King's Research Portal

DOI:

[10.1080/1350486X.2018.1434009](https://doi.org/10.1080/1350486X.2018.1434009)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Cartea, A., Donnelly, R., & Jaimungal, S. (2018). Enhancing trading strategies with order book signals. *Applied Mathematical Finance*, 25(1). <https://doi.org/10.1080/1350486X.2018.1434009>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Enhancing Trading Strategies with Order Book Signals[☆]

Álvaro Cartea^a, Ryan Donnelly^b, Sebastian Jaimungal^c

^a*Department of Mathematics, University of Oxford, Oxford, UK
Oxford-Man Institute of Quantitative Finance, Oxford, UK*

^b*École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland*

^c*University of Toronto, Toronto, Canada*

Abstract

We use high-frequency data from the Nasdaq exchange to build a measure of volume imbalance in the limit order book (LOB). We show that our measure is a good predictor of the sign of the next market order (MO), i.e. buy or sell, and also helps to predict price changes immediately after the arrival of an MO. Based on these empirical findings, we introduce and calibrate a Markov chain modulated pure jump model of price, spread, LO and MO arrivals, and volume imbalance. As an application of the model, we pose and solve a stochastic control problem for an agent who maximizes terminal wealth, subject to inventory penalties, by executing trades using LOs. We use in-sample-data (January to June 2014) to calibrate the model to eleven equities traded in the Nasdaq exchange, and use out-of-sample data (July to December 2014) to test the performance of the strategy. We show that introducing our volume imbalance measure into the optimization problem considerably boosts the profits of the strategy. Profits increase because employing our imbalance measure reduces adverse selection costs and positions LOs in the book to take advantage of favorable price movements.

Keywords: order imbalance, algorithmic trading, high-frequency trading, order flow, market making, adverse selection

[☆]ÁC acknowledges the research support of the Oxford-Man Institute for Quantitative Finance and the hospitality of the Finance Group at Saïd Business School. SJ would like to thank NSERC and GRI for partially funding this work. The authors would like to thank seminar participants at SIAM Financial Mathematics and Engineering (Chicago), TMU Finance Workshop, Institut Louis Bachelier, AMaMeF and Swissquote Conference, Imperial College Finance and Stochastics Seminar, and Morgan Stanley (London) for comments on an earlier version of this article.

Email addresses: `alvaro.cartea@maths.ox.ac.uk` (Álvaro Cartea), `ryan.donnelly@epfl.ch` (Ryan Donnelly), `sebastian.jaimungal@utoronto.ca` (Sebastian Jaimungal)

1. Introduction

The rise in computer power and the dominance of electronic exchanges have paved the way to a surge in computerized trading algorithms. These algorithms are programmed to trade in and out of positions and to handle different types of exposures such as inventory risk and adverse selection costs. The development of many trading algorithms starts with a preconceived strategy which is framed as an optimization problem, and then its implementation is handled by computers. This approach has given way to a vast range of algorithms. These are tailored to different objectives and employed in different trading environments and markets, but common to all is that their performance depends on one of the most basic building blocks in any trading strategy: how are orders scheduled, i.e. how much and when to trade, and what order type should be employed.

How an algorithm determines the timing and type of order to execute a trade depends on what the model predicts about price movements and what type of orders other agents are sending to the market. In order driven markets, the main order types are limit orders (LOs) and market orders (MOs). Agents' LOs show an intention to buy or sell an amount of the asset at a displayed price, and these rest in the limit order book (LOB) until they are filled by an incoming MO or are cancelled by the agent who posted it. MOs, on the other hand, are sent to the market and immediately executed against the LOs resting in the book.

In this paper we use Nasdaq data to show how to employ LOB information in trading algorithms. We use the volume posted on both sides of the LOB for a number of stocks to build a measure of volume imbalance which proxies buying and selling pressures in the market. We show that our measure of imbalance acts as a strong predictor of the rate of incoming MOs as well as the direction and magnitude of price movements following an MO. Given the arrival of an MO when volume imbalance is buy-heavy (sell-heavy), there is a high probability that this MO is a buy (sell) order. Furthermore, immediately following a buy (sell) MO, the magnitude and sign of the midprice is large and positive (negative) when volume imbalance is buy-heavy (sell-heavy).

The ability to form accurate predictions of trade types and price changes is valuable information that an agent can use to optimize her trading strategy. In general, incorporating volume imbalance in algorithmic trading models will improve the performance of strategies. Using signals from the LOB helps to execute directional trades using MOs, and tilt the resting orders in the LOB, to take advantage of favorable price movements and to reduce adverse selection costs.

Furthermore, we use Nasdaq data to analyze the size of MOs relative to the volume posted at the best bid and ask, and to analyze the financial performance of market making strategies that only post at the best bid and ask. First, we find that depending

on the specific equity, between 91.6% and 99.9% of all MOs are filled with the volume posted only at the best price. This means that LOs posted in the book deeper than the best price have less than 0.1% to 8.4% chance of being executed by any given MO. Second, we show that market making strategies that attempt at earning the spread by always placing two-sided LOs at the best quotes suffer greatly from adverse selection to the point where compensation of the spread does not allow positive expected profits.

These two findings shed light into the potential performance of the market making models that decide how deep LOs are posted in the LOB. For large tick stocks, those which regularly trade with a bid-ask spread of one tick, most market making models proposed in the literature result in strategies that mimic a market maker who is always posting at-the-touch, this includes strategies that control exposure to inventory risk, see e.g. Avellaneda and Stoikov (2008), Guéant et al. (2012), Fodra and Labadie (2012), Cartea and Jaimungal (2015), and Cartea et al. (2017). Strategies that optimize expected profits by choosing optimal depths require a level of granularity which is often too fine relative to the tick size of the stock. In this case it is clear that to implement the market making strategy the optimal depths must be rounded to the nearest tick which means that postings are at the best bid or best ask, or at least one tick into the book. However, postings deeper than the best quote are hardly ever filled so these strategies perform very close to one that is always posting at the best quotes, and as we show they are very likely to generate losses.

Our work is related to the microstructure and algorithmic trading literature that studies the information conveyed by the LOB and order flow (i.e. MO activity). The early empirical study of Biais et al. (1995) analyzes the dynamics of the LOB and MOs using French data, and the theoretical work of Foucault (1999) presents a model of price formation and order placement decisions in the LOB. More recently, with access to better quality data, Cont et al. (2013) employ trade and quote data from Nasdaq to perform a statistical study of the price impact of order book events. They show that price changes are driven by order flow imbalances. The impact of volume imbalance on price changes and trade arrivals is also studied in Lipton et al. (2013), and they develop a model for the joint dynamics of the length of the best bid and ask queue and the arrival of MOs. In Huang et al. (2015), the full LOB is modelled as a Markov queuing system where interactions on the dynamics at several price levels are possible. They propose their model as a market simulation tool to test the performance of trading algorithms.

In the context of algorithmic trading, Stoikov and Waeber (2012) consider an asset liquidation problem where they employ LOB information to construct a measure of the instantaneous supply and demand imbalance in the market. More recently, Bechler and Ludkovski (2015) employ order flow information to develop optimal execution models that take into account market impact and informational costs. Finally, Cartea and Jaimungal (2016a) and Cartea and Jaimungal (2016b) employ ultra-high frequency data

to show that there is a positive relation between net order flow, defined as the difference between the volume of buy and sell MOs, and prices of stocks and show how to develop trading algorithms to target VWAP and execute a large number of shares.

The rest of this paper is structured as follows. In Section 2 we introduce our measure of volume imbalance. We investigate the performance of a simple market making strategy using data from Nasdaq, the results of which provide strong motivation for modifying the strategy based on the volume imbalance process. In Section 3 we introduce a dynamic LOB model which reflects the relationship between the volume imbalance process quantities such as arrival and size of MOs, and midprice dynamics. We also propose and solve a trading problem using techniques of stochastic optimal control and provide a verification proof. In Section 4 we examine the performance of the investment strategy on the Nasdaq Exchange. Section 5 concludes, and the appendices contain proofs and parameter estimates used in the paper.

2. Volume Imbalance: Order Arrival and Price Revisions

In this section we introduce and discuss aspects of our measure of volume imbalance. This measure is simple to calculate and easy to incorporate in any algorithmic trading strategy. In this paper we focus on equities that are traded in order driven markets (so the LOB is visible), but in principle, measures based on quantities that show buying and selling pressure can be incorporated into algorithms in a similar way to the trading algorithm that we frame as a stochastic control problem in Section 3.

We define volume imbalance at time t as

$$\rho_t = \frac{V_t^b - V_t^a}{V_t^b + V_t^a} \in [-1, 1], \quad (1)$$

where V_t^b and V_t^a are the volumes at time t of LOs posted at the best bid and best ask respectively. Clearly, when ρ_t is close to 1 there is strong buying pressure and when it is close to -1 there is strong selling pressure.

Volume imbalance is a key quantity because it summarizes agents' willingness to buy or sell assets. This information can be used to predict the arrival and direction of MOs, and help to predict the sign and magnitude of price changes – we return to this point in Subsection 2.2. Before discussing in more detail the attributes and statistical properties of this quantity, we first illustrate how an investment strategy performs in the absence of any measure that accounts for buying and selling pressure.

2.1. An example of trading strategy

A traditional investment strategy is to post on both sides of the LOB with the objective of making roundtrip trades and earning the quoted difference between the LOs. A simple version of this strategy is that of a market maker who is always posted at the best bid and best ask. In this section, we test this strategy for two Nasdaq equities during the period July to December 2014: Intel Corporation (INTC) and Oracle Corporation (ORCL). These stocks exhibit behavior in which the spread, i.e. difference between best offer and best bid, is equal to one cent for a significant portion of the trading day. In Table 2 we report the results for nine other Nasdaq stocks.

We begin by studying the performance of a zero-intelligence strategy that posts LOs at the best bid and ask in an attempt to earn the spread on roundtrip trades. We backtest this strategy on historical data under two scenarios which differ in the volume of the agent's LOs.

- Scenario 1. The agent imposes an inventory constraint such that inventory q always satisfies $|q| \leq Q$. If $q \neq -Q$ ($q \neq Q$), every market buy (sell) order changes the agent's inventory position by -1 ($+1$), and her wealth increases (decreases) by an amount equal to the best ask (bid) price. At the end of the trading period, any remaining inventory of the agent is liquidated
 - (a) using an MO which is filled at the reigning best bid or ask price,
 - (b) at the midprice.
- Scenario 2. As above, the agent imposes the constraint $|q| \leq Q$. When a market buy (sell) order arrives, the agent's inventory position decreases (increases) by $\min\{V, Q + q\}$ ($\min\{V, Q - q\}$), where V is the volume of the MO, and her wealth increases (decreases) by the traded volume multiplied by the best ask (bid) price. At the end of the trading period, any inventory of the agent is liquidated
 - (a) using an MO which is filled at the reigning best bid or ask price,
 - (b) at the midprice.

The difference between the scenarios is that in Scenario 1 the LOs are for one unit of the asset and in Scenario 2 the LOs are for a number of shares large enough to meet the full volume of the incoming MO. We explore two methods of liquidating inventory at the end of the period in order to eliminate the effect of crossing the spread on the performance of the strategy. Moreover, in both scenarios we assume that each trading period is 30 minutes long, but we exclude the first and last trading period in each day. We conduct these tests on each trading day from June to December, 2014, so there are 125 days, each with 11 trading periods during the day.

Figure 1: Performance of best bid/ask LO strategy under Scenario 1. Each point represents a different value of Q . The leftmost point on each curve corresponds to $Q = 1$ with unit increments to $Q = 10$. Increments then increase by 10 up to $Q = 200$. Type a crosses the spread with an MO to liquidate whereas type b) is allowed to liquidate at the midprice.

The left panel of Figure 1 shows the annualized mean and standard deviation of the zero-intelligence strategy under Scenario 1. The leftmost point on each curve corresponds to $Q = 1$ with unit increments to $Q = 10$. Increments then increase by 10 up to $Q = 200$. The right panel shows the Sharpe ratio for a range of values of the inventory constraint. The results for Scenario 2 are shown in Figure 2. Here the values of Q are scaled by a factor of 100 from those of Scenario 1 so the total range is 100 to 20000. The much larger magnitude of earnings under Scenario 2 are due to larger individual trade sizes.

Clearly, the expected earnings of the zero-intelligence strategy are negative under both scenarios. Any agent who provides liquidity to the market in this naive way will be driven out of business due to the strategy's exposure to severe adverse selection costs. These losses are not solely due to liquidating large inventory by crossing the spread at the end of the period. The losses persist even if the agent is allowed to liquidate at the midprice.

An attempt to alleviate adverse selection costs is to post LOs deeper in the LOB, hoping that MOs walk beyond the best quote to fill these LOs. However, this adjustment will have a negligible effect on the expected profits because LOs resting at worse prices than the best quotes are unlikely to be filled. Consider the size of an MO relative to the volume at the best bid or offer. In Table 1, we show the number of market buy and sell orders that only touch the best quotes in the LOB, and the number that touch LOs beyond the best quote. It is clear that a disproportionate number of MOs only involve LOs at the best quotes – for these two stocks, approximately 0.1% of the MOs are large enough to walk beyond the best quote.

Although the strategy employed by the market maker is extremely simple, many of the

Figure 2: Performance of best bid/ask LO strategy under Scenario 2. Each point represents a different value of Q . The leftmost point on each curve corresponds to $Q = 100$ with increments of 100 up to $Q = 1000$. unit increments to $Q = 10$. Increments then increase by 1000 up to $Q = 20000$. Type a) crosses the spread with an MO to liquidate whereas type b) is allowed to liquidate at the midprice.

market making and investments strategies in the extant literature exhibit a very similar behavior to the one discussed here for large tick stocks, see for instance Avellaneda and Stoikov (2008), Guéant et al. (2012), Fodra and Labadie (2012), Cartea and Jaimungal (2015), and Cartea et al. (2017) – large tick stocks are those that generally trade with a bid-ask spread of one tick. The trading strategies developed in these papers use a continuous control variable for the price at which the trader or market maker posts LOs. These optimal depths are the output of a stochastic optimal control problem where the main driver to adjust the depth is inventory position and end of trading horizon. When the strategies are implemented, the optimal depths must be rounded to the nearest tick size, but with the market behavior demonstrated in Table 1 a continuous control price seems unreasonable (see also Table B.8 in the appendix for similar data on other stocks). The effect of this mismatch between a continuous control, and the grid on which asset prices move in exchanges, is more pronounced for large tick stocks because the probability of an MO walking beyond the best quote is essentially zero. When the continuous controls are rounded to an appropriate multiple of the tick size, either all of the values are rounded to the best quote, or they are rounded to a tick which has very little activity in the market. Thus, if the posting strategy based on continuous controls is inventory based, then the resulting strategy essentially becomes one of always posting at the best bid and ask, with an inventory constraint – as the scenarios tested above. On the other hand, models that are not purely inventory based, for instance Cartea et al. (2014), the strategy is modified based on the observation of recent order flow, so rounding the optimal strategy to the nearest tick results in a strategy that posts sometimes at the touch, but also sometimes deeper in the book, and imposes an inventory constraint.

	Best Quote Only		Beyond Best Quote		$\mathbb{P}(V_{MO} \leq V_{LO})$
	Buys	Sells	Buys	Sells	
INTC	35,595	38,451	54	50	0.999
ORCL	30,001	27,502	41	45	0.999

Table 1: Number of MOs that touch or go beyond the best quote. Data are taken from a full month of trading (January, 2014). The column labelled $\mathbb{P}(V_{MO} \leq V_{LO})$ is the probability that an MO has smaller volume than all limit orders posted at the best price, and hence only engages the best quote.

Figure 3: Number of MOs that are buys and sells depending on imbalance level. Imbalance levels correspond to subintervals $[-1, -0.33)$, $[-0.33, 0.33]$, and $(0.33, 1]$. Data are taken from a full month of trading in January 2014.

2.2. Volume imbalance in order driven equity markets

The zero-intelligence strategy described above, as well as those that behave in an essentially identical manner, are a few examples of algorithms whose performance can be enhanced if volume imbalance is employed. In general, if a strategy employs additional information which predicts trade arrivals and price movements then it is able to protect itself from adverse selection costs and take advantage of price movements.

To show the predictive power of our volume imbalance measure we first divide the imbalance measure interval $[-1, 1]$ into three subintervals, referred to as buy-heavy when $\rho_t \in (\frac{1}{3}, 1]$, sell-heavy when $\rho_t \in [-1, -\frac{1}{3})$, and neutral when $\rho_t \in [-\frac{1}{3}, \frac{1}{3}]$. Figure 3 shows that depending on which subinterval imbalance ρ_t lies in, the type of incoming MO, buy or sell, can be established with high accuracy.

Furthermore, the empirical distribution of the midprice change following an MO varies depending on which subinterval was occupied at the time of the MO. In panels (a) and (b) of Figure 4 we show this empirical distribution conditioned on imbalance regime. In

(a) INTC – **all** MOs

(b) ORCL – **all** MOs

(c) INTC – **buy** MOs

(d) ORCL – **buy** MOs

Figure 4: Distribution of midprice change 10 ms after an MO. Subintervals and data are equivalent to those in Figure 3.

this figure, a time lag of 10 ms is used between the arrival of the MO and the calculation of the price change. The positive (negative) bias of price changes after an MO in a buy-heavy (sell-heavy) regime is not simply due to the fact that it is much more likely to observe a market buy (sell) order in that regime. If we compute the price change when only observing market buy orders, we obtain the distribution shown in panels (c) and (d) of Figure 4. These two figures lend support to the idea that the direction of the price change after a market buy order is most of the time positive (see for example ?), and the magnitude of the change has a clear dependence on the level of volume imbalance.

We choose a 10 ms window between MOs and price changes because LO activity increases immediately following an MO to such a large magnitude that this increase in activity can be thought of as having been caused by the MO – the results in the paper are not altered if the window is chosen to be up to 100 ms. We note that our statistical results do not

provide conclusive evidence on the causality between MO and LO activity. However, the framework we develop here does not require a deeper understanding of causality, it suffices to observe that LO activity is greatest right after the arrival of MOs.

To highlight this increased LO activity post MO activity, panels (a) and (b) of Figure 5 shows the proportion of three types of LO events that occur after MO arrivals. The horizontal axis is the time window immediately after the MO event and the vertical axis is proportion of all LO events accounted for in this time window. The three types of LO events counted in the figure are:

- Type 1: placed or cancelled.
- Type 2: placed or cancelled at a price equal to or better than the current best bid or ask.
- Type 3: placed or cancelled and causes the midprice to change.

The difference between a Type 2 order and a Type 3 order is that a Type 2 order does not necessarily have to change the midprice of the asset. If a limit buy order is placed at the best bid, or if an existing order at the best bid is cancelled (but other orders still remain at that price), then this would be a Type 2 order but not a Type 3 order. The inclusions $\text{Type 3} \subset \text{Type 2} \subset \text{Type 1}$ hold.

Panel (c) in Figure 5 shows the proportion of the trading day accounted for immediately following MO arrivals. The calculation of these proportions are performed as follows: let $[0, T]$ be the time interval of interest, panel (c) is all of January 2014 with the first and last 30 minutes of each day removed (we assume it is one connected interval by concatenating trading hours together so that $T = 415800$ seconds). Let $\{t_k\}_{k=1}^N$ be the sequence of times when MOs occur, and let $\{\tau_k^i\}_{k=1}^{N^i}$ be the sequence of times when LO events of type i occur. Let $\Delta t > 0$ represent the lag time after an MO for which its effect on LO activity may be significant. For each value of the lag, Δt , define the union of time intervals $I_{\Delta t} = \cup_{k=1}^N [t_k, t_k + \Delta t)$. The quantity $I_{\Delta t}$ represents the amount of time in $[0, T]$ that takes place less than Δt after an MO. Finally, define $P_{\Delta t}^i = \text{card}(\{\tau_k^i\}_{k=1}^{N^i} \cap I_{\Delta t})$. This quantity represents the number of LO events of type i which occur within Δt after an MO. Panels (a) and (b) plot $P_{\Delta t}^i/N^i$ as a function of Δt for each type of LO event, and panel (c) plots $m(I_{\Delta t})/T$ as a function of Δt for the two equities of interest, where m is the Lebesgue measure.

As Figure 5 shows, a disproportionate amount of all LO activity occurs shortly following an MO. To put in perspective the significance of the disproportion, note that less than 0.5% of the trading day is accounted for in the short time (50 ms) following all MOs.

Figure 5: Proportion of all LO events that occur shortly after an MO (panels a and b). Proportion of trading day as percentage accounted for immediately following MO (panel c). Data are taken from a full month of trading in January 2014.

However, between approximately 40% and 90% of all LO activity (whether defined by type 1, 2 or 3) is accounted for in the same time period.

The combination of using a time lag of 10 ms in computing midprice changes and the behavior shown in Figures 4 produces a notable feature in the distribution of price changes. Most of the mass in these distributions occurs at the 0 and ± 1 tick levels, with less probability of observing a ± 0.5 tick movement. Consider the scenario when the spread is one tick and an incoming MO fills the full volume at the best quote. Soon after, there are three possible events: i) no new LOs are added within the new spread, ii) an LO is added within the new spread to replenish the order that was filled, or iii) an LO is added within the new spread, but on the other side of the book. If the original MO is a buy (sell), the midprice change in each scenario is i) 0.5 (-0.5) ticks, ii) 0 ticks, or iii) 1 (-1) tick, respectively. The drastic increase in LO activity immediately following an MO along with the fact that INTC and ORCL are large tick stocks explains why scenario i) is much more rare than scenarios ii) or iii).

The above empirical results clearly show that volume imbalance conveys important information about arrival rate of MOs and the distribution of prices. Here we have examined a simple way to summarize LOB information on buying and selling pressure, but depending on the particular asset, or how agents post intentions to sell or buy using passive orders in other types of markets, it might be useful to synthesize volume imbalance information in a different way.

Moreover, regardless of the way in which buying and selling pressure is synthesized, trading algorithms could be considerably improved if this information is included. For example, in market making algorithms, or those where the investor aims at profiting from roundtrip trades, it is clear that basing her trading on volume imbalance protects

the strategy from adverse selection costs and increases expected profits as a result of better inventory management and by tilting the LOs to take advantage of expected moves in prices.

3. Trading Algorithm with Volume Imbalance Information

In this section we develop a trading algorithm where the objective is to maximize profits by posting LOs to make roundtrip trades. While the LOB dynamics we propose are appropriate over a wide range of equities, we tailor the control processes to large tick stocks and assume that LOs are posted only at the best bid and/or the best offer – for small tick stocks the strategy needs to be adjusted so that it is possible to post LOs within the spread. In Subsection 3.1 we present model ingredients such as the dynamics for: MOs, midprice, and volume imbalance. In Subsection 3.2 we present and solve the investor’s investment problem. In Subsection 3.4 we discuss the optimal trading strategy. Finally, we devote Section 4 to discussing the performance of the trading algorithm using market data for INTC and ORCL.

3.1. Market orders, midprice dynamics, volume imbalance

We work with a complete filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $E = \mathbb{R}^+ \times \mathbb{R}^3$, $\mathcal{E}$ be the Borel sets of E , and μ^l , μ^+ , and μ^- be three doubly stochastic Poisson random measures (PRMs) on $(E, \mathcal{E})$ with compensators ν^l , ν^+ , and ν^- . More details on their structure are outlined in subsection 3.1.1. Below we describe how the components of the LOB, including LOs, MOs, midprice, and spread evolve.

The PRMs count events in the book due to LO activity (corresponding to μ^l) and MO activity (corresponding to $\mu^\pm$), and thus drive all of the dynamics in the model. The core idea of the model is that the arrival of one of these events (LO or MO) may cause a jump in each of the midprice, imbalance, and spread simultaneously (this corresponds to integrals being taken over three spacial dimensions in what follows). The distribution of each of these random jumps after an event is embedded in the compensators ν^l , ν^+ , and ν^- .

Market Orders: The number of market buy (sell) orders that have occurred up to time t is denoted by M_t^+ (M_t^-). These quantities are

$$M_t^\pm = \int_0^t \int_{\mathbf{y} \in \mathbb{R}^3} \mu^\pm(d\mathbf{y}, du).$$

Limit Order Events: We are only concerned with an LO placement or cancellation

if it causes the midprice, spread, or imbalance regime (defined below) to change. If we refer to these as Type 4 events, then with the description of other types as in Section 2.2 we have $\text{Type 3} \subset \text{Type 4} \subset \text{Type 2}$. We denote the number of Type 4 events as

$$M_t^l = \int_0^t \int_{\mathbf{y} \in \mathbb{R}^3} \mu^l(d\mathbf{y}, du).$$

Midprice: The midprice is modelled by the pure jump process

$$S_t = S_0 + \int_0^t \int_{\mathbf{y} \in \mathbb{R}^3} y_1(\mu^l + \mu^+ - \mu^-)(d\mathbf{y}, du), \quad (2)$$

where the value of $y_1 \in \mathbb{R}$ represents the changes in midprice due to market activity.

Volume Imbalance: Rather than modelling the imbalance process ρ_t directly, which can be rather noisy, we instead consider a finite state imbalance regime process $Z_t \in \{1, \dots, n_Z\}$. This process acts as an approximation, (and smoothing) to the true value of imbalance by dividing the interval $[-1, 1]$ into n_Z subintervals (labelled such that subinterval 1 corresponds to the largest negative values (i.e. the sell heavy regime), and n corresponds to the largest positive values (i.e. the buy heavy regime) and $Z_t = k$ corresponds to ρ_t lying within subinterval k . The process Z_t is given by

$$Z_t = Z_0 + \int_0^t \int_{\mathbf{y} \in \mathbb{R}^3} (y_2 - Z_{u-})(\mu^l + \mu^+ + \mu^-)(d\mathbf{y}, du), \quad (3)$$

where the value of $y_2 \in \mathbb{R}$ represents the state of the imbalance regime process immediately following any market activity.

Spread: The spread between the best bid and best offer is modelled as a finite state process which takes values in $\{1, \dots, n_\Delta\}$ and is given by

$$\Delta_t = \Delta_0 + \int_0^t \int_{\mathbf{y} \in \mathbb{R}^3} (y_3 - \Delta_{u-})(\mu^l + \mu^+ + \mu^-)(d\mathbf{y}, du), \quad (4)$$

where the value of $y_3 \in \mathbb{R}$ represents the value of the spread immediately following any market activity.

3.1.1. Properties of the doubly stochastic PRMs

The compensators of μ^l , μ^+ , and μ^- are denoted by ν^l , ν^+ , and ν^- and are chosen to be of the form $\nu^i(d\mathbf{y}, dt) = \lambda^i(Z_t, \Delta_t) F_{Z_t, \Delta_t}^i(d\mathbf{y}) dt$, where for any (Z, Δ)

$$\int_{\mathbf{y} \in \mathbb{R}^3} F_{Z, \Delta}^i(d\mathbf{y}) = 1. \quad (5)$$

Furthermore, μ^l , μ^+ , and μ^- are independent conditional on (Z_t, Δ_t) , specifically

$$\lim_{h \rightarrow 0} \frac{1}{h} \mathbb{P}(M_{t+h}^i - M_t^i = 1 \cap M_{t+h}^j - M_t^j = 1 \mid Z_t, \Delta_t) = \delta_{i,j} \sqrt{\lambda^i(Z_t, \Delta_t) \lambda^j(Z_t, \Delta_t)}, \quad (6)$$

where $\delta_{i,j}$ is the Kronecker delta and where i and j are any of l , $+$, or $-$. This property ensures that any single change in S_t , Z_t , and Δ_t is a result of only one PRM. If the distributions $F_{Z, \Delta}^i$ are chosen to have support only on $y_2 \in \{1, \dots, n_Z\}$ and $y_3 \in \{1, \dots, n_\Delta\}$, then from the definitions in (3) and (4) we have $Z_t \in \{1, \dots, n_Z\}$ and $\Delta_t \in \{1, \dots, n_\Delta\}$ with probability 1. Furthermore, (6), along with the factored form of the compensators, makes the vector process (Z_t, Δ_t) a continuous-time Markov chain. In the sequel it may occasionally be convenient to refer to the pair imbalance and spread as a single state in which case it is represented by $\mathbf{J} = (Z, \Delta)$.

Proposition 1 (Generator of (Z_t, Δ_t)). *The generator matrix G of the process $\mathbf{J}_t = (Z_t, \Delta_t)$ is given by*

$$G_{\mathbf{J}, \mathbf{K}} = \lambda^l(\mathbf{J}) T_{\mathbf{J}, \mathbf{K}}^l + \lambda^+(\mathbf{J}) T_{\mathbf{J}, \mathbf{K}}^+ + \lambda^-(\mathbf{J}) T_{\mathbf{J}, \mathbf{K}}^-, \quad \text{for } \mathbf{J} \neq \mathbf{K}, \quad (7)$$

$$G_{\mathbf{J}, \mathbf{J}} = -\lambda^l(\mathbf{J}) (1 - T_{\mathbf{J}, \mathbf{J}}^l) - \lambda^+(\mathbf{J}) (1 - T_{\mathbf{J}, \mathbf{J}}^+) - \lambda^-(\mathbf{J}) (1 - T_{\mathbf{J}, \mathbf{J}}^-),$$

for $\mathbf{J} \in \{1, \dots, n_Z\} \times \{1, \dots, n_\Delta\}$, where

$$T_{\mathbf{J}, \mathbf{K}}^i = \int_{\mathbf{y} \in \mathbb{R}^3} \mathbf{1}_{(y_2, y_3) = \mathbf{K}} F_{\mathbf{J}}^i(d\mathbf{y}).$$

Proof See Appendix A.1.

Finally, the distribution $F_{Z, \Delta}^l(d\mathbf{y})$ is chosen to have support only on the two hyperplanes $y_1 = \pm \frac{y_3 - \Delta}{2}$. This is because if a single LO placement or cancellation causes the spread to change, then the midprice must also necessarily change by an amount which is exactly

half of the change in the spread. We do not impose this constraint on $F_{Z,\Delta}^\pm$, even though the change in spread and midprice due to a single MO event must have the same relationship. The reason for not imposing this constraint is that immediately following an MO, there is an increase in the intensity of LO activity which has an effect on the spread and midprice. This extra LO activity (as depicted above in Figure 5) can be viewed as a result of the MO immediately prior, and so the changes in spread and midprice can be attributed to the MO arrival itself.

3.2. Agent's optimization problem

As discussed above, see for example Table 1 (and Table B.8 in Appendix), we observe that MOs hardly ever walk beyond the best quote – this is a common feature in large tick stocks. Thus, here we assume that the agent posts LOs only at-the-touch, i.e. the best bid or best ask. In reality, the agent has LOs posted at many prices in the LOB deeper than the best bid/ask so that if those prices were to become the best bid or ask in the future, the agent's orders at that time are closer to the front of the queue. This is also why we assume that the agent's LOs posted at the best bid or ask are always executed by an incoming MO. It is possible to generalize the results to account for a probability of being filled which is state dependent and less than 1, but the qualitative behavior of the optimal strategy does not change.

The agent begins with a certain amount of wealth x_0 which changes as MOs arrive and fill her LOs and her cash process $X = (X_t)_{0 \leq t \leq T}$ satisfies the SDE

$$dX_t = \gamma_t^+ \left(S_{t^-} + \frac{\Delta_{t^-}}{2} \right) dM_t^+ - \gamma_t^- \left(S_{t^-} - \frac{\Delta_{t^-}}{2} \right) dM_t^-, \quad X_0 = x_0,$$

where the processes $\gamma^\pm = (\gamma_t^\pm)_{0 \leq t \leq T}$, which take on values 0 or 1, are the agent's controls for whether she is posting buy and/or sell orders. Recall that $M_t^\pm$ are the counting processes for market buy and sell orders.

As well, the agent's inventory $q = (q_t)_{0 \leq t \leq T}$ satisfies the SDE

$$dq_t = -\gamma_t^+ dM_t^+ + \gamma_t^- dM_t^-.$$

The dynamics proposed above incorporate some implicit assumptions related to the agent's control processes $(\gamma_t^\pm)_{0 \leq t \leq T}$. First, we do not include a feedback effect of the agent's orders into the dynamics of the LOB. In principle, this could be done by allowing λ^i and F^i (the intensity of order book events and distribution of subsequent price changes) to depend on $\gamma^\pm$. The form of the feedback controls given below and our method of numerically solving for them would be the same, and a feedback effect

would not alter the qualitative behaviour of the optimal strategy. However, an empirical estimation of the magnitude of the feedback effect would be very difficult and beyond the scope of this paper.

Second, we have assumed that the fill probability of the agent's orders is equal to one. One possible approach to include random fills is to draw a Bernoulli variable at the time of each MO with a parameter that depends on other state variables (spread and imbalance). This modification would induce a lower effective λ^+ and λ^- in the HJB equation given below. Once again, the method of solving the equation would be unchanged and the qualitative behavior would be similar to that which we show already. In one set of our historical simulations, we impose a fill probability less than one as a function of the state of imbalance. The results still show significant improvement of the trading strategy over the zero-intelligence case.

Finally, we assume that the agent has zero latency – market changes are followed by immediate repositioning of LO's. In reality, for an agent with colocated servers, a latency on the order of $10ms$ is not unrealistic. After estimating the intensity of MO's conditioned on each state of spread and imbalance, we see that typical arrival rates are slower than $1/10ms$ (see Appendix C.1 for parameters estimated from data).

We consider an agent who aims to find the strategy $\gamma^\pm$ which maximizes expected terminal wealth subject to additional penalties. In particular, her value function is given by

$$H(t, x, q, S, \mathbf{J}) = \sup_{(\gamma_s^\pm)_{t \leq s \leq T} \in \mathcal{A}} \mathbb{E} \left[X_T + q_T (S_T - \ell(q_T, \Delta_T)) - \phi \int_t^T q_u^2 du \middle| \mathcal{F}_t \right], \quad (8)$$

where T is the terminal time of the strategy, $\ell(q_T, \Delta_T)$ is a liquidation penalty, $\phi \geq 0$ is a running inventory penalty parameter, and $\mathcal{A}$ is the set of admissible strategies which are $\mathcal{F}_t$ -predictable such that the inventory remains bounded between finite upper and lower bounds: $\underline{Q} \leq q_t \leq \overline{Q}$. If $q_t = \underline{Q}$ the agent does not post LOs on the sell side of the book, i.e. must choose $\gamma_t^+ = 0$, and similarly, when $q_t = \overline{Q}$ she chooses $\gamma_t^- = 0$.

The liquidation penalty $\ell(q, \Delta)$: is increasing in q , its absolute value is increasing in Δ , and $\ell(0, \Delta) = 0$. This penalty reflects the costs from liquidating terminal inventory with a single MO (which may walk the LOB) and encourages the agent to implement a strategy where q_T is close to 0.

The term $\phi \int_t^T q_u^2 du$ represents a running inventory penalty which discourages the agent from holding non-zero inventory positions for significant lengths of time. The parameter ϕ can be used as a risk control by determining the strength of this penalty. A large value of ϕ makes the agent act in a more conservative manner by inducing the strategy

to unwind positions, long or short, very quickly. See Guilbaud and Pham (2013) and Cartea and Jaimungal (2015) for similar inventory control terms, and Cartea et al. (2017) who show that this running penalty arises in the context of ambiguity aversion when the agent is not confident about the drift of the midprice.

3.3. Feedback control of the optimal strategy

To find the optimal strategy, consider the Hamilton-Jacobi-Bellman (HJB) equation associated with the value function H of the control problem (8)

$$\partial_t H - \phi q^2 + \lambda^l(\mathbf{J}) \mathbb{E}[\mathcal{D}^l H | \mathbf{J}] + \sup_{\gamma^+ \in \{0,1\}} \lambda^+(\mathbf{J}) \mathbb{E}[\mathcal{D}_{\gamma^+}^+ H | \mathbf{J}] + \sup_{\gamma^- \in \{0,1\}} \lambda^-(\mathbf{J}) \mathbb{E}[\mathcal{D}_{\gamma^-}^- H | \mathbf{J}] = 0, \quad (9)$$

subject to the terminal condition $H(T, x, q, S, \mathbf{J}) = x + q(S - \ell(q, \Delta))$. Here the operator $\mathcal{D}^l$ acts as follows:

$$\mathcal{D}^l H(t, x, q, S, \mathbf{J}) = \int_{\mathbf{y} \in \mathbb{R}^3} \left(H(t, x, q, S + y_1, y_2, y_3) - H(t, x, q, S, \mathbf{J}) \right) \mu^l(d\mathbf{y}, \cdot),$$

and the expectation operator conditional on $\mathbf{J}$ acts as

$$\mathbb{E}[\mathcal{D}^l H | \mathbf{J}] = \int_{\mathbf{y} \in \mathbb{R}^3} \left(H(t, x, q, S + y_1, y_2, y_3) - H(t, x, q, S, \mathbf{J}) \right) F_{\mathbf{J}}^l(d\mathbf{y}). \quad (10)$$

Similarly, the operators $\mathcal{D}_{\gamma^\pm}^\pm$ are defined by

$$\mathcal{D}_{\gamma^\pm}^\pm H = \int_{\mathbf{y} \in \mathbb{R}^3} \left(H \left(t, x \pm \gamma^\pm \left(S \pm \frac{\Delta}{2} \right), q \mp \gamma^\pm, S \pm y_1, y_2, y_3 \right) - H \right) \mu^\pm(d\mathbf{y}, \cdot),$$

where we have suppressed explicit dependence on $(t, x, q, S, \mathbf{J})$ for easier readability. These have conditional expectation

$$\mathbb{E}[\mathcal{D}_{\gamma^\pm}^\pm H | \mathbf{J}] = \int_{\mathbf{y} \in \mathbb{R}^3} \left(H \left(t, x \pm \gamma^\pm \left(S \pm \frac{\Delta}{2} \right), q \mp \gamma^\pm, S \pm y_1, y_2, y_3 \right) - H \right) F_{\mathbf{J}}^\pm(d\mathbf{y}). \quad (11)$$

Equation (10) represents the expected change in the value function due to the arrival of an LO event (a new LO or cancellation of an existing order). Similarly, equations (11) represent the expected change in the value function due to the arrival of a buy or sell MO, depending on the agent's strategy $\gamma^\pm$ at the time of the arrival of the MO.

We make the ansatz $H(t, x, q, S, \mathbf{J}) = x + qS + h(t, q, \mathbf{J})$ where each term carries the following interpretation: the cash position x , the book value of the shares which are

marked-to-market at the midprice, and $h(t, q, \mathbf{J})$ which represents the extra value of optimally trading until the terminal date. Inserting this into the HJB (9), after a number of tedious computations, leads to a system of equations for h :

$$\begin{aligned}
& \partial_t h - \phi q^2 + \lambda^l(\mathbf{J}) (q \epsilon^l(\mathbf{J}) + \Sigma^l(t, q, \mathbf{J})) \\
& + \sup_{\gamma^+ \in \{0,1\}} \lambda^+(\mathbf{J}) \left(\gamma^+ \frac{\Delta}{2} + (q - \gamma^+) \epsilon^+(\mathbf{J}) + \Sigma_{\gamma^+}^+(t, q, \mathbf{J}) \right) \\
& + \sup_{\gamma^- \in \{0,1\}} \lambda^-(\mathbf{J}) \left(\gamma^- \frac{\Delta}{2} - (q + \gamma^-) \epsilon^-(\mathbf{J}) + \Sigma_{\gamma^-}^-(t, q, \mathbf{J}) \right) = 0, \\
& h(T, q, \mathbf{J}) = -q \ell(q, \Delta),
\end{aligned} \tag{12}$$

where

$$\begin{aligned}
\epsilon^i(\mathbf{J}) &= \sum_{y_1, y_2, y_3} y_1 F_{\mathbf{J}}^i(y_1, y_2, y_3), \\
\Sigma^l(t, q, \mathbf{J}) &= \sum_{\mathbf{K}} (h(t, q, \mathbf{K}) - h(t, q, \mathbf{J})) T_{\mathbf{J}, \mathbf{K}}^l, \\
\Sigma_{\gamma^\pm}^\pm(t, q, \mathbf{J}) &= \sum_{\mathbf{K}} (h(t, q \mp \gamma^\pm, \mathbf{K}) - h(t, q, \mathbf{J})) T_{\mathbf{J}, \mathbf{K}}^\pm.
\end{aligned}$$

The terms $\Sigma^i(t, q, \mathbf{J})$ represent the expected change in the value function due to state transitions. The terms $\epsilon^i(\mathbf{J})$ represent the expected midprice change after the arrival of an event. This identification, allows us to interpret the remaining contributions to equation (12).

Terms of the form $q \epsilon^i(\mathbf{J})$ represent the instantaneous expected change of the value of the agent's holdings due to market activity of type i (where i can be l , $+$, or $-$). If this product is positive, then the agent is in a state where her inventory holdings are expected to increase in value. If it is negative, then the current market state dictates her inventory holdings are expected to decrease in value. The terms $-\gamma^\pm \epsilon^\pm(\mathbf{J})$ represent the immediate costs incurred due to adverse selection after having an LO filled (note that $\epsilon^\pm(\mathbf{J})$ is generally positive according to Figure 4 and the signs of the PRMs in (2)). Finally, the terms $\gamma^\pm \frac{\Delta}{2}$ are the immediate profits made by the agent due to the arrival of an MO. Consequently, the relative sizes of $\frac{\Delta}{2}$ and $\epsilon^\pm(\mathbf{J})$ are an important indicator of the agent's strategy. We discuss this point below when analyzing the optimal trading strategy.

Proposition 2 (Existence of solution). *Equation (12) together with its terminal condition has a unique classical solution.*

Proof See Appendix A.2.

Recall that the control processes are restricted to values of 0 or 1. From (12), the resulting optimal feedback controls are given by a simple comparison:

$$\gamma^+(t, q, \mathbf{J}) = \begin{cases} 1, & \frac{\Delta}{2} - \epsilon^+(\mathbf{J}) + \Sigma_1^+(t, q, \mathbf{J}) > \Sigma_0^+(t, q, \mathbf{J}), \quad \text{and } q \neq \underline{Q}, \\ 0, & \text{otherwise}, \end{cases} \quad (13a)$$

$$\gamma^-(t, q, \mathbf{J}) = \begin{cases} 1, & \frac{\Delta}{2} - \epsilon^-(\mathbf{J}) + \Sigma_1^-(t, q, \mathbf{J}) > \Sigma_0^-(t, q, \mathbf{J}), \quad \text{and } q \neq \overline{Q}, \\ 0, & \text{otherwise}. \end{cases} \quad (13b)$$

Due to the discrete nature of the controls, there is little hope of finding a closed-form solution for the value function, however, (12) can be easily solved numerically.

The form of the optimal controls in (13a) and (13b), together with some reasonable foresight about the qualitative behavior of the optimal strategy, reveals a simple rule of thumb. It is expected that for a fixed state of imbalance and spread, $\mathbf{J}$, there is a boundary in the space of the variables (t, q) such that if q lies above (below) this boundary, then the agent optimally posts a limit sell (buy) order. The sign of $\frac{\Delta}{2} - \epsilon^\pm(\mathbf{J})$ plays a large part in determining whether this boundary lies above or below the curve $q = 0$. The terms $\Sigma_1^\pm(t, q, \mathbf{J})$ and $\Sigma_0^\pm(t, q, \mathbf{J})$ act as corrections based on the future liquidation penalty and running inventory penalty. If there is no inventory constraint or penalization, and the time to maturity is very large, the terms $\Sigma_1^\pm(t, q, \mathbf{J})$ and $\Sigma_0^\pm(t, q, \mathbf{J})$ are equal and the strategy is determined by the sign of $\frac{\Delta}{2} - \epsilon^\pm(\mathbf{J})$. This last expression is the expected amount of cash that the agent earns if the position is unwound at the midprice.

Theorem 3 (Verification Theorem). *Let h be the solution to (12) and define $\hat{H}(t, x, q, S, \mathbf{J}) = x + qS + h(t, q, \mathbf{J})$. Then the candidate solution $\hat{H}$ equals the value function H as defined in (8).*

Proof See Appendix A.3.

3.4. Optimal Posting Strategy

To show a typical posting strategy, we select a model with 3 states in the volume imbalance regime process, and the stock trades with a spread of 1 or 2 ticks, i.e. $n_Z = 3$ and $n_\Delta = 2$. For a full set of parameters, see the Appendix C.1. These parameters result from the estimation procedure outlined in Section 4.1 when applied to 30 minutes of INTC data starting at 10:00 on January 24, 2014. Parameters of most interest are

shown below. The components of these vectors represent states, $\mathbf{J} = (Z, \Delta)$ with the following convention: the first three components correspond to $\Delta = 1$ and $Z = 1, 2, 3$. The last three components correspond to $\Delta = 2$ and $Z = 1, 2, 3$.

$$\begin{aligned}\hat{\lambda}^+ &= (0.038 \ 0.121 \ 0.600 \ 1.711 \ 1.188 \ 3.042), \\ \hat{\lambda}^- &= (0.600 \ 0.121 \ 0.038 \ 3.042 \ 1.188 \ 1.711), \\ \hat{\epsilon}^+ &= (0.240 \ 0.306 \ 0.713 \ -0.111 \ 0.152 \ 0.469), \\ \hat{\epsilon}^- &= (0.713 \ 0.306 \ 0.240 \ 0.469 \ 0.152 \ -0.111).\end{aligned}$$

These parameters indicate the intensity of MOs and expected midprice change immediately following an MO for each state. Intensities are displayed in units of *second*⁻¹ and midprice changes are in ticks. Note that higher imbalance regimes (i.e. the book is tilted to the buy side) generally indicate a higher frequency of market buy orders and also larger expected (positive) midprice changes after buy MOs. High imbalance also indicates lower frequency of market sell orders and smaller jump sizes after sell orders. This is consistent with market behavior shown in Figures 3 and 4.

Figure 6 shows the optimal strategy corresponding to sell and buy LOs. In the top panel the spread is $\Delta = 1$ tick and in the bottom panel is $\Delta = 2$ ticks. In both panels, the light grey region corresponds to posting a limit sell order, the dark grey region a limit buy order, and the black region indicates an overlap where both a buy and sell order are posted. Finally, the white region indicates that no order should be placed in the book.

There are several features in this figure to note. As imbalance becomes more buy-heavy, the buy (sell) region becomes larger (smaller) in anticipation of the asset price to increase. Limit sell orders become less desirable because they are more likely to incur immediate losses due to adverse selection. Any desire to sell the asset in a buy-heavy imbalance regime is due to the liquidation penalty, to be paid at time T , and the running penalty, which add up to a cost that outweighs the gains from trading in the direction of the expected increase in prices.

In both the buy-heavy and sell-heavy regimes, the boundary between the sell and no sell regions generally lies considerably above or below, respectively, the zero inventory level. Note that in the sell-heavy regime the boundary is significantly below the level of zero inventory. This is consistent with rule of thumb, discussed after Proposition 2, to sell the asset in the sell-heavy regime when $\frac{\Delta}{2} - \epsilon^+(\mathbf{J}) > 0$ because the earnings of the half-spread are larger than the expected price increase immediately following the MO. Similarly, in the buy-heavy regime, the boundary is significantly above the zero inventory level. In this case it is usually not optimal to post a sell order because the gains from having the limit sell order filled are lower than the expected increase in midprice, $\frac{\Delta}{2} - \epsilon^+(\mathbf{J}) < 0$. In this region the agent is better off not posting because on

Figure 6: Dark grey (light grey) region is where buy (sell) LOs are posted. Black region is overlap where both buy and sell orders are be posted. White region indicates no-trade region, where the agent does not have LOs in the book.

average she loses the half-spread she earns to adverse selection costs.

For a fixed imbalance regime, the sell region is larger for larger values of Δ . This implies that within a fixed imbalance regime, the agent is more willing to post orders when the spread is large. This does not always have to be the case, but for the particular set of parameters chosen note that $\frac{\Delta}{2} - \epsilon^+(\mathbf{J})$ is larger for $\Delta = 2$ ticks than for $\Delta = 1$ tick.

Finally, the top panel of the figure shows two large no-trade regions in the sell-heavy and buy-heavy imbalance regimes. In these regions, the agent withdraws from the market. It is optimal to let current inventory appreciate in value because the expected future value of what the agent is holding outweighs the earnings from an additional trade net of any inventory costs (running penalty or terminal liquidation). The agent remains withdrawn from the market until the imbalance or spread change, or until the passage of time takes the strategy close enough to the terminal date T , so the future cost of having to cross the spread to liquidate outweighs the benefit of inventory appreciation. We also observe a small no-trade region in the neutral regime when inventory is close to zero and the strategy approaches the terminal date T .

4. Performance of Strategy on the Nasdaq Exchange

In this section we analyze the performance of the strategy on trade data from the Nasdaq exchange for INTC and ORCL. Our data set consists of all messages sent to Nasdaq’s LOB during all trading days in 2014. We employ the first six months of data to calibrate model parameters and test the strategy’s performance over the following six months. We assume that the agent’s trading window is half-hour intervals. That is, the agent sets a horizon of 30 mins over which she trades and closes all positions at the end of the window. We record the PnL of the strategy for every half-hour and compute the Sharpe ratio for a range of values of the running inventory parameter ϕ .

The choice of a 30 minute trading horizon is arbitrary and may have an effect on the profitability of the strategy. Some stocks are more active than others over that window so the chances of making roundtrip trades varies, thus one could choose the trading window based on the arrival rate of MOs and the average of traded volume. Alternatively, it is possible to sow together the various 30 minute periods to create a strategy that trades all day long and only unwinds at the end of the day. Such a strategy should outperform the one presented here. Nonetheless, the approach taken here allows us to demonstrate the historical value of adding imbalance as a state variable, and also allows us to generate many return periods from which we can compute meaningful estimates of returns and variability of returns.

We organize the rest of this section as follows. In Subsection 4.1 we show how the model parameters are estimated. In Subsection 4.2 we discuss how to forecast and update the model parameters which are used in the execution of the out-of-sample strategy. Finally, in Subsection 4.3 we present the out-of-sample performance of the strategy.

4.1. Estimating Model Parameters

To compute the optimal trading strategy, we solve (12) and then use (13a) and (13b) to find in which regions LOs are posted. To solve these equations we do not require all model parameters. Rather, for each Z and Δ and each $i = +, -, l$, we only require the quantities $\lambda^i(\mathbf{J})$, $\epsilon^i(\mathbf{J}) = \sum_{y_1, y_2, y_3} y_1 F_{\mathbf{J}}^i(y_1, y_2, y_3)$, and $T_{\mathbf{J}, \mathbf{K}}^i = \sum_{y_1} F_{\mathbf{J}}^i(y_1, \mathbf{K})$. The value of $\lambda^i(\mathbf{J})$ is the rate at which an event of type i occurs when the current imbalance and spread are equal to Z and Δ respectively, recall that $\mathbf{J} = (Z, \Delta)$. The value of $\epsilon^i(\mathbf{J})$ is the expected change in midprice immediately after an event of type i when the event occurs in state $\mathbf{J}$. Finally, the value of $T_{\mathbf{J}, \mathbf{K}}^i$ is the probability of transitioning from state $\mathbf{J}$ to state $\mathbf{K}$ after an event of type i . The sum over y_1 represents averaging out all possible midprice changes so the remaining quantity represents transitions between imbalance and spread only.

Estimating the model parameters over a time window is a straightforward procedure. Moreover, to ensure that the model does not contain any long term speculation based on the average growth of the asset midprice, we impose symmetry constraints between complementary states of imbalance. For example if rate of arrival of buy MOs is higher than arrival of sell MOs then a long term strategy is to buy and hold the asset because the model predicts that in the long term prices will increase.

The complementary imbalance state of Z is defined to be $\tilde{Z} = n_Z - Z + 1$. We also use the short-hand notation $\tilde{\mathbf{J}} = (\tilde{Z}, \Delta)$. The constraints we impose are the following:

$$\lambda^l(\mathbf{J}) = \lambda^l(\tilde{\mathbf{J}}), \quad (14a)$$

$$\lambda^\pm(\mathbf{J}) = \lambda^\mp(\tilde{\mathbf{J}}), \quad (14b)$$

$$\epsilon^l(\mathbf{J}) = -\epsilon^l(\tilde{\mathbf{J}}), \quad (14c)$$

$$\epsilon^\pm(\mathbf{J}) = \epsilon^\mp(\tilde{\mathbf{J}}), \quad (14d)$$

$$T_{\mathbf{J},\mathbf{K}}^l = T_{\tilde{\mathbf{J}},\tilde{\mathbf{K}}}^l, \quad (14e)$$

$$T_{\mathbf{J},\mathbf{K}}^\pm = T_{\tilde{\mathbf{J}},\tilde{\mathbf{K}}}^\mp. \quad (14f)$$

We remark that we impose this symmetry to preclude long term speculation strategies and acknowledge that many high-frequency trading strategies are designed to benefit from short-lived trends in the midprice. There is no contradiction in a model where the long term growth of the midprice is symmetric as discussed here, but also exhibits short term deviations around the midprice's long term trend. For example there are models that 'seek alpha', see Cartea et al. (2014), and models that take advantage of how innovations in order flow affect the trend in prices, see for instance ?, ?.

Let $N_{\mathbf{J}}^i$ denote the number of events of type i that occurred from 0 to T when the imbalance and spread are equal to $\mathbf{J}$. Also let $\tau_{\mathbf{J}}$ be the total occupation time of state $\mathbf{J}$ so that $\sum_{\mathbf{J}} \tau_{\mathbf{J}} = T$. Then $\lambda^i(\mathbf{J})$ is estimated by

$$\hat{\lambda}^l(\mathbf{J}) = \frac{N_{\mathbf{J}}^l + N_{\tilde{\mathbf{J}}}^l}{\tau_{\mathbf{J}} + \tau_{\tilde{\mathbf{J}}}}, \quad \text{and} \quad \hat{\lambda}^\pm(\mathbf{J}) = \frac{N_{\mathbf{J}}^\pm + N_{\tilde{\mathbf{J}}}^\mp}{\tau_{\mathbf{J}} + \tau_{\tilde{\mathbf{J}}}}.$$

Further, let $y_{\mathbf{J}}^{i1}, \dots, y_{\mathbf{J}}^{iN_{\mathbf{J}}^i}$ be all the midprice changes that occur following an event of type i from state $\mathbf{J}$. Then

$$\hat{\epsilon}^l(\mathbf{J}) = \frac{\sum_{k=1}^{N_{\mathbf{J}}^l} y_{\mathbf{J}}^{lk} - \sum_{k=1}^{N_{\tilde{\mathbf{J}}}^l} y_{\tilde{\mathbf{J}}}^{lk}}{N_{\mathbf{J}}^l + N_{\tilde{\mathbf{J}}}^l}, \quad \text{and} \quad \hat{\epsilon}^\pm(\mathbf{J}) = \frac{\sum_{k=1}^{N_{\mathbf{J}}^\pm} y_{\mathbf{J}}^{\pm k} + \sum_{k=1}^{N_{\tilde{\mathbf{J}}}^\mp} y_{\tilde{\mathbf{J}}}^{\mp k}}{N_{\mathbf{J}}^\pm + N_{\tilde{\mathbf{J}}}^\mp}.$$

Finally, let $N_{\mathbf{J},\mathbf{K}}^i$ be the number of times that the state transitions from $\mathbf{J}$ to $\mathbf{K}$. Then

$$\hat{T}_{\mathbf{J},\mathbf{K}}^l = \frac{N_{\mathbf{J},\mathbf{K}}^l + N_{\mathbf{J},\widetilde{\mathbf{K}}}^l}{N_{\mathbf{J}}^l + N_{\widetilde{\mathbf{J}}}^l}, \quad \text{and} \quad \hat{T}_{\mathbf{J},\mathbf{K}}^\pm = \frac{N_{\mathbf{J},\mathbf{K}}^\pm + N_{\mathbf{J},\widetilde{\mathbf{K}}}^\pm}{N_{\mathbf{J}}^\pm + N_{\widetilde{\mathbf{J}}}^\pm}.$$

Note that these quantities satisfy (14a) to (14f). These are also the expressions that result from a maximum likelihood estimation with the constraints imposed (see Donnelly (2014)). The parameters used in the numerical example of the previous section were obtained through this procedure, and so by construction they satisfy the constraints which results in the symmetry between buy and sell regions seen in Figure 6.

4.2. Forecasting Model Parameters

Testing the trading strategy on out-of-sample data requires forecasts of model parameters for each half-hour interval based on the estimated parameters from previous intervals – recall that the trading horizon is $T = 30$ mins. We begin by labelling each interval according to the day of the year and the time of day in which it takes place. The data consist of all the messages sent to the Nasdaq exchange in all trading days in 2014, giving a total of 249 days consisting of 13 half-hour intervals, i.e. $T = 30$ mins, in each trading day. We identify each interval as the pair (m, n) where m denotes the day in the year and n denotes the half-hour of the day, thus $m \in \{1, \dots, 249\}$ and $n \in \{1, \dots, 13\}$.

We employ the first 6 months of data, i.e. the first 124 days, to calibrate all model parameters for each half hour in the day. These parameter estimates are used to develop a predictive model so that model parameters in a given 30 minute period (over which we aim to backtest out-of-sample) are estimated from the data we already observed in the last 30 minute period. We do this in two stages described in detail below.

Let $\boldsymbol{\lambda}_{m,n}^i = (\lambda_{m,n}^i(1), \dots, \lambda_{m,n}^i(n_{Z,\Delta}))'$ and $\boldsymbol{\epsilon}_{m,n}^i = (\epsilon_{m,n}^i(1), \dots, \epsilon_{m,n}^i(n_{Z,\Delta}))'$, where $n_{Z,\Delta} = n_Z \times n_\Delta$, and $'$ is the transpose operation. In the first stage, for each trading period in the first six months of the year, January to June ($m = 1, \dots, 124$), we estimate all of the model parameters $\hat{\boldsymbol{\lambda}}_{m,n}^i$ and $\hat{\boldsymbol{\epsilon}}_{m,n}^i$. Thus, we have in-sample estimates of the model parameters (for the first 6 months), and a different model for each day and each trading period.

In Figures 7 and 8 we show motivation for the forecasting method we propose below. When calibrating the model parameters over each 30 minute period, we observe a well known intraday seasonality effect. Figure 7 shows the arrival intensity of MOs, conditional on the imbalance regime, which is generally higher at the beginning and end of the day giving a characteristic ‘U’ shape. Moreover, Figure 8 shows the size of midprice

Figure 7: Intraday seasonality of MO intensity for ORCL. Each curve is a different day in the month of January 2014. Each point on a single curve is the intensity of market buy orders conditional on the state of imbalance over a 30 minute period. The thick black curve is the average intensity within the same window of all days in the sample. Similar behavior is shown for INTC.

changes following an MO. The pattern is generally high at the beginning of the day and low at the end of the day giving a characteristic ‘S’ shape.

In the second stage, we assume a factor model that relates the parameters of one trading interval n to the parameters of the previous trading interval $n - 1$ on the same day m :

$$\hat{\lambda}_{m,n}^i = \alpha_n^{\lambda^i} + \beta_n^{\lambda^i} \hat{\lambda}_{m,n-1}^i + \varepsilon_{m,n}^{\lambda^i}, \quad (15)$$

$$\hat{\epsilon}_{m,n}^i = \alpha_n^{\epsilon^i} + \beta_n^{\epsilon^i} \hat{\epsilon}_{m,n-1}^i + \varepsilon_{m,n}^{\epsilon^i}, \quad (16)$$

where $\alpha_n^{\lambda^i}, \alpha_n^{\epsilon^i} \in \mathbb{R}^{n_Z, \Delta}$, $\beta_n^{\lambda^i}, \beta_n^{\epsilon^i} \in \mathbb{R}^{n_Z, \Delta \times n_Z, \Delta}$, and all of the idiosyncratic error terms $\varepsilon_{m,n}^{\lambda^i}$ and $\varepsilon_{m,n}^{\epsilon^i}$ are independent with mean zero.

Using the model parameters estimated in-sample for the first 6 months, we then perform a multilinear regression to obtain the estimates $\hat{\alpha}_n^{\lambda^i}$ and $\hat{\alpha}_n^{\epsilon^i}$, and the factor loadings $\hat{\beta}_n^{\lambda^i}$ and $\hat{\beta}_n^{\epsilon^i}$.

To test the strategy’s performance, the next step is to forecast the parameters that are used in the out-of-sample period. For day m we employ the data in the first 30 minute period (i.e. $n = 1$ which is the first half hour of the day) to estimate the model parameters $\hat{\lambda}_{m,1}^i$ and $\hat{\epsilon}_{m,1}^i$, but do not trade over this first half-hour of the day. We then use the regression model to predict the model parameters for the next period (in which we trade). Denote these forecasts by $\tilde{\lambda}_{m,2}^i$ and $\tilde{\epsilon}_{m,2}^i$, which are the parameters used to formulate the trading strategy over the second 30 minute period. After the second 30 minute period has passed, we use the observed data for the second period to estimate $\hat{\lambda}_{m,2}^i$ and $\hat{\epsilon}_{m,2}^i$, and once again use the regression model to forecast the parameters of the third period (over which we trade), and so on. Succinctly, we use the estimates resulting

Figure 8: Intraday seasonality of midprice changes following an MO for ORCL. Each curve is a different day in the month of January 2014. Each point on a single curve is the average midprice change following a market buy order (measured in ticks) conditional on the state of imbalance over a 30 minute period. The thick black curve is the average midprice change within the same window of all days in the sample. Similar behavior is shown for INTC.

from the regression (15) and (16) to forecast the parameters for half-hour $n = 2, \dots, 13^1$ in day m using the recursion:²

$$\begin{aligned}\tilde{\lambda}_{m,n}^i &= \hat{\alpha}_n^{\lambda^i} + \hat{\beta}_n^{\lambda^i} \hat{\lambda}_{m,n-1}^i, \\ \tilde{\epsilon}_{m,n}^i &= \hat{\alpha}_n^{\epsilon^i} + \hat{\beta}_n^{\epsilon^i} \hat{\epsilon}_{m,n-1}^i.\end{aligned}$$

We do not perform this type of regression and forecasting for the transition probabilities $T_{\mathbf{J},\mathbf{K},m,n}^i$ because numerical investigation has shown that the trading strategy depends much more heavily on ϵ^i and λ^i . Instead, the value of $T_{\mathbf{J},\mathbf{K},m,n}^i$ is set equal to the mean of $T_{\mathbf{J},\mathbf{K},m,n}^i$ taken over all $m \in \{1, \dots, 124\}$.

4.3. Out-of-sample Performance of Trading Strategy

Now that we have a recursive method to forecast model parameters for the interval (m, n) , based on the parameters of interval $(m, n - 1)$, we test the out-of-sample performance of the trading strategy for all days in July to December 2014. Our forecast

¹Although we are able to use our forecasting method to obtain parameters for $n = 13$, we do not test the trading strategy in the last 30 minute interval as market behaviour is notably different towards the end of the day.

²After computing the forecast of $\tilde{\lambda}_{m,n}^i$ it is possible we obtain a negative value. In this case, we set $\tilde{\lambda}_{m,n}^i = 10^{-2}$. We do not set it equal to zero because we do not want to impose that a state is absorbing. We also impose a maximum value on each $\tilde{\lambda}_{m,n}^i$ of 10. This is solely for purposes of numerical stability when solving equation (12) and does not have a significant qualitative impact on the resulting optimal strategy.

method does not allow us to test the strategy in the first half-hour of each day, and we also exclude testing in the last half-hour of each day as market behavior is notably different towards the end of the trading day. This leaves us with a total of 1375 half-hour intervals to perform an out-of-sample test of the strategy. At the end of each half-hour trading period, any non-zero inventory is liquidated by crossing the spread, consistent with the use of the penalty term $\ell(q, \Delta) = \text{sign}(q) \frac{\Delta}{2}$, which assumes that there is enough liquidity at the best quote to fill the agent's MO.

Here we fix the maximum inventory constraint at $\bar{Q} = -\underline{Q} = 50$ and choose a range of values of the running penalty parameter ϕ to see how it affects the PnL. Moreover, we examine how the strategy performs when volume imbalance is modelled assuming $n_Z = 1, 3, 5$ states, and fixing $n_\Delta = 2$. Recall that volume imbalance $\rho_t \in [-1, 1]$ and that this is divided into n_Z subintervals, and n_Δ is the grid on which the spread lives. When $n_Z = 1$ there is only one volume imbalance state and the midprice can either be in a state where the spread is one or two ticks. The case $n_Z = 1$ can also be considered a benchmark in addition to the zero-intelligence strategy. Comparing the performance of the strategies with $n_Z = 3$ and $n_Z = 5$ reveals the significance of observing imbalance.

Figure 9 shows the annualized mean PnL versus annualized standard deviation of the resulting wealth over all trading intervals for INTC (left panel) and ORCL (right panel). As expected, the worst performance is when $n_Z = 1$ and the results for n_Z equal to 3 and 5 are approximately the same. For both stocks we observe that increasing the value of the running penalty parameter ϕ increases the performance of the strategy: expected PnL increases and its standard deviation decreases.

Figure 10 shows another perspective of the results. It depicts the Sharpe ratio of the strategy as a function of the parameter ϕ , and the risk-free rate is zero. We observe that as the agent enforces stricter controls on inventory by increasing ϕ , the Sharpe ratio of the strategy also increases.

The results in Figures 9 and 10 should be compared to those of the zero-intelligence strategy. It is clear that employing the information impounded in the volume imbalance process, i.e. MO arrival and price innovations, considerably boosts the profits of an investment strategy based on roundtrip trades. The sources of the profits stem from protecting the strategy from adverse selection costs, and positioning the LOs to take advantage of price movements. In addition, we also observe that imposing a running inventory penalty enhances the profitability of the strategy as in Guilbaud and Pham (2013) and Cartea and Jaimungal (2015).

In Table 2 we show the annualized Sharpe ratio for the zero-intelligence strategy when applied to a set of large tick stocks for various values of the inventory constraint Q (these results take place under Scenario 1 as discussed in Section 2.1). These values

(a) INTC

(b) ORCL

Figure 9: Annualized Expectation vs. Standard Deviation of trading strategies based on different number of imbalance states and $\bar{Q} = -\underline{Q} = 50$. Each point represents a different value of ϕ , ranging from 0 to 10^{-5} (larger values of ϕ correspond to smaller values of standard deviation).

(a) INTC

(b) ORCL

Figure 10: Annualized Sharpe ratio vs. ϕ of trading strategies based on different number of imbalance states and $\bar{Q} = -\underline{Q} = 50$, and the risk-free rate is zero.

show similar behaviour to what is seen in Figure 1 indicating the failure to make positive profits through roundtrip trades when always posting LOs on both sides of the LOB.

In contrast, Tables 3, 4, and 5 show the annualized Sharpe ratio over the same set of equities when employing our strategy with volume imbalance and spread information. The results are for $n_Z = 1, 3, 5$, and the spread is 1 or 2 ticks, i.e. $n_\Delta = 2$. We observe a large boost in performance across the board. The results in Table 3 do not incorporate observations of imbalance in the strategy since $n_Z = 1$, but executions are dependent on the spread. This explains the relatively poor performance when compared to the results in Tables 4 and 5. Further, the Sharpe ratio generally increases with an increase in the running inventory parameter ϕ .

	Maximum Inventory Q							
	1	2	4	10	20	40	100	200
AA	-29.62	-33.77	-31.57	-21.89	-15.00	-10.18	- 6.97	- 6.41
AMAT	-34.20	-37.90	-33.66	-22.27	-15.15	-11.60	- 8.81	- 7.47
ARCC	-14.46	-16.07	-14.11	-11.28	- 9.13	- 6.95	- 5.30	- 4.85
BXS	-33.87	-27.24	-22.25	-17.15	-14.49	-12.12	-11.96	-11.96
CSCO	- 8.50	-18.03	-20.56	-15.71	-11.24	- 9.26	- 7.04	- 6.18
EBAY	-17.11	-20.55	-21.21	-18.08	-13.36	- 8.58	- 5.55	- 4.17
FMER	-33.12	-34.11	-28.45	-18.63	-13.44	-11.30	-10.25	-10.22
IMGN	-49.92	-40.68	-29.62	-18.39	-13.70	-10.54	- 8.33	- 7.90
INTC	-20.97	-26.15	-26.53	-23.11	-17.37	-11.98	- 8.24	- 6.90
NTAP	-45.96	-44.68	-37.70	-23.71	-15.08	- 9.75	- 8.06	- 8.04
ORCL	-20.24	-28.23	-28.36	-20.89	-13.59	- 9.17	- 6.71	- 5.53

Table 2: Annualized Sharpe ratio of zero-intelligence trading strategies based on various values of Q .

	Inventory Penalty Parameter ϕ							
	0	10^{-7}	$2 \cdot 10^{-7}$	$4 \cdot 10^{-7}$	10^{-6}	$2 \cdot 10^{-6}$	$4 \cdot 10^{-6}$	10^{-5}
AA	-3.78	- 6.42	- 5.93	- 5.06	- 3.68	- 0.71	3.00	8.75
AMAT	-6.09	- 8.65	- 8.84	- 8.64	- 7.73	- 5.28	- 1.74	4.80
ARCC	-4.94	- 8.41	- 8.52	- 9.06	- 8.73	- 7.16	- 4.09	- 0.53
BXS	1.50	2.68	3.29	3.78	4.49	5.21	5.81	6.01
CSCO	-4.18	- 8.81	- 9.71	-10.23	-10.06	- 8.65	- 5.92	- 0.49
EBAY	0.90	1.27	1.76	2.50	3.86	5.24	7.07	10.28
FMER	1.51	6.40	7.99	9.82	12.90	16.26	20.08	23.81
IMGN	1.40	- 0.91	- 1.08	- 1.25	- 1.22	- 1.20	- 0.85	- 0.07
INTC	-7.52	-16.04	-16.69	-17.19	-17.36	-16.24	-14.43	- 9.42
NTAP	0.50	1.82	1.81	1.51	1.88	2.60	3.17	4.99
ORCL	0.88	5.36	7.13	9.67	13.13	16.65	20.58	26.30

Table 3: Annualized Sharpe ratio of trading strategies based on various values of ϕ , the number of imbalance states here is $n_Z = 1$ and $n_\Delta = 2$ and the maximum inventory constraint is $Q = 50$.

	Inventory Penalty Parameter ϕ							
	0	10^{-7}	$2 \cdot 10^{-7}$	$4 \cdot 10^{-7}$	10^{-6}	$2 \cdot 10^{-6}$	$4 \cdot 10^{-6}$	10^{-5}
AA	6.11	9.24	11.13	12.74	17.82	22.10	24.90	28.32
AMAT	5.32	10.03	12.09	14.62	17.95	20.99	24.56	28.64
ARCC	- 0.85	- 0.41	0.19	0.73	2.10	3.44	5.39	8.18
BXS	2.66	2.86	3.11	3.68	3.93	4.70	6.49	7.96
CSCO	8.19	10.11	11.67	13.86	17.99	21.04	24.82	30.78
EBAY	3.13	5.51	6.28	7.49	9.43	11.17	14.14	18.17
FMER	5.31	7.00	7.92	9.30	12.05	14.45	17.02	22.02
IMGN	2.08	1.65	1.62	1.67	1.98	1.92	2.20	3.60
INTC	1.90	8.04	9.94	12.51	15.96	18.70	21.38	25.13
NTAP	1.98	2.32	2.33	2.73	3.10	4.34	5.88	7.64
ORCL	11.65	16.76	18.74	20.86	23.71	26.94	29.17	34.25

Table 4: Annualized Sharpe ratio of trading strategies based on various values of ϕ , the number of imbalance states here is $n_Z = 3$ and $n_\Delta = 2$ and the maximum inventory constraint is $Q = 50$.

	Inventory Penalty Parameter ϕ							
	0	10^{-7}	$2 \cdot 10^{-7}$	$4 \cdot 10^{-7}$	10^{-6}	$2 \cdot 10^{-6}$	$4 \cdot 10^{-6}$	10^{-5}
AA	6.12	8.35	9.79	12.06	17.34	20.33	23.99	27.93
AMAT	6.81	9.49	10.91	12.79	16.53	19.67	23.72	28.49
ARCC	-0.66	-0.01	0.73	1.40	3.07	4.35	5.85	9.34
BXS	4.25	4.50	4.78	5.08	6.45	6.43	7.57	9.42
CSCO	9.48	11.60	12.91	15.38	18.94	22.73	27.14	32.85
EBAY	3.41	4.59	6.00	7.74	9.96	12.42	16.12	20.91
FMER	5.24	6.57	7.02	8.05	10.48	12.69	15.77	20.71
IMGN	2.83	2.75	2.73	2.43	2.05	2.06	2.49	3.15
INTC	4.49	9.68	11.25	12.85	15.60	17.93	20.88	24.45
NTAP	1.63	2.29	2.63	2.73	2.74	3.79	4.53	7.32
ORCL	11.75	14.28	16.09	18.32	21.90	25.20	28.99	33.29

Table 5: Annualized Sharpe ratio of trading strategies based on various values of ϕ , the number of imbalance states here is $n_Z = 5$ and $n_\Delta = 2$ and the maximum inventory constraint is $Q = 50$.

So far the strategy assumes that the agent's LOs are always at the front of the queue. This assumption is unrealistic and introduces a bias in the profitability of the trading strategy. We address this issue by assigning a fill probability to the LOs. This probability is estimated for each stock separately based on trade data from January 2014 (first and last 30 minutes of data removed from each day). Let $N_Z^\pm$ be the number of market buy (+) and sell (-) orders during this time period which occurred when the imbalance regime is equal to Z . Also let $E_Z^\pm$ be the number of these market orders which are large

enough to fill the entire first price queue. Then fill probabilities are assigned as

$$p_Z^\pm = \frac{E_Z^+ + E_{\tilde{Z}}^-}{N_Z^+ + N_{\tilde{Z}}^-}. \quad (17)$$

This definition imposes the symmetry condition $p_Z^+ = p_{\tilde{Z}}^-$.

The resulting probabilities for the sell side of the LOB (p_Z^+) and $n_Z = 5$ are shown in Table 6. In the case of EBAY for example, if $\rho \in [-\frac{1}{5}, \frac{1}{5}]$ indicating a neutral market in which $Z = 3$, then the agent's LO is filled with probability $p_3^+ = p_3^- = 0.28$. If $\rho \in (\frac{1}{5}, \frac{3}{5}]$ indicating a buy-heavy market in which $Z = 4$, then a buy (sell) MO fills an agent's posted sell (buy) LO with probability $p_4^+ = 0.4$ ($p_4^- = 0.2$). If $\rho \in (\frac{3}{5}, 1]$ indicating a strongly buy-heavy market in which $Z = 5$, then a buy (sell) MO fills an agent's posted sell (buy) LO with probability $p_5^+ = 0.65$ ($p_5^- = 0.18$). The probabilities for imbalance states $Z = 1$ and $Z = 2$ behave similarly.

	Imbalance Regime Z				
	1	2	3	4	5
AA	0.04	0.09	0.16	0.23	0.36
AMAT	0.03	0.04	0.11	0.17	0.37
ARCC	0.05	0.06	0.11	0.18	0.38
BXS	0.34	0.31	0.50	0.69	0.87
CSCO	0.02	0.04	0.08	0.12	0.29
EBAY	0.18	0.20	0.28	0.40	0.65
FMER	0.16	0.17	0.26	0.41	0.72
IMGN	0.35	0.41	0.54	0.71	0.86
INTC	0.05	0.05	0.09	0.15	0.32
NTAP	0.17	0.21	0.31	0.44	0.70
ORCL	0.08	0.09	0.19	0.32	0.52

Table 6: Proportion of MOs that fill entire first ask queue in January 2014. These quantities are interpreted as minimal fill probabilities for an LO posted based on the imbalance regime at the time of the order.

These probabilities are implemented within the historical backtest by drawing a Bernoulli random variable with appropriate success probability at the time of each MO. If this variable is equal to 0, then the MO does not fill the agent's LO. Since these quantities as computed correspond to MOs which fill the entire best price queue, they underestimate the true probability of an LO being filled. These probabilities are only accurate if the LO is at the back of the queue and any position closer to the front would have a larger probability of fill.

The resulting Sharpe ratios of the strategy with $n_Z = 5$ are shown in Table 7. As expected, the Sharpe ratios of the strategy are lower than those shown in Table 5, where

the only difference is the fill rates, but the ratios are still considerable higher than those obtained with the zero-intelligence strategy. We consider this reported performance of the strategy to be a pessimistic view due to the minimal nature of the fill probability estimation, while the true performance should fall somewhere between that of Table 5 and Table 7.

	Inventory Penalty Parameter ϕ							
	0	10^{-7}	$2 \cdot 10^{-7}$	$4 \cdot 10^{-7}$	10^{-6}	$2 \cdot 10^{-6}$	$4 \cdot 10^{-6}$	10^{-5}
AA	1.95	1.68	1.55	-0.34	0.89	2.19	1.38	3.09
AMAT	0.46	1.03	3.74	0.41	0.17	2.73	3.00	4.53
ARCC	-4.43	-3.79	-4.75	-5.02	-2.28	-3.34	-4.40	-3.18
BXS	-0.01	1.43	-0.35	0.14	2.93	1.64	2.67	1.46
CSCO	0.59	1.57	0.94	1.35	2.81	1.99	5.38	5.09
EBAY	-0.31	2.82	0.85	2.89	3.35	4.43	5.16	6.84
FMER	1.10	0.95	1.54	1.46	1.76	1.69	2.47	4.42
IMGN	1.36	1.64	1.61	1.20	0.60	0.54	-0.64	-0.26
INTC	-1.95	-0.90	-0.58	-0.86	0.17	-1.35	1.27	2.40
NTAP	0.44	1.84	2.13	0.27	0.04	-0.43	0.74	1.87
ORCL	3.98	3.52	5.89	5.00	4.59	9.25	10.57	12.17

Table 7: Annualized Sharpe ratio of trading strategies with modified fill probabilities depending on the level of imbalance. The trading strategy here is based on a number of imbalance states equal to $n_Z = 5$ and $n_\Delta = 2$ and the maximum inventory constraint is $Q = 50$.

We finish this section by adding a few remarks about the strategy’s performance. First, for the purposes of this historical test, and indeed for the entire optimization problem as stated, here we focused on an agent who utilizes LOs on both sides of the market. The performance of the strategy is expected to improve if the agent also uses MOs to take advantage of price innovations which are anticipated by the volume imbalance process. Moreover, the use of volume imbalance can be implemented in a vast range of other optimal execution problems. For example, algorithms to acquire/liquidate a large number of shares, as well as ‘Pairs Trading’ algorithms.

Finally, we acknowledge that measures based on the volume of shares resting in the LOB, or number of LOs in the LOB, may be subject to manipulation. In most equity markets it is costless to cancel or amend an LO and although illegal, market participants may artificially tilt the LOB by submitting LOs on one side of the book, only to withdraw them quickly enough whilst taking advantage of traders that employ strategies like the one described, and those suggested, in this paper. This illegal activity is commonly referred to as ‘spoofing’ and very difficult to detect in modern electronic exchanges where a large number of market participants employ different trading strategies and some can trade at extremely low latencies. In the particular cases discussed here, one cannot tell if throughout 2014 the stocks we study are subject to spoofing at any point in time. Even if this is the case, the statistical properties of volume imbalance as a

predictor of MO activity and price innovations seem robust to this manipulation of LOs which would certainly erode the predictive power of volume imbalance.

5. Conclusions

We employ message data from the Nasdaq exchange to build a measure of volume imbalance and show that it predicts type of MO arrival (buy or sell) as well as predict price innovations. This measure captures buying and selling pressure in the LOB. In particular we show that when the LOB is buy-heavy the probability of the next MO being a buy order is much higher than it being a sell order. Moreover, we also show that when the LOB is buy-heavy, price revisions after MO arrival are on average large and positive, i.e. larger than the average price revision seen when the LOB is neither buy-heavy nor sell-heavy. We observe the same empirical behavior when the book is sell-heavy. This relationship between volume imbalance and market order activity is consistent across a large sample of Nasdaq stocks we examined.

This measure is simple to build and can be employed by a wide range of trading algorithms. A clear effect of using this information in a trading model is to reduce adverse selection costs and to take advantage of favorable price movements.

As an example of an algorithmic trading strategy we solve an optimal investment problem of an agent who provides liquidity to the LOB. The agent's objective is to maximize expected terminal wealth by completing roundtrip trades whilst penalizing inventory positions. We use trade data to show the out-of-sample performance of the strategy during the period 1 July to 31 December 2014. We demonstrate that including the volume imbalance process considerably boosts the strategy's process and the Sharpe ratio of the strategy is considerable higher than that obtained by most models proposed in the literature, i.e. Avellaneda and Stoikov (2008), Guéant et al. (2012), Fodra and Labadie (2012), Cartea and Jaimungal (2015), and Cartea et al. (2017).

Appendix A. Proofs

Appendix A.1. Proof of Proposition 1

Proof of Proposition 1 Define the process $F_t = \mathbb{E}[\Phi(\mathbf{J}_T)|\mathcal{F}_t] = f(t, \mathbf{J}_t)$ for some function Φ to be specified later. Using Ito's Lemma and the fact that F_t is a martingale,

we have that f satisfies the following equation

$$\begin{aligned}
& \partial_t f(t, \mathbf{J}) + \int_{\mathbf{y} \in \mathbb{R}^3} \left(f(t, y_2, y_3) - f(t, \mathbf{J}) \right) \left(\sum_i \lambda^i(\mathbf{J}) F_{\mathbf{J}}^i(d\mathbf{y}) \right) = 0 \\
& \partial_t f(t, \mathbf{J}) + \sum_{\mathbf{K}} \int_{\mathbf{y} \in \mathbb{R}^3} \left(f(t, y_2, y_3) - f(t, \mathbf{J}) \right) \mathbb{1}_{(y_2, y_3) = \mathbf{K}} \left(\sum_i \lambda^i(\mathbf{J}) F_{\mathbf{J}}^i(d\mathbf{y}) \right) = 0 \\
& \partial_t f(t, \mathbf{J}) + \sum_{\mathbf{K}} \left(f(t, \mathbf{K}) - f(t, \mathbf{J}) \right) \int_{\mathbf{y} \in \mathbb{R}^3} \mathbb{1}_{(y_2, y_3) = \mathbf{K}} \left(\sum_i \lambda^i(\mathbf{J}) F_{\mathbf{J}}^i(d\mathbf{y}) \right) = 0 \\
& \partial_t f(t, \mathbf{J}) + \sum_{\mathbf{K}} \left(f(t, \mathbf{K}) - f(t, \mathbf{J}) \right) \left(\sum_i \lambda^i(\mathbf{J}) T_{\mathbf{J}, \mathbf{K}}^i \right) = 0.
\end{aligned}$$

Letting $[\mathbf{f}(t)]_{\mathbf{J}} = f(t, \mathbf{J})$ and using the matrix G introduced in the proposition, the matrix form of this equation is

$$\partial_t \mathbf{f} + G \mathbf{f} = 0,$$

which has solution

$$\mathbf{f}(t) = e^{G(T-t)} \Phi.$$

If $\Phi(\mathbf{J})$ is defined as the indicator function for a particular state $\mathbf{J}^0$, then $\mathbf{f}(t)$ is a vector of transition probabilities to state $\mathbf{J}^0$ from each arbitrary state $\mathbf{J}$ over a time of $T - t$. This is the same expression which dictates transition probabilities for a continuous time Markov chain with generator matrix G .

Appendix A.2. Proof of Proposition 2

Proof of Proposition 2 Consider h to be a vector valued function of t , where the components are indexed by $(q, \mathbf{J})$. Then equation (12) can be written in the form:

$$\partial_t \mathbf{h} = \mathbf{F}(\mathbf{h}),$$

where $\mathbf{F}$ is a vector valued function of a vector argument. Component $(q, \mathbf{J})$ of $\mathbf{F}$ is

$$\begin{aligned}
\mathbf{F}_{q, \mathbf{J}}(\mathbf{h}) = & \phi q^2 - \lambda^l(\mathbf{J}) (q \epsilon^l(\mathbf{J}) + \Sigma^l(q, \mathbf{J})) \\
& - \lambda^+(\mathbf{J}) \max \left\{ q \epsilon^+(\mathbf{J}) + \Sigma_0^+(q, \mathbf{J}), \frac{\Delta}{2} + (q-1) \epsilon^+(\mathbf{J}) + \Sigma_1^+(q, \mathbf{J}) \right\} \\
& - \lambda^-(\mathbf{J}) \max \left\{ -q \epsilon^-(\mathbf{J}) + \Sigma_0^-(q, \mathbf{J}), \frac{\Delta}{2} - (q+1) \epsilon^-(\mathbf{J}) + \Sigma_1^-(q, \mathbf{J}) \right\},
\end{aligned} \tag{A.1}$$

with

$$\begin{aligned}
\epsilon^i(\mathbf{J}) &= \sum_{y_1, y_2, y_3} y_1 F_{\mathbf{J}}^i(y_1, y_2, y_3), \\
\Sigma^l(q, \mathbf{J}) &= \sum_{\mathbf{K}} (\mathbf{h}_{q, \mathbf{K}} - \mathbf{h}_{q, \mathbf{J}}) T_{\mathbf{J}, \mathbf{K}}^l, \\
\Sigma_{\gamma}^{-}(q, \mathbf{J}) &= \mathbb{1}_{q < \bar{Q}} \sum_{\mathbf{K}} (\mathbf{h}_{q+\gamma, \mathbf{K}} - \mathbf{h}_{q, \mathbf{J}}) T_{\mathbf{J}, \mathbf{K}}^{-}, \\
\Sigma_{\gamma}^{+}(q, \mathbf{J}) &= \mathbb{1}_{q > \underline{Q}} \sum_{\mathbf{K}} (\mathbf{h}_{q-\gamma, \mathbf{K}} - \mathbf{h}_{q, \mathbf{J}}) T_{\mathbf{J}, \mathbf{K}}^{+}.
\end{aligned}$$

Each component of $\mathbf{F}$ is a piecewise linear function of $\mathbf{h}$, the coefficients depending on the four possible regions corresponding to which terms in (A.1) achieve the maxima. Since inventory, imbalance, and spread are all bounded, there are a finite number of components in $\mathbf{F}$ and the total number of regions for which the coefficients of linear dependence are different is also finite. Therefore, $\mathbf{F}$ is globally Lipschitz with linear growth, and equation (12) has a unique classical solution.

Appendix A.3. Proof of Theorem 3

Proof of Theorem 3 Let h be the solution to equation (12) and define a candidate optimal value function $\hat{H}(t, x, q, S, \mathbf{J}) = x + qS + h(t, q, \mathbf{J})$. From Ito's Lemma we have

$$\begin{aligned}
\hat{H}(T, X_{T-}^{\gamma^{\pm}}, q_{T-}^{\gamma^{\pm}}, S_{T-}, \mathbf{J}_{T-}) &= \hat{H}(t, x, q, S, \mathbf{J}) + \int_t^T \partial_t h(u, X_u^{\gamma^{\pm}}, q_u^{\gamma^{\pm}}, S_u, \mathbf{J}_u) du \\
&+ \int_t^T \int_{\mathbf{y} \in \mathbb{R}^3} q_u^{\gamma^{\pm}} y_1 + h(t, q_u^{\gamma^{\pm}}, y_2, y_3) - h(t, q_u^{\gamma^{\pm}}, \mathbf{J}_u) \mu^l(d\mathbf{y}, du) \\
&+ \int_t^T \int_{\mathbf{y} \in \mathbb{R}^3} \gamma_u^+ \frac{\Delta_u}{2} + (q_u^{\gamma^{\pm}} - \gamma_u^+) y_1 + h(t, q_u^{\gamma^{\pm}} - \gamma_u^+, y_2, y_3) - h(t, q_u^{\gamma^{\pm}}, \mathbf{J}_u) \mu^+(d\mathbf{y}, du) \\
&+ \int_t^T \int_{\mathbf{y} \in \mathbb{R}^3} \gamma_u^- \frac{\Delta_u}{2} - (q_u^{\gamma^{\pm}} + \gamma_u^-) y_1 + h(t, q_u^{\gamma^{\pm}} + \gamma_u^-, y_2, y_3) - h(t, q_u^{\gamma^{\pm}}, \mathbf{J}_u) \mu^-(d\mathbf{y}, du).
\end{aligned}$$

Taking an expectation conditional on $\mathcal{F}_t$ on both sides and rearranging yields

$$\begin{aligned}
\hat{H}(t, x, q, S, \mathbf{J}) &= \mathbb{E} \left[\hat{H}(T, X_{T-}^{\gamma^\pm}, q_{T-}^{\gamma^\pm}, S_{T-}, \mathbf{J}_{T-}) \middle| \mathcal{F}_t \right] - \mathbb{E} \left[\int_t^T \partial_t h(u, q_u^{\gamma^\pm}, \mathbf{J}_u) du \right. \\
&\quad + \int_t^T \int_{\mathbf{y} \in \mathbb{R}^3} q_u^{\gamma^\pm} y_1 + h(t, q_u^{\gamma^\pm}, y_2, y_3) - h(t, q_u^{\gamma^\pm}, \mathbf{J}_u) \nu^l(d\mathbf{y}, du) \\
&\quad + \int_t^T \int_{\mathbf{y} \in \mathbb{R}^3} \gamma_u^+ \frac{\Delta_u}{2} + (q_u^{\gamma^\pm} - \gamma_u^+) y_1 + h(t, q_u^{\gamma^\pm} - \gamma_u^+, y_2, y_3) - h(t, q_u^{\gamma^\pm}, \mathbf{J}_u) \nu^+(d\mathbf{y}, du) \\
&\quad \left. + \int_t^T \int_{\mathbf{y} \in \mathbb{R}^3} \gamma_u^- \frac{\Delta_u}{2} - (q_u^{\gamma^\pm} + \gamma_u^-) y_1 + h(t, q_u^{\gamma^\pm} + \gamma_u^-, y_2, y_3) - h(t, q_u^{\gamma^\pm}, \mathbf{J}_u) \nu^-(d\mathbf{y}, du) \middle| \mathcal{F}_t \right].
\end{aligned} \tag{A.2}$$

Equation (12) then yields the inequality

$$\begin{aligned}
\hat{H}(t, x, q, S, \mathbf{J}) &\geq \mathbb{E} \left[\hat{H}(T, X_{T-}^{\gamma^\pm}, q_{T-}^{\gamma^\pm}, S_{T-}, \mathbf{J}_{T-}) - \phi \int_t^T (q_u^{\gamma^\pm})^2 du \middle| \mathcal{F}_t \right] \\
&= \mathbb{E} \left[\hat{H}(T, X_T^{\gamma^\pm}, q_T^{\gamma^\pm}, S_T, \mathbf{J}_T) - \phi \int_t^T (q_u^{\gamma^\pm})^2 du \middle| \mathcal{F}_t \right] \\
&= \mathbb{E} \left[X_T^{\gamma^\pm} + q_T^{\gamma^\pm} (S_T - \ell(q_T^{\gamma^\pm}, \Delta_T)) - \phi \int_t^T (q_u^{\gamma^\pm})^2 du \middle| \mathcal{F}_t \right].
\end{aligned}$$

Since this inequality holds for arbitrary controls $\gamma^\pm$, we have:

$$\begin{aligned}
\hat{H}(t, x, q, S, \mathbf{J}) &\geq \sup_{(\gamma_s^\pm)_{t \leq s \leq T} \in \mathcal{A}} \mathbb{E} \left[X_T^{\gamma^\pm} + q_T^{\gamma^\pm} (S_T - \ell(q_T^{\gamma^\pm}, \Delta_T)) - \phi \int_t^T (q_u^{\gamma^\pm})^2 du \middle| \mathcal{F}_t \right] \\
&= H(t, x, q, S, \mathbf{J}).
\end{aligned} \tag{A.3}$$

Now, if $\gamma^{\pm*}$ is selected to be of the form in (13a) and (13b), then (A.2) implies:

$$\begin{aligned}
\hat{H}(t, x, q, S, \mathbf{J}) &= \mathbb{E} \left[\hat{H}(T, X_{T-}^{\gamma^{\pm*}}, q_{T-}^{\gamma^{\pm*}}, S_{T-}, \mathbf{J}_{T-}) - \phi \int_t^T (q_u^{\gamma^{\pm*}})^2 du \middle| \mathcal{F}_t \right] \\
&= \mathbb{E} \left[\hat{H}(T, X_T^{\gamma^{\pm*}}, q_T^{\gamma^{\pm*}}, S_T, \mathbf{J}_T) - \phi \int_t^T (q_u^{\gamma^{\pm*}})^2 du \middle| \mathcal{F}_t \right] \\
&= \mathbb{E} \left[X_T^{\gamma^{\pm*}} + q_T^{\gamma^{\pm*}} (S_T - \ell(q_T^{\gamma^{\pm*}}, \Delta_T)) - \phi \int_t^T (q_u^{\gamma^{\pm*}})^2 du \middle| \mathcal{F}_t \right] \\
&\leq \sup_{(\gamma_s^\pm)_{t \leq s \leq T} \in \mathcal{A}} \mathbb{E} \left[X_T^{\gamma^\pm} + q_T^{\gamma^\pm} (S_T - \ell(q_T^{\gamma^\pm}, \Delta_T)) - \phi \int_t^T (q_u^{\gamma^\pm})^2 du \middle| \mathcal{F}_t \right] \\
&= H(t, x, q, S, \mathbf{J}).
\end{aligned} \tag{A.4}$$

Combining (A.3) and (A.4) yields the result.

Appendix B. Market Behavior Across Equities

In this section we show that several equities demonstrate the same behavior which is illustrated in Table 1 and Figures 3 and 4. We begin with Table B.8 which shows how many market orders are filled only at the best posted price versus how many walk through more than one full level of the book. We see that GOOG has the smallest percentage of orders remaining strictly within the best price at 91.6%, but this number is often higher than 99%. Generally, stocks that can be considered large tick exhibit a larger percentage of trades strictly within the best price.

Table B.9 shows the average trade intensities for buy and sell market orders depending on the state of volume imbalance. Generally we see that the frequency of buy (sell) orders increases as imbalance becomes more buy-heavy (sell-heavy).

Finally, in Table B.10 we show the average midprice change 10 ms after an MO that occurred within each state of imbalance. Once again we see a clear tendency for the midprice change to have a larger (smaller) magnitude after a market buy (sell) order when imbalance is more buy-heavy.

	First Tick Only		Beyond First Tick		$\mathbb{P}(V_{MO} \leq V_{LO})$
	Buys	Sells	Buys	Sells	
AAPL	100,362	105,655	4,581	4,527	0.958
FARO	1,745	2,374	64	109	0.960
GOOG	32,096	34,969	3,085	3,075	0.916
INTC	35,595	38,451	54	50	0.999
MMM	22,996	25,745	130	118	0.995
NTAP	28,519	27,118	104	123	0.996
ORCL	30,001	27,502	41	45	0.999
SMH	3,087	3,084	7	4	0.998

Table B.8: Number of MOs that touch only the first tick or go beyond the first tick. Data is taken from a full month of trading in January, 2014 (first and last 30 minutes of each day removed). The column labelled $\mathbb{P}(V_{MO} \leq V_{LO})$ is the probability that an MO has smaller volume than all limit orders posted at the best price, and hence only engages the best quote.

	Average Buy Intensity			Average Sell Intensity		
	Sell-Heavy	Neutral	Buy-Heavy	Sell-Heavy	Neutral	Buy-Heavy
AAPL	0.223	0.257	0.276	0.293	0.271	0.228
FARO	0.005	0.005	0.003	0.008	0.007	0.004
GOOG	0.060	0.090	0.112	0.102	0.095	0.071
INTC	0.024	0.048	0.232	0.288	0.056	0.021
MMM	0.034	0.058	0.080	0.072	0.063	0.046
NTAP	0.031	0.050	0.167	0.133	0.045	0.031
ORCL	0.022	0.041	0.181	0.213	0.037	0.015
SMH	0.002	0.004	0.020	0.017	0.004	0.002

Table B.9: Average trade intensities within each state of imbalance. Data is taken from a full month of trading in January, 2014 (first and last 30 minutes of each day removed).

Appendix C. Parameter Sets

The entire calibrated parameter set is available upon request.

Appendix C.1. Parameters of Section 3.4

The full set of parameters used to compute the optimal trading strategy in Figures 6 are given here. In this example, the number of imbalance regimes is $n_Z = 3$ and the number of spread values is $n_\Delta = 3$ giving a total of $n_Z n_\Delta = 6$ states. These states are ordered in the same convention as used previously: the first three states correspond to $\Delta = 1$ and $Z = 1, 2, 3$ and the last three states correspond to $\Delta = 2$ and $Z = 1, 2, 3$.

	Market Buy Order			Market Sell Order		
	Sell-Heavy	Neutral	Buy-Heavy	Sell-Heavy	Neutral	Buy-Heavy
AAPL	2.212	2.527	3.098	−3.245	−2.751	−2.388
FARO	1.846	3.047	4.251	−4.414	−3.781	−2.602
GOOG	8.115	8.866	10.619	−10.678	−9.191	−8.432
INTC	0.044	0.211	0.457	−0.494	−0.208	−0.057
MMM	1.134	1.507	1.613	−1.716	−1.496	−1.177
NTAP	0.270	0.498	0.723	−0.668	−0.477	−0.291
ORCL	0.031	0.290	0.570	−0.603	−0.376	−0.128
SMH	0.262	0.411	0.759	−0.730	−0.450	−0.217

Table B.10: Average midprice change (in units of ticks) after an MO within each state of imbalance. Data is taken from a full month of trading in January, 2014 (first and last 30 minutes of each day removed).

Due to the form of equation (12), the full form of each compensator is not required. Rather, we only require $\lambda^i(\mathbf{J})$, $\epsilon^i(\mathbf{J}) = \sum_{y_1, \mathbf{K}} y_1 F_{\mathbf{J}}^i(y_1, \mathbf{K})$, and $T_{\mathbf{J}, \mathbf{K}}^i = \sum_{y_1} F_{\mathbf{J}}^i(y_1, \mathbf{K})$.

$$\begin{aligned}
\hat{\lambda}^l &= (\quad 1.257 \quad 0.538 \quad 1.257 \quad 85.187 \quad 23.132 \quad 85.187 \quad), \\
\hat{\lambda}^+ &= (\quad 0.038 \quad 0.121 \quad 0.600 \quad 1.711 \quad 1.188 \quad 3.042 \quad), \\
\hat{\lambda}^- &= (\quad 0.600 \quad 0.121 \quad 0.038 \quad 3.042 \quad 1.188 \quad 1.711 \quad), \\
\hat{\epsilon}^l &= (\quad -0.137 \quad 0.000 \quad 0.137 \quad -0.160 \quad 0.000 \quad 0.160 \quad), \\
\hat{\epsilon}^+ &= (\quad 0.240 \quad 0.306 \quad 0.713 \quad -0.111 \quad 0.152 \quad 0.469 \quad), \\
\hat{\epsilon}^- &= (\quad 0.713 \quad 0.306 \quad 0.240 \quad 0.469 \quad 0.152 \quad -0.111 \quad).
\end{aligned}$$

$$\begin{aligned}
T^l &= \begin{bmatrix} 0.00 & 0.73 & 0.00 & 0.05 & 0.21 & 0.01 \\ 0.50 & 0.00 & 0.50 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.73 & 0.00 & 0.01 & 0.21 & 0.05 \\ 0.04 & 0.01 & 0.33 & 0.00 & 0.61 & 0.01 \\ 0.06 & 0.00 & 0.06 & 0.44 & 0.00 & 0.44 \\ 0.33 & 0.01 & 0.04 & 0.01 & 0.61 & 0.00 \end{bmatrix}, \\
T^+ &= \begin{bmatrix} 0.40 & 0.36 & 0.24 & 0.00 & 0.00 & 0.00 \\ 0.01 & 0.61 & 0.34 & 0.00 & 0.02 & 0.01 \\ 0.02 & 0.40 & 0.53 & 0.00 & 0.03 & 0.03 \\ 0.00 & 0.33 & 0.22 & 0.33 & 0.11 & 0.00 \\ 0.02 & 0.26 & 0.04 & 0.00 & 0.67 & 0.00 \\ 0.00 & 0.69 & 0.25 & 0.00 & 0.00 & 0.06 \end{bmatrix}, \\
T^- &= \begin{bmatrix} 0.53 & 0.40 & 0.02 & 0.03 & 0.03 & 0.00 \\ 0.34 & 0.61 & 0.01 & 0.01 & 0.02 & 0.00 \\ 0.24 & 0.36 & 0.40 & 0.00 & 0.00 & 0.00 \\ 0.25 & 0.69 & 0.00 & 0.06 & 0.00 & 0.00 \\ 0.04 & 0.26 & 0.02 & 0.00 & 0.67 & 0.00 \\ 0.22 & 0.33 & 0.00 & 0.00 & 0.11 & 0.33 \end{bmatrix}.
\end{aligned}$$

References

- Avellaneda, M. and S. Stoikov (2008, November). High-frequency trading in a limit order book. Quantitative Finance 8, 217–224.
- Bechler, K. and M. Ludkovski (2015). Optimal execution with dynamic order flow imbalance. SIAM Journal on Financial Mathematics 6(1), 1123–1151.
- Biais, B., P. Hillion, and C. Spatt (1995). An empirical analysis of the limit order book and the order flow in the paris bourse. The Journal of Finance 50(5), 1655–1689.
- Cartea, A., R. Donnelly, and S. Jaimungal (2017). Algorithmic trading with model uncertainty. SIAM Journal on Financial Mathematics 8(1), 635–671.
- Cartea, Á. and S. Jaimungal (2015). Risk metrics and fine tuning of high frequency trading strategies. Mathematical Finance 25(3), 576–611.
- Cartea, Á. and S. Jaimungal (2016a). A closed-form execution strategy to target VWAP. SIAM Journal on Financial Mathematics 7, 760–785.

- Cartea, Á. and S. Jaimungal (2016b). Incorporating order-flow into optimal execution. Mathematics and Financial Economics 10(3), 339–364.
- Cartea, Á., S. Jaimungal, and J. Ricci (2014). Buy low sell high: A high frequency trading perspective. SIAM Journal on Financial Mathematics 5(1), 415–444.
- Cont, R., A. Kukanov, and S. Stoikov (2013). The price impact of order book events. Journal of Financial Econometrics.
- Donnelly, R. F. (2014). Ambiguity aversion in algorithmic and high frequency trading. PhD thesis. SSRN eLibrary, <http://ssrn.com/abstract=2527808>.
- Fodra, P. and M. Labadie (2012). High-frequency market-making with inventory constraints and directional bets. Working papers, HAL.
- Foucault, T. (1999). Order flow composition and trading costs in a dynamic limit order market. Journal of Financial markets 2(2), 99–134.
- Guéant, O., C.-A. Lehalle, and J. Fernandez-Tapia (2012). Dealing with the inventory risk: a solution to the market making problem. Mathematics and Financial Economics, 1–31.
- Guilbaud, F. and H. Pham (2013). Optimal high frequency trading with limit and market orders. Quantitative Finance 13(1), 79–94.
- Huang, W., C.-A. Lehalle, and M. Rosenbaum (2015). Simulating and analyzing order book data: The queue-reactive model. Journal of the American Statistical Association 110(509), 107–122.
- Lipton, A., U. Pesavento, and M. G. Sotiropoulos (2013). Trade arrival dynamics and quote imbalance in a limit order book. [arXiv:1312.0514](https://arxiv.org/abs/1312.0514).
- Stoikov, S. and R. Waeber (2012). Optimal asset liquidation using limit order book information. Available at SSRN 2113827.