# Buy Low Sell High: a High Frequency Trading Perspective

Álvaro Cartea[a], Sebastian Jaimungal[b], Jason Ricci[b]

[a]*Department of Mathematics, University College London, UK*

[b]*Department of Statistics, University of Toronto, Canada*

---

**Abstract**

We develop a High Frequency (HF) trading strategy where the HF trader uses her superior speed to process information and to post limit sell and buy orders. By introducing a multi-factor self-exciting process we allow for feedback effects in market buy and sell orders and the shape of the limit order book (LOB). Our model accounts for arrival of market orders that influence activity, trigger one-sided and two-sided clustering of trades, and induce temporary changes in the shape of the LOB. We also model the impact that market orders and news have on the short-term drift of the midprice (short-term-alpha). We show that HF traders who do not include predictors of short-term-alpha in their strategies are driven out of the market because they are adversely selected by better informed traders and because they are not able to profit from directional strategies.

*Keywords:* Algorithmic Trading, High Frequency Trading, Short Term Alpha, Adverse Selection, Self-Exciting Processes, Hawkes processes

---

## 1. Introduction

Most of the traditional stock exchanges have converted from open outcry communications between human traders to electronic markets where the activity between participants is handled by computers. In addition to those who have made the conversion, such as the New York Stock Exchange and the London Stock Exchange, new electronic trading platforms have entered the market, for example NASDAQ in the US and Chi-X in Europe. Along with the exchanges, market participants have been increasingly relying on the use of computers to handle their trading needs. Initially, computers were employed to execute trades, but nowadays computers manage inventories and make trading decisions; this modern way of trading in the electronic markets is known as Algorithmic Trading (AT).

Despite the substantial changes that markets have undergone in the recent past, some strategies

*Email addresses:* `alvaro.cartea@uc3m.es` (Álvaro Cartea), `sebastian.jaimungal@utoronto.ca` (Sebastian Jaimungal), `jason.ricci@utoronto.ca` (Jason Ricci)

used by investors remain the same. When asked about how to make money in the stock market, an old adage responds: "Buy low and sell high". Although in principle this sounds like a good strategy, its success relies on spotting opportunities to buy and sell at the right time. Surprisingly, more than ever, due to the incredible growth in computing power, a great deal of the activity in the US and Europe's stock exchanges is based on trying to profit from short-term price predictions by buying low and selling high. The effectiveness of these computerized short-term strategies, a subset of AT known as High Frequency (HF) trading, depends on the ability to process information and send messages to the electronic markets in microseconds, see Cartea and Penalva (2011). In this paper we develop an HF strategy that profits from its superior speed advantage to decide when and how to enter and exit the market over extremely short time intervals.

The increase in computer power has made it easier for market participants to deploy ever more complicated trading strategies to profit from changes in market conditions. Key to the success of HF strategies is the speed at which agents can process information and news events to take trading decisions. A unique characteristic to HF trading is that the strategies are designed to hold close to no inventories over very short periods of time (seconds, minutes, or at most one day) to avoid both exposure to markets after close and to post collateral overnight. Thus, profits are made by turning over positions very quickly to make a very small margin per roundtrip transaction (buy followed by a sell or vice-versa), but repeating it as many times as possible during each trading day.

In the past, markets were quote driven which means that market makers quoted buy and sell prices and investors would trade with them. Nowadays, there are *limit order* markets where all participants can post limit buy or sell orders; i.e. behave as market makers in the old quote driven market. The limit orders show an intention to buy or sell and must indicate the amount of shares and price at which the agent is willing to trade. The limit buy (sell) order with the lowest (highest) price tag is known as the *best bid* (*best offer*). During the trading day, all orders are accumulated in the limit order book (LOB) until they find a counterparty for execution or are canceled by the agent who posted them. The counterparty is a *market order* which is an order to buy or sell an amount of shares, regardless of the price, which is immediately executed against limit orders resting in the LOB at the best execution prices.

As expected, changes in the way trading is conducted in modern electronic markets, coupled with the advent of AT in general and HF trading in particular, is reflected both in the changes of price distributions (see Cvitanic and Kirilenko (2010)) and in the dramatic changes observed in LOB activity. In any one day it is possible to observe up to hundreds of thousands of messages being submitted to the LOB of a single stock.

There is little evidence on the source of HF market making profits, but the picture that is emerging is that price anticipation and short-term price deviations from the fundamental value of the asset are important drivers of profits. On the other hand, we also know that strategies that do not include in their limit orders a buffer to cover adverse selection costs, or that strategically post deeper in the book to avoid being picked off, may see their accumulated profits dwindle as a consequence of trading with other market participants that possess private or better information. In the long term, HF traders (HFTs) who are not able to incorporate short-term price predictability in their optimal

HF market making strategies, as well as account for adverse selection costs, are very likely to be driven out of the market.

The goal of this paper is to develop a particular dynamic HF trading strategy based on optimal postings and cancelations of limit orders to maximize expected terminal wealth over a fixed horizon $T$ whilst penalizing inventories. The HFT we characterize here can be thought of as an ultra-fast market maker where the trading horizon $T$ is at most one trading day, all the limit orders are canceled an instant later if not filled, and inventories are optimally managed and drawn to zero by $T$.[1] Early work on optimal postings by a securities dealer is that of Ho and Stoll (1981) and more recently Avellaneda and Stoikov (2008) study the optimal HF submission strategies of bid and ask limit orders.

Intuitively, the HF dynamic strategy we find maximizes the expected profits resulting from roundtrip trades by specifying how deep on the sell and buy side the limit orders are placed in the LOB. The HF strategy is based on predictable short-term price deviations and managing adverse selection risks that result from trading with counterparties that may possess private or better information. Clearly, the closer the limit orders are to the best bid and best offer, the higher the probability of being executed, but the expected profits from a roundtrip are also lower and adverse selection costs higher.

Accumulated inventories play a key role throughout the entire strategy we develop. Optimal postings control for inventory risks by sending quotes to the LOB which induce mean reversion of inventories to an optimal level and by including a state dependent buffer to cover or avoid expected adverse selection costs. For example, if the probability of the next market order being a buy or sell is the same, and inventories are positive, then the limit sell orders are posted closer to the best ask and the buy orders are posted further away from the best bid so that the probability of the offer being lifted is higher than the bid being hit. Furthermore, as the dynamic trading strategy approaches the terminal date $T$, orders are posted nearer the midquote to induce mean-reversion to zero in inventories which avoids having to post collateral overnight and bearing inventory risks until the market opens the following day. Similarly, if the HF trading algorithm detects that limit orders on one side of the LOB are more likely to be adversely selected, then these limit orders are posted deeper into the book in anticipation of the expected adverse selection costs. An increase in adverse selection risk could be heralded by market orders becoming more one-sided as a consequence of the activity of traders acting on superior or private information who are sending one-directional market orders.

Trade initiation may be motivated by many reasons which have been extensively studied in the literature, see for example Sarkar and Schwartz (2009). Some of these include: asymmetric information, differences in opinion or differential information, and increased proportion of impatient

[1] HFTs closely monitor their exposure to inventories for many reasons. For example, HFTs' own risk controls or regulation do not allow them to build large (long or short) positions; the HFT is capital constrained and needs to post collateral against her inventory position. Moreover, we remark that there is no consensus on characterizing HFTs as market makers because some stakeholders and regulatory authorities point out that their holding periods are too short to consider them as such.

(relative to patient) traders. Likewise, trade clustering can be the result of various market events, see Cartea and Jaimungal (2010). For instance, increases in market activity could be due to shocks to the fundamental value of the asset, or the release of public or private information that generates an increase in trading (two-sided or one-sided) until all information is impounded in stock prices. However, judging by the sharp rise of AT in the recent years and the explosion in volume of submissions and order cancelations it is also plausible to expect that certain AT strategies that generate trade clustering are not necessarily motivated by the reasons mentioned above. An extreme example occurred during the flash crash of May 6 2010 where it is clear that trading between HFTs generated more trading giving rise to the 'hot potato' effect.[2]

The profitability of these low latency AT strategies depends on how they interact with the dynamics of the LOB, and, more importantly, how these AT strategies coexist with each other. The recent increase in the volume of orders shows that fast traders are dominating the market and it is very difficult to link news arrival or other classical ways of explaining motives for trade to the activity one observes in electronic markets. Superfast algorithms make trading decisions in split milliseconds. This speed, and how other superfast traders react, makes it difficult to link trade initiation to private or public information arrival, a particular type of trader, liquidity shock, or any other market event.

Therefore, as part of the model we develop here, we propose a reduced-form model for the intensity of the arrival of market sell and buy orders. The novelty we introduce is to assume that market orders arrive in two types. The first type of orders are influential orders which excite the market and induces other traders to increase the amount of market orders they submit. For instance, the arrival of an influential market sell order increases the probability of observing another market sell order over the next time step and also increases (to a lesser extent) the probability of a market buy order to arrive over the next time step. On the other hand, when non-influential orders arrive the intensity of the arrival of market orders does not change. This reflects the existence of trades that the rest of the market perceives as not conveying any information that would alter their willingness to submit market orders. In our model we also incorporate the arrival of public news as a state variable that increases the intensity of market orders. In this way, our model for the arrival of market orders is able to capture trade clustering which can be one-sided or two-sided and allow for the activity of trading to show the positive feedback that algorithmic trades seem to have brought to the market environment.

In our model the arrival of trades also affects the midprice and the LOB. The arrival of market orders is generally regarded as an informative process because it may convey information about subsequent price moves and adverse selection risks.[3] Here we assume that the dynamics of the midprice of the asset are affected by short-term imbalances in the amount of influential market sell and buy orders – in particular, these imbalances have a temporary effect on the drift of the midprice. The arrival

[2]During that day between 13:45:13 and 13:45:27 CT HFTs traded over 27,000 contracts which accounted for approximately 49% of the total trading volume, while their net position changed by only about 200 contracts (see Kirilenko et al. (2010)).

[3]For instance, periods where the amount of market buy orders is much higher than the amount of market sell orders could be regarded as times where informed traders have a private signal and are adversely selecting market makers who are unaware that they are providing liquidity at a loss, see Easley and O'Hara (1992).

of good and bad news have a similar effect.

Moreover, in our model the arrival of influential orders have a transitory effect on the shape of both sides of the LOB. More specifically, since some market makers anticipate changes in the intensity of both the sell and buy market orders, the shape of the buy and sell side of the book will also undergo a temporary change due to market makers repositioning their limit orders in anticipation of the increased expected market activity and adverse selection risk.

We test our model using simulations where we assume different types of HFTs who are mainly characterized by the quality of the information that they are able to process and incorporate into their optimal postings. We show that only those HFTs who incorporate predictions of short-term price deviations in their strategy will deliver expected positive profits. The other HFTs are driven out of the market because their limit orders are picked off by better informed traders and because they cannot profit from directional strategies which are also based on short-lived predictable trends. We also show that those HFTs who cannot execute profitable directional strategies (and are systematically being picked off) can stay in business if they exert tight controls on their inventories. In our model these controls imply a higher penalty on their inventory position which pushes the optimal limit orders further away from the midprice so the chances of being picked off by other traders are considerably reduced.

## 2. Arrival of Market Orders and Price Dynamics

Very little is known about the details of the strategies that are employed by AT desks or the more specialized proprietary HF trading desks. Algorithms are designed for different purposes and to seek profits in different ways, Bouchard et al. (2011). For example, there are algorithms that are designed to find the best execution prices for investors who wish to minimize the price impact of large buy or sell orders, Bertsimas and Lo (1998), Almgren and Chriss (2000), Kharroubi and Pham (2010) and Bayraktar and Ludkovski (2011), while others are designed to manage inventory risk, Guéant et al. (2011). There are HF strategies that specialize in arbitraging across different trading venues whilst others seek to profit from short-term deviations in stock prices. And finally, there are trading algorithms that seek to profit from providing liquidity by posting bids and offers simultaneously, Guilbaud and Pham (2011). In previous works on algorithmic trading in LOBs, the mid-price is assumed independent of market orders and market orders arrive at Poisson times. Our work differs significantly in that we do not assume independence of market orders, the LOB dynamics and mid-price moves.

A pillar of capital markets is the provision of liquidity to investors when they need it. As compensation for providing immediacy, market makers or liquidity providers earn the *realized spread*, which is the market maker's expected gain from a roundtrip trade.[4] These expected gains depend on, among other things, the architecture of the LOB , and on the ability that market makers have

[4]Some exchanges also pay rebates to liquidity providers, see Cartea and Jaimungal (2010).

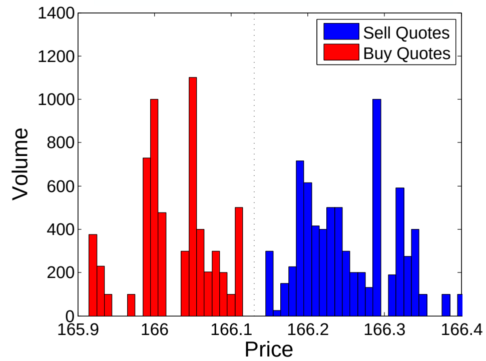


Figure 1: A snapshot of the (NASDAQ) LOB for IBM on June 21, 2011 at 11:45:20.26.

to hold inventories which gives them the opportunity to build strategic long or short positions.

In the LOB, limit orders are prioritized first according to price and then according to time. For example, if two sell (buy) orders are sent to the exchange at the same time, the one with the lowest (highest) price is placed ahead in the queue. Similarly, orders that improve the prices for buy or sell will jump ahead of others regardless of how long they have been resting in the book. Thus, based on the price/time priority rule the LOB stacks on one side all buy orders (also referred to as *bids*) and on the other side all sell orders (also referred to as *offers*). The difference between the best offer and best bid is known as the *spread* and their mean is referred to as the *midquote* price. Another dimension of the book is the quantities on the sell and buy side for each price tick which give 'shape' to the LOB.

The HF trading strategy we develop here is designed to profit from the realized spread where we allow the HFT to build inventories. To this end, before we formalize the HFT's optimization problem, we require a number of building blocks to capture the most important features of the market dynamics.[5] Since the HFT maximizes expected terminal wealth over a finite horizon $T$ and she is continuously repositioning buy and sell limit orders, the success of the strategy depends on optimally picking the 'best places' in the bid and offer queue which requires us to model: (i) The dynamics of the fundamental value of the traded stock. (ii) The arrival of market buy and sell orders. And (iii) how market orders cross the resting orders in the LOB. In this section we focus on (i) and (ii), then in Section 3 we discuss (iii) and after that we present the formal optimal control problem that the HFT solves.

We assume that the **midprice** (or **fundamental price**) of the traded asset follows

$$dS_t = (\upsilon + \alpha_t)\, dt + \sigma\, dW_t\,, \tag{1}$$

[5]Although we focus on a HF trading market making algorithm, the framework we develop here can be adapted for other types of AT strategies.

where $W_t$ is a $\mathbb{P}$-standard Brownian Motion, and $S_0 > 0$ and $\sigma > 0$ are constants.[6] The drift of the midprice is given by a long-term component, $\upsilon$, and by $\alpha_t$ which is a predictable zero-mean reverting component that represents stochastic short-term deviations from $\upsilon$. The recent paper of Gârleanu and Pedersen (2009) develops an optimal dynamic portfolio policy when trading is costly and security returns are predictable by signals with different mean-reversion speeds. Although their model is designed with a 'low frequency' investor in mind, their framework provides a rich and abstract setting for models of predictable returns like the one proposed here. Since we are interested in HF trading, our predictors are based on order flow information where we allow for feedback between market order events and short-term-alpha. In the rest of the paper we assume that $\upsilon = 0$ because the HF strategies we develop are for very short-term intervals.

Below we give details of the dynamics of the process $\alpha_t$, this element of the model plays a key role in the determination of the HF strategies we develop because it captures different features that we observe in the dynamics of the midprice. For instance, it captures the price impact that some orders have on the midprice as a result of: the arrival of news, a burst of activity on one or both sides of the market, orders that eat into the LOB, etc. Furthermore, we also know that a critical component of HF trading is the ability that HFTs have to predict short-term deviations in prices so that they make markets by taking advantage of directional strategies based on short-term predictions (i.e. they are able to predict short-term-alpha) whilst at the same time allowing them to reposition stale or submit new quotes to avoid being picked off by other market participants trading on short-term-alpha – avoid being adversely selected.

### *2.1. Self-exciting incoming market order dynamics*

Markets tend to follow an intraday pattern. Usually after the market opens and before the market closes there is more activity than during the rest of the day. Figure 2 shows the historical intensity of trade arrival, buy and sell, for IBM over a three minute period (starting at 3.30pm, February 1 2008). The historical intensities are calculated by counting the number of buy and sell market orders over the last 1 second. The fitted intensities are computed using our model (see Equation (2)) under the specific assumption that all trades are influential – see Appendix A for more details. From the figures we observe that market orders may arrive in clusters and that there are times when the markets are mostly one-sided (for instance the first 60 seconds of trading is more active on the buy side than on the sell side) and that these bursts of activity die out rather quickly and revert to around 5 events per second.

Why are there bursts of activity on the buy and sell sides? It is difficult to link all these short-lived increases in the levels of activity to the arrival of news. One could argue that trading algorithms, including HF, are also responsible for the sudden changes in the pace of the market activity, including

[6] Unless otherwise stated, all random variables and stochastic processes are on the completed filtered probability space $(\Omega, \mathcal{F}_T, \mathbb{F}, \mathbb{P})$ with filtration $\mathbb{F} = \{\mathcal{F}_t\}_{0 \le t \le T}$ and where $\mathbb{P}$ is the real-world probability measure. What generates the filtration will be defined precisely in Section 5. Simply put, it will be generated by the Brownian motions $W_t$ and $B_t$ (introduced later), counting processes corresponding to buy/sell market and filled limit orders, news events and the indicator of whether a trade is influential or not.

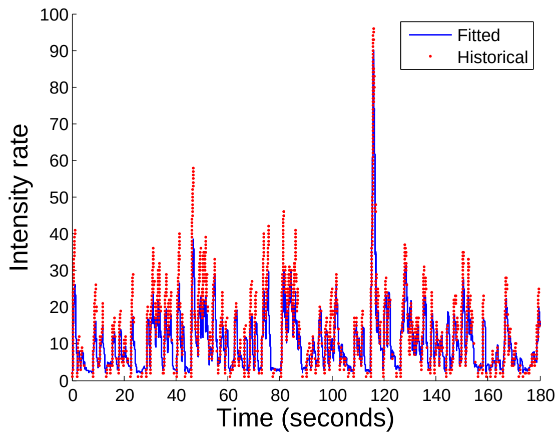

(a) Arrival of IBM market buy orders

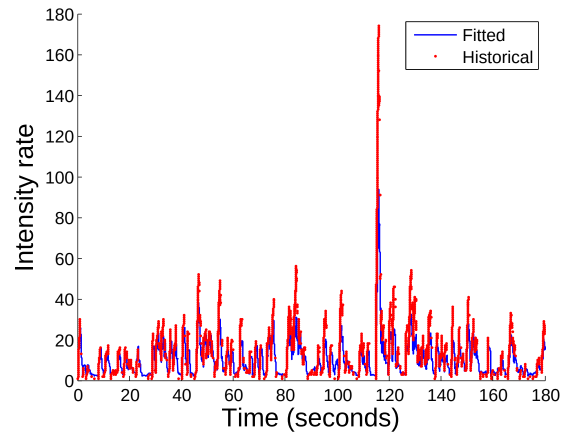

(b) Arrival of IBM market sell orders

Figure 2: IBM market orders. Historical running intensity versus smoothed fitted intensity (restricted to $\rho = 1$) using a 1 second sliding window for IBM for a 3 minute period, between 3.30 and 3.33 pm, February 1 2008.

bursts of activity in the LOB, and most of the times these algorithms act on information which is difficult to link to public news. Thus, here we take the view that some market orders generate more trading activity in addition to the usual effect of news increasing the intensity of market orders.

In our model market orders arrive in two types. The first are *influential* orders which excite the state of the market and induce other traders to increase their trading activity. We denote the total number of arrivals of influential sell/buy market orders up to (and including) time $t$ by the processes $\{\overline{M}_t^-, \overline{M}_t^+\}$. The second type of orders are *non-influential* orders. These are viewed as arising from players who do not excite the state of the market. We denote the total number of arrivals of non-influential sell/buy market orders up to (and including) time $t$ by the processes $\{\widetilde{M}_t^-, \widetilde{M}_t^+\}$. Note that the type indicator of an order is not an observable. Rather all one can observe is whether the market became more active after that trade. Therefore we assume that, conditional on the arrival of a market order, the probability that the trade is influential is a constant $\rho \in [0, 1]$.

Clearly, public bad (good) news increases the sell (buy) activity, but what is not clear is whether market participants always interpret news in the same way. If there is disagreement in how to interpret news or news is ambiguous, then both sides of the market will show an increase in the intensity of buy and sell market orders.[7]

Thus, we model the intensity of sell, $\lambda_t^-$, and buy, $\lambda_t^+$, market orders by assuming that they solve the coupled system of stochastic differential equations:

**Assumption 1.** The rate of arrival of market sell/buy orders $(\lambda_t^-, \lambda_t^+)$ solve the coupled system of

[7]There are automated news feeds designed for AT that already classify the news as good, bad, and neutral.

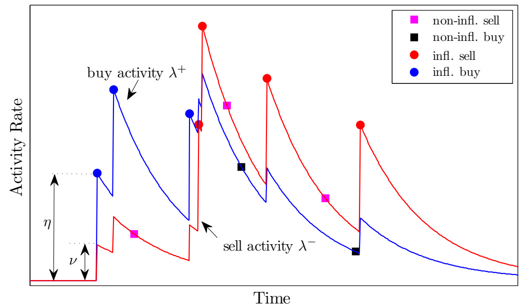


Figure 3: Sample path of market order activity rates. When influential trades arrive, the activity of both buy and sell orders increase but by differing amounts. Circles indicate the arrival of an influential market order, while squares indicate the arrival of non-influential trades.

SDEs

$$\begin{cases} d\lambda_t^- = \beta(\theta - \lambda_t^-)dt + \eta\, d\overline{M}_t^- + \nu\, d\overline{M}_t^+ + \tilde{\eta}\, dZ_t^- + \tilde{\nu}\, dZ_t^+ \,, \\ d\lambda_t^+ = \beta(\theta - \lambda_t^+)dt + \eta\, d\overline{M}_t^+ + \nu\, d\overline{M}_t^- + \tilde{\eta}\, dZ_t^+ + \tilde{\nu}\, dZ_t^- \,, \end{cases} \tag{2}$$

where $Z_t^\pm$ are Poisson processes (independent of all other processes), with constant activity rate $\mu^\pm$, which represent the total amount of good and bad news that have arrived until time $t$, and recall that $\overline{M}_t^+$ and $\overline{M}_t^-$ are the total number of influential buy and sell orders up until time $t$. Moreover, $\beta, \theta, \eta, \nu, \tilde{\eta}, \tilde{\nu}$ are non-negative constants satisfying the constraint $\beta > \rho(\eta + \nu)$.

Market orders are self-exciting because their arrival rates $\lambda^\pm$ jump upon the arrival of influential orders (note that the arrival of non-influential orders do not affect $\lambda^\pm$). If the influential market order was a buy (so that a sell limit order was lifted), the jump activity on the buy side increases by $\eta$ while the jump activity on the sell side increases by $\nu$. On the other hand, if the influential market order was a sell (so that a buy limit order was hit), the jump activity on the sell side increases by $\eta$ while the jump activity on the buy side increases by $\nu$. Typically one would expect $\nu < \eta$ so that jumps on the opposite side of the book are smaller than jumps on the same side of the book (this bears out in the calibration as well as in the moving window activities reported in Figures 2).

News also affects market activity, but does not cause self-excitations. In our model we include cross effects to capture the fact that market participants do not always interpret news in the same way, for example good news increases the intensity of buy market orders by $\tilde{\nu}$ but also affects the intensity of market sell orders by $\tilde{\eta}$.

Trading intensity is mean reverting. Jumps in activity decay back to its long run level of $\theta$ at an exponential rate $\beta$. Figure 3 illustrates a sample path during which no news arrives, but some of the market orders that arrive are influential and induce jumps in the activity level. The lower bound condition on $\beta$ is required for the activity rate to be ergodic. To see this, define the mean future activity rate $m_t^\pm(u) = \mathbb{E}[\lambda_u^\pm | \mathcal{F}_t]$ for $u \geq t$. For the processes $\lambda_t^\pm$ to be ergodic, $m_t^\pm(u)$ must remain bounded as a function of $u$, for each $t$, and the following Lemma provides a justification for the constraint.

**Lemma 1. Lower Bound on Mean-Reversion Rate.** *The mean future rate* $m_t^{\pm}(u)$ *remains bounded for all* $u \geq t$ *if and only if* $\beta > \rho(\eta + \nu)$*. Furthermore,*

$$\lim_{u \to \infty} m_t^{\pm}(u) = \mathbf{A}^{-1}\boldsymbol{\zeta}\,, \quad \textit{where} \quad \mathbf{A} = \begin{pmatrix} \beta - \eta\rho & -\nu\rho \\ -\nu\rho & \beta - \eta\rho \end{pmatrix} \quad \textit{and} \quad \boldsymbol{\zeta} = \begin{pmatrix} \beta\theta + \tilde{\eta}\mu^{-} + \tilde{\nu}\mu^{+} \\ \beta\theta + \tilde{\nu}\mu^{-} + \tilde{\eta}\mu^{+} \end{pmatrix}.$$

Proof. See Appendix B.1. □

The intuition for the constraint is that when a market order arrives the activity will jump either by $\eta$ or by $\nu$ and this occurs with probability $\rho$. Further, since both sell and buy influential orders induce self-excitations, the decay rate $\beta$ must be strong enough to compensate for both jumps to pull the process towards its long-run level of $\theta$.

## 3. Limit Quote Arrival Dynamics and Fill Rates

The LOB can take on a variety of shapes and changes dynamically throughout the day, see Rosu (2009) and Cont et al. (2010). Market orders eat into the LOB until all the volume specified in the order is filled. Limit orders in the tails of the LOB are less likely to be filled than those within a couple of ticks away from the midprice $S_t$. Another important feature of the LOB dynamics is how quickly the book recovers from a large market order; i.e. the quoted spread returns to previous values. This is known as the resilience of the LOB.

Therefore, the decision where to post limit buy and sell orders depends on a number of characteristics of the LOB and on the market orders. Some of the LOB features are: shape of the LOB, resiliency of the LOB, and how the LOB changes in between the arrival of market orders. These features, combined with the size and rate of the incoming market orders, determine the fill rates of the limit orders. The fill rate is the rate of execution of a limit order. Intuitively, a high (low) fill rate indicates that a limit order is more (less) likely to be filled by a market order.

Here we model the fill rate facing the HFT in a general framework where we allow the depth and shape of the book to fluctuate. The fill rate depends on where the HFT posts the limit buy and sell orders, that is at $S_t - \delta_t^-$ and $S_t + \delta_t^+$ respectively, where $\delta^\pm$ denotes how far away from the midprice the orders are posted.

**Assumption 2.** The fill rates are of the form $\Lambda_t^\pm \triangleq \lambda_t^\pm h_\pm(\delta; \boldsymbol{\kappa}_t)$, where the non-increasing function $h_\pm(\delta; \boldsymbol{\kappa}_t) : \mathbb{R} \to [0,1]$ is $C^1$ in $\delta$ (uniformly in $t$ for $\boldsymbol{\kappa}_t \in \mathbb{R}^n$, fixed $\omega \in \Omega$) and $C^3$ in an open neighborhood of the risk-neutral optimal control, and $\lim\limits_{\delta \to \infty} \delta\, h_\pm(\delta; \boldsymbol{\kappa}_t) = 0$ for every $\boldsymbol{\kappa}_t \in \mathbb{R}^n$. Moreover, the functions $h_\pm(\delta; \boldsymbol{\kappa}_t)$ satisfy: $h_\pm(\delta; \boldsymbol{\kappa}_t) = 1$ for $\delta \leq 0, \boldsymbol{\kappa}_t \in \mathbb{R}^n$.

Assumption 2 allows for very general dynamics on the LOB through the dependence of the fill probabilities (FPs) $h_{\pm}(\delta;\boldsymbol{\kappa}_t)$ on the process $\boldsymbol{\kappa}_t$. The FPs can be viewed as a parametric collection with the exponential class $h_{\pm}(\delta;\boldsymbol{\kappa}_t) = e^{-\kappa_t^{\pm}\,\delta^{\pm}}$ and power law class $h_{\pm}(\delta;\boldsymbol{\kappa}_t) = (1 + (\kappa^{\pm}\delta^{\pm})^{\alpha})^{-1}$ being two prime examples. The process $\boldsymbol{\kappa}_t$ introduces dynamics into the collection of FPs reflecting the dynamics in the LOB itself. The differentiability requirements in assumption 2 are necessary for the asymptotic expansions we carry out later on to be correct. The limiting behavior for large $\delta^{\pm}$ implies that the (volume) is the book thins out sufficiently slowly such that the FPs decay sufficiently fast (faster than linear) so that it is not optimal to place orders infinitely far away from the midprice. Finally, the requirement that $h_{\pm}(\delta;\boldsymbol{\kappa}_t) = 1$ for $\delta \leq 0$ and $\forall \boldsymbol{\kappa}_t \in \mathbb{R}^n$ is a financial one. A trader wanting to maximize her chances of being filled the next time a market order arrives, must post the limit orders at the midprice, i.e. $\delta^{\pm} = 0$, or she can also cross the midprice, i.e. $\delta^{\pm} < 0$. In these cases we suppose that the fill rate is $\Lambda_t^{\pm} = \lambda_t^{\pm}$, i.e. it equals the rate of incoming market orders. This assumption makes crossing the midprice a suboptimal decision because the trader cannot improve the arrival rate of market orders, thus she will always post limit orders that are $\delta^{\pm} \geq 0$ away from the midprice. Additionally, this condition is more desirable than explicitly restricting the controls $\delta^{\pm}$ to be non-negative, since it is not necessary to check the boundary condition at $\delta^{\pm} = 0$; it will automatically be satisfied. Moreover, we have the added bonus that the optimal control satisfies the first-order condition.

**Assumption 3.** The dynamics for $\boldsymbol{\kappa}_t$ satisfy

$$\left\{ \begin{aligned} d\kappa_t^- &= \beta_\kappa(\theta_\kappa - \kappa_t^-)\,dt + \eta_\kappa\, d\overline{M}_t^- + \nu_\kappa\, d\overline{M}_t^+ + \tilde{\eta}_\kappa\, dZ_t^- + \tilde{\nu}_\kappa\, dZ_t^+\,, \\ d\kappa_t^+ &= \beta_\kappa(\theta_\kappa - \kappa_t^+)\,dt + \nu_\kappa\, d\overline{M}_t^- + \eta_\kappa\, d\overline{M}_t^+ + \tilde{\nu}_\kappa\, dZ_t^- + \tilde{\eta}_\kappa\, dZ_t^+\,, \end{aligned} \right. \tag{3}$$

where $\eta_\kappa$, $\nu_\kappa$, $\tilde{\eta}_\kappa$ and $\tilde{\nu}_\kappa$ are non-negative constants and $\theta_\kappa$ and $\beta_\kappa$ are strictly positive constants.

Assumption 3 is a specific modeling assumption[8] on $\boldsymbol{\kappa}_t$ which allows for incoming influential market orders and news events to have an impact on the FPs. An increase (decrease) in the fill rate can be due to two main reasons: (i) a decrease (increase) in LOB depth or (ii) an increase (decrease) in the distribution of market order volumes (in a stochastic dominance sense). This is a one-way effect because influential market orders cause jumps in the $\boldsymbol{\kappa}_t$ process, but jumps in the FP do not induce jumps in market order arrivals. While it is possible to allow such feedback, empirical investigations (such as those in Large (2007)) demonstrate that the incoming market orders influence the state of the LOB and not the other way around. The mean-reversion term draws $\kappa_t^{\pm}$ to the long-run mean of $\theta_\kappa$ so that the impact of influential orders on the LOB is only temporary. Typically, we expect that the rate of mean-reversion $\beta_\kappa$ for the LOB to be slower than the rate of mean-reversion $\beta$ of the market order activity. In other words, the impact of influential orders persists in the LOB on a longer time scale compared to their effect on market order activity.

Observe that if market order volumes are iid, then the $\kappa_t^{\pm}$ processes can be interpreted as param-

[8]It is in principle possible to include more general dynamics on $\boldsymbol{\kappa}$; however, we opt to work with this assumption for two reasons (i) the results are easily interpreted and (ii) it reflects realistic behavior of the LOB dynamics.

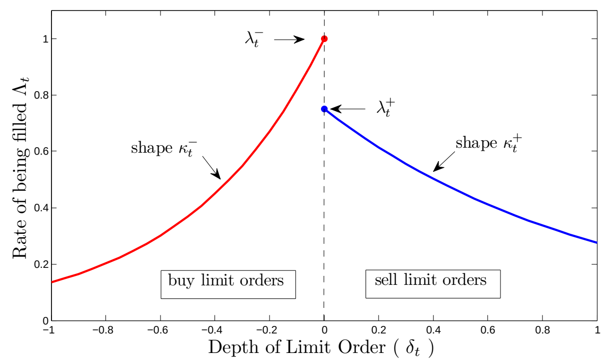


Figure 4: Fill rates $\Lambda^{\pm}$ with $\kappa_t^+ = 2$, $\kappa_t^- = 1$, $\lambda_t^+ = 0.75$ and $\lambda_t^- = 1$. As $\kappa_t^{\pm}$ evolve, the fill rate shape changes, while the market order activity rates $\lambda^{\pm}$ modulate the vertical scale. The maximum fill rate is achieved at $\delta = 0$.

eters directly dictating the shape of the LOB. In particular, if the market order volumes are iid exponentially distributed and the shape of the LOB is flat, then the probability that a limit order at price level $S_t \pm \delta_t^{\pm}$ is executed (given that a market order arrives) is equal to $e^{-\kappa_t^{\pm}\delta_t^{\pm}}$. Consequently, $\kappa_t^{\pm}$ can be interpreted as the exponential decay factor for the fill rate of orders placed away from the midprice. In order to satisfy the $C^1$ condition at $\delta^{\pm} = 0$ and the condition that $h_{\pm}(\delta, \boldsymbol{\kappa}_t) = 1$ for $\delta^{\pm} \leq 0$, it is necessary to smooth the exponential function at $\delta = 0$. This is always possible, since there exists smooth functions for which the sup distance to the target function is less than any positive constant.

Moreover, immediately after an influential market buy/sell order arrives (which eats into the sell/buy side of the book), the probability (given that a market order arrives) that a limit order at price level $S_t \pm \delta_t^{\pm}$ is executed is, for the same $\delta^{\pm}$, smaller than the probability of being filled before the influential order arrived. The intuition is the following. Immediately after an influential market order arrives, market participants react in anticipation of the increase of market activity they will face and decide to send limit orders to the book. Since many market participants react in similar way, the probability of limit orders being filled, conditioned on a market order arriving, decreases.[9]

Figure 4 illustrates the shape of the fill rates at time $t$ describing the rate of arrival of market orders which fill limit orders placed at price levels $S_t \pm \delta_t^{\pm}$. Notice that these rates peak at zero spread at which point they are equal to the arrival rate of market orders. In the figure these rates are asymmetric and decay at differing speeds because we have assumed different parameters for the buy and sell side, $\kappa_t^+ = 2$, $\kappa_t^- = 1$, $\lambda_t^+ = 0.75$ and $\lambda_t^- = 1$. In general these rates will fluctuate throughout the day.

[9] It is also possible to have markets where, conditioned on the arrival of a market order, the probability of a limit order being filled increases immediately after the arrival of an influential order. We can incorporate this feature in our model. Note also that in our general framework, immediately after the influential buy/sell market order arrives, the intensities $\lambda^{\pm}$ increase and the overall effect of an influential order on the fill rates $\Lambda^{\pm} = \lambda_t^{\pm} h_{\pm}(\delta, \boldsymbol{\kappa}_t)$ is ambiguous when $\lambda_t^{\pm}$ and $h_{\pm}(\delta, \boldsymbol{\kappa}_t)$ move in opposite directions after the arrival of an influential order, for example when $h_{\pm}(\delta, \boldsymbol{\kappa}_t) = e^{-\kappa_t^{\pm}\delta_t^{\pm}}$.

## 4. Short-Term-Alpha Dynamics: directional strategies and adverse selection

The actions of market participants affect the dynamics of the midprice via activity in the LOB and/or the execution of market buy and sell orders. For instance, the arrival of public information is impounded in the midprice of the asset as a result of new market orders and the arrival and cancellation of limit orders. Similarly, bursts of activity in buy and/or sell market orders, which are not necessarily the result of the arrival of public information, has market impact by producing momentum in the midprice.

As discussed above, a great deal of the strategies that HFTs employ are directional strategies that take advantage of short-term price deviations. HFTs use their superior speed to gather and process order flow information, as well as public information, to spot fleeting opportunities which in our model are captured by the dynamics of short-term-alpha. Short-term predictability is a key source of HF trading revenues for two reasons. First, it enables HFTs to exploit their superior knowledge of short-term trends in prices to execute profitable roundtrip trades, and second, because it provides key information to update or cancel quotes that can be adversely picked off by other traders.

One can specify the dynamics of the predictable drift $\alpha_t$ in many ways and this depends on the factors that affect the short-term drift which for HF market making are based on order flow and news. Here we assume that $\alpha_t$ is a zero-mean-reverting process and jumps by a random amount at the arrival times of influential trades and news events. If the influential trade was buy initiated (and therefore lifts a sell limit order) the drift will jump up, and if the influential trade was sell initiated (and therefore hit a buy limit order) the drift will jump down; news has a similar effect on $\alpha_t$. As such, we model the predictable drift according to the following assumption.

**Assumption 4.** The dynamics for the predictable component of the midprice's drift, $\alpha_t$, satisfy

$$d\alpha_t = -\zeta\, \alpha_t\, dt + \sigma_\alpha\, dB_t + \epsilon^+\, d\overline{M}_t^+ - \epsilon^-\, d\overline{M}_t^- + \widetilde{\epsilon}^+\, dZ_t^+ - \widetilde{\epsilon}^-\, dZ_t^- \tag{4}$$

where $\epsilon^\pm$ and $\widetilde{\epsilon}^\pm$ are random variables representing the size of the sell/buy influential trade and news impact on the drift of the midprice. Moreover, $B_t$ denotes a Brownian motion independent of all other processes, and $\zeta$, $\sigma_\alpha$ are positive constants.

Moreover, we see how slower traders will be adversely selected by better informed and quicker traders. For example, assume that $\alpha_t = 0$ and an HFT 'detects' that the incoming buy market order is influential. Her optimal directional strategy is to simultaneously send the following orders to the LOB: cancel her sell limit orders, attempt purchase the asset (from a slower market participant), and send new sell limit orders to be able to unwind the transaction. Of course, these types of trades do not guarantee a profit but on average these roundtrips will be profitable because the HFT trades on short-term-alpha and profits from other traders who are not able to update their quotes in time or who submit market sell orders right before prices increase. Finally, even if HFTs who are able to trade on short-term-alpha miss a fleeting opportunity to execute a directional trade, they still benefit from updating their stale quotes in the LOB to avoid being adversely selected by other market participants.

## 5. The High Frequency Trader's Optimization Problem

So far, we have specified counting processes for market orders and dynamics of the LOB through the FPs; however, we also require a counting process for the agent's filled limit orders. To this end, let $N_t^+$ and $N_t^-$ denote the number of the agent's limit sell and buy orders, respectively, that were filled up to and including time $t$ and the process $q_t = N_t^- - N_t^+$ is the agent's total inventory. Note that the arrival rate of these counting processes can be expressed as $\Lambda_t^\pm \triangleq \lambda_t^\pm h_\pm(\delta; \boldsymbol{\kappa}_t)$, as in Assumption 2. Finally, the agent's cash process (i.e., excluding the value of the $q_t$ shares she currently holds) satisfies the SDE

$$dX_t = (S_t + \delta_{t-}^+)\, dN_t^+ - (S_t - \delta_{t-}^-)\, dN_t^- \tag{5}$$

where $\delta_{t-}^\pm$ denotes the left-limit of the spreads since, if the order was filled, the trader receives the spread that was posted an instant prior to the arrival of the market order.[10]

### 5.1. Formulation of the HF investment problem

The HFT wishes to place sell/buy limit orders at the prices $S_t \pm \delta_t^\pm$ at time $t$ such that the expected terminal wealth is maximized whilst penalizing inventories.[11] The HFT is continuously repositioning her limit orders in the book by canceling stale and submitting new limit orders.[12] Specifically, her value function is

$$\Phi(t, X_t, S_t, q_t, \alpha_t, \boldsymbol{\lambda}_t, \boldsymbol{\kappa}_t) = \sup_{(\delta_u^-, \delta_u^+)_{t \le u \le T} \in \mathcal{A}} \mathbb{E}\left[ X_T + q_T S_T - \phi \int_t^T q_s^2\, ds \,\middle|\, \mathcal{F}_t \right], \tag{6}$$

where the supremum is taken over all (bounded) $\mathcal{F}_t$-progressively measurable functions and $\phi$ penalizes deviations of $q_t$ from zero along the entire path of the strategy. Moreover, $\overline{\mathcal{F}}_t$ is the natural (and completed) filtration generated by the collection of processes $\{S_t,\, \alpha_t,\, M_t^\pm = \overline{M}_t^\pm + \widetilde{M}_t^\pm,\, N_t^\pm\}$ and the extended filtration $\mathcal{F}_t = \overline{\mathcal{F}}_t \vee \sigma\{(\overline{M}_u)_{0 \le u \le t}\}$. Note that $\boldsymbol{\lambda}_t$ and $\boldsymbol{\kappa}_t$ are progressively measurable with respect to this expanded filtration. We will often suppress the dependence on many of the variables in $\Phi(\cdot)$ and recall that we assumed $\upsilon = 0$ in the dynamics of the midprice.

The above control problem can be cast into a discrete-time controlled Markov chain as carried out in Bäuerle and Rieder (2009). Classical results from Bertsekas and Shreve (1978) imply that a dynamic

[10] The technical issue is as follows: recall that the driving counting processes, and consequently, the spreads, are right continuous with left-limits (RCLL). However, stochastic integrals must have integrands that are left-continuous with right-limits (LCRL) for the integral w.r.t. a martingale integrator to still be a martingale. Replacing $\delta_t^\pm$ with $\delta_{t-}^\pm$ achieves this goal.

[11] An alternative specification is to assume that the HFT is risk averse so that she maximizes expected utility of terminal wealth. The current approach, however, is more akin to Almgren (2003) where quadratic variation, rather than variance (as is widely incorrectly stated), is penalized which acts on the entire path of the strategy.

[12] In this setup the HFT's limit orders are always of the same size. An interesting extension is to also allow the HFT to choose the amount of shares in each limit order.

programming principle holds and that the value function is the unique viscosity solution of the HJB equation

$$
\begin{aligned}
&(\partial_t+\mathcal{L})\,\Phi+\alpha\,\Phi_s+\tfrac{1}{2}\sigma^2\,\Phi_{ss}\\
&\quad+\lambda^-\sup_{\delta^-}\Big\{\,h_-(\delta^-;\boldsymbol{\kappa})\,\Big[\mathbb{S}^-_{q,\lambda}\Phi(t,x-s+\delta^-)-\Phi\Big]+(1-h_-(\delta^-;\boldsymbol{\kappa}))\,\big[\mathbb{S}^-_{\lambda}\Phi-\Phi\big]\Big\}\\
&\quad+\lambda^+\sup_{\delta^+}\Big\{h_+(\delta^+;\boldsymbol{\kappa})\,\Big[\mathbb{S}^+_{q,\lambda}\Phi(t,x+s+\delta^+)-\Phi\Big]+(1-h_+(\delta^+;\boldsymbol{\kappa}))\,\big[\mathbb{S}^+_{\lambda}\Phi-\Phi\big]\Big\}=\phi\,q^2,
\end{aligned}
\tag{7}
$$

with boundary condition $\Phi(T,\cdot)=x+qs$, and the integro-differential operator $\mathcal{L}$ is the part of the generator of the processes $\alpha_t$, $\boldsymbol{\lambda}_t$, $\boldsymbol{\kappa}_t$, and $Z^{\pm}_t$ which do not depend on the controls $\delta^{\pm}_t$. Explicitly,

$$
\begin{aligned}
\mathcal{L}=\;&\beta(\theta-\lambda^-)\partial_{\lambda^-}+\beta(\theta-\lambda^+)\partial_{\lambda^+}+\beta_\kappa(\theta_\kappa-\kappa^-)\partial_{\kappa^-}+\beta_\kappa(\theta_\kappa-\kappa^+)\partial_{\kappa^+}\\
&-\zeta\,\alpha\partial_\alpha+\tfrac{1}{2}\sigma_\alpha^2\partial_{\alpha\alpha}+\mu^-\left(\widetilde{\mathbb{S}}^-_\lambda-1\right)+\mu^+\left(\widetilde{\mathbb{S}}^+_\lambda-1\right).
\end{aligned}
\tag{8}
$$

Moreover, we have introduced the following shift operators:

$$
\mathbb{S}^{\pm}_{\lambda}\Phi=\rho\,\mathbb{E}\big[\mathcal{S}^{\pm}_{\lambda}\Phi\big]+(1-\rho)\,\Phi\,, \tag{9a}
$$

$$
\mathbb{S}^{\pm}_{q\lambda}\Phi=\rho\,\mathbb{E}\Big[\mathcal{S}^{\pm}_{q\lambda}\Phi\Big]+(1-\rho)\,\mathcal{S}^{\pm}_{q}\Phi\,, \tag{9b}
$$

$$
\mathcal{S}^{\pm}_{q\lambda}=\mathcal{S}^{\pm}_{q}\mathcal{S}^{\pm}_{\lambda}\,, \tag{9c}
$$

$$
\mathcal{S}^{\pm}_{q}\Phi(t,x,s,q,\alpha,\boldsymbol{\lambda},\boldsymbol{\kappa})=\Phi(t,x,s,q\mp 1,\alpha,\boldsymbol{\lambda},\boldsymbol{\kappa})\,, \tag{9d}
$$

$$
\mathcal{S}^{+}_{\lambda}\Phi(t,x,s,q,\alpha,\boldsymbol{\lambda},\boldsymbol{\kappa})=\Phi(t,x,s,q,\alpha+\epsilon^+,\boldsymbol{\lambda}+(\nu,\eta)',\boldsymbol{\kappa}+(\nu_\kappa,\eta_\kappa)')\,, \tag{9e}
$$

$$
\mathcal{S}^{-}_{\lambda}\Phi(t,x,s,q,\alpha,\boldsymbol{\lambda},\boldsymbol{\kappa})=\Phi(t,x,s,q,\alpha-\epsilon^-,\boldsymbol{\lambda}+(\eta,\nu)',\boldsymbol{\kappa}+(\eta_\kappa,\nu_\kappa)')\,, \tag{9f}
$$

$$
\widetilde{\mathbb{S}}^{+}_{\lambda}\Phi(t,x,s,q,\alpha,\boldsymbol{\lambda},\boldsymbol{\kappa})=\mathbb{E}[\Phi(t,x,s,q,\alpha+\widetilde{\epsilon}^+,\boldsymbol{\lambda}+(\tilde{\nu},\tilde{\eta})',\boldsymbol{\kappa}+(\widetilde{\nu}_\kappa,\widetilde{\eta}_\kappa)')]\,, \tag{9g}
$$

$$
\widetilde{\mathbb{S}}^{-}_{\lambda}\Phi(t,x,s,q,\boldsymbol{\lambda},\boldsymbol{\kappa})=\mathbb{E}[\Phi(t,x,s,q,\alpha-\widetilde{\epsilon}^-,\boldsymbol{\lambda}+(\tilde{\eta},\tilde{\nu})',\boldsymbol{\kappa}+(\widetilde{\eta}_\kappa,\widetilde{\nu}_\kappa)')]\,. \tag{9h}
$$

The expectation operator $\mathbb{E}[\cdot]$ in (9a) and (9b) are over the random variables $\epsilon^{\pm}$ and the expectation operator $\mathbb{E}[\cdot]$ in (9g) and (9h) are over the random variables $\widetilde{\epsilon}^{\pm}$.

The terms of the operator $\mathcal{L}$ have the usual interpretations: the first and second terms cause the activity rates $\lambda^{\pm}$ to decay back to the long run level $\theta$. The third and fourth terms pull $\kappa^{\pm}$ to their long run level. The fifth and sixth term causes $\alpha_t$ to diffusive and mean-revert to zero. The seventh and eight terms cause market order activities to jump upon public news arrival. Furthermore, the various terms in the HJB equation represent the jumps in the activity rate and/or a limit order being filled together with the drift and diffusion of the asset price and the effect of penalizing deviations of inventories from zero along the entire path of the strategy is captured by the term $\phi q^2$. More specifically, the sup over $\delta^-$ contain the terms due to the arrival of a market sell order (which are filled by limit buy orders). The first term represents the arrival of a market order (influential or not) which fills the limit order, while the second term represents the arrival of a market order (influential or not) which does not reach the limit order's price level. The sup over $\delta^+$ contain the analogous terms for the market buy orders (which are filled by limit sell orders).

*5.2. The Feedback Control of the optimal trading strategy*

In general, an exact optimal control is not analytically tractable – two exceptions are the cases of an exponential and power FPs where the optimal control admits exact analytical solutions as presented in Appendix C.1. For the general case, we provide an approximate optimal control via an asymptotic expansion which is correct to $o\,(\varsigma)$ where $\varsigma = \max(\phi, \alpha, \mathbb{E}[\epsilon^{\pm}])$. In principle, the expansion can be carried to higher orders if so desired.

**Proposition 2. Optimal Trading Strategy, Feedback Control Form.** *The value function* $\Phi$ *admits the decomposition* $\Phi = x + q\,s + g(t, q, \alpha, \boldsymbol{\lambda}, \boldsymbol{\kappa})$ *with* $g(T, \cdot) = 0$*. Furthermore, assume that* $g(\cdot)$ *can be written as an asymptotic expansion as follows*

$$g(t, q, \alpha, \boldsymbol{\lambda}, \boldsymbol{\kappa}) = g_0(t, q, \boldsymbol{\lambda}, \boldsymbol{\kappa}) + \alpha\, g_\alpha(t, q, \boldsymbol{\lambda}, \boldsymbol{\kappa}) + \varepsilon\, g_\varepsilon(t, q, \boldsymbol{\lambda}, \boldsymbol{\kappa}) + \phi\, g_\phi(t, q, \boldsymbol{\lambda}, \boldsymbol{\kappa}) + o(\varsigma)\,, \tag{10}$$

*with boundary conditions* $g_\cdot(T, \cdot) = 0$ *– note the subscripts on the functions* $g$ *do not denote derivatives, rather they are labels – and we have written* $\epsilon^{\pm} = \varepsilon\, \mathfrak{a}^{\pm}$ *with* $\varepsilon$ *constant. Then, the feedback controls of the optimal trading strategy for the HJB equation* (7) *admit the expansion*

$$\delta_t^{\pm *} = \delta_0^{\pm} + \alpha\, \delta_\alpha^{\pm} + \varepsilon\, \delta_\varepsilon^{\pm} + \phi\, \delta_\phi^{\pm} + o(\varsigma), \tag{11}$$

*where*

$$\delta_\alpha^{\pm} = -\,B(\delta_0^{\pm}; \boldsymbol{\kappa}) \left( \mathbb{S}_{q\lambda}^{\pm} g_\alpha - \mathbb{S}_{\lambda}^{\pm} g_\alpha \right), \tag{12a}$$

$$\delta_\varepsilon^{\pm} = -\,B(\delta_0^{\pm}; \boldsymbol{\kappa}) \left( \mathbb{S}_{q\lambda}^{\pm} g_\varepsilon - \mathbb{S}_{\lambda}^{\pm} g_\varepsilon \pm \rho\, \mathfrak{a}^{\pm} \left( \mathcal{S}_{q\lambda}^{\pm} g_\alpha - \mathcal{S}_{\lambda}^{\pm} g_\alpha \right) \right), \tag{12b}$$

$$\delta_\phi^{\pm} = -\,B(\delta_0^{\pm}; \boldsymbol{\kappa}) \left( \mathbb{S}_{q\lambda}^{\pm} g_\phi - \mathbb{S}_{\lambda}^{\pm} g_\phi \right), \tag{12c}$$

*and the coefficient* $B(\delta_0^{\pm}; \boldsymbol{\kappa}) = \frac{h'_{\pm}(\delta_0^{\pm}; \boldsymbol{\kappa})}{2h'_{\pm}(\delta_0^{\pm}; \boldsymbol{\kappa}) + \delta_0^{\pm} h''_{\pm}(\delta_0^{\pm}; \boldsymbol{\kappa})}$*. Moreover,* $\delta_0^{\pm}$ *is a strictly positive solution to*

$$\delta_0^{\pm}\, h'_{\pm}(\delta_0^{\pm}; \boldsymbol{\kappa}) + h_{\pm}(\delta_0^{\pm}; \boldsymbol{\kappa}) = 0. \tag{13}$$

*A solution to* (13) *always exists. Furthermore, the exact optimal controls are non-negative.*

Proof. See Appendix B.2. □

In the next subsection we use the optimal controls derived here to solve the nonlinear HJB equation and obtain an analytical expression for $g_1$ which is the last step we require to determine $\delta_t^{\pm *}$. Before proceeding we discuss a number of features of the optimal control $\delta_t^{\pm *}$ given by (11). The terms on the right-hand side of equation (11) show how the optimal postings are decomposed into different components: risk-neutral (first term), adverse selection and directional (second and third), and inventory-management (fourth term).

The risk-neutral component, given by $\delta_0^{\pm}$, does not directly depend on: the arrival rate of market orders, short-term-alpha, news arrival, or inventories. It depends on the FPs. To see the intuition behind this result, we note that a risk-neutral HFT, who does not penalize inventories, seeks to maximize the probability of being filled at every instant in time. Therefore, the HFT chooses $\delta^{\pm}$ to maximize the expected spread conditional on a market order hitting or lifting the appropriate side of the book, i.e. maximizes $\delta^{\pm}\, h_{\pm}(\delta^{\pm}; \boldsymbol{\kappa}_t)$. The first order condition of this optimization problem is given by (13) where we see that $\lambda^{\pm}$ plays no role in how the limit orders are calculated.[13]

The optimal halfspreads are adjusted by the impact that influential orders and news have on short-term-alpha through the term $\alpha_t\, \delta_{\alpha}^{\pm} + \varepsilon\, \delta_{\varepsilon}^{\pm}$ to reduce adverse selection costs and to profit from directional strategies. An HFT that is able to process information and estimate the parameters of short-term-alpha will adjust the halfspreads to avoid adverse selection and to profit from short-lived trends in the midprice. For example, if short-term-alpha is positive the HFT's sell halfspread is increased to avoid being picked off, and at the same time the buy halfspread decreases to take advantage of the first leg of a directional strategy by increasing the probability of purchasing the asset in anticipation of a price increase.

Finally, the fourth term is an inventory management component that introduces asymmetry in the postings so that the HFT does not build large long or short inventories. This component of the halfspread is proportional to the penalization parameter $\phi > 0$ which induces mean reversion to the optimal inventory position.

### *5.3. The asymptotic solution of the optimal trading strategy*

Armed with the optimal feedback controls, our remaining task is to solve the resulting non-linear HJB equation to this order in $\varsigma$. The following Theorem contains a stochastic characterization of the asymptotic expansion of the value function. This characterization can be computed explicitly in certain cases and then plugged into the feedback control to provide the optimal strategies.

**Theorem 3. Solving The HJB Equation.** *The solutions for $g_\alpha$, $g_\varepsilon$ and $g_\phi$ can be written as*

[13] If there are multiple solutions to (13) the HFT chooses the $\delta^{\pm}$ that yields the maximum of $\delta^{\pm}\, h_{\pm}(\delta^{\pm}; \boldsymbol{\kappa}_t)$.

$$g_\alpha = a_\alpha(t, \boldsymbol{\lambda}, \boldsymbol{\kappa}) + q\, b_\alpha(t)\,, \tag{14a}$$

$$g_\varepsilon = a_\varepsilon(t, \boldsymbol{\lambda}, \boldsymbol{\kappa}) + q\, b_\varepsilon(t, \boldsymbol{\lambda})\,, \tag{14b}$$

$$g_\phi = a_\phi(t, \boldsymbol{\lambda}, \boldsymbol{\kappa}) + q\, b_\phi(t, \boldsymbol{\lambda}, \boldsymbol{\kappa}) + q^2\, c_\phi(t)\,, \tag{14c}$$

*where*

$$b_\alpha(t) = \tfrac{1}{\zeta}\left(1 - e^{-\zeta(T-t)}\right), \tag{15a}$$

$$b_\varepsilon(t, \boldsymbol{\lambda}) = \mathbb{E}\left[\int_t^T \left\{\rho\left(\mathfrak{a}^+\,\lambda_u^+ - \mathfrak{a}^-\,\lambda_u^-\right) + \left(\widetilde{\mathfrak{a}}^+\mu^+ - \widetilde{\mathfrak{a}}^-\mu^-\right)\right\}\, b_\alpha(u)\,du \,\middle|\, \boldsymbol{\lambda}_t = \boldsymbol{\lambda}\right], \tag{15b}$$

$$b_\phi(t, \boldsymbol{\lambda}, \boldsymbol{\kappa}) = 2\,\mathbb{E}\left[\int_t^T \left\{h_{0,u}^+\,\lambda_u^+ - h_{0,u}^-\,\lambda_u^-\right\}(T-u)\,du \,\middle|\, \boldsymbol{\lambda}_t = \boldsymbol{\lambda},\ \boldsymbol{\kappa}_t = \boldsymbol{\kappa}\right], \quad \textit{and} \tag{15c}$$

$$c_\phi(t) = -(T-t)\,. \tag{15d}$$

*In the above, $h_{0,u}^\pm = h_\pm(\delta_{0,u}^\pm; \boldsymbol{\kappa}_u^\pm)$ and, as before, we have written $\epsilon^\pm = \varepsilon\mathfrak{a}^\pm$ and $\widetilde{\epsilon}^\pm = \varepsilon\widetilde{\mathfrak{a}}^\pm$. Finally, the functions $g_0$, $a_\alpha$, $a_\varepsilon$ and $a_\phi$ do not affect the optimal strategy.*

Proof. See Appendix B.3. □

The asymptotic expansion of the optimal controls now follows as a straightforward corollary to Theorem 3. Note that the functions $b_\varepsilon$ can be computed explicitly and is reported in Appendix C.2. Moreover, under some specific assumptions on the FPs $h_\pm$ (e.g., if $h_\pm$ are exponential or power functions), the function $b_\phi$ can also be computed explicitly. Proposition 7 in Appendix C.3 provides a general class of models (which includes the exponential and power cases) for which simple closed form results are derived and the implications for the optimal limiting order postings have a very natural interpretation.

**Corollary 4. Optimal Limit Orders.** *The asymptotic expansion of the optimal controls to first order in $\varsigma$ is (dependencies on the arguments have been suppressed for clarity)*

$$\delta^{\pm *} = \delta_0^\pm + B(\delta_0^\pm; \boldsymbol{\kappa}^\pm)\left\{\pm\mathbb{S}_\lambda^\pm\left(\mathbb{E}\left[\int_t^T \alpha_u\,du\right]\right) + \phi\left(\pm\mathbb{S}_\lambda^\pm b_\phi + (1 \mp 2q)(T-t)\right)\right\} \tag{16}$$

*where $\delta_0^\pm$ satisfies (13), and we have $|\delta_{opt}^\pm - \delta^{*\pm}| = o(\varsigma)$. Furthermore, the optimal controls $\max\{\delta^{\pm *}, 0\}$ are also of order $o(\varsigma)$.*[14]

[14] Note that the exact solution of the optimal control is non-negative as discussed in Assumption 2, but this is not necessarily the case in the asymptotic solution, thus we write the optimal control as $\max\{\delta^{\pm *}, 0\}$.

Proof. See Appendix B.4. □

The expression for the optimal control warrants some discussion which goes beyond the discussion that followed the general result in Proposition 2. The term $\delta_0^\pm$ represents the action of a risk-neutral agent who is not aware or is not able to estimate the impact that influential market orders and news arrival have on the stochastic drift of the midprice (so she sets it to zero). The first term in the braces accounts for the expected change in midprice due to the potential impact of orders (and news) on the midprice's drift as well as the expected change to the arrival of orders. That is, this term plays a dual role in the optimal strategy: it corrects for the adverse selection effect and positions the quotes to execute directional strategies. If the drift is positive, the agent posts further away from the midprice on the sell side (adverse selection correction) and closer to the midprice on the buy side in anticipation of upward price movements (directional strategy). When the drift is negative the interpretation is similar. The term proportional to $\phi$ contains two terms. The first of these terms accounts for the asymmetry in the arrival rates of market orders on the sell and buy sides and induces mean reversion to an optimal inventory level which is not necessarily zero.

The expression for optimal control simplifies considerably when (i) there are no news events (so that $\tilde{\mu}^+ = \tilde{\mu}^- = 0$); (ii) the impact of influential orders on the stochastic drift is symmetric in the sense that $\varepsilon^+ = \mathbb{E}[\epsilon^+] = \mathbb{E}[\epsilon^-] = \varepsilon^- := \varepsilon$; (iii) the parametric shape of the LOB FPs are symmetric, in the sense that the class of functions $h^+$ and $h^-$ are equal;[15] and (iv) the fill probability at the risk-neutral optimal control is independent of the scale parameters,[16] i.e, $h_\pm(\delta_0^\pm, \boldsymbol{\kappa}) = const$. Under these assumptions, the two important (non-trivial) quantities which appear in the optimal spreads in Equation (16) simplify to

$$\mathbb{E}\left[\int_t^T \alpha_u \, du\right] = \varepsilon \, \frac{\rho}{\zeta} \, (\lambda_t^+ - \lambda_t^-) \left\{ \frac{1 - e^{-\widehat{\beta}(T-t)}}{\widehat{\beta}} - \frac{e^{-\zeta(T-t)} - e^{-\widehat{\beta}(T-t)}}{\widehat{\beta} - \zeta} \right\} + \alpha_t \, \frac{1 - e^{-\zeta(T-t)}}{\zeta}, \tag{17a}$$

$$b_\phi = 2 \, h \, (\lambda_t^+ - \lambda_t^-) \left\{ \frac{1}{\widehat{\beta}}(T - t) - \frac{1 - e^{-\widehat{\beta}(T-t)}}{\widehat{\beta}^2} \right\}, \tag{17b}$$

where $\widehat{\beta} = \beta - \rho(\nu - \eta)$ and $h = h_\pm(\delta_0^\pm, \boldsymbol{\kappa}) = const$. Notice both expressions above contain terms proportional to the difference in the market order activity on the buy and sell sides. If there are no influential orders, these will be equal to their long-run levels and will therefore be zero. However, when influential orders arrive, the buy and sell activities differ and the agent reacts accordingly. Moreover, the contribution of $\alpha$ accounts for the effect of the mean-reverting stochastic drift.

---

[15] This does not imply that the LOB is symmetric because the scale parameters $\kappa_t^\pm$ will differ. For example, exponential FPs $e^{-\kappa^\pm \delta^\pm}$ satisfy this requirement, even though the book may be significantly deeper on one side than the other.

[16] This condition is satisfied by (but not limited to) the exponential and power law FPs as discussed in examples 8 and 9.

## 6. High Frequency Market Making, Short-term-alpha and Directional Strategies

In this section we apply a simulation study of the HF strategy where market orders, buy and sell, are generated over a period of 5 minutes. The HFT is rapidly updating her quotes in the LOB by submitting and canceling limit orders which are filled according to exponential FPs.[17] The optimal postings are calculated using Corollary 4 and the explicit form for $b_\phi$ in Proposition 7. The processes $\boldsymbol{\lambda}_t$, $\boldsymbol{\kappa}_t$ and $\alpha_t$ are updated appropriately and the terminal cash-flows are stored to produce the profit and loss (PnL) generated from these strategies.

We note that, in practice, speed is of utmost importance for HFTs in two ways. First, speed is used to process order flow information (and news) to estimate the parameters and to predict short-term-alpha. These opportunities are short-lived and only ultra-fast traders are able to seize the opportunity. Second, once the HFT has decided what her next move is, speed is important because the orders are sent to the exchange, a notification is sent back to acknowledge that the orders were received, and this must be done in milliseconds (this is latency) otherwise other faster HFTs will get there first and the relatively slower traders might be left providing liquidity at a loss due to their orders arriving 'too late'.

To generate the PnL we assume that the final inventory is liquidated at the midprice with different transactions costs per share: 1 basis point (bp) and 10 bps.[18] In practice the HFT will bear some costs when unwinding a large quantity which could be in the form of a temporary price impact (a consequence of submitting a large market order) and by paying a fee to the exchange for taking liquidity in the form of an aggressive market order. Finally, in each simulation the process is repeated $5,000$ times to obtain the PnLs of the various strategies. More details on the simulation procedure are contained in Appendix D.

We analyze the performance of the HF market making strategy by varying the quality of the information that the HFT is able to employ when calculating the optimal postings. The main difference between our scenarios is whether the HFT is able to calculate the correct $\rho$ which, conditional on the arrival of a market order, is the probability that the trade is influential and whether they are able to estimate the correct dynamics of short-term-alpha – all of them know the equations that determine $\lambda_t^\pm$ and $\kappa_t^\pm$ but do not necessarily know the correct parameters. We contemplate six different types of HFTs:

1. Correct $\rho$. The HFT uses her superior computer power to process information to estimate $\rho$ and the other parameters that determine the dynamics of $\lambda_t^\pm$ and $\kappa_t^\pm$. Furthermore, we assume that the HFT may or may not be able to estimate the correct $\alpha_t$ dynamics.

   (a) Correct $\alpha$ dynamics. This is our **benchmark** because we also assume that the HFT is

[17] The results for power FPs are very similar so in the interest of space we do not show them.

[18] The transaction costs are computed on a percentage basis and since the starting midprice in the simulations is \$100, these correspond to approximately 1 and 10 cents per share, respectively. In particular, $q_T$ shares are liquidated at a value of $q_T(S_T - c_{trans}\text{sng}q_T)$.

able to estimate the parameters of the $\alpha_t$ process and adjust her postings accordingly.

(b) Zero $\alpha$ dynamics. Here we assume that although the HFT is able to estimate the correct $\rho$ she assumes that short-term alpha is zero throughout the entire strategy.

2. High $\rho$. At the other extreme we also have an HFT who cannot distinguish between the type of market order and assumes that all orders are influential, $\rho = 1$. The jump sizes in $\lambda_t^{\pm}$ and $\kappa_t^{\pm}$ are set so that the long-run means are $\lambda_t^{\pm} = m_t^{\pm}(\infty)$ and $\kappa_t^{\pm} = \tilde{m}_t^{\pm}(\infty)$.

   (a) Incorrect $\alpha$ dynamics. Because the HFT assumes that all orders are influential she is not able to correctly predict short-term-alpha – she either overestimates or underestimates the effect that market orders have on short-term-alpha because every time there is an incoming market order the HFT will predict a jump in $\alpha_t$.

   (b) Zero $\alpha$ dynamics. The HFT assumes that short-term-alpha is always zero.

3. Low $\rho$. At one extreme we have an HFT who cannot distinguish between order type and assumes that all orders are non-influential, $\rho = 0$, and assumes that $\lambda_t^{\pm}$, $\kappa_t^{\pm}$ are constant and set at their long-run means: $\lambda_t^{\pm} = m_t^{\pm}(\infty)$, given in Lemma 1, and $\kappa_t^{\pm} = \tilde{m}_t^{\pm}(\infty)$, given in Lemma 10.

   (a) Incorrect $\alpha$ dynamics. Because the HFT assumes that all orders are non-influential she is not able to correctly predict short-term-alpha – she only observes the diffusion components and not the jumps.

   (b) Zero $\alpha$ dynamics. The HFT assumes that short-term-alpha is always zero.

In all six cases, the data generating process is the true process. We assume that news does not arrive during the simulation and assume the following values for the parameters (unless otherwise stated): $\beta = 60$ and $\theta = 1$ (speed and level of mean reversion of intensity of market order arrivals), $\eta = 40$ and $\nu = 10$ (jumps in $\boldsymbol{\lambda}_t$ upon the arrival of influential market orders); $\beta_\kappa = 10$ and $\theta_\kappa = 50$ (speed and level of mean reversion for the $\boldsymbol{\kappa}_t$ process), $\eta_\kappa = 10$ and $\nu_\kappa = 25$ (jumps in $\boldsymbol{\kappa}_t$ upon the arrival of influential market orders); $\upsilon = 0$ (long-term component of the drift of the midprice), $\sigma = 0.01$ (volatility of diffusion component of the midprice), $\zeta = 2$ and $\sigma_\alpha = 0.01$ (speed of mean reversion and volatility of diffusion component of $\alpha_t$ process); and finally $\rho = 0.7$ (probability of the market order being influential). Moreover, $\epsilon^{\pm}$ are both exponentially distributed with the same mean, $\mathbb{E}[\epsilon^{\pm}] = \varepsilon$, for the sell and buy impacts. In the simulations we consider two cases: $\varepsilon = 0.04$ and $\varepsilon = 0.02$.

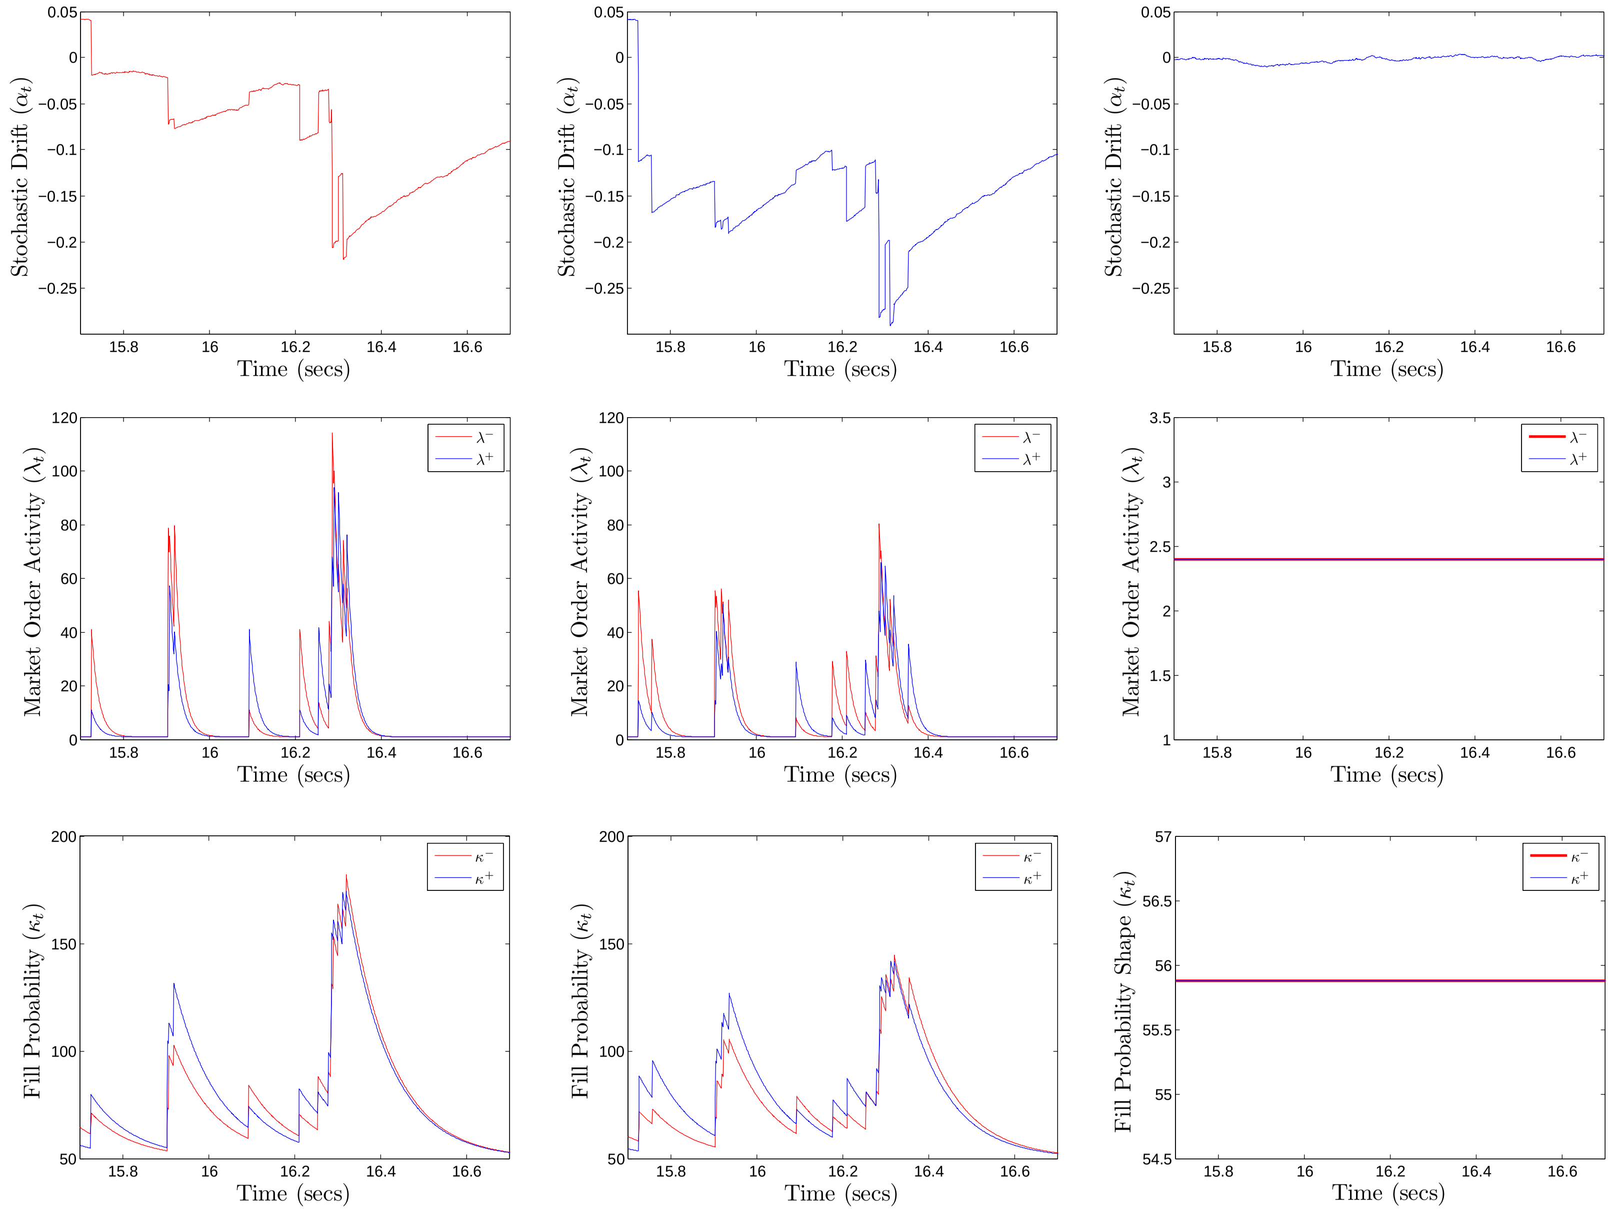


(a) Bench. Model

(b) $\rho = 1$ Model

(c) $\rho = 0$ Model

Figure 5: A sample path of the traders interpretation of the stochastic drift, market order activity, and fill probability parameter. The benchmark model has the correct data generating path, while the others under- or over-estimate the role of trade arrivals.

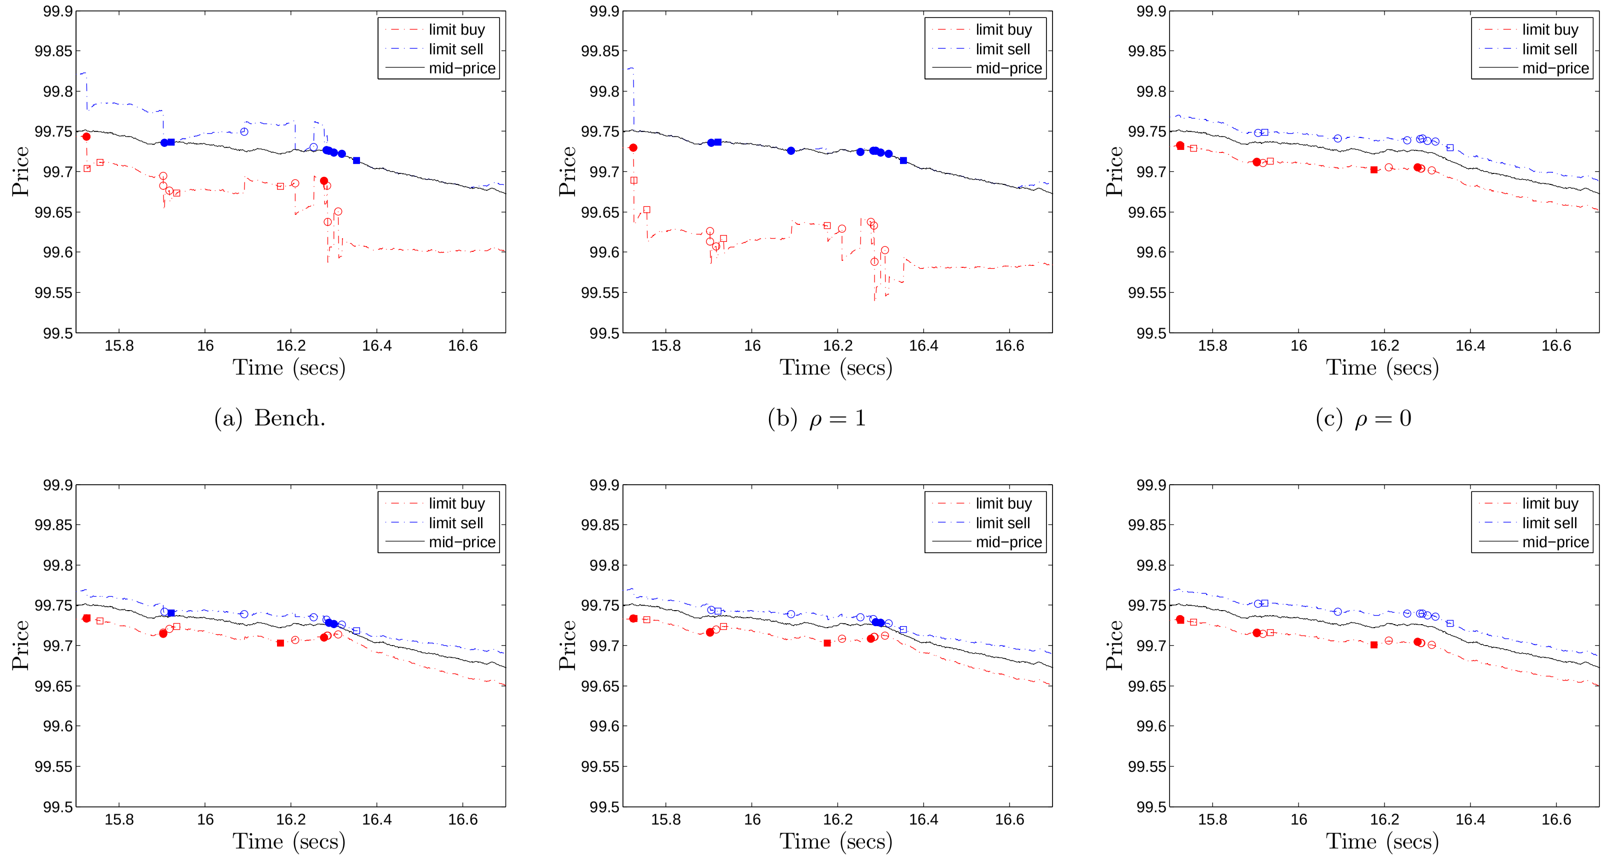


(a) Bench.

(b) $\rho = 1$

(c) $\rho = 0$

(d) Becnh., zero $\alpha$

(e) $\rho = 1$, zero $\alpha$

(f) $\rho = 0$, zero $\alpha$

Figure 6: A sample path generated by the strategy under an exponential fill rate. When influential trades arrive, the activity of both buy and sell orders increase but by differing amounts, the FP parameters also jump, and both decay to their long-run levels. Circles indicate the arrival of an influential market order, while squares indicate the arrival of non-influential trades. Open symbols indicate those market orders which do not fill the traders limit orders, while filled in symbols indicate those market orders which do fill the traders limit orders.

The first column in Figure 5 shows the information that the benchmark HFT employs to calculate her optimal strategy. The top picture shows the dynamics of $\alpha_t$ over approximately a two-second window. We see that for this particular time interval of the simulation, $\alpha_t$ is most of the time negative and every time an influential order arrives the stochastic drift jumps up (buy market order) or down (sell market order) by a random amount which is drawn from an exponential distribution with mean $\varepsilon = 0.04$. In the same column we also show the dynamics of the market order activity $\lambda^{\pm}$ and the FP parameter $\kappa_t^{\pm}$.

In the second column of the same figure we show the information employed by the HFT who incorrectly assumes that all market orders are influential (i.e. assumes $\rho = 1$) when the true parameter is $\rho = 0.7$. In the top picture we observe that by overstating the arrival rate of market orders this HFT also overestimates the impact that market orders have on short-term-alpha. The other two pictures in the column show the HFT's estimate of the arrival rate of market orders and fill rates. Obviously, when compared to those used by the benchmark HFT these are incorrect estimates of the true $\lambda_t^{\pm}, \kappa_t^{\pm}$ because the HFT sets the jump sizes in $\lambda_t^{\pm}$ and $\kappa_t^{\pm}$ so that the long-run means $\lambda_t^{\pm} = m_t^{\pm}(\infty)$ and $\kappa_t^{\pm} = \tilde{m}_t^{\pm}(\infty)$.

Finally, the last column shows the information used by the HFT who incorrectly assumes that all orders are non-influential $\rho = 0$. From the pictures it is clear that this HFT has a poor estimate of $\alpha_t\,, \lambda_t^{\pm}\,, \kappa_t^{\pm}$.

Figure 6 shows how the HFTs post and cancel limit orders during the same two-second window discussed in Figure 5 and the inventory-management parameter is $\phi = 10^{-5}$. The top row of Figure 6, pictures (a), (b), and (c), shows the optimal postings for: the benchmark HFT with correct $\rho = 0.7$, the HFT with $\rho = 1$, and the HFT with $\rho = 0$ respectively, all of which employ the information shown in Figure 5. The bottom row of Figure 6 shows the postings of the HFTs with the correct and incorrect $\rho$s, where all of these assume that $\alpha_t = 0$ and $\varepsilon = 0$ along the entire strategy.

To understand the intuition behind the optimal postings of the benchmark HFT, let us focus on picture (a) of Figure 6. The solid line shows the midprice and the dash-dot lines show the buy and sell limit orders. Circles denote the arrival of influential market orders and squares the arrival of non-influential orders. When the circles and squares are colored in, it shows that the market order was filled by the benchmark HFT. Otherwise, when the circles and squares are not colored in, it represents market orders that arrived but were filled by other more competitive resting orders in the LOB. One can observe that a key driver of the optimal postings is $\alpha_t$. At the beginning of the window (see (a) in Figure 6) we see that an influential sell market order arrived (hitting the buy side of the book), thus short-term-alpha drops (see first picture in Figure 5) and the benchmark HFT cancels her existing limit orders and reposts limit orders that see an increase in the halfspread on the buy side and a decrease in the halfspread on the sell side. These changes in the halfspread are due to adverse selection and directional strategies. Because the benchmark HFT observes a very low $\alpha_t$ at the beginning of the window, she knows that negative short-term-alpha over a very short time interval (recall that $\alpha_t$ is quickly mean reverting to zero) indicates that if she does not adjust her buy halfspread upward (i.e. adjust the buy quote downward) she will be picked off by traders that might have anticipated a decline in the midprice (or picked off by a noise trader that sends a sell market order) – thus the benchmark HFT's strategy is to avoid buying the asset right before the price drops. Similarly, the downward adjustment in the sell spread (i.e. adjust the sell quote downward) is part of a directional strategy whereby the benchmark HFT wants to sell the asset before its price drops and then purchase it back at a lower price.[19]

### *6.1. Profit and Loss from High Frequency Market Making*

Here we show the Profit and Losses (PnL) that the HFTs face when executing the optimal strategy. We report the results in Tables 1 and 2 where the difference between the two tables is the impact that influential orders have ($\varepsilon = 0.04$ and $\varepsilon = 0.02$) on short-term-alpha. In both tables terminal inventories $q_T$ are liquidated at the midprice $S_T$ and pick up a penalty of 1bps and 10bps per share. The tables show the results for different values of the inventory-management parameter $\phi = \{1 \times 10^{-5}, 2 \times 10^{-5}, 4 \times 10^{-5}\}$. For each value of $\phi$ we show the mean and standard deviation of the six PnLs where the top row, for each $\phi$, reports the three PnLs resulting from: the benchmark HFT (who uses the correct $\rho = 0.7$), and the other two HFTs who incorrectly specify the arrival of influential and non-influential market orders. For each $\phi$ the bottom row shows the other three PnLs that result from assuming that the HFTs set $\alpha_t = 0$ throughout the entire strategy.

Case I: $\varepsilon = 0.04$, $\rho = 0.7$ and liquidation costs = 1bp

| $\phi$ | $\alpha_t$ | | Bench. | $\rho = 1$ | $\rho = 0$ |
|---|---|---|---|---|---|
| $1 \times 10^{-5}$ | Yes | mean | 14.16 | 12.97 | 9.96 |
| | | (std.) | (6.84) | (7.31) | (5.88) |
| | No | mean | -3.82 | -3.86 | -4.35 |
| | | (std.) | (2.85) | (2.85) | (2.99) |
| $2 \times 10^{-5}$ | Yes | mean | 13.42 | 12.36 | 9.68 |
| | | (std.) | (5.35) | (5.64) | (4.75) |
| | No | mean | -1.68 | -1.79 | -2.82 |
| | | (std.) | (2.09) | (2.08) | (2.19) |
| $4 \times 10^{-5}$ | Yes | mean | 12.27 | 11.34 | 9.19 |
| | | (std.) | (4.23) | (4.41) | (3.86) |
| | No | mean | 0.07 | -0.06 | -1.27 |
| | | (std.) | (1.50) | (1.49) | (1.60) |

Case II: $\varepsilon = 0.04$, $\rho = 0.7$ and liquidation costs = 10bp

| $\phi$ | $\alpha_t$ | | Bench. | $\rho = 1$ | $\rho = 0$ |
|---|---|---|---|---|---|
| $1 \times 10^{-5}$ | Yes | mean | 13.39 | 12.18 | 9.25 |
| | | (std) | (6.68) | (7.16) | (5.80) |
| | No | mean | -4.32 | -4.35 | -4.84 |
| | | std | (3.01) | (3.00) | (3.14) |
| $2 \times 10^{-5}$ | Yes | mean | 12.79 | 11.70 | 9.09 |
| | | (std) | (5.22) | (5.51) | (4.68) |
| | No | mean | -2.08 | -2.18 | -3.21 |
| | | std | (2.22) | (2.21) | (2.31) |
| $4 \times 10^{-5}$ | Yes | mean | 11.74 | 10.79 | 8.71 |
| | | (std) | (4.13) | (4.31) | (3.80) |
| | No | mean | -0.26 | -0.38 | -1.60 |
| | | std | (1.60) | (1.60) | (1.70) |

Table 1: The mean and standard deviation of the PnL from the various strategies as the inventory-management parameter $\phi$ increases, $\varepsilon = 0.04$, and final inventory liquidation costs are 1bps and 10bps per share. Recall that only the benchmark HFT, who uses $\rho = 0.7$, is able to correctly specify the dynamics of short-term-alpha.

The tables clearly show that market making is more profitable if the HFTs incorporate in their optimal strategies predictions of short-term-alpha – this is true even if the HFTs incorrectly specify the short-term-alpha parameters. Moreover, when the mean impact of influential orders on $\alpha_t$ is $\varepsilon = 0.04$, Table 1 clearly shows that HFTs who are not able to execute market making strategies based on predictable trends in the midprice will be driven out of the market because their trades are being adversely selected and because they are unable to profit from directional strategies – HFTs

[19] The drop in the sell halfspread is also due to inventory management. If the HFT holds a large inventory she would prefer to sell some of it because of the expected decline in the value of the inventory as a consequence of a drop in the midprice.

| Case III $\varepsilon = 0.02$, $\rho = 0.7$ and liquidation costs = 1bp | | | | | |
|---|---|---|---|---|---|
| $\phi$ | | | Bench. | $\rho = 1$ | $\rho = 0$ |
| $1 \times 10^{-5}$ | Yes | mean | 8.79 | 8.18 | 7.30 |
| | | (std.) | (2.81) | (2.92) | (2.58) |
| | No | mean | 1.65 | 1.61 | 1.21 |
| | | (std.) | (1.47) | (1.47) | (1.54) |
| $2 \times 10^{-5}$ | Yes | mean | 8.20 | 7.66 | 6.95 |
| | | (std.) | (2.20) | (2.27) | (2.04) |
| | No | mean | 2.50 | 2.43 | 1.79 |
| | | (std.) | (1.09) | (1.09) | (1.14) |
| $4 \times 10^{-5}$ | Yes | mean | 7.27 | 6.82 | 6.39 |
| | | (std.) | (1.71) | (1.77) | (1.62) |
| | No | mean | 3.00 | 2.93 | 2.28 |
| | | (std.) | (0.81) | (0.80) | (0.85) |

| Case IV $\varepsilon = 0.02$, $\rho = 0.7$ and liquidation costs = 10bps | | | | | |
|---|---|---|---|---|---|
| $\phi$ | | | Bench. | $\rho = 1$ | $\rho = 0$ |
| $1 \times 10^{-5}$ | Yes | mean | 8.16 | 7.54 | 6.68 |
| | | (std) | (2.73) | (2.84) | (2.55) |
| | No | mean | 1.17 | 1.14 | 0.74 |
| | | std | (1.63) | (1.63) | (1.69) |
| $2 \times 10^{-5}$ | Yes | mean | 7.68 | 7.12 | 6.44 |
| | | (std) | (2.14) | (2.21) | (2.02) |
| | No | mean | 2.10 | 2.04 | 1.40 |
| | | std | (1.23) | (1.23) | (1.27) |
| $4 \times 10^{-5}$ | Yes | mean | 6.84 | 6.38 | 5.97 |
| | | (std) | (1.67) | (1.72) | (1.60) |
| | No | mean | 2.67 | 2.60 | 1.95 |
| | | std | (0.92) | (0.92) | (0.96) |

Table 2: The mean and standard deviation of the PnL from the various strategies as the inventory-management parameter $\phi$ increases, $\varepsilon = 0.02$, and final inventory liquidation costs are 1bps and 10bps per share. Recall that only the benchmark HFT, who uses $\rho = 0.7$, is able to correctly specify the dynamics of short-term-alpha.

who omit short-term-alpha face negative, or at best close to zero, mean PnLs. Table 2 shows that if the mean impact of influential orders decreases to $\varepsilon = 0.02$, HFTs are able to subsist even if they do not use predictors of short-term-alpha when making markets; however we believe that in practice HFTs will not survive if they are not able to trade on short-term-alpha to profit from directional strategies and to reduce the effects of adverse selection.[20]

The inventory-management parameter $\phi$ plays an important role in the performance of the HFT strategies. Although the HFTs are maximizing expected terminal wealth (and not expected utility of terminal wealth), they are capital constrained and their own internal risk-measures require them to penalize building large positions. HFTs that wish to, or are required to, exert a tight control on their exposure to inventories will prefer a high $\phi$. Tables 1 and 2 show an interesting effect of $\phi$ on the PnL of the different strategies that we study. If the HFT uses her predictions of short-term-alpha to make markets, increasing $\phi$ reduces both the mean and standard deviation of the PnL. Thus, in these cases the tradeoff between mean and standard deviation of profits is clear: those HFTs who trade on short-term-alpha are able to trade off mean against standard deviation of PnL.

On the other hand, the effect of increasing $\phi$ on the PnL of HFTs that do not take into account short-term-alpha is to increase the mean and to decrease the standard deviation of the PnL. The intuition behind this result is the following. As we have shown, HFTs that do not trade using predictions of short-term-alpha suffer from being picked off by better informed traders and are unable to boost their profits using directional strategies. However, increasing $\phi$ makes their postings more conservative because, everything else equal, the limit orders are posted deeper in the LOB and this makes it more difficult for other traders to pick off their quotes. Thus, by increasing $\phi$ the HFT reduces her exposure to adverse selection and this explains why the mean PnL increases in $\phi$. Finally, the standard deviation of the PnL decreases because when $\phi$ increases the strategy induces very quick mean reversion of inventories to zero.

[20] We are grateful to an anonymous referee for pointing this out.

Finally, we repeat the simulations by assuming that influential orders arrive with probability $\rho = 0.3$. Table 3 shows the results when we assume that $\varepsilon = 0.04$ and $\varepsilon = 0.02$ and that final inventory liquidation costs are 1bp. The results are qualitatively the same as those discussed above. We also run simulations with different parameter choices and the benchmark HFT always performs better than the other HFTs.

| $\varepsilon = 0.04$, $\rho = 0.3$ and liquidation costs = 1bp | | | | | |
|---|---|---|---|---|---|
| $\phi$ | | | Bench. | $\rho = 1$ | $\rho = 0$ |
| $1 \times 10^{-5}$ | Yes | mean | 4.89 | 3.65 | 3.79 |
| | | (std) | (1.39) | (1.41) | (1.33) |
| | No | mean | 2.61 | 2.63 | 2.59 |
| | | std | (1.24) | (1.24) | (1.25) |
| $2 \times 10^{-5}$ | Yes | mean | 4.71 | 3.52 | 3.69 |
| | | (std) | (1.09) | (1.10) | (1.02) |
| | No | mean | 2.77 | 2.76 | 2.69 |
| | | std | (0.92) | (0.92) | (0.94) |
| $4 \times 10^{-5}$ | Yes | mean | 4.37 | 3.27 | 3.51 |
| | | (std) | (0.86) | (0.86) | (0.78) |
| | No | mean | 2.86 | 2.81 | 2.73 |
| | | std | (0.69) | (0.69) | (0.70) |

| $\varepsilon = 0.02$, $\rho = 0.3$ and liquidation costs = 1bp | | | | | |
|---|---|---|---|---|---|
| $\phi$ | | | Bench. | $\rho = 1$ | $\rho = 0$ |
| $1 \times 10^{-5}$ | Yes | mean | 4.67 | 4.07 | 4.24 |
| | | (std) | (0.81) | (0.79) | (0.77) |
| | No | mean | 3.88 | 3.88 | 3.85 |
| | | std | (0.77) | (0.77) | (0.78) |
| $2 \times 10^{-5}$ | Yes | mean | 4.48 | 3.91 | 4.10 |
| | | (std) | (0.65) | (0.63) | (0.60) |
| | No | mean | 3.85 | 3.84 | 3.79 |
| | | std | (0.60) | (0.60) | (0.60) |
| $4 \times 10^{-5}$ | Yes | mean | 4.16 | 3.61 | 3.85 |
| | | (std) | (0.53) | (0.52) | (0.49) |
| | No | mean | 3.71 | 3.67 | 3.63 |
| | | std | (0.48) | (0.48) | (0.48) |

Table 3: The mean and standard deviation of the PnL from the various strategies as the inventory-management parameter $\phi$ increases, $\varepsilon = 0.04$ and 0.02, and final inventory liquidation costs are 1bps per share. Recall that only the benchmark HFT, who uses $\rho = 0.3$, is able to correctly specify the dynamics of short-term-alpha.

## 7. Conclusions

We develop an HF trading strategy where the HFT uses her superior speed advantage to process information and to send orders to the LOB to profit from roundtrip trades over very short-time scales. One of our contributions is to differentiate market orders between influential and non-influential. The arrival of influential market orders increases market order activity and also affects the shape and dynamics of the LOB. On the other hand, when non-influential market orders arrive they eat into the LOB but have no effect on the demand or supply of shares in the market.

Another contribution is to model short-term-alpha in the drift of the midprice as a zero-mean reverting process which jumps by a random amount upon the arrival of influential market orders and news. Influential buy and sell market orders induce a short-lived upward and downward trend in the midprice of the asset (good and bad news have a similar effect). This specification allows us to capture the essence of HF trading: to exploit short-lived predictable opportunities by the way of directional strategies, and to supply liquidity to the market taking into account adverse selection costs.

The trading strategy that the HFT employs is given by the solution of an optimal control problem where the trader is constantly submitting and canceling limit orders to maximize expected terminal wealth, whilst managing inventories, over a short time interval $T$. The strategy shows how to optimally post (and cancel) buy and sell orders and is continuously updated to incorporate information of the arrival of market orders, news (good, bad and ambiguous), size and sign of inventories, and short-term-alpha. The optimal strategy captures many of the key characteristics that differentiate HFTs from other algorithmic traders: profit from directional strategies based on predicting short-term-alpha; reduce exposure to limit orders being picked off by better informed traders; and strong mean reversion of inventories to an optimal level throughout the entire strategy and to zero at the terminal date.

Our framework allows us to derive asymptotic solutions of the optimal control problem under very general assumptions of the dynamics of the LOB. We test our model using simulations where we assume different types of HFTs who are mainly characterized by the quality of the information that they are able to process and incorporate into their optimal postings. We show that only those HFTs who incorporate predictions of short-term price deviations in their strategy will deliver expected positive profits. The other HFTs are driven out of the market because their limit orders are picked off by better informed traders and because they cannot profit from directional strategies which are also based on short-lived predictable trends. We also show that those HFTs who cannot execute profitable directional strategies and are systematically being picked off can stay in business if they exert tight controls on their inventories. In our model, these controls imply a higher penalty on their inventory position which pushes the optimal limit orders further away from the midprice so the chances of being picked off by other traders are considerably reduced.

One aspect that we have left unmodeled is when is it optimal for the HFT to submit market orders. We know that HFTs submit both aggressive and passive orders. Depending on short-term-alpha it might be optimal for the HFT to submit aggressive orders (for one or both legs of the trade) to complete a directional strategy. In our stochastic optimal control problem the HFT does not execute market orders, the best she can do is send limit orders at the midprice (zero spread) but this is no guarantee that the limit order will be filled in time for the HF strategy to be profitable. We leave for future research the optimal control problem where HFTs can submit both passive and aggressive orders.

Finally, the self-exciting nature of our model captures other important features of strategic behavior which include 'market manipulation'. For example, algorithms could be designed to send market orders, in the hope of being perceived as influential, to trigger other algorithms into action and then profit from anticipating the temporary changes in the LOB and short-term-alpha. Market manipulation strategies are not new to the marketplace, they have been used by some market participants for decades, perhaps what has changed is the speed at which these techniques are executed and the question is whether speed enhances the ability to go undetected. Analyzing such strategies are beyond the scope of this paper.

**Acknowledgements**

SJ and JR would like to thank the Natural Sciences and Engineering Research Council and Mprime for partially funding this work. JR would additionally like to thank the Ontario Graduate Scholarship for funding support. The authors also thank participants at the SAFI meeting at U. Michigan

(May 2011), ICIAM2011 meeting in Vancouver (July 2011), 3C Risk Forum at the Fields Institute for Mathematical Sciences, Toronto (Oct 2011), Universidad Carlos III de Madrid (Oct 2011), University College London (Nov 2011) and R. Almgren, B. Dupire, D. Kramkov, U. Horst, P. Bank, M. Ludkovski, J. Penalva and S. Stoikov for helpful comments on various stages of this project. Finally, we would like to thank two anonymous referees who lead us to include the effect of stochastic drift into the model and ultimately improved the paper.

Almgren, R. (2003, January). Optimal execution with nonlinear impact functions and trading-enhanced risk. Applied Mathematical Finance 10(1), 1–18.

Almgren, R. and N. Chriss (2000). Optimal execution of portfolio transactions. Journal of Risk 3.

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

Bäuerle, N. and U. Rieder (2009). MDP algorithms for portfolio optimization problems in pure jump markets. Finance and Stochastics 13, 591–611.

Bayraktar, E. and M. Ludkovski (2011). Liquidation in limit order books with controlled intensity. Working Paper.

Bertsekas, D. P. and S. Shreve (1978). Stochastic Optimal Control, Volume 139 of Mathematics in Science and Engineering. New York: Academic Press.

Bertsimas, D. and A. W. Lo (1998, April). Optimal control of execution costs. Journal of Financial Markets 1(1), 1–50.

Bouchard, B., N.-M. Dang, and C.-A. Lehalle (2011). Optimal control of trading algorithms: A general impulse control approach. SIAM Journal on Financial Mathematics 2, 404–438.

Cartea, Á. and S. Jaimungal (2010). Modeling asset prices for algorithmic and high frequency trading. SSRN eLibrary.

Cartea, Á. and J. Penalva (2011). Where is the value in high frequency trading? SSRN eLibrary.

Cont, R., R. Talreja, and S. Stoikov (2010). A stochastic model for order book dynamics. Operations Research 58, 217–224.

Cvitanic, J. and A. A. Kirilenko (2010). High Frequency Traders and Asset Prices. SSRN eLibrary.

Easley, D. and M. O'Hara (1992, June). Time and the process of security price adjustment. The Journal of Finance XLVII(2), 577–605.

Gârleanu, N. and L. H. Pedersen (2009). Dynamic trading with predictable returns and transaction costs. SSRN eLibrary.

Guéant, O., C.-A. Lehalle, and J. Fernandez Tapia (2011, May). Dealing with the inventory risk. Quantitative Finance Papers 1105.3115, arXiv.org.

Guilbaud, F. and H. Pham (2011). Optimal high frequency trading with limit and market orders. SSRN eLibrary.

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

Kharroubi, I. and H. Pham (2010). Optimal portfolio liquidation with execution cost and risk. SIAM Journal on Financial Mathematics 1, 897–931.

Kirilenko, A. A., A. P. S. Kyle, M. Samadi, and T. Tuzun (2010). The Flash Crash: The Impact of High Frequency Trading on an Electronic Market. SSRN eLibrary.

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

Rosu, I. (2009, November). A dynamic model of the limit order book. Review of Financial Studies 22(11), 4601–4641.

Sarkar, A. and R. A. Schwartz (2009). Market sidedness: Insights into motives for trade initiation. The Journal of Finance 64(1), 375–423.

## Appendix A. Fitting the Model

When all market orders are influential (i.e., when $\rho = 1$), the path of the intensity process is fully specified once the times at which the buy and sell trades are specified. Consequently, the likelihood can be written explicitly, and a straightforward maximum likelihood estimation (MLE) can be used (albeit it must be maximized numerically). To be specific, suppose $\{t_1, t_2, \dots, t_n\}$ are a set of observed trade times (with $t_n \le T$ the time of the last trade) and $\{B_1, B_2, \dots, B_n\}$ are buy/sell indicators, i.e., 0 if the trade is a market sell and 1 if the trade is a market buy. Then the hazard rates and their integral at an arbitrary time $t$ are

$$\lambda_t^{\pm} = \theta + \sum_{i=1}^{n} H_i^{\pm}\, e^{-\beta(t-t_i)} \quad \text{and} \quad \int_0^t \lambda_u^{\pm}\, du = \theta\, t + \sum_{i=1}^{n} H_i^{\pm}\, \frac{1-e^{-\beta(t-t_i)}}{\beta}\,, \tag{A.1}$$

where $H_i^{\pm} = (B_i\,\eta + (1-B_i)\,\nu,\ B_i\,\nu + (1-B_i)\,\eta)$. Finally, the log-likelihood

$$\mathscr{L} = -2\theta\, T - \sum_{i=1}^{n} \left\{ B_i\, \log(\lambda_{t_i}^{+}) + (1-B_i)\log(\lambda_{t_i}^{-}) + (\eta+\nu)\frac{1-e^{-\beta(T-t_i)}}{\beta} \right\}. \tag{A.2}$$

Maximizing this log-likelihood results in the MLE estimates of the model parameters, and upon back substitution into Equation (A.1) provides the estimated path of activity. Integrating this activity over the last second, i.e., $\int_{t-1}^{t} \lambda_u^{\pm}\, du$ provides us with a smoothed version of the intensity and shown in Figure 2 as the path labeled "Fitted". This is directly comparable to the one second historical intensity in Figure 2 labeled as "Historical".

For the time window 3:30pm to 4:00pm on Feb 1, 2008 for IBM the estimated parameters are as follows:

$$\widehat{\beta} = 180.05, \quad \widehat{\theta} = 2.16, \quad \widehat{\eta} = 64.16, \quad \text{and} \quad \widehat{\nu} = 55.73\,.$$

Notice that the spikes in the historical intensity are often above the fitted intensities. The reason for this difference is that, here, the fitted intensities assume that all trades are influential (i.e., $\rho = 1$). Consequently, the size of the jump in intensities must be smaller than the true jump size to preserve total mean activity of trades. When a full calibration is carried out – in which $\rho$ is not necessarily 1 and the influential/non-influential nature of the event must be filtered – the jump sizes are indeed larger. We will report on this more involved estimation procedure and filtering problem in a forthcoming paper.

## Appendix B. Proof of Results

### *Appendix B.1. Proof of Lemma 1*

Integrating both sides of (2), taking conditional expectation, applying Fubini's Theorem, and then taking derivative gives the following coupled system of coupled ODEs for $m_t^\pm(u)$

$$\frac{d}{du}\begin{pmatrix} m_t^-(u) \\ m_t^+(u) \end{pmatrix} + \begin{pmatrix} \beta-\eta\rho & -\nu\rho \\ -\nu\rho & \beta-\eta\rho \end{pmatrix}\begin{pmatrix} m_t^-(u) \\ m_t^+(u) \end{pmatrix} - \begin{pmatrix} \beta\theta+\tilde\eta\mu^- + \tilde\nu\mu^+ \\ \beta\theta+\tilde\nu\mu^- + \tilde\eta\mu^+ \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \tag{B.1}$$

with initial conditions $m_t^\pm(t) = \lambda_t^\pm$. This is a standard matrix equation and, if $\mathbf{A}$ has no zero eigenvalues, admits the unique solution

$$\begin{pmatrix} m_t^-(u) \\ m_t^+(u) \end{pmatrix} = e^{-\mathbf{A}(u-t)}\left[\begin{pmatrix} \lambda_t^- \\ \lambda_t^+ \end{pmatrix} - \mathbf{A}^{-1}\zeta\right] + \mathbf{A}^{-1}\zeta\,. \tag{B.2}$$

Since $\mathbf{A}$ is symmetric, it is diagonalizable by an orthonormal matrix $\mathbf{U}$. Furthermore, its eigenvalues are $\beta-(\eta\pm\nu)\rho$. Clearly, in the limit $u\to\infty$, $m_t(u)$ converges if and only if $\beta-(\eta\pm\nu)\rho>0$ which implies $\beta>(\eta+\nu)\rho$ since $\eta,\nu,\rho\geq 0$.

The remaining case is if $\mathbf{A}$ has at least one zero eigenvalue. However, it is easy to see that in this case, the solution to (B.1) has at least one of $m_t^\pm(u)$ growing linearly as a function of $u$. Furthermore, if one eigenvalue is zero, then either $\beta=(\eta-\nu)\rho$ or $\beta=(\eta+\nu)\rho$, which lie outside the stated the bounds. Finally, if both eigenvalues are zero, then we must have $\beta=\nu=\eta=0$. Once again outside of the stated bounds. □

### *Appendix B.2. Proof of Proposition 2*

Applying the ansatz on the form on $\Phi$, differentiating inside the supremum in (7) with respect to $\delta^\pm$, then expanding $g$ using the specified ansatz, writing $\delta^{\pm *} = \delta_0^\pm + \alpha\delta_\alpha^\pm + \varepsilon\delta_\varepsilon^\pm + \phi\delta_\phi^\pm + o(\varsigma)$, and setting the resulting equation to 0 gives our first-order optimality condition. To this order, the first-order conditions imply that

$$\begin{aligned} h_\pm(\delta_0^\pm) + \delta_0^\pm\, h_\pm'(\delta_0^\pm) + \alpha\Big\{\delta_\alpha\left(h_\pm''(\delta_0^\pm) + 2\,h_\pm'(\delta_0^\pm)\right) + h'(\delta_0^\pm)\left(\mathbb{S}_{q\lambda}^\pm g_\alpha - \mathbb{S}_\lambda^\pm g_\alpha\right)\Big\} \\ +\varepsilon\Big\{\delta_\varepsilon\left(h_\pm''(\delta_0^\pm) + 2\,h_\pm'(\delta_0^\pm)\right) + h_\pm'(\delta_0^\pm)\left(\mathbb{S}_{q\lambda}^\pm g_\varepsilon - \mathbb{S}_\lambda^\pm g_\varepsilon + \pm\rho\,\mathfrak{a}^\pm\,\mathcal{S}_{q\lambda}^\pm g_\alpha\right)\Big\} \\ +\phi\Big\{\delta_\phi\left(h_\pm''(\delta_0^\pm) + 2\,h_\pm'(\delta_0^\pm)\right) + h_\pm'(\delta_0^\pm)\left(\mathbb{S}_{q\lambda}^\pm g_\phi - \mathbb{S}_\lambda^\pm g_\phi\right)\Big\} \quad = o(\varsigma)\,. \end{aligned} \tag{B.3}$$

Observe that the Taylor expansion of $h(\delta)$ about $\delta_0$ requires the $C^3$ regularity condition to keep the error of the correct order. The $C^1$ regularity condition ensures that the global maximizer satisfies (B.3). Setting the constant term in (B.3) to zero yields (13). Setting the coefficients of $\alpha$, $\varepsilon$ and $\phi$ each separately to zero and solving for $\delta_\alpha$, $\delta_\varepsilon$ and $\delta_\phi$ results in (11). The finiteness of the optimal control correct to this order is ensured by the last condition in Assumption 2.

The existence of a solution to (13) is clear by noticing that (13) is a critical point of the function $\delta h(\delta)$. The critical point exists since $\delta h(\delta)$ is non-positive for $\delta\leq 0$, strictly positive on an open interval of the form $(0,d)$ due to $C^1$, and goes to 0 in the limit by Assumption 2.

To see that the exact values of the optimal controls are non-negative, observe that the value function is increasing in $x$. Therefore, $\Phi(t,x+\delta,\cdot) < \Phi(t,x,\cdot)$ for any $\delta<0$. Since the shift operators appearing in the argument of the supremum are linear operators, and $h(\delta;\kappa)$ is bounded above by 1 and attains this maxima at $\delta = 0$, the $\delta = 0$ strategy dominates all strategies which have $\delta < 0$. $\square$

### *Appendix B.3. Proof of Theorem 3*

Inserting the expansion for $g$ and the feedback controls (11) for $\delta$ into the HJB equation (7), and carrying out tedious but ultimately straightforward expansions, to order $\varsigma$, equation (7) reduces to

$$\begin{aligned} o(\varsigma) = \mathcal{D}g_0 &+ (\lambda^+ \delta_0^+ \, h_+(\delta_0^+) + \lambda^- \delta_0^- \, h_-(\delta_0^-)) \\ &+ \alpha \left\{ q + (\mathcal{D} - \zeta)\, g_\alpha + \lambda^+ \, h_+(\delta_0) \left[ \mathbb{S}_{q\lambda}^+ - \mathbb{S}_\lambda^+ \right] g_\alpha + \lambda^- \, h_-(\delta_0) \left[ \mathbb{S}_{q\lambda}^- - \mathbb{S}_\lambda^- \right] g_\alpha \right\} \\ &+ \varepsilon \left\{ \mathcal{D}g_\varepsilon + \lambda^+ \, h_+(\delta_0) \left( \left[ \mathbb{S}_{q\lambda}^+ - \mathbb{S}_\lambda^+ \right] g_\varepsilon + \rho \, \mathfrak{a}^+ \, \mathcal{S}_{q\lambda}^+ g_\alpha \right) + \lambda^- \, h_-(\delta_0) \left( \left[ \mathbb{S}_{q\lambda}^- - \mathbb{S}_\lambda^- \right] g_\varepsilon - \rho \, \mathfrak{a}^- \, \mathcal{S}_{q\lambda}^- g_\alpha \right) \right\} \\ &+ \phi \left\{ -q^2 + \mathcal{D}g_\phi + \lambda^+ \, h_+(\delta_0) \left[ \mathbb{S}_{q\lambda}^+ - \mathbb{S}_\lambda^+ \right] g_\phi + \lambda^- \, h_-(\delta_0) \left[ \mathbb{S}_{q\lambda}^- - \mathbb{S}_\lambda^- \right] g_\phi \right\} \,, \end{aligned} \tag{B.4}$$

where $\mathcal{D} = \partial_t + \mathcal{L}$ and the boundary conditions $g_0(T,\cdot) = g_\alpha(T,\cdot) = g_\varepsilon(T,\cdot) = g_\phi(T,\cdot) = 0$ apply. Clearly, $g_0$ is a independent of $q$ and, as seen in Proposition 2, does not affect the optimal strategy. Next, perform the following steps (i) set the coefficients of $\alpha$, $\varepsilon$ and $\phi$ to zero separately; (ii) write $g_\alpha$, $g_\varepsilon$ and $g_\phi$ as in (14); and (iii) collect powers of $q$ and set them individually to zero;[21] then one finds the following equations for the functions $b_\alpha(t)$, $b_\varepsilon(t,\boldsymbol{\lambda})$, $b_\phi(t,\boldsymbol{\lambda},\boldsymbol{\kappa})$ and $c_\phi(t)$:

$$0 = \mathcal{D}b_\alpha - \zeta \, b_\alpha + \lambda^+ \left[ \mathbb{S}_\lambda^+ - 1 \right] b_\alpha + \lambda^- \left[ \mathbb{S}_\lambda^- - 1 \right] b_\alpha + 1, \tag{B.5a}$$

$$0 = \mathcal{D}b_\varepsilon + \lambda^+ \left[ \mathbb{S}_\lambda^+ - 1 \right] b_\varepsilon + \lambda^- \left[ \mathbb{S}_\lambda^- - 1 \right] b_\varepsilon + \left\{ \rho \left( \lambda^+ \mathfrak{a}^+ - \lambda^- \mathfrak{a}^- \right) + \left( \mu^+ \tilde{\mathfrak{a}}^+ - \mu^- \tilde{\mathfrak{a}}^- \right) \right\} b_\alpha, \tag{B.5b}$$

$$0 = \mathcal{D}b_\phi + \lambda^+ \left[ \mathbb{S}_\lambda^+ - 1 \right] b_\phi + \lambda^- \left[ \mathbb{S}_\lambda^- - 1 \right] b_\phi - 2\, h(\delta_0) \, (\lambda^+ - \lambda^-) \, c_\phi, \tag{B.5c}$$

$$0 = \mathcal{D}c_\phi + \lambda^+ \left[ \mathbb{S}_\lambda^+ - 1 \right] c_\phi + \lambda^- \left[ \mathbb{S}_\lambda^- - 1 \right] c_\phi - 1\,. \tag{B.5d}$$

These equations, together with the boundary conditions that $b_\alpha(T,\cdot) = b_\varepsilon(T,\cdot) = b_\phi(T,\cdot) = c_\phi(T,\cdot) = 0$, admit, through a Feynman-Kac argument, the solutions presented in (15).

The functions $a_\alpha$, $a_\varepsilon$ and $a_\phi$ are independent of $q$ and, since the optimal spreads given in (12) contain difference operators in $q$ which vanish when the difference operators act on functions independent of $q$, do not influence the optimal strategy. $\square$

### *Appendix B.4. Proof of Corollary 4*

Applying Equation (14) for $g_\alpha$, $g_\varepsilon$ and $g_\phi$ in Theorem 3 to Equations (11) and (12) of Proposition 2 and using the fact that the $a$, $b$ and $c$ functions are all independent of $q$, after some tedious computations, $\delta^{*\pm}$ reduces to

$$\delta^{\pm *} = \delta_0^\pm + B(\delta_0^\pm; \boldsymbol{\kappa}^\pm) \left\{ \pm \alpha \, b_\alpha + \varepsilon \left( \pm \mathbb{S}_\lambda^\pm b_\varepsilon + \rho \, \mathfrak{a}^\pm \, b_\alpha \right) + \phi \left( \pm \mathbb{S}_\lambda^\pm b_\phi + (1 \mp 2q)(T - t) \right) \right\} \,.$$

Next, observing that $\pm \alpha \, b_\alpha + \varepsilon \left( \pm \mathbb{S}_\lambda^\pm b_\varepsilon + \rho \, \mathfrak{a}^\pm \, b_\alpha \right) = \pm \mathbb{S}_\lambda^\pm \left( \mathbb{E} \left[ \int_t^T \alpha_u \, du \right] \right)$, we find (16). Finally, let $\delta_{\text{opt}}^\pm$ denote the exact optimal controls. Using Proposition 2 we have that $\delta_{\text{opt}}^\pm$ is non-negative and $|\delta^* - \delta_{\text{opt}}^\pm| = o(\varsigma)$ therefore,

$$\left| \max\{\delta^{\pm *}, 0\} - \delta_{\text{opt}}^\pm \right| \leq \left| \delta^{\pm *} - \delta_{\text{opt}}^\pm \right| = o(\varsigma)\,,$$

and we are done.

[21] Note that this step is not an asymptotic expansion in $q$, rather it is exact given the prescribed expansion in the other parameters.

## Appendix C. Some Explicit Formulae

This appendix contains several explicit formulae for the optimal spreads as well as quantities that feed into the optimal spreads.

### *Appendix C.1. Exact Optimal Trading Strategy*

Although an exact optimal control is not analytically tractable in general, the feedback control form for the cases of exponential and power-law FPs can be obtained within our modeling framework.

**Proposition 5. Exact Optimal Controls for Exponential and Power-Law.** *Suppose that the scale parameter process $\kappa_t^{\pm}$ is strictly positive almost surely, more specifically, assume that $\mathbb{P}\left[\inf_{t\in[0,T]}\kappa_t^{\pm} > 0\right] = 1$.*

1. *If $h^{\pm}(\delta;\boldsymbol{\kappa}) = e^{-\kappa^{\pm}\delta}$ for $\delta > 0$, then the feedback control of the optimal trading strategy for the HJB equation (7) is given by*
$$\delta_t^{\pm} = \frac{1}{\kappa^{\pm}} - \left\{\mathbb{S}_{q\lambda}^{\pm} g - \mathbb{S}_{\lambda}^{\pm} g\right\}. \tag{C.1a}$$
2. *If $h^{\pm}(\delta;\boldsymbol{\kappa}) = (1+\kappa^{\pm}\delta)^{\alpha^{\pm}}$ for $\delta > 0$, then the feedback control of the optimal trading strategy for the HJB equation (7) is given by*
$$\delta_t^{\pm} = \frac{\alpha}{\alpha - 1}\left(\frac{1}{\kappa^{\pm}} - \left\{\mathbb{S}_{q\lambda}^{\pm} g - \mathbb{S}_{\lambda}^{\pm} g\right\}\right). \tag{C.1b}$$

*Here, the ansatz $\Phi = x + q\,s + g(t,q,\alpha,\boldsymbol{\lambda},\boldsymbol{\kappa})$ with boundary condition $g(T,\cdot) = 0$ has been applied. Furthermore, the solutions in (2) are unique.*

Proof. Applying the first order conditions to the supremum terms and using the specified ansatz leads, after some simplifications, to stated result. Uniqueness is trivial. □

### *Appendix C.2. Explicit Computation of $b_\varepsilon$*

Rather than computing $b_\varepsilon$ directly, it is more convenient to compute the expected integrated drift and then identify the appropriate terms. To this end we have the following result.

**Proposition 6. Expected Integrated Drift.** *The expected integrated drift is given by the expression*

$$\mathbb{E}\left[\int_t^T \alpha_s\,ds \,\middle|\, \boldsymbol{\lambda}_t = \boldsymbol{\lambda},\ \alpha_t = \alpha\right] = \varepsilon\, b_\epsilon(t,\boldsymbol{\lambda}) + \alpha\, b_\alpha(t) \tag{C.2}$$

*where* $\varepsilon\, b_\varepsilon(t,\boldsymbol{\lambda}) = A(t) + \boldsymbol{\lambda}\cdot\mathbf{C}(t)$ *and*

$$A(t) = \tfrac{1}{\zeta}(\mu^+\widetilde{\varepsilon}^+ - \mu^-\widetilde{\varepsilon}^-)\left((T-t) - b_\alpha(t)\right) + \boldsymbol{\chi}\cdot\mathbf{B}(t)\,, \tag{C.3a}$$

$$\mathbf{B}(t) = \frac{\rho}{\zeta}\left\{\mathbf{A}^{-1}\left((T-t)\,\mathbf{I} - \mathbf{A}^{-1}\left(\mathbf{I} - e^{-\mathbf{A}(T-t)}\right)\right) - (\mathbf{A}-\zeta\,\mathbf{I})^{-1}\left(b_\alpha(t)\,\mathbf{I} - \mathbf{A}^{-1}\left(\mathbf{I} - e^{-\mathbf{A}(T-t)}\right)\right)\right\}\boldsymbol{\varepsilon}\,, \tag{C.3b}$$

$$\mathbf{C}(t) = \frac{\rho}{\zeta}\left\{\mathbf{A}^{-1}\left(\mathbf{I} - e^{-\mathbf{A}(T-t)}\right) - (\mathbf{A}-\zeta\,\mathbf{I})^{-1}\left(e^{-\zeta(T-t)}\,\mathbf{I} - e^{-\mathbf{A}(T-t)}\right)\right\}\boldsymbol{\varepsilon}\,. \tag{C.3c}$$

*Moreover,* $\boldsymbol{\varepsilon} = (-\mathbb{E}[\epsilon^-],\, \mathbb{E}[\epsilon^+])'$, $\widetilde{\varepsilon}^+ = \mathbb{E}[\widetilde{\epsilon}^+]$, $\widetilde{\varepsilon}^- = \mathbb{E}[\widetilde{\epsilon}^-]$ *and* $\boldsymbol{\chi} = (\beta\theta + \mu^+\widetilde{\nu} + \mu^-\widetilde{\eta},\ \beta\theta + \mu^+\widetilde{\eta} + \mu^-\widetilde{\nu})'$.

Proof. Denoting $f(t,\alpha,\boldsymbol{\lambda}) = \mathbb{E}\left[\int_t^T \alpha_s\,ds \,\middle|\, \boldsymbol{\lambda}_t = \boldsymbol{\lambda},\ \alpha_t = \alpha\right]$, we have, through a Feynman-Kac theorem, that $f$ satisfies the PDE

$$(\partial_t + \mathcal{L})f + \lambda^+\left(\mathbb{S}^+_\lambda f - f\right) + \lambda^-\left(\mathbb{S}^-_\lambda f - f\right) + \alpha = 0, \tag{C.4}$$

where the infinitesimal generator of $\alpha$ and $\boldsymbol{\lambda}$ is

$$\mathcal{L} = \beta(\theta - \lambda^-)\partial_{\lambda^-} + \beta(\theta - \lambda^+)\partial_{\lambda^+} - \zeta\,\alpha\partial_\alpha + \tfrac{1}{2}\sigma^2\partial_{\alpha\alpha} + \mu^-\left(\widetilde{\mathcal{S}}^-_\lambda - 1\right) + \mu^+\left(\widetilde{\mathcal{S}}^+_\lambda - 1\right).$$

Substituting the affine ansatz $f = A(t) + \boldsymbol{\lambda}\mathbf{C}(t) + \alpha\, b_\alpha(t)$ into the PDE, subject to the boundary conditions $A(T) = \mathbf{C}(t) = 0$, leads to the system of coupled ODEs:

$$\begin{cases} \partial_t A(t) + (\mu^+\widetilde{\varepsilon}^+ - \mu^-\widetilde{\varepsilon}^-)\, b_\alpha(t) + \boldsymbol{\chi}\cdot\mathbf{C}(t) &= 0, \\ \partial_t\mathbf{C}(t) - \mathbf{A}\mathbf{C}(t) + \rho\, b_\alpha(t)\,\boldsymbol{\varepsilon} &= 0. \end{cases} \tag{C.5}$$

The solution of this coupled system is given by (C.3). The assertion that $A(t) + \boldsymbol{\lambda}\cdot\mathbf{C}(t) = \varepsilon\, b_\varepsilon$ with $b_\varepsilon$ provided in (15b) can be confirmed by (i) writing down the PDE which the function $b_\varepsilon$ satisfies, (ii) note that it admits an affine ansatz $A_\varepsilon(t) + \boldsymbol{\lambda}\cdot\mathbf{C}_\varepsilon(t)$, and (iii) the ODEs that $A_\varepsilon(t)$ and $\mathbf{C}_\varepsilon(t)$ satisfy are the same ODEs as $A(t)$ and $\mathbf{C}(t)$ with the same boundary conditions. Uniqueness then implies they are equal. □

### *Appendix C.3. Computing $b_\phi$ When Risk-Neutral Fill Probabilities are Constants*

Closed form expressions for the function $b_\phi$ can only be derived under further assumptions on the FPs $h_\pm(\delta;\boldsymbol{\kappa})$. As a motivating factor, note that both exponential and power law FPs have the property that $h_\pm(\delta_0^\pm;\boldsymbol{\kappa}) = const.$, irrespective of the dynamics on the shape parameter $\kappa^\pm$. This leads us to investigate the larger class of models for which $h_\pm(\delta_0^\pm;\boldsymbol{\kappa})$ are constant. Under these assumptions, the following proposition provides an explicit form for the function $b_\phi$.

**Proposition 7. Explicit Solution for $b_\phi(t,\boldsymbol{\lambda},\boldsymbol{\kappa})$.** *If $h_\pm(\delta_0^\pm;\boldsymbol{\kappa}) = h_\pm$ are constants $\mathbb{P}$-a.s., then the function $b_\phi(t,\boldsymbol{\lambda},\boldsymbol{\kappa})$ is independent of $\boldsymbol{\kappa}$, and is explicitly*

$$b_\phi(t,\boldsymbol{\lambda}) = 2\boldsymbol{\xi}'\left\{\left(\mathbf{A}^{-1}(T-t) - \mathbf{A}^{-2}\left(\mathbf{I} - e^{-\mathbf{A}(T-t)}\right)\right)\left[\boldsymbol{\lambda} - \mathbf{A}^{-1}\boldsymbol{\zeta}\right] + \tfrac{1}{2}(T-t)^2\,\mathbf{A}^{-1}\boldsymbol{\zeta}\right\}, \tag{C.6}$$

*where $\mathbf{I}$ is the $2\times 2$ identity matrix and $\boldsymbol{\xi} = (-h_-,\ h_+)'$.*

Proof. Note that $\mathbb{E}\left[\int_t^T \lambda_u^{\pm}\,(T-u)\,du|\mathcal{F}_t\right] = \int_t^T \mathbb{E}\left[\lambda_u^{\pm}|\mathcal{F}_t\right]\,(T-u)\,du = \int_t^T m_t^{\pm}(u)\,(T-u)\,du$. Using the form of $m_t^{\pm}(u)$ provided in (B.2) and integrating over $u$ implies that

$$\int_t^T \boldsymbol{m}_t(u)(T-u)\,du = \left(\mathbf{A}^{-1}(T-t) - \mathbf{A}^{-2}\left(\mathbf{I} - e^{-\mathbf{A}(T-t)}\right)\right)\left(\boldsymbol{\lambda}_t - \mathbf{A}^{-1}\boldsymbol{\zeta}\right) + \mathbf{A}^{-1}\boldsymbol{\zeta}\,\tfrac{1}{2}(T-t)^2 . \tag{C.7}$$

This result is valid under the restriction that $\mathbf{A}$ is invertible, which is implied by the arrival rate of market orders (2) and Lemma 1. Moreover, when $h_{\pm}(\delta_0^{\pm};\boldsymbol{\kappa}) = h_{\pm}$ we have $b_\phi(t,\boldsymbol{\lambda}) = 2\int\limits_t^T \{h_+ \cdot m_t^+(u) - h_- \cdot m_t^-(u)\}(T-u)\,du$ and (C.6) follows immediately. ☐

As already mentioned, studying the class of models for which $h_{\pm}(\delta_0^{\pm};\boldsymbol{\kappa}) = h_{\pm}$ are constant was motivated by the exponential and power-law cases which we formalize in the two examples below.

**Example 8. Exponential Fill Rate.** *Take $\kappa^{\pm} = f^{\pm}(\boldsymbol{\kappa})$, where $f^{\pm} : \mathbb{R}^k \mapsto \mathbb{R}^+$ are continuous functions. If $h^{\pm}(\delta;\boldsymbol{\kappa}) = e^{-\kappa^{\pm}\delta}$ for $\delta > 0$ and $\mathbb{P}\left[\inf\limits_{t\in[0,T]} \kappa_t^{\pm} > 0\right] = 1$, then $h_{\pm}(\delta_0^{\pm};\boldsymbol{\kappa}) = e^{-1}$ is constant and Proposition 7 applies.*

**Example 9. Power Fill Rate.** *Take $\kappa^{\pm} = f^{\pm}(\boldsymbol{\kappa})$, where $f^{\pm} : \mathbb{R}^k \mapsto \mathbb{R}^+$ are continuous functions, and $\alpha^{\pm} > 1$ as fixed constants. If $h_{\pm}(\delta;\boldsymbol{\kappa}) = \left[1 + (\kappa^{\pm}\delta)^{\alpha^{\pm}}\right]^{-1}$ for $\delta > 0$ and $\mathbb{P}\left[\inf\limits_{t\in[0,T]} \kappa_t^{\pm} > 0\right] = 1$, then $\delta_0^{\pm} = (\alpha^{\pm} - 1)^{-\frac{1}{\alpha^{\pm}}}(\kappa^{\pm})^{-1}$ and $h_{\pm}(\delta_0^{\pm};\boldsymbol{\kappa}) = \frac{\alpha^{\pm}-1}{\alpha^{\pm}}$ is constant and Proposition 7 applies.*

Notice that the Poisson model of trade arrivals can be recovered by setting $\rho = 0$. Furthermore, if the initial states $\lambda_0^{\pm}$ are equal and $\kappa^{\pm}$ are equal then $b_\phi \equiv 0$.

### *Appendix C.4. Conditional Mean of Fill Probability process*

**Lemma 10. Conditional Mean of $\boldsymbol{\kappa}_t$.** *Under the dynamics given in* (3), *the conditional mean $\tilde{m}_t^{\pm}(u) := \mathbb{E}[\kappa_u^{\pm}|\mathcal{F}_t]$ is*

$$\tilde{m}_t^{\pm}(u) = \theta_\kappa + \frac{\rho}{\beta_\kappa}\left[\eta_\kappa m_t^{\pm}(u) + \nu_\kappa m_t^{\mp}(u)\right] + \left[\kappa_t^{\pm} - \theta_\kappa - \frac{\rho}{\beta_\kappa}(\eta_\kappa \lambda_t^{\pm} + \nu_\kappa \lambda_t^{\mp})\right] e^{-\beta_\kappa(u-t)} \tag{C.8}$$

*where $m_t^{\pm}(u)$ are given in Appendix B.1, and $\mathbf{A}$, $\boldsymbol{\zeta}$ are given in Lemma 1.*

Proof. Proceeding as in the proof of Lemma 1 in Appendix B.1, $\tilde{m}_t^{\pm}(u)$ satisfies the (uncoupled) system of ODEs

$$\frac{d\tilde{m}_t^{\pm}(u)}{du} + \beta_\kappa \tilde{m}_t^{\pm}(u) = \beta_\kappa\theta_\kappa + \rho\left[\eta_\kappa m_t^{\pm}(u) + \nu_\kappa m_t^{\mp}(u)\right] \tag{C.9}$$

where $m_t^{\pm}(u)$ is given by (B.2). Solving (C.9) with the boundary condition $\tilde{m}_t^{\pm}(t) = \kappa_t^{\pm}$ gives the stated result. ☐

## Appendix D. Simulation Procedure

Here we describe in more detail the approach to simulating the PnL distribution of the HF strategy. Note that this produces an exact simulation – specifically, there are no discretization errors.

1. Generate the duration until the next market order given the current level of activity $\lambda^{\pm}_{t_n}$.
   - In between orders, the total rate of order arrival is $\lambda_t = 2\theta + (\lambda^+_{t_n} + \lambda^-_{t_n} - 2\theta)e^{-\beta(t-t_n)}$. To obtain a random draw of the time of the next trade, draw a uniform $u \sim U(0,1)$ and find the root[22] of the equation $\tau e^{\tau} = \frac{1}{2\theta}(\lambda_{t_n} - 2\theta)\, e^{\varsigma}$ where $\varsigma = \frac{\lambda_{t_n}-2\theta}{2\theta} + \frac{\beta}{2\theta}\ln u$. Then, $T_{n+1} = \frac{1}{\beta}(\tau - \varsigma)$ is a sample for the next duration and $t_{n+1} = t_n + T_{n+1}$.
2. Decide if the trade is a buy or sell market order.
   - The probability that the market order is a buy order is $p_{buy} = \frac{\theta + (\lambda^+_{t_n} - \theta)\, e^{-\beta\, T_{n+1}}}{2\theta + (\lambda^+_{t_n} + \lambda^-_{t_n} - 2\theta)\, e^{-\beta\, T_{n+1}}}$. Therefore, draw a uniform $u \sim U(0,1)$ and if $u < p_{buy}$ the order is a buy order, otherwise it is a sell order.
   - Set the buy/sell indicator $B_{n+1} = 1$ if it is a buy market order and $B_{n+1} = -1$ if it is a sell market order.
3. Decide whether the market order filled the agent's posted limit order.
   - Compute the posted limit order at the time of the market order $\lambda^{\pm}_{t_{n+1}} = \theta + (\lambda^{\pm}_{t_n} - \theta)\, e^{-\beta\, T_{n+1}}$.
   - Draw a uniform $u \sim U(0,1)$.
   - If the market order was a sell (buy) order, then if $u < e^{-\kappa^-_t\, \delta^-_t}$ ($u < e^{-\kappa^+_t\, \delta^+_t}$) the agent's buy (sell) limit order was lifted (hit).
4. Update the midprice and drift of the asset.
   - Generate two correlated normals $Z_1$ and $Z_2$ with zero mean and covariances:

     $\mathbb{C}(Z_1, Z_1) = \frac{\sigma^2}{\zeta^2}\left(T_{n+1} - 2\frac{1-e^{-\zeta T_{n+1}}}{\zeta} + \frac{1-e^{-2\zeta T_{n+1}}}{2\zeta}\right)$, $\mathbb{C}(Z_2, Z_2) = \frac{\sigma^2}{2\zeta}(1 - e^{-2\zeta T_{n+1}})$, and $\mathbb{C}(Z_1, Z_2) = \frac{\sigma^2}{2\zeta^2}\left(1 - 2e^{-\zeta T_{n+1}} + e^{-2\zeta T_{n+1}}\right)$.

     Generate a third independent standard normal $Z$.
   - Update price and drift. $S_{t_{n+1}} = S_{t_n} + \alpha_{t_n}\frac{1}{\zeta}(1 - e^{-\zeta_{n+1}}) + Z_1 + \sigma\sqrt{T_{n+1}}\, Z$ and $\alpha_{t_{n+1}} = e^{-\zeta T_{n+1}}\alpha_{t_n} + Z_2$.
5. Update the inventory and agent's cash: $X_{t_{n+1}} = X_{t_n} + B_{n+1}S_{t_{n+1}} + \delta^+_{t_{n+1}}$ and $q_{t_{n+1}} = q_{t_n} - B_{n+1}$.
6. Decide if the trade is influential and update activities, FPs and drift.
   - Draw a uniform $u \sim (0,1)$, if $u < \rho$ the trade is influential, set $H_{n+1} = 1$, otherwise set $H_{n+1} = 0$. Finally,

$$
\begin{aligned}
\lambda^{\pm}_{t_{n+1}} &= \theta + (\lambda^{\pm}_{t_n} - \theta)\, e^{-\beta\, T_{n+1}} + \left(\tfrac{1}{2}(1 \pm B_{n+1})\nu + \tfrac{1}{2}(1 \mp B_{n+1})\eta\right) H_{n+1}\,, \\
\kappa^{\pm}_{t_{n+1}} &= \theta_\kappa + (\kappa^{\pm}_{t_n} - \theta_\kappa)\, e^{-\beta_\kappa\, T_{n+1}} + \left(\tfrac{1}{2}(1 \pm B_{n+1})\nu_\kappa + \tfrac{1}{2}(1 \mp B_{n+1})\eta_\kappa\right) H_{n+1}\,, \\
\alpha_{t_{n+1}} &= \tfrac{1}{2}(1 + B_{n+1})\epsilon^+ - \tfrac{1}{2}(1 - B_{n+1})\epsilon^- + \alpha_{t_{n+1}}\,,
\end{aligned}
$$

7. Repeat from step 1 until $t_{n+1} \geq T$.
8. Flow the diffusion from the last time prior to maturity until maturity using step 4 with $t_{n+1} = T$.
9. Compute the terminal $PnL = X_T + q_T\, S_T\, (1 - c_{trans}\, \mathrm{sgn}(q_T))$ where $c_{trans}$ is the liquidation cost.

The PnLs for the other types of HFTs employed in the simulation are obtained similarly.

[22] This is efficiently computed using the Lambert-W function since $A_0$ is typically small.