King's Research Portal

DOI:

[10.1080/1350486X.2018.1506257](https://doi.org/10.1080/1350486X.2018.1506257)

Document Version

Peer reviewed version

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

Citation for published version (APA):

Donnelly, R., & Gan, L. (2018). Optimal Decisions in a Time Priority Queue. *Applied Mathematical Finance*, 107-147. <https://doi.org/10.1080/1350486X.2018.1506257>

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.

Optimal Decisions in a Time Priority Queue[☆]

Ryan Donnelly^a, Luhui Gan^b

^a*Department of Applied Mathematics, University of Washington, Seattle, WA 98195, USA*

^b*Cubist Systematic Strategies, Singapore*

Abstract

We show how the position of a limit order in the queue influences the decision of whether to cancel the order or let it rest. Using ultra high-frequency data from the Nasdaq exchange, we perform empirical analysis on various limit order book events and propose novel ways for modelling some of these events, including cancellation of limit orders in various positions and size of market orders. Based on our empirical findings, we develop a queuing model that captures stylized facts on the data. This model includes a distinct feature which allows for a potentially random effect due to the agent's impulse control. We apply the queuing model in an algorithmic trading setting by considering an agent maximizing her expected utility through placing and cancelling of limit orders. The agent's optimal strategy is presented after calibrating the model to real data. A simulation study shows that for the same level of standard deviation of terminal wealth, the optimal strategy has a 2.5% higher mean compared to a strategy which ignores the effect of position; or a 8.8% lower standard deviation for the same level of mean. This extra gain stems from posting a limit order during adverse conditions and obtaining a good queue position before conditions become favourable.

Keywords: algorithmic trading, high frequency trading, limit order book, queuing model, order flow, impulse control, adverse selection

1. Introduction

Most major modern stock exchanges have switched to electronic limit order books (LOBs) as the primary matching mechanism of trades. Along with this evolution comes a surge in computerized trading algorithms. These algorithms are developed to perform a variety of tasks (for example,

[☆]The authors would like to thank Sebastian Jaimungal (University of Toronto) and Yaroslav Melnyk (EPFL) for their comments and suggestions, as well as participants at the Research in Options 2016 conference and the Conference on Mathematical Modelling in Finance 2017.

Email addresses: rdon@uw.edu (Ryan Donnelly), luke.gan@mail.utoronto.ca (Luhui Gan)

executing a large order or providing liquidity to the market) by quickly digesting information from the market and sending orders to the exchange. The success of a trading algorithm relies on two fundamental questions: what type of orders should be sent and what is the optimal time to do so?

Many LOBs adopt a price-time priority rule, that is, the priorities of limit orders (LOs) facing execution by a market order (MO) are based first on their price, and then on their time of submission. Within a queue at a fixed price, the first LO which is matched to an incoming MO is the one which was placed at the earliest time. When an LO is placed, it resides at the back of the queue and only moves forward if other LOs in front of it are cancelled or if an incoming MO lifts other orders from the queue.

If an agent has an LO active in the LOB at the best price, there are a number of factors to consider when trying to determine the best course of action with respect to cancelling the order or letting it rest. One factor is the probability that an incoming MO would transact with the agent's LO. This will depend on the position of the LO in the queue, since if it is farther from the front then it will have a smaller chance of being filled by a MO. Another factor is the likelihood that the agent may want to place the order back in the queue a short time after cancelling it. This will be affected by both the position of the LO as well as the total queue length. If the queue is very long and the agent's LO is near the front, then it will take a long time for a new order to reach the current position should it be cancelled and replaced later.

In this paper, we show how to form the optimal decisions to place and cancel LOs in a single queue based on the queue length and position of the LO. To achieve this goal, our first step is to perform an empirical study on LOB dynamics. Using Nasdaq data, we study various LOB features including the rate of addition and cancellation of LOs, rate of MO arrival, distribution of MO size, and distribution of replenished queue length. In particular, we wish to investigate how these events depend on the length of the best price queue. We also attempt to gauge the profitability of an LO execution based on an observable trade indicator. We propose a statistical model to describe the dynamics of these features. To the best of our knowledge, some of these empirical findings are novel, in particular the dependence of these events on the length of the queue. For example, we find that the intensity of cancellation of a limit order at a particular position in the queue has significant dependence on the position of the order, but little dependence on the total length of the queue. We also find that though market orders seldom go beyond the best price level, a large proportion of them deplete the entire queue at the best price level; the size of market orders is well described as a random variable right-censored by the volume at the best price level.

Based on our empirical study, we construct a queuing model for a single price level on one side of the LOB. In our set-up, the dynamics of this queue are partially dictated by a stochastic time dependent regime which can be interpreted as an abstract trade signal (TS). The gain or loss of a filled LO depends on the regime of the TS: during a gainful regime, a filled LO is likely to be profitable, but the rate of MO arrivals is low so it takes longer for a LO to be executed; during an adverse regime, a filled LO is likely to constitute a loss, and the rate of MO arrivals is high so it

takes less time for a LO to be executed. This behaviour is consistent with what we observe in the data.

An important feature of our model is that the agent’s impulse control may have a random effect on the state of the queue. Impulse control problems generally feature the ability for the agent to dictate the precise state of the system after their impulse. Mathematically this is proposed by imposing that the state of the system due to the action be a random variable which is measurable at the time of the impulse. The model proposed here allows for the resulting state to be random from the perspective of the agent without introducing issues related to measurable random variables.

We then consider the problem of an agent that maximizes her expected utility by optimally choosing the timing of placing and cancelling her LOs. For simplicity, we restrict the agent to have at most one LO active in the queue at a single time, and the agent incurs a small fixed cost when she places or cancels the LO. In order to form her optimal decision, she keeps track of three state variables: the TS, the queue length, and the queue position of her LO. When she has an LO active in the queue, she stays in the queue as long as the TS is in a gainful regime. However, it is worthwhile mentioning that even in an adverse regime the agent may keep an active LO in the queue, or place a new LO if she does not have an active one. This is because in an adverse regime, the agent may be willing to wait for her active LO to move forward in the queue so that it is in a good position by the time the TS becomes gainful. In an adverse regime, the agent optimally balances between the loss from an executed order and the gain from a good queue position when the TS changes to a favourable one.

1.1. Literature Review

Our work is related to much of the literature on optimal execution. In the seminal work of Almgren and Chriss (2001), the authors consider the problem of executing a large portfolio using MOs. Their work has been generalized in a number of ways, see for example, Alfonsi et al. (2010) and Gatheral et al. (2012). An alternative approach using LOs is taken by Avellaneda and Stoikov (2008), and has been extended by Bayraktar and Ludkovski (2011), Guéant et al. (2012), Guilbaud and Pham (2013), and Cartea and Jaimungal (2015). In all the above studies, the actual queuing dynamics of the LOB are abstracted away. For example, in Almgren and Chriss (2001), Alfonsi et al. (2010), and Gatheral et al. (2012), the agent does not interact directly with the LOB; instead, her trading activities generate price impacts, which are assumed to be deterministic functions of the trading rate. In Avellaneda and Stoikov (2008) and subsequent studies, it is assumed that the agent’s filled LOs follow a counting process with rate independent of the state of the LOB. In both streams of literature, the state of the LOB becomes irrelevant. More recent developments incorporate order flow information as a state variable for the LOB. In Bechler and Ludkovski (2015), order flow imbalance is assumed to be an Ornstein-Uhlenbeck process, whereas Cartea et al. (2015a) uses volume imbalance as a measure for order flow imbalance and model it as a discrete state Markov chain. In quantifying the profitability of filled LOs, for the purposes of our work we choose the TS

to be volume imbalance. In general, the TS can be chosen to be any observable quantity which the agent believes will offer information about the profitability of trades and the dynamics of the LOB. Our contribution is to quantify when it is beneficial to cancel or place an LO, based on its queue position and the TS.

Our work is also related to the literature on models for LOB events and queues. This stream of the literature started with Cont et al. (2010). A simplified version which models only the best bid and ask can be found in Cont and De Larrard (2013), where its diffusive limit is also derived. Diffusive limits of queuing model for LOBs are also studied in Lakner et al. (2013), Jaisson et al. (2015), and Guo et al. (2015). In terms of empirical features, our model is closest to Huang et al. (2015). Our modeling approach to order book events also has similarities to Bacry et al. (2016) in which the authors use a multivariate Hawkes process to drive the dynamics of the best bid and ask queues. The distinction between our work and these previous works is that we investigate a control problem within the queuing system rather than only investigating the dynamics and other related quantities.

A closely related work is Lehalle and Mounjid (2016) where the authors also consider an optimal trading problem in a queuing system, however there are some key differences in our work. First, we consider both the optimal timing of placing LOs as well as cancelling them, whereas in the other work they only investigate when it is optimal to cancel an order. Second, we do not model MO size as a constant, electing instead to model the size as a censored distribution. As suggested by our empirical study, a large proportion of MOs deplete the queue at the best price. Therefore, if an agent models MO size as a constant, then she will significantly underestimate the probability of her LO being executed if its position is far from the front of the queue. This has significant effects on the agent’s desire to cancel an order even if it is not at the front of the queue.

It is also worthwhile mentioning Maglaras et al. (2014) and Maglaras et al. (2015), who consider order placement problems under fluid (deterministic) queuing models. Our work is different in that we have a stochastic queuing model where the queue length is discrete-valued. Moreover, we incorporate adverse selection risk, so the agent’s executed LOs are not always considered a gain by earning the spread. Rather, the asset price might move in the direction such that the agent’s executed LO becomes an instantaneous loss.

The rest of the paper is organized as follows. Section 2 presents some empirical findings on LOB dynamics. In Section 3 we describe a queuing model which reflects many of our empirical findings and the agent’s stochastic control problem. In Section 4 we derive and simplify the dynamic programming equation for the agent’s control problem. Section 5 contains a numerical example using model parameters calibrated from real data. In Section 6 we prove that the numerical method we employ to solve the associated system of differential equations converges. Section 7 concludes.

2. Empirical Analysis

In this section, we present empirical findings that motivate our queuing model in Section 3, including rate of addition and cancellation of other LOs, rate of MO arrivals, distribution of MO size, and distribution of replenished queue length. We also choose to use the value of volume order imbalance as the TS and investigate the magnitude of gain or loss of a filled LO as a function of imbalance. In general, other TSs can be used if a dependence structure between the signal and other queue dynamics can be specified.

For each quantity that we investigate, we present our empirical findings and also introduce a parametric fit of a model to the data. When we demonstrate the optimal strategy in Section 5 we use the estimated parametric model for all numerical values of parameters. The exact form of the parametric model we use is essentially arbitrary and in general any form can be used within the model presented in Section 3, but we have found that our choices capture the essential features from the data in addition to being relatively straightforward to estimate. For the purposes of this analysis, we primarily consider only events on the sell side of the LOB, and as such we are concerned only with market buy orders (MBOs).

The estimation of empirical quantities, parametric fits, and creation of all figures in this section was performed in 41.5 seconds in Matlab on an Intel Core i7 2.7GHz CPU.

2.1. Data

The data that we analyze is the Nasdaq Historical Total View (ITCH) for the ticker INTC on Oct 1, 2014, with the first and last half hours of trading removed. The remaining data consists of 772,166 events in total, 280,989 of which occur at the best ask price or improve the best ask price. A further breakdown of the types of events in the data is presented in Table 1.

Event Type	Count	Average Frequency (per second)
Best Ask Event	280,989	14.19
Order Placement	142,455	7.19
Order Cancellation	128,973	6.51
Market Buy Order	9,561	0.48

Table 1: A best ask event is any order cancellation that occurs at the best ask price, any sell order placement at the best ask price or lower, or any market buy order.

We divide all volumes and order sizes by 100, round them to the nearest integer and discard results that are equal to zero. In other words, we assume that all buy and sell volumes are multiples of 100. Our choice is based on the observation that a large amount of LOs and MOs have a “round-lot” size, i.e, sizes that are of multiples of 100. For other equities this typical order size may be different and should be modified accordingly. See, for example, Cartea et al. (2015b) for more details. Henceforth it should be understood that any volume quantity is to be interpreted as representing

a multiple of this “round-lot” size. We denote by L_t^b (L_t^a) $\in \{1, 2, \dots\}$ the observed volumes at time t of LOs posted at the best buy (sell) price. For most of this section, only the sell side of the LOB is examined. We have removed all events when $L_t^a > 200$ (0.7% of all events).

2.2. Volume Imbalance as Trade Signal

As many researchers have pointed out, the imbalance between buy and sell order flows exhibits a dependence structure with future price changes. There are various ways of measuring order flow imbalance: Cont et al. (2014) constructs a measure based on net aggregated volumes of buy and sell orders over a time interval; Cartea et al. (2015a) uses another measure, the volume imbalance, computed from volumes of LOs posted at the best buy and sell price. Both works show that their measures of order flow imbalance were significant in predicting price movement, and so we construct our TS regime from observations of the imbalance process. In this paper, we consider the volume imbalance as in Cartea et al. (2015a), defined as:

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

We construct a three-state regime process Z_t which we will interpret as our TS by dividing the imbalance measure interval $[-1, 1]$ into three subintervals and assigning different values as follows:

1. sell-heavy: $\rho_t \in [-1, 0.2] \Rightarrow Z_t = 1$;
2. neutral: $\rho_t \in (0.2, 0.6] \Rightarrow Z_t = 2$;
3. buy-heavy: $\rho_t \in (0.6, 1] \Rightarrow Z_t = 3$.

As will be shown later, $Z_t = 1$ is considered by the agent a gainful regime, where a filled order is likely to be profitable; $Z_t = 3$ is considered by the agent an adverse regime where a filled order is likely to be considered a loss. In interpreting Z_t as a TS we shall henceforth refer to these regimes as gainful, neutral, and adverse.

The choice of the boundaries which define different TS regimes will affect the estimation of other model parameters which subsequently will dictate the optimal order placement and cancellation strategy. For example, if the boundary between the sell-heavy and neutral regimes were selected as 0 instead of 0.2, then some changes to the empirical results below will take place. Since the sell-heavy regime has decreased in size, the transition rates (as in Figure 1) from the sell-heavy regime to the neutral and buy-heavy regimes will increase. In addition, both sell-heavy and neutral regimes would experience increased selling pressure with this redefined boundary, so the distribution of price changes (as in Figure 5) would exhibit higher probabilities of lower outcomes. This has the effect of making the sell-heavy and neutral regimes more attractive states to the agent. As such,

she would be less likely to cancel her order in the neutral regime (cancellation region in Figure 7 shrinks) and more likely to place the order in the sell-heavy regime.

For the purposes of the present work, these boundaries are selected so that when we illustrate the optimal strategy there will be a clear effect based on the regime state. If this model were used in practice, other methods could be used to determine these boundaries either subjectively (for example, to reflect a desired level of trading aggression or passivity based on the risk tolerances of the agent) or objectively (for example by choosing the bounds symmetrically around 0 such that the TS is expected to occupy each regime for an equal length of time). In addition it may be desirable to employ a finer discretization of the TS for practical purposes, but we limit the number to three in this work in order investigate highlight the effects of queue position on the optimal strategy.

2.2.1. Empirical Estimation

From the definition in (1) we expect dependence between the distribution of queue length L_t^a and the regime of the TS Z_t . To capture this pattern, we propose letting the rate of transition between regimes dependent on the queue length. Let $\lambda^{z,\bar{z},\ell}$ be the rate of transition from regime z to regime $\bar{z}$ ($z \neq \bar{z}$), when the queue length is $\ell > 0$ (we use the convention $\lambda^{z,z,\ell} = 0$ for all z). Our empirical estimate of $\lambda^{z,\bar{z},\ell}$ is found through the following procedure: we denote by $N^Z(z, \bar{z}, \ell)$ the total number of transitions from z to $\bar{z}$ when $L_t^a = \ell$ and by $T(z, \ell)$ the occupation time for which $L_t^a = \ell$ and $Z_t = z$. Then the empirical estimate of $\lambda^{z,\bar{z},\ell}$ is

$$\frac{N^Z(z, \bar{z}, \ell)}{T(z, \ell)}.$$

2.2.2. Parametric Model

We run a Poisson regression $N^Z(z, \bar{z}, \ell) \sim \text{Poisson}(\lambda^{z,\bar{z},\ell} T(z, \ell))$ with the canonical link function

$$\lambda^{z,\bar{z},\ell} = \exp(\beta_{0,z,\bar{z}}^Z + \beta_{1,z,\bar{z}}^Z \ell), \quad (2)$$

for constants $\beta_{0,z,\bar{z}}^Z$ and $\beta_{1,z,\bar{z}}^Z$ and set these coefficients equal to their maximum likelihood estimates, and the estimated functions $\lambda^{z,\bar{z},\ell}$ are truncated for the highest 1% and the lowest 1% values of ℓ weighted by occupations time¹ and replaced by linear extrapolation. Estimated coefficients are reported in Table B.6 and B.7.

Figures 1(a) 1(d) show the rates of transition from the gainful to the neutral and adverse regimes

¹For each ℓ , we compute $f_{z,\ell} = \frac{T_{z,\ell}}{\sum_l T_{z,\ell}}$, the fraction of time that the queue length is equal to ℓ when the LOB is in regime z . The upper and lower truncation thresholds are defined as: $\bar{\ell}_z = \min\{\ell : \sum_{m \leq \ell} f_{z,m} \geq 0.01\}$ and $\underline{\ell}_z = \max\{\ell : \sum_{m \leq \ell} f_{z,m} \leq 0.99\}$.

Figure 1: Estimated rate of transition. Points represent the empirically observed transition rate in each regime for each value of queue length. Curves represent the estimated fit corresponding to equation (2). Note that $\lambda^{z,z,\ell} = 0$ for all z , so it is not shown.

(empirical estimates are shown as well as the rates which result from the Poisson regression). Both rates increase as the queue becomes shorter, an indication that there is an increasing buy pressure, thus it is more and more likely that we transition to an increasingly adverse regime. Also note that the transition rate from gainful to neutral is higher than that from gainful to adverse (the jump which can be interpreted as being of smaller size is more likely to occur). Figures 1(b) and 1(e) show the transition rates to gainful and adverse regimes from the neutral regime. When the queue is short, the regime is more likely to become adverse, whereas when the queue becomes long, it is more likely to become gainful. Similar patterns can be found in Figures 1(c) and 1(f), where the rates of transition out of the adverse regime are presented.

Figure 2: The rate of cancellation for different queue length and queue position. For each point (ℓ, y) on the grid, we calculate the total number of cancellation divided by the total time that the best sell queue length is equal to ℓ . We then apply Gaussian smoothing to obtain this figure.

2.3. Rate of Cancellation and Addition

2.3.1. Empirical Estimation

We assume that all cancellations and additions are of a unit size². We denote by $\tilde{\lambda}^{c,z,\ell,y}$ the rate of cancellation of an order at position y ($0 < y \leq \ell$) when the queue length is ℓ and the TS regime is z (the c in the notation stands for “cancellation”). Queue position $y = 1$ corresponds to the front of the queue, and $y = \ell$ represents the back. Figure 2 shows the empirical estimates of $\tilde{\lambda}^{c,z,\ell,y}$. Each panel corresponds to a different regime z , and the colour at each point (ℓ, y) represents the corresponding rate of cancellation $\tilde{\lambda}^{c,z,\ell,y}$ (lighter colours indicating a higher rate of cancellation). For each regime, we see that the rate of cancellations depends much more heavily on the position in the queue, y , than on the total length of the queue, ℓ . Therefore, a reasonable model for LO cancellation is to assume that $\tilde{\lambda}^{c,z,\ell,y}$ does not depend on ℓ and henceforth we use the notation $\tilde{\lambda}^{c,z,y}$.

The agent is not concerned with the rate of cancellation of any particular order, but she is concerned with the total rate of cancellations in front of her order and behind her order. Therefore, we denote by $\lambda^{c,z,y}$ the total rate of cancellation at or before position y , i.e., $\lambda^{c,z,y} = \sum_{k=1}^y \tilde{\lambda}^{c,z,k}$. For instance, when the LOB is in regime z , the rate of cancellations happening within the interval $(a, b]$ is $\lambda^{c,z,b} - \lambda^{c,z,a}$. We do not attempt to estimate $\tilde{\lambda}^{c,z,y}$, but below we describe how we estimate the function $\lambda^{c,z,y}$.

We begin by taking all events corresponding to LO cancellations in the best sell queue, marked as $(Q_{z,i}^c, L_{z,i}^c)_{i=1,2,\dots}$, where $Q_{z,i}^c$ is the volume of the i -th cancellation that occurs during regime z and $L_{z,i}^c$ is the total volume at the best sell price immediately prior to this cancellation. For each level

²The details given here show that the number of additions and cancellations are computed in a volume weighted fashion. Thus, even though we assume they have unit volume, the resulting intensities of the events will be larger to compensate for this discrepancy producing roughly equal expected volumes of added and cancelled orders.

$L_{z,i}^c = \ell$ ($\ell = 1, 2, \dots$) and each regime we denote by $N^c(z, \ell)$ the aggregate volume of cancelled LOs, so that we have

$$N^c(z, \ell) = \sum_{i: L_{z,i}^c = \ell} Q_{z,i}^c.$$

Recall that $T(z, \ell)$ denotes the occupation time of the state $(Z_t, L_t^a) = (z, \ell)$. The empirical estimate of the cancellation rate is then

$$\frac{N^c(z, \ell)}{T(z, \ell)}.$$

The empirical estimate of the addition rate is computed similarly and equal to

$$\frac{N^a(z, \ell)}{T(z, \ell)},$$

where $N^a(z, \ell)$ is the cumulative number of additions defined in the same fashion as $N^c(z, \ell)$.

2.3.2. Parametric Model

We conduct a Poisson regression $N^c(z, \ell) \sim \text{Poisson}(\lambda^{c,z,\ell} T(z, \ell))$ with the canonical link function

$$\lambda^{c,z,\ell} = \exp(\beta_{0,z}^c + \beta_{1,z}^c \ell), \quad (3)$$

where $\beta_{0,z}^c$ and $\beta_{1,z}^c$ are constants. The above form gives $\lambda^{c,z,0} = 0$ which will be an important restriction when we write our full queue model to ensure that when the agent's order is at the front of the queue, the rate of cancellations in front of her order is zero. The values of $\beta_{0,z}^c$ and $\beta_{1,z}^c$ are obtained via maximum likelihood estimation.

The rate of addition $\lambda^{a,z,\ell}$ is estimated through a similar procedure. We assume that

$$\lambda^{a,z,\ell} = \exp(\beta_{0,z}^a + \beta_{1,z}^a \ell), \quad (4)$$

for constants $\beta_{0,z}^a$ and $\beta_{1,z}^a$.

Figure 3 shows the estimated rate of addition and cancellation as well as the rates which result from the Poisson regression. Coefficients are reported in Table B.8.

2.4. Rate of Market Order Arrival

We denote by $\lambda^{m,z,\ell}$ the rate of MO arrival when the TS is in regime z and the length of the queue is ℓ . In principle, $\lambda^{m,z,\ell}$ can depend on both z and ℓ , as in the case for the rate of addition and

Figure 3: Estimated (volume-weighted) rate of addition/cancellation for different queue length. Points represent empirical estimates and curves represent the fit of the parametric model.

cancellation. However, given the fact that the number of MO events is significantly smaller than the number of LO additions and cancellations, it is difficult to estimate the dependence of $\lambda^{m,z,\ell}$ on ℓ . For example, in the data set we consider here there are only 2,911 MBOs. In particular, there are only 556 MBOs when the TS is in the gainful regime. Therefore, to keep the model parsimonious, we assume that $\lambda^{m,z,\ell}$ depends on the regime z only and henceforth we denote the rate of MOs by $\lambda^{m,z}$. The quantity $\lambda^{m,z}$ can be estimated as

$$\lambda^{m,z} = \frac{M(z)}{T(z)},$$

where $M(z)$ is the total number of MBOs when the TS is in regime z and $T(z)$ is the occupation time of regime z . In our numerical examples we do not use a parametric model as in the previous sections, but rather use this empirical estimate as the market order arrival rate. Estimated results are present in Table B.9.

2.5. Distribution of Market Order Size

2.5.1. Empirical Estimation

As documented in a number of studies, it is rare for an MO to walk through the first level of the LOB (see for example Cartea et al. (2015a)). Here, however, we present another empirical finding: a non-negligible fraction of MOs deplete the entire queue at the best price. Table 2 shows some summary statistics on MO size. The column “Average MO size” contains the average size of MBOs when the queue length is at different intervals. We can see that as the queue length increases, the average MO size increases as well. The column “% of depletion” contains the proportion of MOs that deplete the entire queue at the best sell price. Surprisingly, these proportions are very high: for example, even when the queue length is between $[40, 60)$, which is about 5 times the average MO size, the probability that an MO depletes the whole queue is 12.1%. Therefore, a realistic model for MO size should incorporate the probability of an LO being executed even it is far away

from the front of the queue.

Another observation from Table 2 is that the percentage of MOs depleting the entire queue decreases as the queue length increases. This suggests that we model the MO size as a random variable right-censored by the queue length. To validate this hypothesis, we conduct the following bootstrap exercise. First we estimate the empirical distribution of MO size using the non-parametric Kaplan-Meier estimator (Kaplan and Meier (1958)), assuming that MO size is indeed right-censored by the queue length. We then draw independent samples $(\bar{Q}_i^b)_{i=1,2,\dots}$ from the estimated empirical distribution of the MO size and $(L_i^b)_{i=1,2,\dots}$ from the empirical distribution of queue length. Next we compute the censored MO size $Q_i^b = \min\{\bar{Q}_i^b, L_i^b\}$ and the summary statistics based on the bootstrapped dataset (Q_i^b, L_i^b) . The results are shown in Table 3. It can be seen that the summary statistics in Table 3 are very close to those in Table 2.

Queue Length	Average MO size	% of depletion
[0, 20)	4.5	39.8%
[20, 40)	9.2	14.3%
[40, 60)	10.0	12.1%
[60, 80)	13.6	5.8%
[80, 100)	16.1	4.1%

Table 2: Size of market sell orders (empirical).

Queue Length	Average MO size	% of depletion
[0, 20)	4.3	40.6%
[20, 40)	8.7	14.8%
[40, 60)	11.2	9.5%
[60, 80)	12.8	6.2%
[80, 100)	13.6	3.7%

Table 3: Size of market sell orders (bootstrapped).

2.5.2. Parametric Model

The parametric model we employ in our numerical examples is to model MO size by a negative binomial distribution which depends on the regime of the TS. In particular, for a fixed regime z and for each pair of observed MO size and queue length $(Q_{z,i}^m, L_{z,i}^m)_{i=1,2,\dots}$, we assume that $Q_{z,i}^m = \min\{\bar{Q}_{z,i}^m, L_{z,i}^m\}$, where $(\bar{Q}_{z,i}^m)_{i=1,2,\dots}$ are i.i.d. negative binomial random variables. The parameters are estimated via maximum likelihood. Figure 4 shows the empirical and the estimated distributions for the MO size. Estimated parameters are presented in Table B.10.

2.6. Gain/Loss of a Filled Limit Order

The agent's LO faces different levels of adverse selection risk in different regimes. In our study, we use $\theta = \Upsilon - \Delta_{dt}$ as a proxy for the agent's wealth change when her order is filled, where Υ is equal

Figure 4: Estimated and empirical distribution of MO size. The empirical distributions are estimated using Kaplan-Meier estimator and the curves are the result of fitting a negative binomial distribution.

Figure 5: Histogram of best sell price change 100ms after MO arrival conditional on TS regime at the time of the order.

to the size of a half-tick plus the rebate, multiplied by the size of the agent's order, and Δ_{dt} is the change in the best sell price dt seconds after an MBO arrival, multiplied by the size of the agent's order. Υ can be interpreted as the agent's gain from an executed limit order if the fundamental asset price does not change upon execution. In our numerical example in Section 5, we will use the empirical distribution of Δ_{dt} for $dt = 100$ milliseconds.

Figure 5 shows the histogram of best sell price changes 100 milliseconds after the arrival of an MBO. It is clear that when the TS is in the adverse regime, the best sell price is more likely to move up after an MBO arrival and the agent suffers more from adverse selection. The time interval of 100 milliseconds is to allow for significant activity of LO additions and cancellations to occur shortly after the MBO, and this interval may need to be adjusted depending on the asset in question.

In our numerical examples below, we do not use a parametric model to describe this distribution. In a similar fashion to MO intensity, we use the empirical distribution in Figure 4 as the distribution of the price change after an MO.

Figure 6: Empirical distribution of the replenished queue length, and estimated fit of distribution.

2.7. Distribution of Replenished Queue Length

2.7.1. Empirical Estimation

Since we will model only one queue, we must make an assumption of what happens when the entire queue is depleted, either due to an MO filling the entire queue, or the cancellation of the last LO in the queue. We will assume that the queue is immediately replenished to a random quantity with a distribution depending on the regime of the TS. Here we show the empirical distribution of the volume of the best ask queue immediately after an event depletes the previously best ask queue, conditional on the TS regime immediately before the depletion.

2.7.2. Parametric Model

We will assume that the replenished queue length is drawn from a distribution $\mu^{\zeta, z}$, which depends on the regime z . To estimate $\mu^{\zeta, z}$, we fit negative binomial distribution to the empirical distribution considered in the previous section. Both the empirical distribution and the estimated fit are shown in Figure 6. Estimated parameters can be found in Table B.11. This method of modelling the replenished queue is similar to that of Cont and De Larrard (2013). In that work, new queue lengths are drawn from a distribution immediately after a price change. Another approach is taken in Lehalle and Mounjid (2016) where the new queue length is fully determined by the state of the LOB immediately before the queue depletes. Both of these studies, as well as ours, take this approach because it allows us to avoid the necessity of modelling the volume process of the queue at the next price level, significantly reducing the dimensionality of the problem.

3. Model

In this section we propose a model for the queue which encompasses all of the possible events which may alter its state. Once the set of possible events has been established, the rate at which these

events occur and the distributions of the changes they cause are chosen to reflect the behaviour observed in Section 2.

Our construction is non-standard in the sense that our stochastic processes will be indexed by $\mathbb{R} \times \{1, 2\}$ rather than by $\mathbb{R}$ as usual. We introduce the notation $\bar{t} = (t, k)$ where $t \in \mathbb{R}$ and $k \in \{1, 2\}$. We equip the set $\mathbb{R} \times \{1, 2\}$ with the lexicographical order, i.e. $\bar{t}_1 \prec \bar{t}_2$ if $t_1 < t_2$ or if $t_1 = t_2$ and $k_1 < k_2$. We require this construction because at a single instant in time, there may occur multiple events which must follow a logical sequence. Thus, the first index denoted t is to be interpreted as the usual flow of continuous time, and the second index denoted k will be used to construct the proper sequencing of events which occur simultaneously.

The need for this type of construction stems from two sources. First, we model a replenishing of the queue to a random length when it is depleted by some event, and second, the agent's action of cancelling the last LO in the queue will cause the queue to be replenished. The cancellation and replenishing occur at the same instant in time, but within this particular time we would still like to interpret the cancellation as occurring first, and we do not want the agent to have foreknowledge of the replenished length when she cancels her order.

A specific example which illustrates this sequencing of simultaneous events is the following: suppose there are exactly two LOs active in the queue, the order at the back being the one belonging to the agent. At time t :

1. The order at the front of the queue is cancelled.
2. Based on the cancellation of the first order, the agent decides to cancel her order.
3. The cancellation of the last order causes the queue to be replenished at a random length.
4. Based on the new length of the queue, the agent decides to place a new LO at the back of the queue.

This example demonstrates how we would like the decisions of the agent to be based on information in the queue. The agent's decision to cancel her order in step 2 is not allowed to depend on the length of the replenished queue in step 3, but the agent's decision to replace the order in step 4 is allowed to depend on this quantity. The typical set-up of an impulse control problem gives the agent perfect knowledge of the effect of her control on the state of the system by stating that the controlled state is measurable with respect to the stopping time of the impulse. See for example Bensoussan et al. (1982), Baccarin and Sanfelici (2006), and Ly Vath and Pham (2007). Our extended indices allow us to circumvent this issue.

Fix a complete filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F} = (\mathcal{F}_{\bar{t}})_{(0,1) \preceq \bar{t} \preceq (T,2)})$. In developing the model it may be useful to consider the filtration $\mathbb{G}$ defined by $\mathcal{G}_t = \mathcal{F}_{(t,1)}$.

3.1. State Space

The state of the system will be denoted by $\mathcal{S} = (X, Z, L, Y)$, where X represents the agent's wealth, Z is the regime of the TS, L is the total length of the queue, and Y is the position of the agent's order (should it exist). We set a maximum possible queue length denoted by N^L so that $L \in \{0, 1, \dots, N^L\}$. We allow the length of the queue to be zero, but we will specify the dynamics such that if $L_{(t,1)} = 0$, then $L_{(t,2)} > 0$ due to instantaneous replenishing. In addition, we will always have $Y \in \{1, \dots, N^L, \Xi\}$ where $Y = \Xi$ is a placeholder state which indicates that the agent does not have an active order. Otherwise we must enforce the constraint $Y \leq L$ which indicates that the position of the agent's order is at most equal to the length of the queue. We denote the set of feasible states by

$$\mathbb{S} = \left\{ \mathcal{S} = (X, Z, L, Y) : X \in \mathbb{R}, Z \in \{1, \dots, N^Z\}, \right. \\ \left. L \in \{0, 1, \dots, N^L\}, Y \in \{1, \dots, N^L, \Xi\}, Y = \Xi \text{ or } Y \leq L \right\}. \quad (5)$$

Future sections will involve many conditional statements with respect to the state variable Y . In order to simplify notation slightly, we adopt the convention that $\mathbb{1}_{k=\Xi} = \mathbb{1}_{k<\Xi} = \mathbb{1}_{k>\Xi} = 0$ whenever $k \in \mathbb{R}$. That is, statements of the form $\Xi = k$, $\Xi < k$, and $\Xi > k$ are interpreted to be false for $k \in \mathbb{R}$.

3.2. Uncontrolled Queue Dynamics

Here we indicate the dynamics of the state given that the agent does not place or cancel any orders. To this end, for each value of $z, \bar{z} \in \{1, \dots, N^Z\}$ and $\ell \in \{1, \dots, N^L\}$, we let $N_{(t,1)}^{z,\bar{z},\ell}$, $N_{(t,1)}^{c,z,\ell}$, $N_{(t,1)}^{a,z,\ell}$, and $M_{(t,1)}^z$ be independent Poisson processes adapted to $\mathcal{G}_t$. An arrival of the process $N_{(t,1)}^{z,\bar{z},\ell}$ indicates a change in the TS from z to $\bar{z}$ at time t . An arrival of $N_{(t,1)}^{c,z,\ell}$ represents the cancellation of an order at position ℓ while the TS is equal to z . An arrival of $N_{(t,1)}^{a,z,\ell}$ represents the addition of an LO when the length of the queue is ℓ in regime z . Lastly, an arrival of $M_{(t,1)}^z$ indicates an MO in regime z . We extend these processes to be defined at all $\bar{t}$ by setting $N_{(t,2)}^{z,\bar{z},\ell} = N_{(t,1)}^{z,\bar{z},\ell}$, and similarly for the others. The interpretation is that these processes may jump at an index of the form $(t, 1)$, but not at one of the form $(t, 2)$. Staying consistent with the notation of Section 2, the intensities of these processes are denoted by $\lambda^{z,\bar{z},\ell}$, $\tilde{\lambda}^{c,z,\ell}$, $\lambda^{a,z,\ell}$, and $\lambda^{m,z}$. In addition, as before, we have $\lambda^{c,z,\ell} = \sum_{k=1}^{\ell} \tilde{\lambda}^{c,z,k}$. The quantity $\tilde{\lambda}^{c,z,\ell}$ represents the rate of cancellation of an order at position ℓ (if it exists), so $\lambda^{c,z,\ell}$ is the rate of cancellation of any order in front of (and including) the one in position ℓ .

We also suppose that for each z and each t we have three random variables $\epsilon_{(t,1)}^z$, $\theta_{(t,1)}^z$, and $\zeta_{(t,2)}^z$, all of which are independent from all other variables and with distributions $\mu^{\epsilon,z}$, $\mu^{\theta,z}$, and $\mu^{\zeta,z}$ respectively. The distribution $\mu^{\epsilon,z}$ is chosen so that $\epsilon_{(t,1)}^z$ is a positive integer, $\mu^{\zeta,z}$ is chosen so that $\zeta_{(t,2)}^z$ is an integer between 1 and N^L , and $\mu^{\theta,z}$ is chosen so that $\mathbb{E}[e^{-\gamma\theta_{(t,1)}^z}] < \infty$ (γ is the

agent's risk aversion parameter which will be introduced later). Note that two of these variables are defined at time $(t, 1)$, and the third is defined at time $(t, 2)$. This is to emphasize that $\epsilon_{(t,1)}^z$ and $\theta_{(t,1)}^z$ are $\mathcal{F}_{(t,1)}$ -measurable, but $\zeta_{(t,2)}^z$ is $\mathcal{F}_{(t,2)}$ -measurable. The need for this will be clear below when we specify the dynamics of the queue. The interpretation of these quantities is the following: the value of $\epsilon_{(t,1)}^z$ is the size of an MO that arrives at time t while in regime z . If the agent has an LO which is filled by an incoming MO, then the change in the agent's wealth is equal to $\theta_{(t,1)}^z$. The value of $\zeta_{(t,2)}^z$ will be used as the size of the replenished queue if there is an event at time t which depletes it.

Let $\bar{\tau} = (\tau, \kappa)$ be a stopping time in the filtration $\mathcal{F}_{\bar{t}}$ and let φ be an $\mathcal{F}_{\bar{\tau}}$ -measurable random variable taking values in $\mathbb{S}$ with the restriction that if $L(\varphi) = 0$, then $\kappa = 1$. We define a sequence of stopping times $\bar{\rho}_j = (\rho_j, k_j)$ and states φ_j as follows: let $\bar{\rho}_0 = \bar{\tau}$ and $\varphi_0 = \varphi$, and set

$$\begin{aligned} X_{\bar{\rho}_0}^{u, \bar{\tau}, \varphi} &= X(\varphi) & Z_{\bar{\rho}_0}^{u, \bar{\tau}, \varphi} &= Z(\varphi) \\ L_{\bar{\rho}_0}^{u, \bar{\tau}, \varphi} &= L(\varphi) & Y_{\bar{\rho}_0}^{u, \bar{\tau}, \varphi} &= Y(\varphi) \end{aligned}$$

via the appropriate projection mappings. The superscript u in these quantities is to indicate that these represent the uncontrolled versions of each process. For the remainder of this subsection we will suppress the dependence on $\bar{\tau}$ and φ of the state variables to reduce clutter in notation. Now recursively define stopping times by

$$\bar{\rho}_{j+1} = \begin{cases} (\rho_{j+1}, 1), & \text{if } L_{\bar{\rho}_j}^u \neq 0, \\ (\rho_j, 2), & \text{if } L_{\bar{\rho}_j}^u = 0, \end{cases} \quad (6)$$

where if $L_{\bar{\rho}_j}^u \neq 0$,

$$\rho_{j+1} = \inf_{\substack{\bar{z} \neq Z_{\bar{\rho}_j}^u \\ \ell \neq Y_{\bar{\rho}_j}^u, \ell \leq L_{\bar{\rho}_j}^u}} \left\{ t > \rho_j : \max\{\Delta N_t^{Z_{\bar{\rho}_j}^u, \bar{z}, \ell}, \Delta N_t^{c, Z_{\bar{\rho}_j}^u, \ell}, \Delta N_t^{a, Z_{\bar{\rho}_j}^u, L_{\bar{\rho}_j}^u}, \Delta M_t^{Z_{\bar{\rho}_j}^u}\} = 1 \right\}, \quad (7)$$

and define the state transitioned into at these times by

$$\varphi_{j+1} = \begin{cases} \left(X_{\bar{\rho}_j}^u, \bar{z}, L_{\bar{\rho}_j}^u, Y_{\bar{\rho}_j}^u \right), & \text{if } \Delta N_{\bar{\rho}_{j+1}}^{Z_{\bar{\rho}_j}^u, \bar{z}, \ell} = 1, L_{\bar{\rho}_j}^u \neq 0, \\ \left(X_{\bar{\rho}_j}^u, Z_{\bar{\rho}_j}^u, L_{\bar{\rho}_j}^u + 1, Y_{\bar{\rho}_j}^u \right), & \text{if } \Delta N_{\bar{\rho}_{j+1}}^{a, Z_{\bar{\rho}_j}^u, L_{\bar{\rho}_j}^u} = 1, L_{\bar{\rho}_j}^u \neq 0, \\ \left(X_{\bar{\rho}_j}^u, Z_{\bar{\rho}_j}^u, L_{\bar{\rho}_j}^u - 1, Y_{\bar{\rho}_j}^u - \mathbb{1}_{\ell < Y_{\bar{\rho}_j}^u} \right), & \text{if } \Delta N_{\bar{\rho}_{j+1}}^{c, Z_{\bar{\rho}_j}^u, \ell} = 1, L_{\bar{\rho}_j}^u \neq 0, \\ \left(X_{\bar{\rho}_j}^u + \theta \mathbb{1}_{\epsilon \geq Y_{\bar{\rho}_j}^u}, Z_{\bar{\rho}_j}^u, (L_{\bar{\rho}_j}^u - \epsilon) \mathbb{1}_{\epsilon < L_{\bar{\rho}_j}^u}, \right. \\ \quad \left. (Y_{\bar{\rho}_j}^u - \epsilon) \mathbb{1}_{\epsilon < Y_{\bar{\rho}_j}^u} + \Xi(\mathbb{1}_{\epsilon \geq Y_{\bar{\rho}_j}^u} + \mathbb{1}_{Y_{\bar{\rho}_j}^u = \Xi}) \right), & \text{if } \Delta M_{\bar{\rho}_{j+1}}^{Z_{\bar{\rho}_j}^u} = 1, L_{\bar{\rho}_j}^u \neq 0, \\ \left(X_{\bar{\rho}_j}^u, Z_{\bar{\rho}_j}^u, \zeta, \Xi \right), & \text{if } L_{\bar{\rho}_j}^u = 0 \end{cases} \quad (8)$$

where for the sake of brevity in (8) we have replaced $\theta_{\bar{\rho}_{j+1}}^{Z_{\bar{\rho}_j}^u}$, $\epsilon_{\bar{\rho}_{j+1}}^{Z_{\bar{\rho}_j}^u}$, and $\zeta_{\bar{\rho}_{j+1}}^{Z_{\bar{\rho}_j}^u}$ by θ , ϵ , and ζ respectively. We now set:

$$\begin{aligned} X_{\bar{\rho}_{j+1}}^u &= X(\varphi_{j+1}) & Z_{\bar{\rho}_{j+1}}^u &= Z(\varphi_{j+1}) \\ L_{\bar{\rho}_{j+1}}^u &= L(\varphi_{j+1}) & Y_{\bar{\rho}_{j+1}}^u &= Y(\varphi_{j+1}) \end{aligned}$$

This construction requires some explanation. The definition given by (6) and (7) corresponds to the arrival index of the next feasible event. If there is an event that sends L to zero (which must happen at an index of the form $(t, 1)$ by construction), then the next stopping time is immediate but occurs at an index of the form $(t, 2)$. If the most recent previous event does not result in L being sent to zero, then the next event is due to a regime change, MO arrival, LO arrival, or LO cancellation, and must occur at an index of the form $(t, 1)$. Note that the infimum is taken only over values of ℓ corresponding to order positions which are within the queue and not equal to the agent's order should it exist, hence the agent's order can never be exogenously cancelled. Also, by independence of the Poisson processes, with probability 1 the increments in (7) never coincide, so the definition in (8) is well defined.

The first two types of events in (8) are rather straightforward. The first type corresponds to a change in the TS from its previous value of Z_{ρ_k} to the new value of $\bar{z}$ (by convention we set the intensity $\lambda^{z, z, \ell} = 0$ for all z). The second type corresponds to the addition of a single LO (we set the intensity of $\lambda^{a, z, N^L} = 0$ to enforce $L_t^u \leq N^L$).

The third type of event corresponds to a cancellation of the order in position ℓ . This will reduce the length of the queue by 1 (note that if this results in the queue being depleted, then by (6) the next stopping time will be $\bar{\rho}_{k+2} = (\rho_{k+1}, 2)$ to allow the queue to be randomly replenished immediately). If the agent does not have an order in the queue, then this status is unchanged due

to the cancellation. If the agent does have an active order, then its position decreases by one if the cancellation occurs in front of the agent's order. Finally, note that if there is exactly one order in the queue and the agent has an active order, then an exogenous cancellation cannot occur due to the definition in (7).

The fourth event type is the arrival of an MO. If the size of the MO (denoted here by ϵ) is large enough to fill the agent's order, then the wealth changes by the random value θ . The length of the queue either decreases by ϵ , or if the MO is large enough to deplete the queue then it sends L to zero (and it will be immediately replenished as described in the third type of event). Finally, if the MO is not large enough to fill the agent's LO, then the agent's position decreases by the size of the MO. Otherwise, the status of the agent's order changes to Ξ as the order is filled and removed from the queue.

The model enforces a constraint that the agent always places orders with volume equal to that of a single "round-lot" order. In addition, when a market order arrives the agent's order is either fully executed or not executed. We have avoided the inclusion of partial executions in this model in order to focus on the aspects of queuing in the system. It is possible to include partial executions in the model but this comes with a significant increase in complexity due to the addition of more state variables. In particular, the volume of the agent's order becomes a state variable, and the agent needs to decide not only on the timing of orders but also on their volume. We do not believe that the inclusion of this feature would significantly change the qualitative behaviour of the optimal strategy so we have decided to consider only full executions.

The state of the uncontrolled system starting at index $\bar{\tau}$ in state φ is then defined as

$$\mathcal{S}_{\bar{t}}^{u, \bar{\tau}, \varphi} = (X_{\bar{t}}^{u, \bar{\tau}, \varphi}, Z_{\bar{t}}^{u, \bar{\tau}, \varphi}, L_{\bar{t}}^{u, \bar{\tau}, \varphi}, Y_{\bar{t}}^{u, \bar{\tau}, \varphi}) = \varphi_j, \quad \text{for } \bar{\rho}_j \preceq \bar{t} \prec \bar{\rho}_{j+1} \quad (9)$$

3.3. Controlled Queue Dynamics

The previous subsection outlined how the state of the queue changes if the agent does not intervene. The agent's control will be defined by a sequence of stopping times whose effect on the state of the queue will be outlined here. The full controlled process will then consist of stopping the uncontrolled queue at times corresponding to the agent's strategy, then restarting the uncontrolled queue in a new state defined appropriately by the effect of the agent's action. We will also describe the restrictions put on the stopping times to ensure that the agent's action is feasible.

Let $\{\bar{\tau}_j\}_{j \geq 1} = \{(\tau_j, \kappa_j)\}_{j \geq 1}$ be a sequence of non-decreasing stopping times which represent the times at which the agent sends a signal to place or cancel an order. We denote the controlled state by $\mathcal{S}_t^c$ and define it iteratively by setting $\bar{\tau}_0 = (0, 1)$ and $\phi_0 = \mathcal{S}_{(0,1)}^c = (x, z, \ell, y) \in \mathbb{S}$, and for

$\bar{\tau}_j \preceq \bar{t} \prec \bar{\tau}_{j+1}$ we set

$$\mathcal{S}_t^c = (X_t^c, Z_t^c, L_t^c, Y_t^c) = \mathcal{S}_t^{u, \bar{\tau}_j, \phi_j}. \quad (10)$$

The superscript c in the processes above indicate that they represent the state of the system when taking into account the effects of the agent's control. It remains to define the states ϕ_j which represent the states at which we start the uncontrolled queue immediately after the agent exercises her control. Let $\eta > 0$ be constant which is to reflect a small incremental cost incurred by the agent for sending a signal. Define the operator $A : \mathbb{S} \rightarrow \mathbb{S}$ by

$$A(x, z, \ell, y) = \begin{cases} (x - \eta, z, \ell + 1, \ell + 1), & y = \Xi \\ (x - \eta, z, \ell - 1, \Xi), & y \neq \Xi \end{cases} \quad (11)$$

and then set

$$\phi_{j+1} = A(\mathcal{S}_{\bar{\tau}_{j+1}}^{u, \bar{\tau}_j, \phi_j}) \quad (12)$$

The value of $A(x, z, \ell, y)$ represents the state of the queue once the signal is sent given that the state is (x, z, ℓ, y) immediately before the signal.

Finally, we impose the following restrictions on $\{\bar{\tau}_j\}_{j \geq 1} = \{(\tau_j, \kappa_j)\}_{j \geq 1}$:

$$\lim_{j \rightarrow \infty} \tau_j > T, \quad (13)$$

$$Y_{\bar{\tau}_j}^{u, \bar{\tau}_{j-1}, \phi_{j-1}} = \Xi \Rightarrow L_{\bar{\tau}_j}^{u, \bar{\tau}_{j-1}, \phi_{j-1}} < N^L \quad (14)$$

$$\kappa_j = \begin{cases} 1 & \text{if } L_{\bar{\tau}_j}^{u, \bar{\tau}_{j-1}, \phi_{j-1}} = 1 \text{ and } Y_{\bar{\tau}_j}^{u, \bar{\tau}_{j-1}, \phi_{j-1}} = 1, \\ 2 & \text{otherwise} \end{cases} \quad (15)$$

We denote the set of sequences which satisfy these conditions by $\mathcal{A}$. Condition (13) ensures that the number of signals sent by the agent over the trading period is finite almost surely. This simplifying assumption will be satisfied automatically by the optimal strategy anyway since the cost satisfies $\eta > 0$. Condition (14) simply states that if the agent does not have an order in the queue, then she may only place one if the length of the queue is strictly less than the maximum allowable length. This simplifying assumption could be relaxed in a more generalized framework and is also convenient for numerical purposes.

Condition (15) is related to the information that the agent has at the time of executing her strategy. If $L_t^c = Y_t^c = 1$, then the agent's order is the only one in the queue, and cancelling it would send the length of the queue to zero to be replenished immediately after. We do not wish for the agent's strategy to consist of a cancellation when she has foreknowledge of the length of the replenished queue. Thus, when cancelling at an index of the form $\bar{t} = (t, 1)$, the queue is replenished at $(t, 2)$ to the level $\zeta_{(t, 2)}^{Z_{(t, 1)}^c}$, which is an $\mathcal{F}_{(t, 2)}$ measurable random variable. The agent is then allowed to replace the order at index $(t, 2)$, depending on the value of $\zeta_{(t, 2)}^{Z_{(t, 1)}^c}$.

3.4. Performance Criteria and Optimization

The agent attempts to choose a sequence of stopping times satisfying conditions (13) to (15), denoted by $\{\bar{\tau}_j\}_{j \geq 1} \in \mathcal{A}$, which maximizes the utility of her terminal wealth. We define the value function as

$$H(\bar{t}, x, z, \ell, y) = \sup_{\{\bar{\tau}_j\}_{j \geq 1} \in \mathcal{A}} \mathbb{E} \left[-e^{-\gamma X_{\bar{t}}^c} \middle| X_{\bar{t}}^c = x, Z_{\bar{t}}^c = z, L_{\bar{t}}^c = \ell, Y_{\bar{t}}^c = y \right], \quad (16)$$

where γ is a risk aversion parameter. This utility function will be convenient because it will allow us to make an ansatz for the form of H and eliminate the variable x from the equations that are to be solved, reducing the dimensionality of the system. In addition this will have the effect that the optimal strategy does not depend on the agent's wealth. We suppose that the dynamic programming principle holds for any stopping time $\bar{\tau} \preceq (T, 2)$:

$$H(\bar{t}, x, z, \ell, y) = \sup_{\{\bar{\tau}_j\}_{j \geq 1} \in \mathcal{A}} \mathbb{E} \left[H(\bar{\tau}, X_{\bar{\tau}}^c, Z_{\bar{\tau}}^c, L_{\bar{\tau}}^c, Y_{\bar{\tau}}^c) \middle| X_{\bar{t}}^c = x, Z_{\bar{t}}^c = z, L_{\bar{t}}^c = \ell, Y_{\bar{t}}^c = y \right]. \quad (17)$$

At this point we would like to eliminate the use of the time index notation of the form $\bar{t} = (t, k)$ in order to write a more easily readable dynamic programming equation. To this end, suppose $\bar{t} = (t, 1)$ on the left hand side of (17) and $\bar{\tau} = (t, 2)$ on the right hand side. If $\ell > 1$ or $y = \Xi$ then

$$H((t, 1), x, z, \ell, y) = H((t, 2), x, z, \ell, y), \quad (18)$$

due to (15) (the agent is not allowed to take action at this particular time index, so the uncontrolled dynamics proceed). On the other hand, if $\ell = y = 1$, then

$$H((t, 1), x, z, \ell, y) = \max \left\{ \mathbb{E} \left[H((t, 2), x - \eta, z, \zeta_{(t,2)}^z, \Xi) \right], H((t, 2), x, z, \ell, y) \right\} \quad (19)$$

where the expectation is taken over the random variable $\zeta_{(t,2)}^z$. Equation (19) carries the intuition that if the agent is the only one with an order in the queue, then the optimal choice is based on the relative value of proceeding to the next instant in time with no action versus the average value to the agent if they were to cancel and generate a random queue length. The form of the first term within the maximum of (19) is due to the last lines of (6) and (8). By considering both (18) and (19) we are able to eliminate any use of time index of the form $(t, 1)$ and write the value function in terms of only $(t, 2)$. With a slight abuse of notation, from this point forward we now denote the value function by $H(t, x, z, \ell, y)$ where we are implicitly referring only to the value function at an expanded time index of the form $(t, 2)$.

4. Solving for the Value Function and Optimal Strategy

4.1. Dynamic Programming Equation

In this section we write the dynamic programming equation corresponding to the agent's optimization problem as described above. Due to the drastic difference in the dynamics depending on whether $y = \Xi$ or $y \neq \Xi$ as well as the restriction $\ell \leq N^L$, we will separate the writing of this equation into three cases. We first write the equation for $y = \Xi$ and $\ell < N^L$. Second, we consider $y = \Xi$ and $\ell = N^L$. Finally, we have the case $y \neq \Xi$. Recall that we have introduced the notation $\lambda^{c,z,\ell} = \sum_{k=1}^{\ell} \tilde{\lambda}^{c,z,k}$. This notation will be convenient because any cancellation of an order in front of the agent's order, regardless of its position, has the same effect on the agent's value function. Similarly for cancellations behind the agent's order. For the case $y = \Xi$ and $\ell < N^L$ the dynamic programming equation is

$$\begin{aligned} \max \Big\{ & \partial_t H + \sum_{\bar{z} \neq z} \lambda^{z,\bar{z},\ell} \left(H(t, x, \bar{z}, \ell, \Xi) - H \right) \\ & + \lambda^{m,z} \mathbb{E} \left[\mathbb{1}_{\epsilon < \ell} H(t, x, z, \ell - \epsilon, \Xi) + \mathbb{1}_{\epsilon \geq \ell} H(t, x, z, \zeta, \Xi) - H \right] \\ & + \lambda^{c,z,\ell} \left(\mathbb{1}_{\ell > 1} H(t, x, z, \ell - 1, \Xi) + \mathbb{1}_{\ell = 1} \mathbb{E} \left[H(t, x, z, \zeta, \Xi) \right] - H \right) \\ & + \lambda^{a,z,\ell} \left(H(t, x, z, \ell + 1, \Xi) - H \right); \\ & H(t, x - \eta, z, \ell + 1, \ell + 1) - H \Big\} = 0. \end{aligned} \quad (20)$$

For the case $y = \Xi$ and $\ell = N^L$, the equation takes the form

$$\begin{aligned} & \partial_t H + \sum_{\bar{z} \neq z} \lambda^{z,\bar{z},N^L} \left[H(t, x, \bar{z}, N^L, \Xi) - H \right] \\ & + \lambda^{m,z} \mathbb{E} \left[\mathbb{1}_{\epsilon < N^L} H(t, x, z, N^L - \epsilon, \Xi) + \mathbb{1}_{\epsilon \geq N^L} H(t, x, z, \zeta, \Xi) - H \right] \\ & + \lambda^{c,z,N^L} \left(H(t, x, z, N^L - 1, \Xi) - H \right) = 0. \end{aligned} \quad (21)$$

Finally, when $y \neq \Xi$ the equation is

$$\begin{aligned}
& \max \left\{ \partial_t H + \sum_{\bar{z} \neq z} \lambda^{z, \bar{z}, \ell} \left(H(t, x, \bar{z}, \ell, y) - H \right) \right. \\
& \quad + \lambda^{m, z, \ell} \mathbb{E} \left[\mathbf{1}_{\epsilon < y} H(t, x, z, \ell - \epsilon, y - \epsilon) + \mathbf{1}_{y \leq \epsilon < \ell} H(t, x + \theta, z, \ell - \epsilon, \Xi) \right. \\
& \quad \quad \left. + \mathbf{1}_{\epsilon \geq \ell} H(t, x + \theta, z, \zeta, \Xi) - H \right] \\
& \quad + \lambda^{c, z, y-1} \left(H(t, x, z, \ell - 1, y - 1) - H \right) \\
& \quad + (\lambda^{c, z, \ell} - \lambda^{c, z, y}) \left(H(t, x, z, \ell - 1, y) - H \right) \\
& \quad + \lambda^{a, z, \ell} \left(H(t, x, z, \ell + 1, y) - H \right); \\
& \quad \left. \mathbf{1}_{\ell=1} \mathbb{E} \left[H(t, x - \eta, z, \zeta, \Xi) \right] + \mathbf{1}_{\ell > 1} H(t, x - \eta, z, \ell - 1, \Xi) - H \right\} = 0.
\end{aligned} \tag{22}$$

The terminal condition does not depend on the restrictions on the state variables and is given by

$$H(T, x, z, \ell, y) = -e^{-\gamma x}. \tag{23}$$

The expectations in (20) to (22) are taken with respect to the independent random variables ϵ , θ , and ζ with marginal distributions $\mu^{\epsilon, z}$, $\mu^{\theta, z}$, and $\mu^{\zeta, z}$ respectively. Equation (21) does not take the form of a quasi-variational inequality because the agent is restricted from adding an order to the queue when $\ell = N^L$.

Each term in the dynamic programming equation has a rather straightforward explanation. We will focus on giving this for equation (22), the others follow similarly. The max between two terms represents the optimal decision by the agent whether to cancel her order based on the current value of the state variables, or to take no action and let the uncontrolled dynamics of the state variables proceed.

Within the first quantity being compared via the max operator, the summation term represents the average rate of change of the value function due to regime changes of the TS. The next term represents the expected rate of change due to the arrival of an MO. Note that there are three possible outcomes: the MO does not fill the agent's order (in which case the order moves forward by the MO size and the queue length decreases), the MO fills the agent's order (wealth changes by random quantity and the agent's order status becomes Ξ), or the MO depletes the entire queue (the agent's order is filled and the queue replenished at the random size ζ).

The next two terms represent the expected rate of change due to exogenous LO cancellations. These may happen either ahead of the agent's order (with total rate $\lambda^{c, z, y-1}$) or behind the agent's

order (with total rate $\lambda^{c,z,\ell} - \lambda^{c,z,y}$). The difference between the two changes in value is that in one case the agent's order moves closer to the front of the queue and in the other case the position stays the same but total queue length still decreases. The final rate term is similar but corresponds to the addition of a single order to the back of the queue.

The second quantity within the max comparison represents the change in the value function given that the agent cancels her order. Note that there are again two cases depending on the state of the queue. If the agent's order is the only one present ($\ell = 1$) then a cancellation will cause the queue to be replenished at the random length ζ . Otherwise if $\ell > 1$ the queue simply decreases by size 1.

The form of equations (20) to (22) along with the terminal conditions (23) allow us to make the ansatz $H(t, x, z, \ell, y) = -e^{-\gamma(x+h(t,z,\ell,y))}$. Making this substitution in the above equations results in the following: for $y = \Xi$ and $\ell < N^L$

$$\begin{aligned} \min \Big\{ & -\gamma \partial_t h + \sum_{\bar{z} \neq z} \lambda^{z,\bar{z},\ell} \left(e^{-\gamma(h(t,\bar{z},\ell,\Xi)-h)} - 1 \right) \\ & + \lambda^{m,z} \mathbb{E} \left[\mathbb{1}_{\epsilon < \ell} e^{-\gamma(h(t,z,\ell-\epsilon,\Xi)-h)} + \mathbb{1}_{\epsilon \geq \ell} e^{-\gamma(h(t,z,\zeta,\Xi)-h)} - 1 \right] \\ & + \lambda^{c,z,\ell} \left(\mathbb{1}_{\ell > 1} e^{-\gamma(h(t,z,\ell-1,\Xi)-h)} + \mathbb{1}_{\ell=1} \mathbb{E} \left[e^{-\gamma(h(t,z,\zeta,\Xi)-h)} \right] - 1 \right) \\ & + \lambda^{a,z,\ell} \left(e^{-\gamma(h(t,z,\ell+1,\Xi)-h)} - 1 \right); \\ & e^{-\gamma(h(t,z,\ell+1,\ell+1)-h-\eta)} - 1 \Big\} = 0, \end{aligned} \quad (24)$$

for $y = \Xi$ and $\ell = N^L$ we have

$$\begin{aligned} & -\gamma \partial_t h + \sum_{\bar{z} \neq z} \lambda^{z,\bar{z},N^L} \left(e^{-\gamma(h(t,\bar{z},N^L,\Xi)-h)} - 1 \right) \\ & + \lambda^{m,z} \mathbb{E} \left[\mathbb{1}_{\epsilon < N^L} e^{-\gamma(h(t,z,N^L-\epsilon,\Xi)-h)} + \mathbb{1}_{\epsilon \geq N^L} e^{-\gamma(h(t,z,\zeta,\Xi)-h)} - 1 \right] \\ & + \lambda^{c,z,N^L} \left(e^{-\gamma(h(t,z,N^L-1,\Xi)-h)} - 1 \right), \end{aligned} \quad (25)$$

and finally for $y \neq \Xi$ we have

$$\begin{aligned}
& \min \left\{ -\gamma \partial_t h + \sum_{\bar{z} \neq z} \lambda^{z, \bar{z}, \ell} \left(e^{-\gamma(h(t, \bar{z}, \ell, y) - h)} - 1 \right) \right. \\
& \quad + \lambda^{m, z} \mathbb{E} \left[\mathbb{1}_{\epsilon < y} e^{-\gamma(h(t, z, \ell - \epsilon, y - \epsilon) - h)} + \mathbb{1}_{y \leq \epsilon < \ell} e^{-\gamma(\theta + h(t, z, \ell - \epsilon, \Xi) - h)} \right. \\
& \quad \quad \left. \left. + \mathbb{1}_{\epsilon \geq \ell} e^{-\gamma(\theta + h(t, z, \zeta, \Xi) - h)} - 1 \right] \right. \\
& \quad + \lambda^{c, z, y-1} \left(e^{-\gamma(h(t, z, \ell-1, y-1) - h)} - 1 \right) \\
& \quad + (\lambda^{c, z, \ell} - \lambda^{c, z, y}) \left(e^{-\gamma(h(t, z, \ell-1, y) - h)} - 1 \right) \\
& \quad + \lambda^{a, z, \ell} \left(e^{-\gamma(h(t, z, \ell+1, y) - h)} - 1 \right); \\
& \quad \left. \mathbb{1}_{\ell=1} \mathbb{E} \left[e^{-\gamma(h(t, z, \zeta, \Xi) - h - \eta)} \right] + \mathbb{1}_{\ell > 1} e^{-\gamma(h(t, z, \ell-1, \Xi) - h - \eta)} - 1 \right\} = 0.
\end{aligned} \tag{26}$$

These equations are subject to terminal conditions

$$h(T, z, \ell, y) = 0. \tag{27}$$

The form of equations (24) and (26) due to the ansatz allow us to conclude that the agent's decision to place and cancel orders does not depend on how much wealth she has accumulated over time (the equations are independent of x). The decision is only based on the current state of the queue and market parameters.

4.2. Optimal Strategy

Construction of the optimal strategy follows along the same lines as in other optimal impulse control problems. We simply need to make some simple modifications due to the form of the indexing set of our stochastic processes. By inspecting the system of QVI's (24) to (26), we see that they for an appropriate choice of functions F_i and G_i they can be written in the form

$$\min \left\{ F_i(t, h(t), \partial_t h_i(t)); G_i(t, h(t)) \right\} = 0, \tag{28}$$

where $i = (z, \ell, y)$ indexes the system of equations (the exact form of F_i and G_i will be given later in Section 6). We then use the system to define a continuation region and an impulse region:

$$\mathcal{C} = \left\{ (t, z, \ell, y) : t \leq T \text{ and } G_i(t, h(t)) > 0 \right\} \quad (29)$$

$$\mathcal{I} = \left\{ (t, z, \ell, y) : t \leq T \text{ and } G_i(t, h(t)) = 0 \right\}. \quad (30)$$

The optimal strategy consists of a sequence of stopping times $\{\bar{\tau}_j\}_{j=1}^\infty = \{(\tau_j, \kappa_j)\}_{j=1}^\infty$ which will be constructed recursively. By convention we set $\bar{\tau}_0 = (0, 1)$, then given $\bar{\tau}_{j-1}$ we first define τ_j as:

$$\tau_j = \inf \left\{ t \geq \tau_{j-1} : (t, Z_{(t,2)}^{u, \bar{\tau}_{j-1}, \phi_{j-1}}, L_{(t,2)}^{u, \bar{\tau}_{j-1}, \phi_{j-1}}, Y_{(t,2)}^{u, \bar{\tau}_{j-1}, \phi_{j-1}}) \in \mathcal{I} \right\}. \quad (31)$$

Recall that $Z^{u, \bar{\tau}_{j-1}, \phi_{j-1}}$, $L^{u, \bar{\tau}_{j-1}, \phi_{j-1}}$, and $Y^{u, \bar{\tau}_{j-1}, \phi_{j-1}}$ are the uncontrolled processes starting from index $\bar{\tau}_{j-1}$ in state ϕ_{j-1} . Once τ_j is defined, κ_j is chosen based on the state of the queue at time τ_j as

$$\kappa_j = \begin{cases} 1 & \text{if } L_{(\tau_j, 1)}^{u, \bar{\tau}_{j-1}, \phi_{j-1}} = 1 \text{ and } Y_{(\tau_j, 1)}^{u, \bar{\tau}_{j-1}, \phi_{j-1}} = 1, \\ 2 & \text{otherwise} \end{cases} \quad (32)$$

This strategy is easily seen to be admissible: $\{\bar{\tau}_j\}_{j=1}^\infty$ satisfies (13) since with probability 1 there are a finite number of jumps of each of the Poisson processes $N^{z, \bar{z}, \ell}$, $N^{c, z, \ell}$, $N^{a, z, \ell}$, and M^z , so $\tau_j = \infty$ for all sufficiently large j ; it satisfies (14) because $(t, z, N^L, \Xi) \notin \mathcal{I}$ by the definition $G_{z, N^L, \Xi} = \infty$ (see Section 6); and it satisfies (15) by virtue of (32).

In the next section we graphically illustrate the continuation and impulse regions.

5. Numerical Examples

The system of equations needed to be solved for the value function are highly coupled and complex. As is often the case with impulse control problems, we are not able to present closed form expressions for the value function or the optimal strategy. Thus, in order to investigate the qualitative behaviour of the optimal strategy as it depends on the state of the queue, we must resort to numerical solutions. In this section, we present numerical results that demonstrates the agent's optimal strategy (in discussing the intuition behind the results we will always interpret the queue as being on the sell side of the LOB).

5.1. Parameters

We use volume imbalance to construct the trade signal Z_t (see Subsection 2.2). Most parameters, including the rate of addition and cancellation ($\lambda^{a, z, \ell}$ and $\lambda^{c, z, \ell}$), the rate of transition between

regimes $(\lambda^{z,\bar{z},\ell})$, the rate of MO arrival $(\lambda^{m,z})$, the distribution of MO size $(\mu^{\epsilon,z})$, the distribution of gain/loss of a filled LO $(\mu^{\theta,z})$, and the distribution of replenished queue length $(\mu^{\zeta,z})$ are described in Section 2.

In all examples below we select the constant $T = 50 \text{ seconds}$ as the length of the trading period. This is done because the boundary between the continuation region and impulse region reaches a steady state when t is far from T . In our numerical examples we find that 50 seconds is sufficiently large. Using a larger value of T would not change the optimal strategy at time $t = 0$.

We also set the constant $\Upsilon = 0.5$. Recall that Υ represents the instantaneous profit of a filled LO if the underlying fundamental price does not change after execution. If the midprice is seen as a representation of the fundamental price of the stock, then the instantaneous profit is the half-spread. The smallest possible value of the spread is 1, and for large tick stocks (such as INTC) nearly all of the trading time during the day observes a spread of 1.

The parameters γ and η are subjective and can change between different agents. For this reason we briefly investigate how the optimal strategy changes depending on these two parameters in this section. For a single set of parameters the computation of the optimal strategy was performed in 2,046 seconds in Matlab on an Intel Core i7 2.7GHz CPU.

5.2. The Optimal Strategy

Figure 7 shows the optimal switching regions when the agent has an active LO ($y \neq \Xi$). The left panel shows both the switching region and the continuation region when the TS is in the neutral regime ($Z_t = 2$) and the right panel shows these regions when the TS is in the adverse regime ($Z_t = 3$). The switching region (represented by yellow colouring in both panels) indicates when it is optimal to cancel an existing LO rather than let it remain in the queue. When the TS is in the gainful regime ($Z_t = 1$), the agent will never cancel an existing LO, so the figure is omitted.

There are several features of these figures which are worthy of discussion. First, for both panels, we can see that if we fix any queue position Y_t , as the queue length L_t decreases, the agent eventually changes her strategy from keeping the order to cancelling once a threshold for L_t is crossed. There are two intuitive reasons for this. First, as the length of the queue decreases, the rate of transition to the gainful regime decreases as well, making the choice to remain in the queue more risky (see Figure 1). Second, the agent has an incentive to cancel the order to avoid adverse selection when not in the gainful regime, but if she does and the TS becomes favourable shortly after, then she would like to put the order back in the queue and she has needlessly sacrificed her queue position. However, if the total queue length is shorter, than this sacrifice becomes less significant to the point where it is worth giving up the position to avoid adverse selection, but not difficult to replace the order in the case that the TS becomes gainful again after cancelling.

Second, we see that the switching regions in Figure 7 are not monotone in the queue position Y_t .

Figure 7: Optimal strategy when $t = 0$ and $Y_t \neq \Xi$. Other parameters are $\eta = 10^{-4}$ and $\gamma = 0.4$. The yellow region indicates where the agent's optimal policy is to cancel her order. The blue region indicates where she takes no action. The red curve represents the position in the queue as a function of its total length which maximizes the agent's value function.

For example, on the right panel, when the queue length L_t is around 40, the agent will cancel her LO when Y_t is close to 0 and when Y_t is close to L_t ; in between, she stays in the queue. The reasoning for this is similar to the second line of intuition above. When the agent's order is at the back of the queue and there is a small level of risk to being filled by an MO, the agent may simply cancel the order since there is essentially no sacrifice to queue position, and replace the order if the TS becomes favourable shortly after. As the order moves slightly closer to the front, the sacrifice of queue position eventually outweighs the risk of adverse selection, so the agent retains her order in the queue. The value of this queue position will increase if the TS changes to a favourable state. With a further move towards the front of the queue, the risk of adverse selection becomes imminent as most MOs are certain to fill her LO.

Although the switching regions in the two panels of Figure 7 look similar, the actual frequency of switching can be quite different. This is because the distribution of queue length is different for the different regimes. From Figure 8, we can see that when the regime is neutral, the distribution of queue length peaks at around 40; whereas when the regime is adverse, the distribution peaks at somewhere between 10 and 20. The implication of this is that in the adverse regime, the agent almost always cancels an LO and in the neutral regime, there is some proportion of time that the agent stays in the queue.

Lastly, we also plot the position in the queue in terms of its total length which maximizes the agent's value function. Precisely stated, the red curve corresponds to:

$$y^*(t, Z_t, L_t) = \arg \max_{y \leq L_t} h(t, Z_t, L_t, y) .$$

We only plot the curve within the continuation region, since every position within the cancellation region would have the same value to the agent. In the gainful regime, the most valuable position to the agent is the front of the queue, so this is also not displayed. Of particular significance is

Figure 8: Distribution of queue length in different regimes. The solid lines are the simulation results from our queuing model. The dashed lines are the empirical distributions.

Figure 9: Optimal strategy when $Y_t = \Xi$ and $t = 0$. Other parameters are $\eta = 10^{-4}$ and $\gamma = 0.4$.

the fact that there are certain states in which the front of the queue is not considered the most valuable. This is another illustration of the trade-off between the adverse selection suffered by being filled by an MO and the desire to be near the front of the queue in preparation for when the regime becomes gainful. When the queue length becomes sufficiently long the front position is always the most valuable. This stems from the fact that a longer queue indicates a higher rate of regime change into a gainful regime to the point where this is more likely to occur before the next MO (see Figure 1).

Figure 9 shows the optimal strategy as a function of L_t when the agent does not have an active LO ($Y_t = \Xi$). When the LOB is in the gainful regime, the agent always places a new LO immediately. When the LOB is in the neutral or adverse regime, the agent will place an order when the length of the queue is above a certain threshold. The agent does not enter the queue when it is short because she wants to avoid the risk of adverse selection, but when the queue is long this risk is reduced and it is worth placing the order to ensure a good position before the TS changes to a favourable regime. As expected, and the threshold for placing the order in the adverse regime is higher than that in the neutral regime.

Figure 10: Optimal strategy when $t = 0$ and $Y_t \neq \Xi$. Other parameters are $\eta = 10^{-3}$ and $\gamma = 0.4$. The difference in the parameters from Figure 7 is that η is 10 times larger.

Figure 11: Optimal strategy when $t = 0$ and $Y_t = \Xi$. Other parameters are $\eta = 10^{-3}$ and $\gamma = 0.4$. The difference in the parameters from Figure 9 is that η is 10 times larger.

5.3. Increasing Transaction Cost

Figures 10 and 11 show the optimal strategy when the transaction cost η is increased to 10^{-3} . Figure 10 shows the switching (yellow) region when the agent has an active LO. We can see that comparing to Figure 7, the yellow regions are smaller and the blue regions are larger. This is expected as the agent has less incentive to cancel her existing LO. Similarly, Figure 11 shows the agent's action when she does not have an active LO. When comparing to Figure 9 in all regimes, particularly the gainful regime, we see that the agent is less likely to place a new LO. Of particular interest is the fact that the agent will not post an order in the gainful regime when the queue is short. This may seem counter-intuitive at first, but it is due to the fact that the rate of transition to the adverse and neutral regimes increase as the queue length decreases. If the agent places an LO and the queue switches to one of the other two regimes, she will likely be in the cancellation region (yellow region in Figure 10). When she cancels, she incurs another transaction cost. To avoid this situation, her optimal strategy is not to enter the queue when it is very short in the gainful regime. As the queue becomes longer she becomes willing to place her order because even in a longer queue there is a significant change that an MO fills her order, but less chance that the regime changes before this occurs. In summary, when the transaction cost η increases, the agent is less likely to either cancel or place an LO.

Figure 12: Optimal strategy and most valuable queue position when $t = 0$ and $Y_t \neq \Xi$. Other parameters are $\eta = 10^{-3}$ and $\gamma = 1.2$. The difference in the parameters from Figure 7 is a larger value of γ .

Figure 13: Optimal strategy when $t = 0$ and $Y_t = \Xi$. Other parameters are $\eta = 10^{-3}$ and $\gamma = 1.2$. The difference in the parameters from Figure 9 is a larger value of γ .

5.4. Increasing Risk-Aversion

Figure 12 and Figure 13 demonstrate the optimal strategy when the agent is more risk averse ($\gamma = 1.2$ increased from $\gamma = 0.4$ in the previous examples). Comparing Figure 12 to Figure 7, the yellow region is much larger because the agent is more inclined to cancel her LO to avoid the risk associated with adverse selection. Similarly, comparing Figure 13 to Figure 9, we see that the thresholds to place a new LO have increased to avoid the risk of having LOs active during the presence of an undesirable TS.

5.5. Decreasing Risk-Aversion

Figure 14 and Figure 15 show the optimal strategy for the agent when she is less risk-averse ($\gamma = 0.1$). Comparing Figures 14 and 15 to Figures 7 and 9 shows typical expected behaviour through the fact that the agent is in all cases more willing to have active LOs, hence the strategy is more aggressive. Note again the presence of non-monotonicity of the strategy with respect to queue position Y_t when the TS is adverse.

Figure 14: Optimal strategy and most valuable queue position when $t = 0$ and $Y_t \neq \Xi$. Other parameters are $\eta = 10^{-3}$ and $\gamma = 0.1$. The difference in the parameters from Figure 7 is a smaller value of γ .

Figure 15: Optimal strategy when $t = 0$ and $Y_t = \Xi$. Other parameters are $\eta = 10^{-3}$ and $\gamma = 0.1$. The difference in the parameters from Figure 9 is a smaller value of γ .

5.6. Simulation

In this section we attempt to quantify how the incorporation of queue length and order position into the optimal decision policy benefits the agent when compared to the performance of various strategies which only consider the value of the TS. We do this by simulating the dynamics of the queue according to the model outlined in Section 3. We simulate 10^6 sample paths using the parameters $\eta = 10^{-4}$, $\Upsilon = 0.5$ and $T = 50$ seconds. The initial state of the regime Z_0 is randomly drawn from the empirical distribution of Z_t , and the initial queue length L_0 is drawn from the empirical distribution $\mu^{\zeta, z}$, conditioning on $z = Z_0$. The agent starts with zero wealth and no active LO, i.e., $X_0 = 0$ and $Y_0 = \Xi$. The four different order placement strategies which we test in these simulations are outlined below.

- *Optimal*: the optimal strategy (as defined by cancel regions, continue regions, and new order thresholds) with risk aversion parameter γ varying from 0.1 to 1.27.
- *Passive Benchmark*: the agent places a new order whenever the TS becomes gainful, and cancels the order whenever the TS becomes neutral or adverse.

Figure 16: Risk-reward plot. Along the curve indicating performance of the optimal strategy, the right-most point represents $\gamma = 0.1$. Increasing γ lowers both the mean and standard deviation of terminal wealth.

- *Aggressive Benchmark*: the agent places a new order whenever the TS becomes gainful or neutral, and cancels the order whenever the TS becomes adverse.
- *Always-post*: the agent never cancels her order and always replaces it when it is filled by an MO.

Figure 16 shows the risk-reward plot for all the strategies. The horizontal axis is the standard deviation of the agent’s terminal wealth and the vertical axis is her mean terminal wealth. The blue curve represents the performance of the optimal strategy for various values of the risk-aversion parameter. From the figure, we see the always-post strategy clearly shows poor performance compared to the other strategies. Even though the always-post action is likely to occupy order positions which are at the front of the queue for a greater proportion of time than the other strategies, previous analysis and reasoning indicates that such a position can have negative value to the agent if there is a high chance of being filled by an MO and suffering adverse selection.

Table 4 compares some of the optimal strategies for different risk aversion parameters with aggressive benchmark. The optimal strategy when $\gamma = 0.1$ and aggressive benchmark have very similar standard deviations for the terminal wealth; but the optimal strategy has a mean terminal wealth of 2.23, which is 2.5% higher than aggressive benchmark. When $\gamma = 0.4$, the optimal strategy and aggressive benchmark have very similar means, but the standard deviation of the optimal strategy is 1.825, which is 8.8% lower.

Strategy	Mean	Std
$\gamma = 0.1$	2.230	1.976
$\gamma = 0.4$	2.182	1.825
Aggressive Benchmark	2.175	1.979

Table 4: Mean and standard deviation of terminal wealth.

To better understand the differences between these strategies, let us have a look at Figure 17, which shows the histograms of terminal wealth for all strategies. Panel (d) shows the case for

Figure 17: Histogram of terminal wealth.

the naive always-post strategy. Its terminal wealth has a lot of variation, with a large number of samples being negative. Panel (c) shows the case for aggressive benchmark. Compared to the always-post strategy, aggressive benchmark is more conservative: the variance is smaller and there are fewer samples ending up with negative wealth. Panel (b) shows the case for passive benchmark, which seems to be much more conservative than the rest. Panel (a) shows the case for the optimal strategy ($\gamma = 0.4$). We can see that its variance lies in between passive benchmark and aggressive benchmark. Table 5 shows the proportions of LOs executed as three regimes for the four strategies. The always-post strategy has a large fraction of LOs (21.1%) executed at the adverse regime. Passive benchmark goes the other extreme, with all LOs being executed at the gainful regime. For the optimal strategy, we see that there is only a tiny fraction of LOs (0.2%) executed at the adverse regime. Compared to aggressive benchmark, there are relatively more LOs executed at the gainful regime.

Strategy	$z = 1$	$z = 2$	$z = 3$	number of filled orders
optimal ($\gamma = 0.4$)	76.5	23.3	0.2	9.85
passive benchmark	100.0	0.0	0.0	2.43
aggressive benchmark	64.0	36.0	0.0	11.91
always post	45.1	33.7	21.1	17.64

Table 5: Percentages of orders executed in each regime.

6. Convergence Results

In this section, we show that the system of QVIs (24) to (26) admits a unique solution and provide a numerical scheme that converges to this solution.

We denote by $\bar{\mathbb{S}}$ the reduced state space, equivalent to that of (5) without the X component:

$$\bar{\mathbb{S}} = \left\{ i = (Z, L, Y) : Z \in \{1, \dots, N^Z\}, L \in \{0, 1, \dots, N^L\}, \right. \\ \left. Y \in \{1, \dots, N^L, \Xi\}, Y = \Xi \text{ or } Y \leq L \right\}. \quad (33)$$

This is the set of state variables of concern with respect to equations (24) to (26). We also write $\mathbb{R}^{\bar{\mathbb{S}}} = \{r_i : r_i \in \mathbb{R}, i \in \bar{\mathbb{S}}\}$. With the above notations, the original system of QVIs can be restated as

$$J_i(t, u, \partial_t u_i) = \min \{F_i(t, u, \partial_t u_i); G_i(t, u)\} = 0, \\ u_i(T) = 0, \quad (34)$$

for $i \in \bar{\mathbb{S}}$ and $u : [0, T] \rightarrow \mathbb{R}^{\bar{\mathbb{S}}}$. In other words, we are viewing the original system of QVIs as a system indexed by $i \in \bar{\mathbb{S}}$. The mappings F_i and G_i in (34) are defined as

$$F_{z,\ell,\Xi}(t, r, p) = -\gamma p + \sum_{\bar{z} \neq z} \lambda^{z,\bar{z},\ell} \left(e^{-\gamma(r_{\bar{z},\ell,\Xi} - r_{z,\ell,\Xi})} - 1 \right) \\ + \lambda^{m,z,\ell} \left(\mathbb{E} \left[\mathbf{1}_{\epsilon < \ell} e^{-\gamma(r_{z,\ell-\epsilon,\Xi} - r_{z,\ell,\Xi})} + \mathbf{1}_{\epsilon \geq \ell} e^{-\gamma(r_{z,\zeta,\Xi} - r_{z,\ell,\Xi})} \right] - 1 \right) \\ + \lambda^{c,z,\ell} \left(\mathbf{1}_{\ell > 1} e^{-\gamma(r_{z,\ell-1,\Xi} - r_{z,\ell,\Xi})} + \mathbf{1}_{\ell=1} \mathbb{E} \left[e^{-\gamma(r_{z,\zeta,\Xi} - r_{z,\ell,\Xi})} \right] - 1 \right) \\ + \lambda^{a,z,\ell} \left(e^{-\gamma(r_{z,\ell+1,\Xi} - r_{z,\ell,\Xi})} - 1 \right), \\ G_{z,\ell,\Xi}(t, r) = e^{-\gamma(r_{z,\ell+1,\ell+1} - r_{z,\ell,\Xi} - \eta)} - 1,$$

for $i = (z, \ell, \Xi)$ with $\ell < N^L$, as

$$F_{z,N^L,\Xi}(t, r, p) = -\gamma p + \sum_{\bar{z} \neq z} \lambda^{z,\bar{z},N^L} \left(e^{-\gamma(r_{\bar{z},N^L,\Xi} - r_{z,N^L,\Xi})} - 1 \right) \\ + \lambda^{m,z,N^L} \left(\mathbb{E} \left[\mathbf{1}_{\epsilon < N^L} e^{-\gamma(r_{z,N^L-\epsilon,\Xi} - r_{z,N^L,\Xi})} + \mathbf{1}_{\epsilon \geq N^L} e^{-\gamma(r_{z,\zeta,\Xi} - r_{z,N^L,\Xi})} \right] - 1 \right) \\ + \lambda^{c,z,N^L} \left(e^{-\gamma(r_{z,N^L-1,\Xi} - r_{z,N^L,\Xi})} - 1 \right), \\ G_{z,N^L,\Xi}(t, r) = \infty,$$

for $i = (z, N^L, \Xi)$, and as

$$\begin{aligned}
F_{z,\ell,y}(t, r, p) &= -\gamma p + \sum_{\bar{z} \neq z} \lambda^{z,\bar{z},\ell} \left(e^{-\gamma(r_{\bar{z},\ell,y} - r_{z,\ell,y})} - 1 \right) \\
&\quad + \lambda^{m,z,\ell} \left(\mathbb{E} \left[\mathbb{1}_{\epsilon < y} e^{-\gamma(r_{z,\ell-\epsilon,y} - r_{z,\ell,y})} + \mathbb{1}_{y \leq \epsilon < \ell} e^{-\gamma(\theta + r_{z,\ell-\epsilon,\Xi} - r_{z,\ell,y})} \right. \right. \\
&\quad \left. \left. + \mathbb{1}_{\epsilon \geq \ell} e^{-\gamma(\theta + r_{z,\zeta,\Xi} - r_{z,\ell,y})} \right] - 1 \right) \\
&\quad + \lambda^{c,z,y-1} \left(e^{-\gamma(r_{z,\ell-1,y-1} - r_{z,\ell,y})} - 1 \right) \\
&\quad + (\lambda^{c,z,\ell} - \lambda^{c,z,y}) \left(e^{-\gamma(r_{z,\ell-1,y} - r_{z,\ell,y})} - 1 \right) \\
&\quad + \lambda^{a,z,\ell} \left(e^{-\gamma(r_{z,\ell+1,y} - r_{z,\ell,y})} - 1 \right), \\
G_{z,\ell,y}(t, r) &= \mathbb{1}_{\ell=1} \mathbb{E} \left[e^{-\gamma(r_{z,\zeta,\Xi} - r_{z,\ell,y} - \eta)} \right] + \mathbb{1}_{\ell > 1} e^{-\gamma(r_{z,\ell-1,\Xi} - r_{z,\ell,y} - \eta)} - 1.
\end{aligned}$$

for $i = (z, \ell, y)$ with $y \neq \Xi$. We define $G_{z,N^L,\Xi}(t, r) = \infty$ so that the QVI (34) reduces to $F_{z,N^L,\Xi}(t, u, \partial_t u_i) = 0$.

Our first result is the following comparison principle for (34), which ensures the uniqueness of the solution.

Theorem 1. *Let u be a bounded upper-semi-continuous subsolution and v a bounded lower-semi-continuous supersolution of (34), then $u \leq v$.*

PROOF. See Appendix A.1.

For $(z, \ell, \Xi) \in \bar{\mathbb{S}}$ with $\ell < N^L$ and $\phi \in C^1([0, T])$ we provide the following finite difference scheme:

$$\begin{aligned}
\mathbf{F}_{z,\ell,\Xi}^\Delta(t, r, \phi) &= \gamma \frac{\phi(t) - r_{z,\ell,\Xi}}{\Delta} + \sum_{\bar{z} \neq z} \lambda^{z,\bar{z},\ell} \left(e^{-\gamma(r_{\bar{z},\ell,\Xi} - r_{z,\ell,\Xi})} - 1 \right) \\
&\quad + \lambda^{m,z,\ell} \left(\mathbb{E} \left[\mathbb{1}_{\epsilon < \ell} e^{-\gamma(r_{z,\ell-\epsilon,\Xi} - r_{z,\ell,\Xi})} + \mathbb{1}_{\epsilon \geq \ell} e^{-\gamma(r_{z,\zeta,\Xi} - r_{z,\ell,\Xi})} \right] - 1 \right) \\
&\quad + \lambda^{c,z,\ell} \left(\mathbb{1}_{\ell > 1} e^{-\gamma(r_{z,\ell-1,\Xi} - r_{z,\ell,\Xi})} + \mathbb{1}_{\ell=1} \mathbb{E} \left[e^{-\gamma(r_{z,\zeta,\Xi} - r_{z,\ell,\Xi})} \right] - 1 \right) \\
&\quad + \lambda^{a,z,\ell} \left(e^{-\gamma(r_{z,\ell+1,\Xi} - r_{z,\ell,\Xi})} - 1 \right), \\
\mathbf{G}_{z,\ell,\Xi}(t, r, \phi) &= e^{-\gamma(r_{z,\ell+1,\ell+1} - \phi(t) - \eta)} - 1.
\end{aligned}$$

The obvious analogous expressions are used for all cases in $\bar{\mathbb{S}}$ that remain.

Here $\Delta = T/N$ is the size of time steps where N is a positive integer. We denote by $\mathbf{T}^\Delta = \{t_j := \Delta j : 0 \leq j \leq N\}$ the set of grid points for the finite difference scheme. Consider the following discrete problem:

$$\mathbf{J}_i^\Delta(t_j, u^\Delta(t_{j+1}), u_i^\Delta) = \min \{ \mathbf{F}_i^\Delta(t_j, u^\Delta(t_{j+1}), u_i^\Delta); \mathbf{G}_i(t_j, u^\Delta(t_{j+1}), u_i^\Delta) \} = 0, \quad (35a)$$

$$u_i^\Delta(T) = 0, \quad (35b)$$

for $i \in \bar{\mathbb{S}}$, $1 \leq j \leq N-1$ and $u^\Delta : [0, T] \rightarrow \mathbb{R}^{\mathcal{S}}$. We will show that the solution to (35) can be used to approximate the solution to (34). To proceed we first need the following results.

Proposition 2. (*Stability*)

Let Δ be sufficiently small and u^Δ be the solution to (35). For any $i = (z, \ell, y) \in \bar{\mathbb{S}}$, we have

$$u_i(t_j) \leq A(T - t_j), \quad (36)$$

$$u_i(t_j) \geq 0, \quad y = \Xi, \quad (37)$$

$$u_i(t_j) \geq -\eta, \quad y \neq \Xi, \quad (38)$$

where $A > 0$ is a constant.

PROOF. See Appendix A.2.

Proposition 3. (*Monotonicity*)

Consider functions $u : [0, T] \rightarrow \mathbb{R}^{\bar{\mathbb{S}}}$ and $v : [0, T] \rightarrow \mathbb{R}^{\bar{\mathbb{S}}}$, satisfying $C^{\min} \leq u \leq v \leq C^{\max}$, for some constants $C^{\min}$ and $C^{\max}$. Suppose for some $i \in \bar{\mathbb{S}}$, we have $u_i \leq v_i$ and $u_j = v_j$ for $j \neq i$; and for some $t \in [0, T]$, we have $u_i(t) = v_i(t)$. Then

$$\mathbf{J}_i^\Delta(t, u(t'), u_i) - \mathbf{J}_i^\Delta(t, v(t'), v_i) \geq 0 \quad (39)$$

for all $t' \in [0, T]$ and Δ sufficiently small.

PROOF. See Appendix A.3

Proposition 4. (*Consistency*)

Fix $i \in \bar{\mathbb{S}}$. For $\delta > 0$, $t \in [0, T]$ and any family of continuous functions $\Phi \in C^0([0, T], \mathbb{R})$ with continuously differentiable i -th component Φ_i , we have

$$|\mathbf{J}_i^\Delta(t, \Phi(t + \Delta), \Phi_i) - J_i(t, \Phi(t), \Phi_i(t))| \leq o(\Delta), \quad (40)$$

where $o(\Delta) \rightarrow 0$ as $\Delta \rightarrow 0$.

PROOF. Immediately follows from the fact that Φ_i is continuous and

$$\frac{\Phi_i(t) - \Phi(t + \Delta)}{\Delta} \rightarrow \partial_t \Phi_i(t),$$

as $\Delta \rightarrow 0$. □

With the above results, we are now ready to prove the convergence of our numerical scheme.

Proposition 5. Let u^Δ be the solution to (35) and let u be the solution to (34). Then $u^\Delta \rightarrow u$ as $\Delta \rightarrow 0$.

PROOF. We define the following limits.

$$\underline{u}(t) = \liminf_{\Delta \rightarrow 0, t' \rightarrow t} u^\Delta(t'), \quad \bar{u}(t) = \limsup_{\Delta \rightarrow 0, t' \rightarrow t} u^\Delta(t').$$

Proposition 2 ensures the above limits are well-defined. Following Barles and Souganidis (1991) and Briani et al. (2012), it suffices to show that $\underline{u}$ is a supersolution and $\bar{u}$ is a subsolution to (34). Then by comparison principle (Theorem 1), $\bar{u} \leq \underline{u}$. The opposite inequality is obviously true and we have $\bar{u} = \underline{u} = \lim_{\Delta \rightarrow 0} u^\Delta$.

We will prove that $\bar{u}$ is a subsolution (the proof for $\underline{u}$ being a supersolution is similar). Let ϕ be a smooth function and $i \in \bar{\mathbb{S}}$ such that $\bar{u}_i - \phi$ has a strict global maximum at t_0 and $\bar{u}_i(t_0) = \phi(t_0)$.

Standard arguments imply that there exists sequences $\Delta_n \rightarrow 0$ and $t_n \rightarrow 0$ such that $u_i^{\Delta_n} - \phi$ has a local maximum at t_n and $u^{\Delta_n}(t_n) \rightarrow \bar{u}_i(t_0)$ as $n \rightarrow \infty$.

Set $\delta_n := (u_i^{\Delta_n} - \phi)(t_n)$ and define the function $\Phi^n(t) = \{\Phi_j^n(t)\}_{j \in \bar{\mathbb{S}}}$

$$\Phi_j^n(t) = \begin{cases} \phi(t) + \delta_n & j = i, \\ u_j^{\Delta_n}(t) & j \neq i. \end{cases}$$

We have

$$\begin{aligned} 0 &= \mathbf{J}_i^{\Delta_n}(t_n, u^{\Delta_n}(t_n + \Delta_n), u_i^{\Delta_n}) \\ &\geq \mathbf{J}_i^{\Delta_n}(t_n, \Phi^n(t_n + \Delta_n), \Phi_i^n) \\ &= J_i(t_n, \phi(t_n + \Delta_n) + \delta_n, \partial_t \phi(t_n)) + o(\Delta_n), \end{aligned}$$

where the second line follows from Proposition 3 and the third line follows from Proposition 4. Sending $n \rightarrow \infty$, we have the required result. $\square$

7. Conclusion

Using tick level data from the Nasdaq exchange, we conduct an empirical study on various limit order book events, including dynamics of volume imbalance, gain/loss of a filled limit order, rate of addition and cancellation, rate of market order arrival, distribution of market order size and distribution of replenished queue length. We propose novel ways of modeling cancellation position of limit orders and distribution of market order size. Based on our empirical findings, we develop a queuing model that captures stylized facts on limit order book data. One important feature of our model is that the dynamics of the limit order book depends on the regime, which is a function of volume imbalance. Moreover, an executed limit order is more likely to be considered a gain (loss) by the agent when the limit order book is at a gainful (adverse) regime.

We show how to apply our proposed queuing model in the context of algorithmic trading. We consider the problem of an agent maximizes her expected utility by placing and canceling limit orders in such queuing model. In our set up, the agent tries to let more of her limit orders executed at a gainful regime and less of her limit orders executed at an adverse regime. We demonstrate our result using a numerical example, in which parameters are calibrated from real data. Our result shows that even at an adverse regime, the agent might still be willing to stay in the queue

if she already has an active limit order, or place a new limit order if she does not have one. This is because the agent tries to obtain a good queue position when the regime switches to a gainful one. It also implies that strategies that only look at regimes are sub-optimal. Simulation study shows that the optimal strategy achieve a 2.4% higher mean terminal wealth than the benchmark strategy that only looks at regimes, given the same level of standard deviation of terminal wealth; or 11.4% lower standard deviation given the same level of mean terminal wealth.

Appendix A. Proofs

Appendix A.1. Proof of Theorem 1

PROOF. Consider $\tilde{u} = (1 + \delta)u - \delta A(T - t)$, where $0 < \delta < 1$ and $A > 0$ is a constant to be fixed later. We first show that $\tilde{u}$ is a subsolution to

$$J_i(t, \tilde{u}, \partial_t \tilde{u}_i) \leq -\delta C \quad (\text{A.1})$$

for $i \in \bar{\mathbb{S}}$ and some constant $C > 0$, independent of the choice of δ .

We will show (A.1) for the case $i = (z, \ell, \Xi)$. Showing the case $i = (z, \ell, y)$ with $y \neq \Xi$ is similar. Let $\phi(t)$ be a smooth test function and $t_0 < T$ be such that $\tilde{u}_i(t) - \phi(t)$ attains maximum at t_0 . The subsolution property of u implies either of the following

$$F_i\left(t_0, u(t_0), \frac{\partial_t \phi(t_0) - A\delta}{1 + \delta}\right) \leq 0, \quad (\text{A.2})$$

$$G_i(t_0, u(t_0)) \leq 0. \quad (\text{A.3})$$

Suppose (A.2) holds. Without loss of generality, we assume that $\lambda^{z, \bar{z}, \ell} = \lambda^{m, z, \ell} = \lambda^{c, z, \ell} = 0$. The case where these are strictly positive can be treated similarly. Multiplying (A.2) by $1 + \delta$ and re-arranging the terms, we have

$$-\gamma \partial_t \phi(t_0) + \lambda^{a, z, \ell} \left(e^{-\gamma(u_{z, \ell+1, \Xi}(t_0) - u_{z, \ell, \Xi}(t_0))} - 1 \right) \leq -\lambda^{a, z, \ell} \left(e^{-\gamma(u_{z, \ell+1, \Xi}(t_0) - u_{z, \ell, \Xi}(t_0))} - 1 \right) \delta - \gamma A \delta.$$

Using the fact that u is bounded, we can conclude that there exists a constant C_1 such that

$$-\gamma \partial_t \phi(t_0) + \lambda^{a,z,\ell} (e^{-\gamma(u_{z,\ell+1,\Xi}(t_0) - u_{z,\ell,\Xi}(t_0))} - 1) \leq (C_1 - \gamma A) \delta. \quad (\text{A.4})$$

From (A.4) we have

$$\begin{aligned} F_{z,\ell,\Xi}(t_0, \tilde{u}(t_0), \partial_t \phi(t_0)) &\leq (C_1 - \gamma A) \delta + \lambda^{a,z,\ell} e^{-\gamma(\tilde{u}_{z,\ell+1,\Xi}(t_0) - \tilde{u}_{z,\ell,\Xi}(t_0))} (1 - e^{\gamma \delta(u_{z,\ell+1,\Xi}(t_0) - u_{z,\ell,\Xi}(t_0))}) \\ &\leq (C_1 - \gamma A) \delta - \lambda^{a,z,\ell} e^{-\gamma(1+\delta)(u_{z,\ell+1,\Xi}(t_0) - u_{z,\ell,\Xi}(t_0))} (u_{z,\ell+1,\Xi}(t_0) - u_{z,\ell,\Xi}(t_0)) \gamma \delta, \\ &\leq (C_2 - \gamma A) \delta, \end{aligned} \quad (\text{A.5})$$

for some constant C_2 . In the second inequality we use the fact $e^x \geq x + 1$ for all x . In the third inequality we use the fact that u and δ are bounded. Choosing $A > \frac{C_2}{\gamma}$, we have

$$F_{z,\ell,\Xi}(t_0, \tilde{u}(t_0), \partial_t \phi(t_0)) \leq -C_3 \delta, \quad (\text{A.6})$$

for some constant $C_3 > 0$.

For the case where $\lambda^{z,\bar{z},\ell}$, $\lambda^{m,z,\ell}$ and $\lambda^{c,z,\ell}$ are nonzero, the same argument applies: on the right-hand-side of inequalities (A.4) and (A.5), we would have different constants C_1 and C_2 . The rest of the proof is the same.

If (A.3) holds we have $u_{z,\ell+1,\ell+1}(t_0) - u_{z,\ell,\Xi}(t_0) \geq \eta$ and

$$\begin{aligned} G_{z,\ell,\Xi}(t_0, \tilde{u}(t_0)) &= e^{-\gamma(u_{z,\ell+1,\ell+1}(t_0) - u_{z,\ell,\Xi}(t_0) - \eta)} e^{-\gamma \delta(u_{z,\ell+1,\ell+1}(t_0) - u_{z,\ell,\Xi}(t_0))} - 1 \\ &= G_{z,\ell,\Xi}(t_0, u(t_0)) e^{-\gamma \delta(u_{z,\ell+1,\ell+1}(t_0) - u_{z,\ell,\Xi}(t_0))} + e^{-\gamma \delta(u_{z,\ell+1,\ell+1}(t_0) - u_{z,\ell,\Xi}(t_0))} - 1 \\ &\leq e^{-\gamma \delta(u_{z,\ell+1,\ell+1}(t_0) - u_{z,\ell,\Xi}(t_0))} - 1 \\ &\leq -\delta e^{-\gamma \eta} \gamma \eta. \end{aligned} \quad (\text{A.7})$$

Since either (A.6) or (A.7) holds, by choosing $C = \min\{C_3, \gamma \eta e^{-\gamma \eta}\}$, we have shown that $\tilde{u}$ is a subsolution to (A.1).

In order to prove $u \leq v$, it suffices to show

$$\max_{i \in \bar{\mathbb{S}}} \sup(\tilde{u}_i - v_i) \leq 0, \quad (\text{A.8})$$

for all $\delta > 0$. The required result can be obtained by sending δ to 0.

Let us show (A.8) by contradiction. Suppose there exists some $\delta > 0$ and $i^* \in \bar{\mathbb{S}}$ such that

$$\mathcal{M} := \max_{i \in \bar{\mathbb{S}}} \sup_t (\tilde{u}_i - v_i^0) = \sup_t \{\tilde{u}_{i^*}(t) - v_{i^*}(t)\} > 0.$$

Clearly, $\tilde{u}_{i^*} - v_{i^*}$ attains its maximum $\mathcal{M}$ at some $\bar{t} \in [0, T)$, and we can assume without loss of generality that $\bar{t} > 0$. Consider the following class of test functions

$$\Phi_\alpha(t, s) = \tilde{u}_{i^*}(t) - v_{i^*}(s) - \alpha|t - s|^2,$$

with $\alpha > 0$. By standard argument, we have

$$\begin{aligned} t_\alpha, s_\alpha &\rightarrow \bar{t}, \\ \mathcal{M}_\alpha = \sup \Phi_\alpha &= \Phi_\alpha(t_\alpha, s_\alpha) \rightarrow \mathcal{M}, \end{aligned}$$

as $\alpha \rightarrow \infty$, where (t_α, s_α) is the maximizer of Φ_α .

Without loss of generality, we assume that $i^* = (z^*, \ell^*, \Xi)$. The case for $i^* = (z^*, \ell^*, y^*)$ with $y^* \neq \Xi$ can be shown using a similar argument. By the viscosity solution property of $\tilde{u}$ and v , we have

$$\begin{aligned} \min \{F_{z^*, \ell^*, \Xi}(t_\alpha, \tilde{u}(t_\alpha), -2\alpha(t_\alpha - s_\alpha)); G_{z^*, \ell^*, \Xi}(t_\alpha, \tilde{u}(t_\alpha))\} &\leq -C\delta, \\ \min \{F_{z^*, \ell^*, \Xi}(s_\alpha, v(s_\alpha), -2\alpha(t_\alpha - s_\alpha)); G_{z^*, \ell^*, \Xi}(s_\alpha, v(s_\alpha))\} &\geq 0. \end{aligned}$$

The above two inequalities imply either of the following

$$F_{z^*, \ell^*, \Xi}(t_\alpha, \tilde{u}(t_\alpha), -2\alpha(t_\alpha - s_\alpha)) - F_{z^*, \ell^*, \Xi}(s_\alpha, v(s_\alpha), -2\alpha(t_\alpha - s_\alpha)) \leq -\delta C, \quad (\text{A.9})$$

$$G_{z^*, \ell^*, \Xi}(t_\alpha, \tilde{u}(t_\alpha)) - G_{z^*, \ell^*, \Xi}(s_\alpha, v(s_\alpha)) \leq -\delta C. \quad (\text{A.10})$$

Note that for any (z, ℓ, Ξ) , we have

$$\begin{aligned}
& \liminf_{\alpha \rightarrow \infty} e^{-\gamma(\tilde{u}_{z,\ell,\Xi}(t_\alpha) - \tilde{u}_{z^*,\ell^*,\Xi}(t_\alpha))} - e^{-\gamma(v_{z,\ell,\Xi}(s_\alpha) - v_{z^*,\ell^*,\Xi}(s_\alpha))} \\
&= \liminf_{\alpha \rightarrow \infty} \left(e^{-\gamma(\tilde{u}_{z,\ell,\Xi}(t_\alpha) - v_{z,\ell,\Xi}(s_\alpha))} - e^{-\gamma(\tilde{u}_{z^*,\ell^*,\Xi}(t_\alpha) - v_{z^*,\ell^*,\Xi}(s_\alpha))} \right) e^{-\gamma(v_{z,\ell,\Xi}(s_\alpha) - \tilde{u}_{z^*,\ell^*,\Xi}(t_\alpha))} \\
&\geq \liminf_{\alpha \rightarrow \infty} \left(e^{-\gamma(\tilde{u}_{z,\ell,\Xi}(t_\alpha) - v_{z,\ell,\Xi}(s_\alpha))} - e^{-\gamma(\tilde{u}_{z^*,\ell^*,\Xi}(t_\alpha) - v_{z^*,\ell^*,\Xi}(s_\alpha))} \right) C_4 \\
&\geq \left(e^{-\gamma(\tilde{u}_{z,\ell,\Xi}(\bar{t}) - v_{z,\ell,\Xi}(\bar{t}))} - e^{-\gamma(\tilde{u}_{z^*,\ell^*,\Xi}(\bar{t}) - v_{z^*,\ell^*,\Xi}(\bar{t}))} \right) C_4 \\
&\geq 0,
\end{aligned}$$

where $C_4 > 0$ is a constant. (A similar result hold for (z, ℓ, y) with $y \neq \Xi$.) Using the above inequality, we can conclude that as $\alpha \rightarrow \infty$, the $\liminf$ of the l.h.s of (A.9) and the $\liminf$ of the l.h.s. of (A.10) are both nonnegative, a contradiction. $\square$

Appendix A.2. Proof of Proposition 2

PROOF. We prove by induction backwards in t_j . Without loss of generality, we assume that $\lambda^{z,\bar{z},\ell} = \lambda^{m,z,\ell} = \lambda^{c,z,\ell} = 0$. Equation (35a) can be re-written as

$$u_i(t_j) = \max \left\{ u_i(t_{j+1}) - \frac{\Delta \lambda^a}{\gamma} (e^{-\gamma(u_k(t_{j+1}) - u_i(t_{j+1}))} - 1) ; \quad u_m(t_{j+1}) - \eta \right\}, \quad (\text{A.11})$$

for $i, k, m \in \bar{\mathbb{S}}$ and $\lambda^a = \lambda^{a,z,\ell}$ for some z and ℓ . Clearly (36), (37) and (38) holds when $t_j = T$. Suppose they also hold for t_{j+1} where $j \leq N - 1$. Using the fact that $e^y > 0$ for all y and $u(t_{j+1}) \leq A(T - t_{j+1})$, we have

$$u_i(t_j) \leq A(T - t_j - \Delta) + \frac{\Delta \lambda^a}{\gamma} \leq A(T - t_j),$$

for $A > \lambda^a/\gamma$ and we have shown (36).

On the other hand, for $i = (z, \ell, \Xi)$, (A.11) implies

$$\begin{aligned}
u_i(t_j) &\geq u_i(t_{j+1}) - \frac{\Delta \lambda^a}{\gamma} (e^{-\gamma(u_k(t_{j+1}) - u_i(t_{j+1}))} - 1) \\
&= u_i(t_{j+1}) + \frac{\Delta \lambda^a}{\gamma} e^{-\gamma(u_k(t_{j+1}) - u_i(t_{j+1}))} (e^{\gamma(u_k(t_{j+1}) - u_i(t_{j+1}))} - 1) \\
&\geq u_i(t_{j+1}) + \Delta \lambda^a e^{-\gamma(u_k(t_{j+1}) - u_i(t_{j+1}))} (u_k(t_{j+1}) - u_i(t_{j+1})) \\
&\geq (1 - \Delta \lambda^a e^{\gamma AT}) u_i(t_{j+1}) \\
&\geq 0,
\end{aligned}$$

for $\Delta < e^{-\gamma AT}/\lambda^a$. In the forth line (third inequality) we use the fact that $0 \leq u_i(t_{j+1}) \leq A(T - t_{j+1}) \leq AT$ and $u_k(t_{j+1}) \geq 0$. For $i = (z, \ell, y)$ with $y \neq \Xi$, (A.11) implies

$$\begin{aligned}
u_{z,\ell,y}(t_j) &\geq u_{z,\ell-1,\Xi}(t_{j+1}) - \eta \\
&\geq -\eta.
\end{aligned}$$

□

Appendix A.3. Proof of Proposition 3

PROOF. It suffices to show

$$\mathbf{F}_i^\Delta(t, u(t'), u_i) - \mathbf{F}_i^\Delta(t, v(t'), v_i) \geq 0, \quad (\text{A.12})$$

$$\mathbf{G}_i^\Delta(t, u(t'), u_i) - \mathbf{G}_i^\Delta(t, v(t'), v_i) \geq 0. \quad (\text{A.13})$$

The inequality in (A.13) immediately follows from $u_i \leq v_i$. To show (A.12), we assume with out loss of generality that $i = (z, \ell, \Xi)$ and $\lambda^{z,\bar{z},\ell} = \lambda^{m,z,\ell} = \lambda^{c,z,\ell} = 0$. Let $w = u - v$ and $i^* = (z, \ell + 1, \Xi)$.

Note that $w_i(t) = 0$, $w_i(t') \leq 0$ and $u_{i*} = v_{i*}$. The left-hand-side of (A.12) is

$$\begin{aligned}
& \gamma \frac{w_i(t)}{\Delta} - \gamma \frac{w_i(t')}{\Delta} + \lambda^{a,z,\ell} e^{-\gamma(u_{i*}(t') - v_i(t'))} (e^{\gamma w_i(t')} - 1) \\
& \geq -\gamma \frac{w_i(t')}{\Delta} + \lambda^{a,z,\ell} e^{-\gamma(u_{i*}(t') - v_i(t'))} \gamma w_i(t') \\
& \geq \left(-\frac{1}{\Delta} + \lambda^{a,z,\ell} e^{-\gamma(C^{\min} - C^{\max})} \right) \gamma w_i(t') \\
& \geq 0,
\end{aligned}$$

for Δ sufficiently small. □

Appendix B. Estimated Parameters

	$z = 1$	$z = 2$	$z = 3$
$\bar{z} = 1$		7.769 (0.145)	17.962 (0.891)
$\bar{z} = 2$	-14.79 (0.193)		7.939 (0.14)
$\bar{z} = 3$	-19.614 (1.095)	-8.6 (0.15)	

Table B.6: Estimated $\beta_{0,z,\bar{z}}^Z$ in (2). Numbers in brackets represent standard errors.

	$z = 1$	$z = 2$	$z = 3$
$\bar{z} = 1$		-2.024 (0.037)	-6.621 (0.336)
$\bar{z} = 2$	3.932 (0.049)		-2.718 (0.046)
$\bar{z} = 3$	5.033 (0.292)	2.931 (0.048)	

Table B.7: Estimated $\beta_{1,z,\bar{z}}^Z$ in (2).

	$\beta_{0,z}^c$	$\beta_{1,z}^c$	$\beta_{0,z}^a$	$\beta_{1,z}^a$
$z = 1$	0.273 (0.005)	0.316×10^{-3} (0.043×10^{-5})	0.695 (0.005)	0.293×10^{-3} (0.004×10^{-4})
$z = 2$	0.54 (0.014)	0.411×10^{-3} (0.288×10^{-5})	1.945 (0.014)	0.092×10^{-3} (0.036×10^{-4})
$z = 3$	0.23 (0.023)	0.52×10^{-3} (1.13×10^{-5})	2.932 (0.018)	-0.714×10^{-3} (0.135×10^{-4})

Table B.8: Estimated rate of addition and cancellation.

$z = 1$	$z = 2$	$z = 3$
0.044 (3.92×10^{-4})	0.175 (0.003)	0.596 (0.012)

Table B.9: Estimated rate of MO arrival.

	$z = 1$	$z = 2$	$z = 3$
r	0.477 (0.002)	0.630 (0.002)	0.866 (0.003)
p	3.339×10^{-4} (1.961×10^{-6})	5.408×10^{-4} (2.968×10^{-6})	9.097×10^{-4} (4.697×10^{-6})

Table B.10: Estimated parameter of MO size distribution (negative binomial) with probability mass function $p(k) = \binom{k+r-1}{k}(1-p)^r p^k$.

	$z = 1$	$z = 2$	$z = 3$
r	6.031 (0.294)	7.530 (0.513)	3.618 (0.452)
p	8.441×10^{-4} (4.285×10^{-5})	1.9×10^{-3} (1.345×10^{-4})	2.3×10^{-3} (3.041×10^{-4})

Table B.11: Estimated parameter of replenished queue length distribution (negative binomial) with probability mass function $p(k) = \binom{k+r-1}{k}(1-p)^r p^k$.

References

- Alfonsi, A., Fruth, A., and Schied, A. (2010). Optimal execution strategies in limit order books with general shape functions. Quantitative Finance, 10(2):143–157.
- Almgren, R. and Chriss, N. (2001). Optimal execution of portfolio transactions. Journal of Risk, 3:5–40.
- Avellaneda, M. and Stoikov, S. (2008). High-frequency trading in a limit order book. Quantitative Finance, 8(3):217–224.
- Baccarin, S. and Sanfelici, S. (2006). Optimal impulse control on an unbounded domain with nonlinear cost functions. Computational Management Science, 3(1):81–100.
- Bacry, E., Jaisson, T., and Muzy, J.-F. (2016). Estimation of slowly decreasing hawkes kernels: application to high-frequency order book dynamics. Quantitative Finance, 16(8):1179–1201.
- Barles, G. and Souganidis, P. E. (1991). Convergence of approximation schemes for fully nonlinear second order equations. Asymptotic analysis, 4(3):271–283.
- Bayraktar, E. and Ludkovski, M. (2011). Optimal trade execution in illiquid markets. Mathematical Finance, 21(4):681–701.
- Bechler, K. and Ludkovski, M. (2015). Optimal execution with dynamic order flow imbalance. SIAM Journal on Financial Mathematics, 6(1):1123–1151.
- Bensoussan, A., Lions, J.-L., and Lions, J.-L. (1982). Contrôle impulsionnel et inéquations quasi variationnelles, volume 1. Dunod Paris.
- Briani, A., Camilli, F., and Zidani, H. (2012). Approximation schemes for monotone systems of nonlinear second order partial differential equations: convergence result and error estimate. Differential Equations and Applications, 4:297–317.

- Cartea, Á., Donnelly, R. F., and Jaimungal, S. (2015a). Enhancing trading strategies with order book signals. Available at SSRN 2668277.
- Cartea, Á. and Jaimungal, S. (2015). Risk metrics and fine tuning of high frequency trading strategies. Mathematical Finance, 25(3):576–611.
- Cartea, Á., Jaimungal, S., and Penalva, J. (2015b). Algorithmic and high-frequency trading. Cambridge University Press.
- Cont, R. and De Larrard, A. (2013). Price dynamics in a markovian limit order market. SIAM Journal on Financial Mathematics, 4(1):1–25.
- Cont, R., Kukanov, A., and Stoikov, S. (2014). The price impact of order book events. Journal of financial econometrics, 12(1):47–88.
- Cont, R., Stoikov, S., and Talreja, R. (2010). A stochastic model for order book dynamics. Operations research, 58(3):549–563.
- Gatheral, J., Schied, A., and Slynko, A. (2012). Transient linear price impact and fredholm integral equations. Mathematical Finance, 22(3):445–474.
- Guéant, O., Lehalle, C.-A., and Fernandez-Tapia, J. (2012). Optimal portfolio liquidation with limit orders. SIAM Journal on Financial Mathematics, 3(1):740–764.
- Guilbaud, F. and Pham, H. (2013). Optimal high-frequency trading with limit and market orders. Quantitative Finance, 13(1):79–94.
- Guo, X., Ruan, Z., and Zhu, L. (2015). Dynamics of order positions and related queues in a limit order book. Available at SSRN 2607702.
- Huang, W., Lehalle, C.-A., and Rosenbaum, M. (2015). Simulating and analyzing order book data: The queue-reactive model. Journal of the American Statistical Association, 110(509):107–122.
- Jaisson, T., Rosenbaum, M., et al. (2015). Limit theorems for nearly unstable hawkes processes. The Annals of Applied Probability, 25(2):600–631.
- Kaplan, E. L. and Meier, P. (1958). Nonparametric estimation from incomplete observations. Journal of the American statistical association, 53(282):457–481.
- Lakner, P., Reed, J., and Stoikov, S. (2013). High frequency asymptotics for the limit order book. Preprint.
- Lehalle, C.-A. and Mounjid, O. (2016). Limit order strategic placement with adverse selection risk and the role of latency. arXiv preprint arXiv:1610.00261.
- Ly Vath, V. and Pham, H. (2007). Explicit solution to an optimal switching problem in the two-regime case. SIAM Journal on Control and Optimization, 46(2):395–426.

- Maglaras, C., Moallemi, C. C., and Zheng, H. (2014). Queueing dynamics and state space collapse in fragmented limit order book markets. Columbia Business School Research Paper, (14-13).
- Maglaras, C., Moallemi, C. C., and Zheng, H. (2015). Optimal execution in a limit order book and an associated microstructure market impact model. Columbia Business School Research Paper, (15-60).