{x})) + \mathfrak{c} s \right\} = 0, \quad (\text{C.28})$$

where $\mathcal{D} = [0, T] \times \mathbb{R}^+ \times (-\infty, \mathfrak{M}] \times \mathbb{R}^+$, \mathcal{L} is an infinitesimal generator given by

$$\mathcal{L} \cdot = \frac{\partial \cdot}{\partial s} b(s) + \frac{1}{2} \frac{\partial^2 \cdot}{\partial s^2} \sigma^2(s), \quad (\text{C.29})$$

and the terminal condition is

$$v(T, \mathbf{x}) = c + q(s + \zeta - a(q - 1))(1 - \mathfrak{c}), \quad \forall l \in L. \quad (\text{C.30})$$

Next, employ the ansatz $v(t, s, q, c) = c + h(t, s, q)$ and write

$$\min \left\{ -\frac{\partial h}{\partial t}(t, s, q) - \mu(s) \frac{\partial h}{\partial s}(t, s, q) - \frac{1}{2} \sigma^2(s) \frac{\partial^2 h}{\partial s^2}(t, s, q) + \phi q^2, \right. \\ \left. h(t, s, q) - s - h(t, s - \kappa, q - 1) + \mathfrak{c} s \right\} = 0 \quad (\text{C.31})$$

on $[0, T] \times \mathbb{R}^+ \times (-\infty, \mathfrak{M}]$, with terminal condition $h(T, s, q) = q (s + \zeta - a(q - 1)) (1 - \mathfrak{c})$.