

ALGORITHMIC TRADING, STOCHASTIC CONTROL, AND MUTUALLY-EXCITING PROCESSES*

ÁLVARO CARTEA[†], SEBASTIAN JAIMUNGAL[‡], AND JASON RICCI[‡]

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 multifactor mutually 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 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.

Key words. algorithmic trading, high frequency trading, short-term-alpha, adverse selection, mutually exciting processes, Hawkes processes

AMS subject classifications. 91G80, 49L20, 60G55, 60G99

DOI. 10.1137/130911196

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 today computers manage inventories and make trading decisions; this modern way of trading in the electronic markets is known as algorithmic trading (AT), see [18][27].

Despite the substantial changes that markets have undergone in the recent past, some strategies 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 European stock exchanges is based on trying to profit from short-term price predictions by buying low and selling high.¹ 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 [19]. 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

*XXXX XXXXXXXXXXXX February 26, 2013; accepted for publication (in revised form) April 15, 2014; published electronically July 22, 2014.

<http://www.siam.org/journals/sifin/5/91119.html>

[†]Mathematical Institute, University of Oxford, Oxford, UK (alvaro.carte@maths.ox.ac.uk).

[‡]Department of Statistics, University of Toronto, Toronto M5S3G3, Canada (sebastian.jaimungal@utoronto.ca, jason.ricci@utoronto.ca). These authors would like to thank the Natural Sciences and Engineering Research Council and Mprime for partially funding this work. The third author would additionally like to thank the Ontario Graduate Scholarship for funding support.

¹See, for example, [22] and [43].

time intervals. A unique characteristic of HF trading is that the strategies are designed to hold almost no inventories over very short periods of time (seconds, minutes, or at most one day) to avoid exposure both 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. Today, 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 (LOs) show an intention to buy or sell and must indicate the number of shares and the price at which the agent is willing to trade. The limit buy (sell) order with the highest (lowest) 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* (MO), which is an order to buy or sell a number of shares, regardless of the price, which is immediately executed against LOs resting in the LOB at the best execution prices.

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, from the classical microstructure literature on adverse selection (see, e.g., [40]), we also know that strategies that do not include in their LOs 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 who 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 cancellations of LOs to maximize expected terminal wealth over a fixed horizon T while penalizing inventories, which will be made mathematically precise in section 5.1. The HFT we characterize here can be thought of as an ultrafast market maker where the trading horizon T is at most one trading day, all the LOs are canceled an instant later if not filled, and inventories are optimally managed (to maximize expected penalized terminal wealth) and drawn to zero by T .² Early work on optimal postings by a securities dealer is that of Ho and Stoll [32], and more recently Avellaneda and Stoikov [3] studied the optimal HF submission strategies of bid and ask LOs.

Intuitively, our HF dynamic strategy maximizes the expected profits resulting from roundtrip trades by specifying how deep on the sell and buy sides the LOs 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 LOs are to the best bid and best offer, the higher the probability of being executed, but

²HFTs 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; see, for example, [25].

the expected profits from a roundtrip are also lower, and adverse selection costs are 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, as will be discussed in sections 5.2, 5.3, and 6. For example, if the probability of the next MO 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 to 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 LOs on one side of the LOB are more likely to be adversely selected, then these LOs are posted deeper into the book in anticipation of the expected adverse selection costs. An increase in adverse selection risk could be heralded by MOs becoming more one-sided as a consequence of the activity of traders acting on superior or private information who are sending one-directional MOs.

Trade initiation may be motivated by many factors which have been extensively studied in the literature; see, for example, [42]. Some of these include asymmetric information, differences in opinion or differential information, and increased proportion of impatient (relative to patient) traders. Likewise, trade clustering can be the result of various market events; see [12]. 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 over the last ten years and the explosion in the volume of submissions and order cancellations, it is also plausible to expect that certain AT strategies that generate trade clustering are not necessarily motivated by the factors mentioned above.

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. The recent increase in the number of orders and in the frequency of LOB updates shows that fast traders are responsible for most of the market activity, 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 MOs arrive in two types. The first type of orders are *influential* orders, which excite the market and induce other traders to increase the number of MOs 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 *noninfluential* orders arrive, the intensity of the arrival of MOs 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 MOs. In this way, our model for the arrival of MOs 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. Multivariate Hawkes processes have recently been used in the financial econometrics literature to model clustering in trade arrival and changes in the LOB; see, e.g., [38], [9], and [44]. However, this paper is the first to incorporate such effects into optimal control problems related to AT.

In our model the arrival of trades also affects the midprice and the LOB. The arrival of MOs 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 number of influential market sell and buy orders; in particular, these imbalances have a temporary effect on the drift of the midprice.

Moreover, in our model the arrival of influential orders has 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 MOs, the shape of the buy and sell sides of the book will also undergo a temporary change due to market makers repositioning their LOs 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 those HFTs who incorporate predictions of short-term price deviations in their strategy will deliver positive expected profits. At the other extreme we have the HFTs who are driven out of the market because their LOs 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 in between these two cases, 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 LOs further away from the midprice so the chances of being picked off by other traders are considerably reduced.

2. Arrival of MOs 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 [8]. 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 [2], [37], [6], [35], [36], [13], while others are designed to manage inventory risk [14], [28]. There are HF strategies that specialize in arbitraging across different trading venues, while 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 [30], [17]. In previous works on algorithmic trading in LOBs, the midprice is assumed to be independent of MOs, and MOs arrive at Poisson times. Our work differs significantly in that we allow for dependence between MOs, the LOB dynamics, and midprice moves.

In the LOB, LOs are prioritized first according to price and then according to

³For instance, periods where the number of market buy orders is much higher than the number 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 [24].

time.⁴ 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 sides 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 salient features of the market dynamics.⁵ Since the HFT maximizes expected terminal wealth over a finite horizon T , while being penalized for holding large inventories, and she is continuously repositioning buy and sell LOs, 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 MOs cross the resting orders in the LOB. In this section we focus on (i) and (ii), in section 3 we discuss (iii), and after that we present the formal optimal control problem that the HFT solves.

2.1. Price dynamics. We assume that the *midprice* (or *fundamental price*) of the traded asset follows

$$(1) \quad dS_t = (v + \alpha_t) dt + \sigma dW_t,$$

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 v and by a short-term component α_t , which is a predictable zero-mean reverting process. Since we are interested in HF trading, our predictors are based on order flow information where we allow for feedback between MO events and short-term-alpha. In the rest of the paper we assume that $v = 0$ because the HF strategies we develop are for very short-term intervals.

Section 4 gives details of the dynamics of the process α_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 a burst of activity on one or both sides of the market, orders that walk 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) while at the same time allowing them to reposition stale quotes or submit new quotes to avoid being picked off by other market participants trading on short-term-alpha, i.e., avoid being adversely selected.

⁴This is the case for most exchanges. Some exchanges use pro rata order books, where MOs are matched with all traders posting at the touch proportional their posted volume (see, e.g., [31]), or other alternatives such as Broker priority in Scandinavian markets.

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

⁶Unless 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 $\mathbb{F} = \{\mathcal{F}_t\}_{0 \leq t \leq T}$ 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 LOs, news events, and the indicator of whether a trade is influential or not.




FIG. 1. *INTC market orders. Historical running intensity versus smoothed fitted intensity (restricted to $\rho = 1$) using a 2 second sliding window for INTC for the period between 9:30am and 9:33am, January 29, 2018.*

2.2. Mutually exciting incoming market order dynamics. Order flow tends to fluctuate throughout the day; indeed, as a motivating example, Figure 1 shows the historical intensity of trade arrival for INTC over a three minute period (starting at 9:30am, January 29, 2018). The historical intensities are calculated by counting the number of buy and sell MOs over the last two seconds. The fitted intensities are computed using our model (see (2)) under the specific assumption that all trades are influential; see Appendix A for more details. From the figures we observe that MOs may arrive in clusters, that there are times when the markets are mostly one-sided (for instance, the period around 100 seconds), and that these bursts of activity die out rather quickly and revert to around two 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 bursts of activity in the LOB, and most of the time these algorithms act on information which is difficult to link to public news. Thus, here we take the view that some MOs generate more trading activity in addition to the usual effect of news increasing the intensity of MOs.

In our model MOs 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 MOs up to (and including) time t by the processes $\{\bar{M}_t^-, \bar{M}_t^+\}$. The second type of orders are *noninfluential* orders. These orders do not excite the state of the market. We denote the total number of arrivals of noninfluential sell/buy MOs 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 an MO, the probability that the trade is influential is a constant $\rho \in [0, 1]$.

Thus, we model the intensity of sell, λ_t^- , and buy, λ_t^+ , MOs by assuming that they solve the coupled system of stochastic differential equations (SDEs).

Assumption 1. The market sell/buy order arrival rates $(\lambda_t^-, \lambda_t^+)$ solve the coupled system of SDEs

$$(2a) \quad d\lambda_t^- = \beta(\theta - \lambda_t^-)dt + \eta d\bar{M}_t^- + \nu d\bar{M}_t^+,$$

$$(2b) \quad d\lambda_t^+ = \beta(\theta - \lambda_t^+)dt + \eta d\bar{M}_t^+ + \nu d\bar{M}_t^-,$$




FIG. 2. Sample path of MO 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 MO, while squares indicate the arrival of noninfluential trades.

where, as previously stated, \overline{M}_t^+ and \overline{M}_t^- are the total number of influential buy and sell orders up until time t . Moreover, β, θ, η, ν are nonnegative constants satisfying the constraint $\beta > \rho(\eta + \nu)$.

MOs are mutually exciting since their arrival rates λ^\pm jump upon the arrival of influential orders (note that the arrival of noninfluential orders does not affect λ^\pm). If the influential MO is a buy (so that a sell LO was lifted), the jump activity on the buy side increases by η while the jump activity on the sell side increases by ν , and similarly when the MO is a sell. 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 Figure 1).

Trading intensity is mean-reverting. Jumps in activity decay back to its long run level of θ at an exponential rate β . Figure 2 illustrates an intensity sample path. The lower constraint $\beta > \rho(\eta + \nu)$ is required for the intensity processes 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 λ_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,*

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

The intuition for the constraint is that when an MO arrives the activity will jump either by η or by ν and this occurs with probability ρ . Further, since both sell and buy influential orders induce mutual excitations, the decay rate β must be strong enough to compensate for both jumps to pull the process toward its long-run level of θ .

News events may also induce increases in trading activity, and incorporating them into the analysis is straightforward, e.g., by adding an exogenous counting process which induces jumps in activity at the time of news arrivals. To keep the modeling to a minimum, however, we opt to exclude them in our analysis.

3. Limit quote arrival dynamics and fill rates. The LOB can take on a variety of shapes, and it changes dynamically throughout the day; see [41] and [23]. MOs walk the book until all the volume specified in the order is filled. LOs in the tails of the LOB are less likely to be filled than those within a couple of ticks from the midprice S_t . The decision on where to post limit buy and sell orders depends on a number of characteristics of the LOB and on the MOs. Some of the LOB features are shape of the LOB, resiliency of the LOB, and how the LOB changes in between the arrival of MOs. These features, combined with the size and rate of the incoming MOs, determine the fill rates of the LOs. The fill rate is the rate of execution of an LO. Intuitively, a high (low) fill rate indicates that an LO is more (less) likely to be filled by an MO.

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 δ^\pm denotes how far from the midprice the orders are posted. Note that the agent continuously adjusts their posting relative to the midprice; hence, it is not possible for the midprice to move through the agent's posts. Rather, fills occur when MOs arrive and reach the level at which the agent is posted. This is in line with how a number of other authors have modeled optimal postings and fill probabilities; see [32], [3], [6], [14], and [28]. This approach can be viewed as a reduced-form one, in contrast to models which focus on modeling the dynamics of each level of the LOB, together with MO arrivals (see, e.g., [41] and [23]). In all of the reduced-form approaches, when an MO arrives, it walks through the LOB, and the probability that the HFT's limit order is filled is a (static) function of the posted depth δ^\pm , and it is also assumed that filled MOs do not affect the shape of the LOB. In our approach, however, we allow fill probabilities to be stochastic due to changes in the LOB which result from the arrival of MOs.

Assumption 2. The fill rates are of the form $\Lambda_t^\pm \triangleq \lambda_t^\pm h_\pm(\delta; \kappa_t)$, where the non-increasing function $h_\pm(\delta; \kappa_t) : \mathbb{R} \rightarrow [0, 1]$ is C^2 in δ (uniformly in t for $\kappa_t \in \mathbb{R}^n$, fixed $\omega \in \Omega$), and $\lim_{\delta \rightarrow \infty} \delta h_\pm(\delta; \kappa_t) = 0$ for every $\kappa_t \in \mathbb{R}^n$. Moreover, the functions $h_\pm(\delta; \kappa_t)$ satisfy $h_\pm(\delta; \kappa_t) = 1$ for $\delta \leq 0$, $\kappa_t \in \mathbb{R}^n$.

Assumption 2 allows for very general dynamics on the LOB through the dependence of the fill probabilities (FPs) $h_\pm(\delta; \kappa_t)$ on the process κ_t . The FPs can be viewed as a parametric collection, with the exponential class $h_\pm(\delta; \kappa_t) = e^{-\kappa_t^\pm \delta^\pm}$ and power law class $h_\pm(\delta; \kappa_t) = (1 + (\kappa_t^\pm \delta^\pm)^\alpha)^{-1}$ being two prime examples. The process κ_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 δ^\pm implies that the book (volume) 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 $h_\pm(\delta; \kappa_t) = 1$ for $\delta \leq 0$ and $\forall \kappa_t \in \mathbb{R}^n$ is a financial one. A trader wanting to maximize her chances of being filled the next time an MO arrives must post the LOs 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 MOs. This assumption makes crossing the midprice a suboptimal decision because the trader cannot improve the arrival rate of MOs (since Λ_t^\pm is constant when $\delta^\pm \leq 0$); thus she will always post LOs that are $\delta^\pm \geq 0$ away from the midprice. Furthermore, this condition is more desirable than explicitly restricting the controls δ^\pm to be nonnegative, since it

is not necessary to check the boundary condition at $\delta^\pm = 0$; it will automatically be satisfied. Finally, we have the added bonus that the optimal control satisfies the first-order condition.

Observe that if MO volumes are independent and identically distributed (i.i.d.), then the κ_t^\pm processes can be interpreted as parameters directly dictating the shape of the LOB. In particular, if the MO volumes are i.i.d. and exponentially distributed and the shape of the LOB is flat, then the probability that an LO at price level $S_t \pm \delta_t^\pm$ is executed (given that an MO arrives) is equal to $e^{-\kappa_t^\pm \delta_t^\pm}$. Consequently, κ_t^\pm can be interpreted as depth of the LOB at each price level. In order to satisfy the C^1 condition at $\delta^\pm = 0$ and the condition that $h_\pm(\delta, \kappa_t) = 1$ for $\delta^\pm \leq 0$, it is necessary to smooth the exponential function at $\delta = 0$. This is always possible, though, since there exist C^2 functions for which the L^2 distance to the target function is less than any positive constant.

Assumption 3. The dynamics for κ_t satisfy

$$(3a) \quad d\kappa_t^- = \xi(\vartheta - \kappa_t^-) dt + \eta_\kappa d\overline{M}_t^- + \nu_\kappa d\overline{M}_t^+,$$

$$(3b) \quad d\kappa_t^+ = \xi(\vartheta - \kappa_t^+) dt + \nu_\kappa d\overline{M}_t^- + \eta_\kappa d\overline{M}_t^+,$$

where η_κ and ν_κ are nonnegative constants and ϑ and ξ are strictly positive constants.

Assumption 3 is a specific modeling assumption on κ_t which allows for incoming influential MOs to have an impact on the FPs. An increase (decrease) in the fill rate can be due to two main factors: (i) a decrease (increase) in LOB depth and/or (ii) an increase (decrease) in the distribution of MO volumes (in a stochastic dominance sense). This is a one-way effect because influential MOs cause jumps in the κ_t process, but jumps in the FP do not induce jumps in MO arrivals. While it is possible to allow such feedback, empirical investigations (such as those in [38]) demonstrate that the incoming MOs influence the state of the LOB and not the other way around. The mean-reversion term draws κ_t^\pm to the long-run mean of ϑ so that the impact of influential orders on the LOB is only temporary. Typically, we expect the rate of mean-reversion ξ for the LOB to be slower than the rate of mean-reversion β of the MO activity. In other words, the impact of influential orders persists in the LOB on a longer time scale compared to their effect on MO activity.

Moreover, immediately after an influential market buy/sell order arrives, the probability that an LO at price level $S_t \pm \delta_t^\pm$ is executed is, for the same δ^\pm , smaller than the probability of it being filled before the influential order arrives. The intuition is the following. Immediately after an influential MO arrives, market participants react in anticipation of the increase of market activity they will face and decide to send LOs to the book. Since many market participants react in a similar way, the probability of LOs being filled, conditional on an MO arriving, decreases.⁷

Figure 3 illustrates the shape of the fill rates at time t describing the rate of arrival of MOs which fill LOs 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 MOs. In the figure these rates are asymmetric and decay at differing speeds because we have

⁷It is also possible to have markets where, conditional on the arrival of an MO, the probability of an LO being filled increases immediately after the arrival of an influential order. We can incorporate this feature into our model. Note also that in our general framework, immediately after the influential buy/sell MO arrives, the intensities λ^\pm increase and the overall effect of an influential order on the fill rates $\Lambda^\pm = \lambda_t^\pm h_\pm(\delta, \kappa_t)$ is ambiguous when λ_t^\pm and $h_\pm(\delta, \kappa_t)$ move in opposite directions after the arrival of an influential order—for example, when $h_\pm(\delta, \kappa_t) = e^{-\kappa_t^\pm \delta_t^\pm}$.




FIG. 3. Fill rates Λ^\pm with $\kappa_t^+ = 2$, $\kappa_t^- = 1$, $\lambda_t^+ = 0.75$, and $\lambda_t^- = 1$. As κ_t^\pm evolve, the fill rate shape changes, while the MO activity rates λ_t^\pm modulate the vertical scale. The maximum fill rate is achieved at $\delta = 0$.

assumed different parameters for the buy and sell sides, $\kappa_t^+ = 2$, $\kappa_t^- = 1$, $\lambda_t^+ = 0.75$, and $\lambda_t^- = 1$. In general, these curves will fluctuate throughout the day.

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 MOs and the arrival and cancellation of LOs. Similarly, bursts of activity in buy and/or sell MOs, which are not necessarily the result of the arrival of public information, have market impact by producing momentum in the midprice, see also [16] who show how order-flow affects midprice dynamics.

As discussed above, a great deal of the strategies that HFTs employ are directional strategies that take advantage of short-term price deviations in two ways. First, the strategies enable HFTs to exploit their superior knowledge of short-term trends in prices to execute profitable roundtrip trades, and second, they provide 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 α_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. Here we assume that α_t is a zero-mean reverting process and jumps by a random amount at the arrival times of influential trades. If the influential trade was buy-initiated, the drift will jump up, and if the influential trade was sell-initiated, the drift will jump down. As such, we model the predictable drift as below.

Assumption 4. The dynamics for the predictable component of the midprice's drift, α_t , satisfy

$$(4) \quad d\alpha_t = -\zeta \alpha_t dt + \sigma_\alpha dB_t + \epsilon^+ d\bar{M}_t^+ - \epsilon^- d\bar{M}_t^-,$$

where ϵ^\pm are random variables representing the size of the sell/buy influential trade's impact on the drift of the midprice. Moreover, B_t denotes a Brownian motion independent of all other processes, and ζ, σ_α are positive constants.

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 MO

is influential. Her optimal directional strategy is to simultaneously send the following orders to the LOB: cancel her sell LOs, attempt to purchase the asset (from a slower market participant), and send new sell LOs 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. Given our chosen dynamics on the fill probability driving process κ_t^\pm in (3), the aforementioned effect can be modeled by taking $\eta_\kappa < \nu_\kappa$, which will induce more arrivals of limit buy (sell) quotes when an influential market buy (sell) order arrives.

An alternative approach to adverse selection was introduced in [14], whereby MOs may induce an immediate jump in the midprice. The result of such direct adverse selection effects was that the agent increases her optimal postings by the expected jump size. In this work, we will see a similar, but distinct, result whereby the agent adjusts her posting to protect herself against the potential change in the midprice drift.

5. The HFT's optimization problem. So far, we have specified counting processes for MOs and dynamics of the LOB through the FPs; however, we also require a counting process for the agent's filled LOs. 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 $\Lambda_t^\pm \triangleq \lambda_t^\pm h_\pm(\delta; \kappa_t)$, as in Assumption 2. Finally, the agent's cash process X_t (excluding the shares she currently holds) satisfies the SDE

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

where δ_{t-}^\pm denotes the left-limit of the LO's distance from the midprice, i.e., if the LO was filled, the agent receives the quote that was posted an instant prior to the arrival of the MO.

5.1. Formulation of the HF investment problem. The HFT wishes to place sell/buy LOs at the prices $S_t \pm \delta_t^\pm$ at time t such that the expected terminal wealth is maximized while penalizing inventories.⁸ The HFT is continuously repositioning her LOs in the book by canceling stale LOs and submitting new LOs.⁹ Specifically, her value function is

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

where the supremum is taken over all (bounded) \mathcal{F}_t -progressively measurable functions and ϕ penalizes deviations of q_t from zero along the entire path of the strategy.

⁸An 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 [1], where quadratic variation, rather than variance, is penalized, which acts on the entire path of the strategy.

⁹In this setup the HFT's LOs are always of the same size. An interesting extension is to also allow the HFT to choose the number of shares in each LO.

$\overline{\mathcal{F}}_t$ 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 $\mathcal{F}_t = \overline{\mathcal{F}}_t \vee \sigma\{(\overline{M}_u)_{0 \leq u \leq t}\}$. Note that λ_t and κ_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 $v = 0$ in the dynamics of the midprice. Note that [11] shows that the running penalty term in (6) can be interpreted as arising from the agent's ambiguity aversion with respect to the asset's midprice.

The above control problem can be cast into a discrete-time controlled Markov chain as carried out in [5]. Classical results from [7] imply that a dynamic programming principle holds and that the value function is the unique viscosity solution of the HJB equation

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

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

$$(8) \quad \mathcal{L} = \beta(\theta - \lambda^-) \partial_{\lambda^-} + \beta(\theta - \lambda^+) \partial_{\lambda^+} + \xi(\vartheta - \kappa^-) \partial_{\kappa^-} + \xi(\vartheta - \kappa^+) \partial_{\kappa^+} - \zeta \alpha \partial_\alpha + \frac{1}{2} \sigma_\alpha^2 \partial_{\alpha\alpha}.$$

Moreover, we have introduced the following shift operators:

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

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

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

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

where the expectation operators $\mathbb{E}[\cdot]$ in (9a) are over the random variables ϵ^\pm .

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 section 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 Φ admits the decomposition $\Phi = x + qs + g(t, q, \alpha, \lambda, \kappa)$ with $g(T, \cdot) = 0$. Furthermore, assume that $g(\cdot)$ can be written as an asymptotic expansion as follows:*

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

with boundary conditions $g(T, \cdot) = 0$. Note that the subscripts on the functions g do not denote derivatives; rather they are labels, and we have written $\mathbb{E}[\epsilon^\pm] = \varepsilon \mathfrak{a}^\pm$ with

ε constant. Then, the feedback controls of the optimal trading strategy for the HJB equation (7) admit the expansion

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

where

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

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

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

and the coefficient $B(\delta_0^{\pm}; \kappa) = \frac{h'_{\pm}(\delta_0^{\pm}; \kappa)}{2h'_{\pm}(\delta_0^{\pm}; \kappa) + \delta_0^{\pm} h''_{\pm}(\delta_0^{\pm}; \kappa)}$. Moreover, δ_0^{\pm} is a strictly positive solution to

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

A solution to (13) always exists. Furthermore, the exact optimal controls are nonnegative.

In the next subsection we use the optimal controls derived here to solve the nonlinear HJB equation and obtain an analytical expression for g_{α} , g_{ε} , and g_{ϕ} to obtain explicit expressions for the optimal postings. 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 (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 δ_0^{\pm} , does not directly depend on the arrival rate of MOs, short-term-alpha, 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 δ^{\pm} to maximize the expected spread conditional on an MO hitting or lifting the appropriate side of the book, i.e., maximizes $\delta^{\pm} h_{\pm}(\delta^{\pm}; \kappa_t)$. The first order condition of this optimization problem is given by (13), where we see that λ^{\pm} plays no role in how the LOs are calculated.¹⁰

The optimal halfspreads are adjusted by the impact that influential orders 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.

¹⁰If there are multiple solutions to (13), the HFT chooses the δ^{\pm} that yields the maximum of $\delta^{\pm} h_{\pm}(\delta^{\pm}; \kappa_t)$.

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 ς . 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_α , g_ε , and g_ϕ can be written as*

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

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

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

where

$$(15a) \quad b_\alpha(t) = \frac{1}{\varsigma} \left(1 - e^{-\varsigma(T-t)} \right),$$

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

$$(15c)$$

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

$$(15d) \quad c_\phi(t) = -(T-t).$$

In the above, $h_{0,u}^\pm = h_\pm(\delta_{0,u}^\pm; \boldsymbol{\kappa}_u^\pm)$, and we have written $E[\epsilon^\pm] = \varepsilon \mathbf{a}^\pm$. Finally, the functions g_0 , a_α , a_ε , and a_ϕ do not affect the optimal strategy.

The asymptotic expansion of the optimal controls now follows as a straightforward corollary to Theorem 3. Note that the functions b_ε can be computed explicitly, as is reported in section C.2. Moreover, under some specific assumptions on the fill probabilities h_\pm (e.g., if h_\pm are exponential or power functions), the function b_ϕ can also be computed explicitly. Proposition 7 in section 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 LOs). *The asymptotic expansion of the optimal controls to first order in ς is (dependencies on the arguments have been suppressed for clarity)*

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

where δ_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 $o(\varsigma)$.¹¹

The expression for the optimal control warrants some discussion which goes beyond the discussion that followed the general result in Proposition 2. The term δ_0^\pm

¹¹Note that the exact solution of the optimal control is nonnegative, as discussed in Assumption 2, but this is not necessarily the case in the asymptotic solution; thus we write the optimal control as $\max\{\delta^{\pm*}, 0\}$.

represents the action of a risk-neutral agent who is not aware of or is not able to estimate the impact that influential MOs 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 on the midprice's drift, the expected change to the arrival of orders, and the Brownian component in the short-term-alpha dynamics; see (4). 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 ϕ contains two terms. The first of these terms accounts for the asymmetry in the arrival rates of MOs on the sell and buy sides, while the second term controls for inventories. Both terms together help to induce mean reversion to an optimal inventory level which is not necessarily zero.

There are a number of special cases that are interesting to analyze. For instance, if $\mathbb{E}[\epsilon^\pm] = \epsilon = 0$ influential trades do not affect the short-term-alpha dynamics, then in the optimal control only the fact that α_t reverts to zero, at the exponential speed ζ , is taken into account. Clearly, if $\alpha_t > 0$, the HFT's sell halfspread is increased to avoid selling an asset which is trending up in price, and for the same reason the optimal buy halfspread decreases to increase the probability of purchasing the asset in anticipation of a price increase.

In addition, the expression for optimal control simplifies considerably when (i) the impact of influential orders on the stochastic drift is symmetric in the sense that $\epsilon^+ = \mathbb{E}[\epsilon^+] = \mathbb{E}[\epsilon^-] = \epsilon^- := \epsilon$; (ii) the parametric shape of the LOB FPs are symmetric in the sense that the class of functions h^+ and h^- are equal;¹² and (iii) the fill probability at the risk-neutral optimal control is independent of the scale parameters,¹³ i.e., $h_\pm(\delta_0^\pm, \kappa) = \text{const.}$ Under these assumptions, the two important (nontrivial) quantities which appear in the optimal spreads in (16) simplify to

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

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

where $\hat{\beta} = \beta - \rho(\eta - \nu)$ and $h = h_\pm(\delta_0^\pm, \kappa) = \text{const.}$ Both expressions contain terms proportional to the difference in the MO 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 to this order flow imbalance. Finally, it is straightforward to see in this symmetric case ($\lambda_t^+ = \lambda_t^-$) that if $\epsilon = 0$, the short-term-alpha component

¹²This does not imply that the LOB is symmetric because the scale parameters κ_t^\pm will differ. For example, exponential FPs $e^{-\kappa_t^\pm \delta^\pm}$ satisfy this requirement even though the book may be significantly deeper on one side than the other.

¹³This condition is satisfied by (but not limited to) the exponential and power law FPs, as discussed in Examples 9 and 10.

affects the optimal posting only via the effect of the last term on the right-hand side of (17a).

6. HF market making, short-term-alpha, and directional strategies. In this section we apply a simulation study of the HF strategy where MOs are generated over a period of five minutes. The HFT is rapidly updating her quotes in the LOB by submitting and canceling LOs which are filled according to exponential FPs.¹⁴ The optimal postings are calculated using Corollary 4 and the explicit form for b_ϕ in Proposition 7. The processes λ_t , κ_t , and α_t are updated appropriately, and the terminal cash-flows are stored to produce the profit and loss (PnL) generated from these strategies.

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.¹⁵ 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 MO) and by paying a fee to the exchange for taking liquidity in the form of an aggressive MO. 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 has when calculating the optimal postings. The main difference between our scenarios is whether the HFT is able to calculate the correct ρ which, conditional on the arrival of an MO, is the probability that the trade is influential and whether they are able to estimate the correct dynamics of short-term-alpha; all of the HFTs know the equations that determine λ_t^\pm and κ_t^\pm but do not necessarily know the correct parameters. We contemplate six different types of HFTs:

1. *Correct probability of influential event (ρ).* The HFT uses her superior computing power to process information to estimate ρ and the other parameters that determine the dynamics of λ_t^\pm and κ_t^\pm . Furthermore, we assume that the HFT may or may not be able to estimate the correct α_t dynamics.
 - (a) *Correct midprice drift (α) dynamics.* This is our *benchmark* because we also assume that the HFT is able to estimate the parameters of the α_t process and adjust her postings accordingly.
 - (b) *Zero midprice drift (α) dynamics.* Here we assume that, although the HFT is able to estimate the correct ρ , she assumes that short-term-alpha is zero throughout the entire strategy.
2. *High probability of influential event (ρ).* At the other extreme we also have an HFT who cannot distinguish between the type of MO and assumes that all orders are influential, $\rho = 1$. The jump sizes in λ_t^\pm and κ_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 midprice drift (α) 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 MOs have on short-term-alpha because every time there is an incoming MO, the HFT will predict a jump in α_t . The mean jump size parameter, ε ,

¹⁴The results for power FPs are very similar, and so in the interest of space we do not show them.

¹⁵The 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 $q_T(S_T - c_{\text{trans}} \text{sgn}(q_T))$.

is also rescaled by the correct ρ .

- (b) *Zero midprice drift (α) dynamics.* The HFT assumes that short-term-alpha is always zero.
- 3. *Low probability of influential event (ρ).* At one extreme we have an HFT who cannot distinguish between order type, assumes that all orders are noninfluential, $\rho = 0$, and assumes that λ_t^\pm , κ_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 11.
 - (a) *Incorrect midprice drift (α) dynamics.* Because the HFT assumes that all orders are noninfluential, she is not able to correctly predict short-term-alpha; she observes only the diffusion components and not the jumps.
 - (b) *Zero midprice drift (α) dynamics.* The HFT assumes that short-term-alpha is always zero.

In all six cases, the data generating processes (DPGs) are identical and are given by the full model, where we assume the following values for the parameters (unless otherwise stated): $\beta = 60$ and $\theta = 1$ (speed and level of mean reversion of intensity of MO arrivals); $\eta = 40$ and $\nu = 10$ (jumps in λ_t upon the arrival of influential MOs); $\xi = 10$ and $\vartheta = 50$ (speed and level of mean reversion for the κ_t process); $\eta_\kappa = 10$ and $\nu_\kappa = 25$ (jumps in κ_t upon the arrival of influential MOs); $v = 0$ (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 α_t process); and, finally, $\rho = 0.7$ (probability of the MO being influential). Moreover, ϵ^\pm are both exponentially distributed with the same mean, $\mathbb{E}[\epsilon^\pm] = \varepsilon$, for the sell and buy impacts, with $\varepsilon = 0.04$.

In Table 1 we show the PnLs that the HFTs face when executing the optimal strategy. 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 ϕ we show the mean and standard deviation of the six PnLs where the top row, for each ϕ , 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 noninfluential MOs). For each ϕ the bottom row shows the other three PnLs that result from assuming that the HFTs set $\alpha_t = 0$ throughout the entire strategy.

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 α_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 who omit short-term-alpha face negative, or at best close to zero, mean PnLs.

The inventory-management parameter ϕ 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 ϕ . Table 1 shows an interesting effect of ϕ on the PnL of the different strategies that we study. If the HFT uses her predictions of

TABLE 1

The mean and standard deviation of the PnL from the various strategies as the inventory-management parameter ϕ 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.

Case I: $\varepsilon = 0.04$, $\rho = 0.7$ and liquidation costs = 1bp					
ϕ	α_t		Bench.	$\rho = 1$	$\rho = 0$
1×10^{-5}	Yes	mean	14.09	12.77	-4.34
		(std)	(6.98)	(6.44)	(3.00)
	No	mean	-3.81	-3.88	-4.32
		(std)	(2.83)	(2.83)	(2.94)
2×10^{-5}	Yes	mean	13.52	12.15	-2.80
		(std)	(5.44)	(5.01)	(2.21)
	No	mean	-1.57	-1.71	-2.80
		(std)	(2.07)	(2.06)	(2.17)
4×10^{-5}	Yes	mean	12.49	11.08	-1.24
		(std)	(4.28)	(3.94)	(1.60)
	No	mean	0.24	0.06	-1.25
		(std)	(1.48)	(1.48)	(1.58)

Case II: $\varepsilon = 0.04$, $\rho = 0.7$ and liquidation costs = 10bp					
ϕ	α_t		Bench.	$\rho = 1$	$\rho = 0$
1×10^{-5}	Yes	mean	13.32	12.05	-4.81
		(std)	(6.80)	(6.29)	(3.12)
	No	mean	-4.28	-4.34	-4.78
		(std)	(2.98)	(2.98)	(3.07)
2×10^{-5}	Yes	mean	12.87	11.55	-3.19
		(std)	(5.30)	(4.90)	(2.32)
	No	mean	-1.96	-2.10	-3.19
		(std)	(2.20)	(2.19)	(2.28)
4×10^{-5}	Yes	mean	11.94	10.58	-1.56
		(std)	(4.17)	(3.85)	(1.69)
	No	mean	-0.09	-0.27	-1.57
		(std)	(1.59)	(1.59)	(1.67)

short-term-alpha to make markets, increasing ϕ 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 the PnL.

On the other hand, the effect of increasing ϕ 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 ϕ makes their postings more conservative because, everything else being equal, the LOs are posted deeper in the LOB, and this makes it more difficult for other traders to pick off their quotes. Thus, by increasing ϕ , the HFT reduces her exposure to adverse selection, and this explains why the mean PnL increases in ϕ . Finally, the standard deviation of the PnL decreases because, when ϕ increases, the strategy induces very quick mean reversion of inventories to zero.

7. New Developments. Since the original version of this paper, there have been a number of exciting developments in stochastic control and stochastic game formulations of algorithmic trading problems. The problem setup here, and in all other works at the time, ignore the optimization of other agents. While [16] address how order-flow from other agents modifies the optimal strategy of an individual agent, the authors do not address the optimisation problem that the other agents are simultaneously pursuing. Doing so requires studying large stochastic games. [34] is the earliest work to address the problem where there are multiple agents trading, some of whom are also optimizing their trades, while others add noise to the order-flow – so-called uninformed traders – and there is a large liquidating agent. [10] also address a multiple agent setting where all agents are homogeneous. In both cases, the trading problem becomes a multi-player stochastic game, and in general, such problems cannot be solved in closed-form. Instead, the authors make use of the theory of mean-field games (MFGs), developed simultaneously (with different approaches) in [33] and [39], to obtain Nash equilibria for the infinite player game, and demonstrate

that the MFG solution leads to ϵ -Nash equilibria¹⁶ for the finite player version of the game.

Trading problems where the agent aims to track a particular price path, such as volume weighted asset price (VWAP) [29, 26, 15], or tracking a particular trading rate [4], such as when the trader is hedging an option taking into account the microstructure impact, have also received attention in the recent literature.

Another important innovation in the field stems from the observation that financial markets, and in particular intra-day markets, have factors that drive the price and order-flow dynamics which cannot be observed – so-called latent factors. [20] demonstrate how to perform statistical arbitrage in the presence of these latent factors, and derive closed-form optimal strategies under a general class of Markov-modulated diffusion and counting process models. They utilize tools from optimal control with partial information, and while the mathematical results are in closed-form, estimation and implementation of the strategy requires use of machine-learning tools, which the authors also address. This integration of machine learning and stochastic control is a very exciting and promising direction of research.

A final, but not exhaustive, innovation is to incorporate both latent factors and the optimal actions of multiple heterogeneous agents. [21] analyzes a very general model where asset prices are adapted to a large filtration $\mathcal{G} = \{\mathcal{G}_t\}_{t \in [0, T]}$, while agents have access to a coarser filtration $\mathcal{F} = \{\mathcal{F}_t\}_{t \in [0, T]}$, s.t. $\mathcal{F}_t \subseteq \mathcal{G}_t$, and various sub-populations of agents have differing performance criteria. The authors succeed, by using methods of convex analysis (rather than using the dynamic programming principle, or the stochastic Pontryagin maximum principle), in obtaining forward-backward stochastic differential equations which the latent, heterogeneous, multiple agent trading problem solves. Moreover, they solve this FBSDE system, and demonstrate that the MFG leads to a ϵ -Nash equilibria for the finite player game.

8. 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 between influential and noninfluential MOs. The arrival of influential MOs increases MO activity and also affects the shape and dynamics of the LOB. On the other hand, when noninfluential MOs arrive, they walk 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 MOs and news. Influential buy and sell MOs induce short-lived upward and downward trends in the midprice of the asset. This specification allows us to capture the essence of HF trading—to exploit short-lived predictable opportunities by 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 LOs to maximize expected terminal wealth, while 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 on the arrival of MOs, size and sign of inventories, and short-term-alpha. The optimal strategy captures many of the key

¹⁶An ϵ -Nash equilibria is one where agents can improve their performance criteria by deviating from the equilibria, but only by an amount ϵ which tends to zero as the population size tends to infinity.

characteristics that differentiate HFTs from other algorithmic traders: profit from directional strategies based on predicting short-term-alpha; reduced exposure to LOs 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 positive expected profits. The other HFTs are driven out of the market because their LOs 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 LOs 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 it is optimal for the HFT to submit MOs. 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 MOs; the best she can do is send LOs at the midprice (zero spread), but this is no guarantee that the LO 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 mutually exciting nature of our model captures other important features of strategic behavior which include “market manipulation.” For example, algorithms could be designed to send MOs, 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 is beyond the scope of this paper.

Appendix A. Fitting the model. Here we focus on the case when all MOs are influential. Online calibration and estimation of the current state of activity in the general model are beyond the scope of this work and will be reported on elsewhere.

When all MOs 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 $t_n \leq T$ 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 can be found by integrating (2a)–(2b) and are

explicitly given by

$$(18) \quad \lambda_t^\pm = \theta + \sum_{i=1}^n H_i^\pm e^{-\beta(t-t_i)} \quad \text{and} \quad \int_0^t \lambda_u^\pm du = \theta t + \sum_{i=1}^n H_i^\pm \frac{1 - e^{-\beta(t-t_i)}}{\beta},$$

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

$$(19) \quad \mathcal{L} = -2\theta T + \sum_{i=1}^n \left\{ B_i \log(\lambda_{t_i}^+) + (1 - B_i) \log(\lambda_{t_i}^-) - (\eta + \nu) \frac{1 - e^{-\beta(T-t_i)}}{\beta} \right\}.$$

Maximizing this log-likelihood results in the MLE of the model parameters, and upon back-substitution into (18) provides the estimated path of activity. Integrating this activity over the last two seconds, i.e., $\frac{1}{2} \int_{t-2}^t \lambda_u^\pm du$, provides us with a smoothed version of the intensity, shown in Figure 1 as the path labeled “Fitted.” This is directly comparable to the one second historical intensity in Figure 1 labeled “Historical.”

For the time window 10:00am to 11:00 am on January 29, 2018, for INTC the estimated parameters are as follows:

$$\hat{\beta} = 3,774, \quad \hat{\theta} = 0.22, \quad \hat{\eta} = 2,766, \quad \text{and} \quad \hat{\nu} = 54.$$

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 ρ is not necessarily 1 and the influential/noninfluential nature of the event must be filtered, the jump sizes are indeed larger.

Appendix B. Proof of results.

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 ODEs for $m_t^\pm(u)$:

$$(20) \quad \frac{d}{du} \begin{pmatrix} m_t^-(u) \\ m_t^+(u) \end{pmatrix} + \begin{pmatrix} \beta - \eta\rho & -\nu\rho \\ -\nu\rho & \beta - \eta\rho \end{pmatrix} \begin{pmatrix} m_t^-(u) \\ m_t^+(u) \end{pmatrix} - \begin{pmatrix} \beta\theta \\ \beta\theta \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

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

$$(21) \quad \begin{pmatrix} m_t^-(u) \\ m_t^+(u) \end{pmatrix} = e^{-\mathbf{A}(u-t)} \left[\begin{pmatrix} \lambda_t^- \\ \lambda_t^+ \end{pmatrix} - \mathbf{A}^{-1} \zeta \right] + \mathbf{A}^{-1} \zeta.$$

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 $u \rightarrow \infty$, $m_t(u)$ 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 (20) 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 bounds. Finally, if both eigenvalues are zero, then we must have $\beta = \nu = \eta = 0$, which is, once again, outside of the stated bounds. None

B.2. Proof of Proposition 2. Applying the ansatz on the form on Φ , differentiating inside the supremum in (7) with respect to δ^\pm , 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

$$(22) \quad \begin{aligned} & h_\pm(\delta_0^\pm) + \delta_0^\pm h'_\pm(\delta_0^\pm) + \alpha \left\{ \delta_\alpha (h''_\pm(\delta_0^\pm) + 2h'_\pm(\delta_0^\pm)) + h'_\pm(\delta_0^\pm) \left(\mathbb{S}_{q\lambda}^\pm g_\alpha - \mathbb{S}_\lambda^\pm g_\alpha \right) \right\} \\ & + \varepsilon \left\{ \delta_\varepsilon (h''_\pm(\delta_0^\pm) + 2h'_\pm(\delta_0^\pm)) + h'_\pm(\delta_0^\pm) \left(\mathbb{S}_{q\lambda}^\pm g_\varepsilon - \mathbb{S}_\lambda^\pm g_\varepsilon + \pm \rho \mathbf{a}^\pm \mathcal{S}_{q\lambda}^\pm g_\alpha \right) \right\} \\ & + \phi \left\{ \delta_\phi (h''_\pm(\delta_0^\pm) + 2h'_\pm(\delta_0^\pm)) + h'_\pm(\delta_0^\pm) \left(\mathbb{S}_{q\lambda}^\pm g_\phi - \mathbb{S}_\lambda^\pm g_\phi \right) \right\} = o(\varsigma). \end{aligned}$$

Observe that the Taylor expansion of $h(\delta)$ about δ_0 requires the C^2 regularity condition to keep the error of the correct order. The C^1 regularity condition ensures that the global maximizer satisfies (22). Setting the constant term in (22) to zero yields (13). Setting the coefficients of α , ε , and ϕ each separately to zero and solving for δ_α , δ_ε , and δ_ϕ 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 the solution to (13) is a critical point of the function $\delta h(\delta)$. The critical point exists since $\delta h(\delta)$ is nonpositive for $\delta \leq 0$, is strictly positive on an open interval of the form $(0, d)$ due to $h \in C^1$ (since $h > 0$ in an open neighborhood of $\delta = 0$), and goes to 0 in the limit (when $\delta \rightarrow \infty$) by Assumption 2.

To see that the exact values of the optimal controls are nonnegative, 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 $\delta < 0$. None

B.3. Proof of Theorem 3. Inserting the expansion for g and the feedback controls (11) for δ into the HJB equation (7), and carrying out tedious but ultimately straightforward expansions, to order ς , (7) reduces to

$$(23) \quad \begin{aligned} o(\varsigma) = & \mathcal{D}g_0 + (\lambda^+ \delta_0^+ h_+(\delta_0^+) + \lambda^- \delta_0^- h_-(\delta_0^-)) \\ & + \alpha \left\{ q + (\mathcal{D} - \zeta) g_\alpha + \lambda^+ h_+(\delta_0) \left[\mathbb{S}_{q\lambda}^+ - \mathbb{S}_\lambda^+ \right] g_\alpha + \lambda^- h_-(\delta_0) \left[\mathbb{S}_{q\lambda}^- - \mathbb{S}_\lambda^- \right] g_\alpha \right\} \\ & + \varepsilon \left\{ \mathcal{D}g_\varepsilon + \lambda^+ h_+(\delta_0) \left(\left[\mathbb{S}_{q\lambda}^+ - \mathbb{S}_\lambda^+ \right] g_\varepsilon + \rho \mathbf{a}^+ \mathcal{S}_{q\lambda}^+ g_\alpha \right) \right. \\ & \quad \left. + \lambda^- h_-(\delta_0) \left(\left[\mathbb{S}_{q\lambda}^- - \mathbb{S}_\lambda^- \right] g_\varepsilon - \rho \mathbf{a}^- \mathcal{S}_{q\lambda}^- g_\alpha \right) \right\} \\ & + \phi \left\{ -q^2 + \mathcal{D}g_\phi + \lambda^+ h_+(\delta_0) \left[\mathbb{S}_{q\lambda}^+ - \mathbb{S}_\lambda^+ \right] g_\phi + \lambda^- h_-(\delta_0) \left[\mathbb{S}_{q\lambda}^- - \mathbb{S}_\lambda^- \right] g_\phi \right\}, \end{aligned}$$

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 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 α , ε , and ϕ to zero separately; (ii) write g_α , g_ε , and g_ϕ as in (14); and (iii) collect powers of q , and set them individually to zero.¹⁷ Then one finds the

¹⁷Note that this step is not an asymptotic expansion in q ; rather it is exact given the prescribed expansion in the other parameters.

following equations for the functions $b_\alpha(t)$, $b_\varepsilon(t, \boldsymbol{\lambda})$, $b_\phi(t, \boldsymbol{\lambda}, \boldsymbol{\kappa})$, and $c_\phi(t)$:

$$\begin{aligned} (24a) \quad & 0 = \mathcal{D}b_\alpha - \zeta b_\alpha + \lambda^+ [\mathbb{S}_\lambda^+ - 1] b_\alpha + \lambda^- [\mathbb{S}_\lambda^- - 1] b_\alpha + 1, \\ (24b) \quad & 0 = \mathcal{D}b_\varepsilon + \lambda^+ [\mathbb{S}_\lambda^+ - 1] b_\varepsilon + \lambda^- [\mathbb{S}_\lambda^- - 1] b_\varepsilon + \rho (\lambda^+ \mathbf{a}^+ - \lambda^- \mathbf{a}^-) b_\alpha, \\ (24c) \quad & 0 = \mathcal{D}b_\phi + \lambda^+ [\mathbb{S}_\lambda^+ - 1] b_\phi + \lambda^- [\mathbb{S}_\lambda^- - 1] b_\phi - 2h(\delta_0) (\lambda^+ - \lambda^-) c_\phi, \\ (24d) \quad & 0 = \mathcal{D}c_\phi + \lambda^+ [\mathbb{S}_\lambda^+ - 1] c_\phi + \lambda^- [\mathbb{S}_\lambda^- - 1] c_\phi - 1. \end{aligned}$$

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). More specifically, we apply the Feynman–Kac formula in Lemma 5 to link the solution of the derived partial integro-differential equation (PIDE) back to its stochastic representation, as presented in (15).

The functions a_α , a_ε , and a_ϕ 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. None

B.4. Proof of Corollary 4. Applying (14) for g_α , g_ε , and g_ϕ in Theorem 3 to (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

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

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

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

and we are done. None

B.5. Feynman–Kac formula for jump diffusions. Let X_t be a jump diffusion on \mathbb{R}^k with \mathbb{P} -generator \mathcal{L} . Suppose that some function $u(t, x)$ satisfies the following PIDE with boundary condition:

$$(25) \quad \begin{cases} (\partial_t + \mathcal{L})u(t, x) + f(t, x) &= V(t, x)u(t, x), \\ u(x, T) &= \varphi(x). \end{cases}$$

LEMMA 5 (Feynman–Kac formula). *The solution to (25) has the following stochastic representation:*

$$(26) \quad u(t, x) = \mathbb{E}^\mathbb{P} \left[\int_t^T e^{-\int_t^s V(z, X_z) dz} f(s, X_s) ds + e^{-\int_t^T V(z, X_z) dz} \varphi(X_T) \middle| X_t = x \right].$$

Proof. Suppose $u(t, x)$ is a solution to (25), and consider the process

$$Y(r, X_r) := \int_t^r e^{-\int_t^s V(z, X_z) dz} f(s, X_s) ds + e^{-\int_t^r V(z, X_z) dz} u(r, X_r).$$

It then follows that

$$\mathbb{E}_{t,x} [Y(T, X_T) - Y(t, X_t)] = \mathbb{E}_{t,x} \left[\int_t^T e^{-\int_t^r V_z dz} \{f_r + (\partial_r + \mathcal{L})u_r - V_r u_r\} dr \right] = 0.$$

The first equality is true by Dynkin's formula. The second equality is due to the condition in (25). Hence, we have

$$\begin{aligned} u(t, x) &= Y(t, x) = \mathbb{E}_{t,x} [Y(t, X_t)] = \mathbb{E}_{t,x} [Y(T, X_T)] \\ &= \mathbb{E}_{t,x} \left[\int_t^T e^{-\int_t^s V(z, X_z) dz} f(s, X_s) ds + e^{-\int_t^T V(z, X_z) dz} u(T, X_T) \right] \\ &= \mathbb{E}_{t,x} \left[\int_t^T e^{-\int_t^s V(z, X_z) dz} f(s, X_s) ds + e^{-\int_t^T V(z, X_z) dz} \varphi(X_T) \right]. \quad \square \end{aligned}$$

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.

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 6 (exact optimal controls for exponential and power-law). *Suppose that the scale parameter process κ_t^\pm is bounded away from zero almost surely. More specifically, assume that $\mathbb{P}[\inf_{t \in [0, T]} \kappa_t^\pm > 0] = 1$.*

1. *If $h^\pm(\delta; \kappa) = e^{-\kappa^\pm \delta}$ for $\delta > 0$, then the feedback control of the optimal trading strategy for the HJB equation (7) is given by*

$$(27a) \quad \delta_t^\pm = \max \left\{ \frac{1}{\kappa^\pm} - \left\{ \mathbb{S}_{q\lambda}^\pm g - \mathbb{S}_\lambda^\pm g \right\}, 0 \right\}.$$

2. *If $h^\pm(\delta; \kappa) = (1 + \kappa^\pm \delta)^\alpha$ for $\delta > 0$, then the feedback control of the optimal trading strategy for the HJB equation (7) is given by*

$$(27b) \quad \delta_t^\pm = \max \left\{ \frac{\alpha}{\alpha - 1} \left(\frac{1}{\kappa^\pm} - \left\{ \mathbb{S}_{q\lambda}^\pm g - \mathbb{S}_\lambda^\pm g \right\} \right), 0 \right\}.$$

Here, the ansatz $\Phi = x + q s + g(t, q, \alpha, \lambda, \kappa)$ with boundary condition $g(T, \cdot) = 0$ has been applied. Furthermore, the solutions in (27a)–(27b) are unique.

Proof. Applying the first order conditions to the supremum terms and using the specified ansatz leads, after some simplifications, to the stated result. We show this in detail for δ^- in (27a) only as the other cases are analogous. The relevant supremum term in the HJB equation in (7) simplifies to, after applying the ansatz $\Phi = x + q s + g(t, q, \alpha, \lambda, \kappa)$,

$$(28) \quad e^{-\delta^- \kappa^-} [\mathbb{S}_\lambda^- g - g + \delta^-] + (1 - e^{-\delta^- \kappa^-}) [\mathbb{S}_\lambda^- g - g].$$

Differentiating (28) with respect to δ^- and setting the resulting expression equal to zero yields (27a). Checking the second derivative of (28) verifies that this point is in fact a local maximum. If this point is negative, then the optimal δ^- is $\delta^- = 0$ by Assumption 2.

Uniqueness is trivial. \square

C.2. Explicit computation of b_ε . Rather than computing b_ε 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 7 (expected integrated drift). *The expected integrated drift is given by the expression*

$$(29) \quad \mathbb{E} \left[\int_t^T \alpha_s ds \middle| \lambda_t = \lambda, \alpha_t = \alpha \right] = \varepsilon b_\varepsilon(t, \lambda) + \alpha b_\alpha(t),$$

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

$$(30a) \quad A(t) = \zeta' \cdot \mathbf{B}(t),$$

$$(30b) \quad \mathbf{B}(t) = \frac{\rho}{\zeta} \left\{ \mathbf{A}^{-1} \left((T-t) \mathbf{I} - \mathbf{A}^{-1} \left(\mathbf{I} - e^{-\mathbf{A}(T-t)} \right) \right) - (\mathbf{A} - \zeta \mathbf{I})^{-1} \left(b_\alpha(t) \mathbf{I} - \mathbf{A}^{-1} \left(\mathbf{I} - e^{-\mathbf{A}(T-t)} \right) \right) \right\} \varepsilon \mathbf{a},$$

$$(30c) \quad \mathbf{C}(t) = \frac{\rho}{\zeta} \left\{ \mathbf{A}^{-1} \left(\mathbf{I} - e^{-\mathbf{A}(T-t)} \right) - (\mathbf{A} - \zeta \mathbf{I})^{-1} \left(e^{-\zeta(T-t)} \mathbf{I} - e^{-\mathbf{A}(T-t)} \right) \right\} \varepsilon \mathbf{a}.$$

Moreover, $\mathbf{a} = (-\mathbf{a}^-, \mathbf{a}^+)'$ and $\zeta = (\beta\theta, \beta\theta)'$.

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

$$(31) \quad (\partial_t + \mathcal{L})f + \lambda^+ (\mathbb{S}_\lambda^+ f - f) + \lambda^- (\mathbb{S}_\lambda^- f - f) + \alpha = 0,$$

where the infinitesimal generator of α and λ is

$$\mathcal{L} = \beta(\theta - \lambda^-) \partial_{\lambda^-} + \beta(\theta - \lambda^+) \partial_{\lambda^+} - \zeta \alpha \partial_\alpha + \frac{1}{2} \sigma^2 \partial_{\alpha\alpha}.$$

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

$$(32) \quad \begin{cases} \partial_t A(t) + \zeta' \cdot \mathbf{C}(t) = 0, \\ \partial_t \mathbf{C}(t) - \mathbf{A} \mathbf{C}(t) + \rho b_\alpha(t) \varepsilon \mathbf{a} = 0. \end{cases}$$

The solution of this coupled system is given by (30a)–(30c). The assertion that $A(t) + \lambda \cdot \mathbf{C}(t) = \varepsilon b_\varepsilon$ with b_ε provided in (15b) can be confirmed by (i) writing down the PDE which the function b_ε satisfies, (ii) noting that it admits an affine ansatz $A_\varepsilon(t) + \lambda \cdot \mathbf{C}_\varepsilon(t)$, and (iii) noting that 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 that they are equal. \square

C.3. Computing b_ϕ when risk-neutral fill probabilities are constants.

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

PROPOSITION 8 (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 given by*

$$(33) \quad b_\phi(t, \boldsymbol{\lambda}) = 2\xi' \left\{ \left(\mathbf{A}^{-1}(T-t) - \mathbf{A}^{-2} \left(\mathbf{I} - e^{-\mathbf{A}(T-t)} \right) \right) [\boldsymbol{\lambda} - \mathbf{A}^{-1}\boldsymbol{\zeta}] + \frac{1}{2}(T-t)^2 \mathbf{A}^{-1}\boldsymbol{\zeta} \right\},$$

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

Proof. Note that $\mathbb{E} \left[\int_t^T \lambda_u^\pm (T-u) du | \mathcal{F}_t \right] = \int_t^T \mathbb{E} [\lambda_u^\pm | \mathcal{F}_t] (T-u) du = \int_t^T m_t^\pm(u) (T-u) du$. Using the form of $m_t^\pm(u)$ provided in (21) and integrating over u implies that

$$(34) \quad \int_t^T \mathbf{m}_t(u) (T-u) du = \left(\mathbf{A}^{-1}(T-t) - \mathbf{A}^{-2} \left(\mathbf{I} - e^{-\mathbf{A}(T-t)} \right) \right) (\boldsymbol{\lambda}_t - \mathbf{A}^{-1}\boldsymbol{\zeta}) + \mathbf{A}^{-1}\boldsymbol{\zeta} \frac{1}{2}(T-t)^2.$$

This result is valid under the restriction that \mathbf{A} is invertible, which is implied by the arrival rate of MOs (2) and Lemma 1. Moreover, when $h_\pm(\delta_0^\pm; \boldsymbol{\kappa}) = h_\pm$ we have $b_\phi(t, \boldsymbol{\lambda}) = 2 \int_t^T \{h_+ \cdot m_t^+(u) - h_- \cdot m_t^-(u)\} (T-u) du$, and (33) follows immediately. \square

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 9 (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 $\mathbb{P}[\inf_{t \in [0, T]} \kappa_t^\pm > 0] = 1$, then $h_\pm(\delta_0^\pm; \boldsymbol{\kappa}) = e^{-1}$ is constant and Proposition 7 applies.

Example 10 (power fill rate). Take $\kappa^\pm = f^\pm(\boldsymbol{\kappa})$, where $f^\pm : \mathbb{R}^k \mapsto \mathbb{R}^+$ are continuous functions, and $\alpha^\pm > 1$ are fixed constants. If $h_\pm(\delta; \boldsymbol{\kappa}) = [1 + (\kappa^\pm \delta)^{\alpha^\pm}]^{-1}$ for $\delta > 0$ and $\mathbb{P}[\inf_{t \in [0, T]} \kappa_t^\pm > 0] = 1$, 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 λ_0^\pm are equal and κ^\pm are equal, then $b_\phi \equiv 0$.

C.4. Conditional mean of fill probability process.

LEMMA 11 (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

$$(35) \quad \tilde{m}_t^\pm(u) = \vartheta + \frac{\rho}{\xi} [\eta_\kappa m_t^\pm(u) + \nu_\kappa m_t^\mp(u)] + \left[\kappa_t^\pm - \vartheta - \frac{\rho}{\xi} (\eta_\kappa \lambda_t^\pm + \nu_\kappa \lambda_t^\mp) \right] e^{-\xi(u-t)},$$

where $m_t^\pm(u)$ are given in section B.1.

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

$$(36) \quad \frac{d\tilde{m}_t^\pm(u)}{du} + \xi \tilde{m}_t^\pm(u) = \xi \vartheta + \rho [\eta_\kappa m_t^\pm(u) + \nu_\kappa m_t^\mp(u)],$$

where $m_t^\pm(u)$ is given by (21). Solving (36) with the initial condition $\tilde{m}_t^\pm(t) = \kappa_t^\pm$ gives the stated result. \square

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 that would be associated with approximating a continuous time process by a discrete one (i.e., simulated interarrival times are correct up to machine precision).

1. Generate the duration until the next MO given the current level of activity $\lambda_{t_n}^\pm$.
 - 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 root¹⁸ 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 MO.
 - The probability that the MO 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 MO and $B_{n+1} = -1$ if it is a sell MO.
3. Decide whether the MO filled the agent's posted LO.
 - Compute the posted LO at the time of the MO $\lambda_{t_{n+1}}^\pm = \theta + (\lambda_{t_n}^\pm - \theta)e^{-\beta T_{n+1}}$.
 - Draw a uniform $u \sim U(0, 1)$.
 - If the MO was a sell (buy) order, then if $u < e^{-\kappa_t^- \delta_t^-}$ ($u < e^{-\kappa_t^+ \delta_t^+}$), the agent's buy (sell) LO 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:

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

Generate a third independent standard normal Z .

- Update price and drift. $S_{t_{n+1}} = S_{t_n} + \alpha_{t_n} \frac{1}{\zeta} (1 - e^{-\zeta T_{n+1}}) + Z_1 + \sigma \sqrt{T_{n+1}} Z$ 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, so set $H_{n+1} = 1$; otherwise set $H_{n+1} = 0$. Finally,

$$\begin{aligned}\lambda_{t_{n+1}}^\pm &= \theta + (\lambda_{t_n}^\pm - \theta) e^{-\beta T_{n+1}} + \left(\frac{1}{2}(1 \pm B_{n+1})\nu + \frac{1}{2}(1 \mp B_{n+1})\eta \right) H_{n+1}, \\ \kappa_{t_{n+1}}^\pm &= \vartheta + (\kappa_{t_n}^\pm - \vartheta) e^{-\xi T_{n+1}} + \left(\frac{1}{2}(1 \pm B_{n+1})\nu_\kappa + \frac{1}{2}(1 \mp B_{n+1})\eta_\kappa \right) H_{n+1}, \\ \alpha_{t_{n+1}} &= \frac{1}{2}(1 + B_{n+1})\epsilon^+ - \frac{1}{2}(1 - B_{n+1})\epsilon^- + \alpha_{t_{n+1}}.\end{aligned}$$

7. Repeat from step 1 until $t_{n+1} \geq T$.

¹⁸This is efficiently computed using the Lambert-W function since A_0 is typically small.

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 $PnL = X_T + q_T S_T (1 - c_{trans} \operatorname{sgn}(q_T))$, where c_{trans} is the *liquidation cost* (additional fee charged to the trader for taking liquidity from the market at time T due to forced liquidation of inventory).

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

Acknowledgments. The authors thank participants at the SAFI meeting at U. Michigan, the ICIAM meeting in Vancouver, and the 3C Risk Forum at the Fields Institute for Mathematical Sciences in Toronto; Universidad Carlos III de Madrid, University College London, and Imperial College; 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.

Bibliography.

REFERENCES

- [1] R. ALMGREN, Optimal execution with nonlinear impact functions and trading-enhanced risk, *Applied Mathematical Finance*, 10 (2003), pp. 1–18.
- [2] R. ALMGREN AND N. CHRISS, Optimal execution of portfolio transactions, *Journal of Risk*, 3 (2000).
- [3] M. AVELLANEDA AND S. STOIKOV, High-frequency trading in a limit order book, *Quantitative Finance*, 8 (2008), pp. 217–224.
- [4] P. BANK, H. M. SONER, AND M. VOSS, Hedging with temporary price impact, *Mathematics and Financial Economics*, 11 (2017), pp. 215–239.
- [5] N. BÄUERLE AND U. RIEDER, MDP algorithms for portfolio optimization problems in pure jump markets, *Finance and Stochastics*, 13 (2009), pp. 591–611.
- [6] E. BAYRAKTAR AND M. LUDKOVSKI, Liquidation in limit order books with controlled intensity, *Mathematical Finance*, 24 (2014), pp. 627–650.
- [7] D. P. BERTSEKAS AND S. SHREVE, Stochastic Optimal Control, vol. 139 of *Mathematics in Science and Engineering*, Academic Press, New York, 1978.
- [8] B. BOUCHARD, N.-M. DANG, AND C.-A. LEHALLE, Optimal control of trading algorithms: A general impulse control approach, *SIAM Journal on Financial Mathematics*, 2 (2011), pp. 404–438.
- [9] C. G. BOWSHER, Modelling security market events in continuous time: Intensity based, multivariate point process models, *Journal of Econometrics*, 141 (2007), pp. 876 – 912.
- [10] P. CARDALIAGUET AND C.-A. LEHALLE, Mean field game of controls and an application to trade crowding, *Mathematics and Financial Economics*, (2016), pp. 1–29.
- [11] Á. CARTEA, R. DONNELLY, AND S. JAIMUNGAL, Algorithmic trading with model uncertainty, *SIAM Journal on Financial Mathematics*, 8 (2017), pp. 635–671.
- [12] Á. CARTEA AND S. JAIMUNGAL, Modeling asset prices for algorithmic and high-frequency trading, *Applied Mathematical Finance*, 20(6) (2013), pp. 512–547.
- [13] —, Optimal execution with limit and market orders, *Quantitative Finance*, 15 (2015), pp. 1279–1291.
- [14] —, Risk metrics and fine tuning of high-frequency trading strategies, *Mathematical Finance*, 25 (2015), pp. 576–611.
- [15] —, A closed-form execution strategy to target volume weighted average price, *SIAM Journal on Financial Mathematics*, 7 (2016), pp. 760–785.
- [16] —, Incorporating order-flow into optimal execution, *Mathematics and Financial Economics*, 10 (2016), pp. 339–364.
- [17] Á. CARTEA, S. JAIMUNGAL, AND D. KINZEBULATOV, Algorithmic trading with learning, *International Journal of Theoretical and Applied Finance*, 19 (2016), p. 1650028.
- [18] Á. CARTEA, S. JAIMUNGAL, AND J. PENALVA, Algorithmic and high-frequency trading, Cambridge University Press, 2015.
- [19] Á. CARTEA AND J. PENALVA, Where is the value in high frequency trading?, *Quarterly Journal of Finance*, 2 (2012), pp. 1–46.

- [20] P. CASGRAIN AND S. JAIMUNGAL, Trading algorithms with learning in latent alpha models, Available at <https://ssrn.com/abstract=2871403>, (2016).
- [21] ———, Algorithmic trading with partial information: A mean field game approach, Available at <http://arxiv.org/abs/1803.04094>, (2018).
- [22] CFTC AND SEC, Findings regarding the market events of May 6, 2010, report, SEC, September 2010.
- [23] R. CONT, R. TALREJA, AND S. STOIKOV, A stochastic model for order book dynamics, Operations Research, 58 (2010), pp. 217–224.
- [24] D. EASLEY AND M. O'HARA, Time and the process of security price adjustment, The Journal of Finance, XLVII (1992), pp. 577–605.
- [25] EUROPEAN COMMISSION, Public consultation: Review of the markets in financial instruments directive (MiFID), 2010.
- [26] C. FREI AND N. WESTRAY, Optimal execution of a VWAP order: a stochastic control approach, Mathematical Finance, 25 (2015), pp. 612–639.
- [27] O. GUÉANT, The financial mathematics of market liquidity: From optimal execution to market making, vol. 33, CRC Press, 2016.
- [28] O. GUÉANT, C.-A. LEHALLE, AND J. FERNANDEZ TAPIA, Dealing with the inventory risk: a solution to the market making problem, Mathematics and Financial Economics, 7(4) (2013), pp. 477–507.
- [29] O. GUÉANT AND G. ROYER, VWAP execution and guaranteed VWAP, SIAM Journal on Financial Mathematics, 5 (2014), pp. 445–471.
- [30] F. GUILBAUD AND H. PHAM, Optimal high frequency trading with limit and market orders, Quantitative Finance, 13(1) (2013), pp. 79–94.
- [31] F. GUILBAUD AND H. PHAM, Optimal high-frequency trading in a pro rata microstructure with predictive information, Mathematical Finance, 25 (2015), pp. 545–575.
- [32] T. HO AND H. R. STOLL, Optimal dealer pricing under transactions and return uncertainty, Journal of Financial Economics, 9 (1981), pp. 47–73.
- [33] M. HUANG, P. E. CAINES, AND R. P. MALHAMÉ, Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ϵ -Nash equilibria, IEEE Trans. Autom. Control, 52 (2007), pp. 1560–1571.
- [34] X. HUANG, S. JAIMUNGAL, AND M. NOURIAN, Mean-field game strategies for optimal execution, Available at <http://ssrn.com/abstract=2578733>, (2015).
- [35] S. JAIMUNGAL AND D. KINZEBULATOV, Optimal Execution with a Price Limiter, SSRN eLibrary, <http://ssrn.com/abstract=2199889>, (2013).
- [36] S. JAIMUNGAL, D. KINZEBULATOV, AND D. RUBISOV, Optimal accelerated share repurchases, Applied Mathematical Finance, 24 (2017), pp. 216–245.
- [37] I. KHARROUBI AND H. PHAM, Optimal portfolio liquidation with execution cost and risk, SIAM Journal on Financial Mathematics, 1 (2010), pp. 897–931.
- [38] J. LARGE, Measuring the resiliency of an electronic limit order book, Journal of Financial Markets, 10 (2007), pp. 1–25.
- [39] J.-M. LASRY AND P.-L. LIONS, Mean field games, Jap. J. Math., 2 (2007), pp. 229–260.
- [40] M. O'HARA, Market Microstructure Theory, Wiley, 1998.
- [41] I. ROȘU, A dynamic model of the limit order book, Review of Financial Studies, 22 (2009), pp. 4601–4641.
- [42] A. SARKAR AND R. A. SCHWARTZ, Market sidedness: Insights into motives for trade initiation, The Journal of Finance, 64 (2009), pp. 375–423.
- [43] THE GOVERNMENT OFFICE FOR SCIENCE, LONDON, Foresight: The future of computer trading in financial markets, 2012.
- [44] I. TOKE, market making in an order book model and its impact on the spread, in Econophysics of Order-driven Markets, F. Abergel, B. Chakrabarti, A. Chakraborti, and M. Mitra, eds., New Economic Windows, Springer Milan, 2011, pp. 49–64.