MATHEMATICAL FINANCE
An International Journal of Mathematics, Statistics and Financial Economics

Mathematical Finance, Vol. 25, No. 3 (July 2015), 457–495

GENERAL INTENSITY SHAPES IN OPTIMAL LIQUIDATION

OLIVIER GUÉANT

UFR de Mathématiques, Laboratoire Jacques-Louis Lions, Université Paris-Diderot

CHARLES-ALBERT LEHALLE

Capital Fund Management

The classical literature on optimal liquidation, rooted in Almgren–Chriss models, tackles the optimal liquidation problem using a trade-off between market impact and price risk. It answers the general question of optimal scheduling but the very question of the actual way to proceed with liquidation is rarely dealt with. Our model, which incorporates both price risk and nonexecution risk, is an attempt to tackle this question using limit orders. The very general framework we propose to model liquidation with limit orders generalizes existing ones in two ways. We consider a risk-averse agent, whereas the model of Bayraktar and Ludkovski only tackles the case of a risk-neutral one. We consider very general functional forms for the execution process intensity, whereas Guéant, Lehalle and Fernandez-Tapia are restricted to exponential intensity. Eventually, we link the execution cost function of Almgren–Chriss models to the intensity function in our model, providing then a way to see Almgren–Chriss models as a limit of ours.

KEY WORDS: optimal liquidation, limit orders, stochastic optimal control, viscosity solutions

1. INTRODUCTION

Since the late nineties and the first papers on the impact of execution costs on trading strategies (e.g., Bertsimas and Lo 1998), an important literature has developed to tackle the problem of optimal liquidation. This literature, often rooted in the seminal papers by Almgren and Chriss (1999, 2001), has long been characterized by a trade-off between, on the one hand, market impact that encourages to trade slowly and, on the other hand, price risk that provides an incentive to trade fast.

The first family of models, following Almgren and Chriss, considered general instantaneous price impact (sometimes called execution cost) and linear permanent price impact. Several generalizations have been proposed such as an extension to random execution costs Almgren (2003), or stochastic volatility and stochastic liquidity

This research was partially supported by the Research Initiative "Microstructure des marchés financiers" under the aegis of the Europlace Institute of Finance. The authors wish to acknowledge the helpful conversations with Yves Achdou (Université Paris-Diderot), Guy Barles (Universitéde Tours), Erhan Bayraktar (University of Michigan), Bruno Bouchard (Université Paris-Dauphine), Jean-Michel Lasry (Université Paris-Dauphine), Mike Ludkovski (UC Santa Barbara), and Nizar Touzi (Ecole Polytechnique). The anonymous referees also need to be warmly thanked for their thorough reading and the numerous improvements following their remarks.

Manuscript received July 2012; final revision received June 2013.

Address correspondence to Olivier Guéant, UFR de Mathématiques, Laboratoire Jacques-Louis Lions, Université Paris-Diderot, 5 rue Thomas Mann, 75013 Paris, France; e-mail: gueant@ljll.univ-paris-diderot.fr.

DOI: 10.1111/mafi.12052

(Almgren 2009, 2012). Also, many objective functions to design optimal strategies have been proposed and discussed in order to understand the assumptions under which optimal strategies are deterministic (as opposed to adaptive). The initial mean-variance framework has been expressed in an expected utility setting (using a CARA utility function) in Schied, Schöneborn, and Tehranchi (2010), Schied and Schöneborn (2007), and Guéant (2012), a mean-quadratic-variation framework has been considered (Forsyth et al. 2012), an initial-time mean variance criterion has been discussed (Almgren and Lorenz 2007; Lorenz and Almgren 2011), and the very interesting case of a general utility function has recently been considered in Schied and Schöneborn (2009) to justify aggressive-in-the-money or passive-in-the-money strategies. Slightly different approaches have been proposed in this first generation of models (see, e.g., He and Mamaysky 2005; Huberman and Stanzl 2005; Bouchard et al. 2011; and Kharroubi and Pham 2010). They all derive from the initial models by Almgren and Chriss since market impact is either permanent or instantaneous. In other words, they do not take into account explicitly the resilience of the underlying order book.

Another family of models appeared following a paper by Obizhaeva and Wang (2005). In these models, the limit order book is directly modeled and the authors consider its resilient dynamic after each trade. This second generation of optimal liquidation models, based on transient market impact, has developed in recent years (Alfonsi, Fruth, and Schied 2008; Alfonsi, Fruth, and Schied 2010; Alfonsi and Schied 2010; and Predoiu, Shaikhet, and Shreve 2011). It raises the theoretical question of the functional forms for the transient market impact that are compatible with the absence of price manipulation (see Alfonsi, Schied, and Slynko 2012; Gatheral 2010; and Gatheral, Schied, and Slynko 2012).

All these models only make use of market orders, and hence only consider liquidity-taking strategies. They do not consider the possible use of limit orders that provide liquidity, nor the possible use of dark pools. Notwithstanding the preceding criticism, models *à la* Almgren–Chriss provide a rather acceptable answer to the macroscopic question of the optimal scheduling of liquidation—at least once the instantaneous market impact function has been replaced by an execution cost function modeling the ability to trade over short periods of time, with all possible means including limit orders, dark pools, and market orders. However, they do not answer the question of the optimal way to proceed in practice and the methods currently used in the industry are seldom based on optimal control models at the microscopic level. This paper provides such a model of optimal liquidation using limit orders, and can be used, either to liquidate a portfolio as a whole over a few hours, or on shorter periods of time to follow a trading curve, be it a TWAP curve, a VWAP curve, or an Almgren–Chriss (Implementation Shortfall) trading curve.

In our approach, a trader posts limit orders (thus providing liquidity instead of taking it) and does not know when his orders are going to be executed, if at all. As a consequence, the classical trade-off between market impact/execution cost and price risk is not central in our model. In our setting, a new risk is borne by the trader because execution is now a random process. This nonexecution risk is very different, in its nature, from price risk. This new risk characterizes the recent literature on optimal liquidation, which focuses on the optimal way to liquidate rather than on optimal scheduling. The recent literature on optimal liquidation focuses indeed on alternatives to the use of market orders. Kratz and Schöneborn (2014) proposed an approach inspired from models of the first family, but with both market orders and access to dark pools. Although they did not consider risk aversion with respect to the new risk borne by the trader, their model is one of the first

in this new family of models. The optimal split of large orders across liquidity pools has then been studied by Laruelle, Lehalle, and Pagès (2011a). Liquidation with limit orders has been developed by Bayraktar and Ludkovski (2014) for general intensity functions but only in a risk-neutral framework. Guéant, Lehalle, and Fernandez-Tapia (2012) considered in parallel the specific case of an exponential intensity for a risk-averse agent (see also Laruelle, Lehalle, and Pagès 2011b). More recently, Huitema (2012) considered liquidation involving market orders and limit orders, and Guilbaud and Pham (2015) also propose a liquidation model in a pro-rata microstructure.

One should also note that many models dealing with high-frequency market making have been developed that can be adapted to deal with optimal liquidation. Building on the model proposed by Ho and Stoll (1981) and then modified by Avellaneda and Stoikov (2008),[1] Cartea, Jaimungal, and Ricci (2014) considered a model with exponential intensity, market impact on the limit order book, adverse selection effects, and predictable α. Cartea and Jaimungal (2015) use a similar model to introduce risk measures for high-frequency trading. Earlier, the same authors proposed a model Cartea and Jaimungal (2013) in which the reference price is modeled by a hidden Markov model. Eventually, Guilbaud and Pham (2013) also used a model including both market orders and limit orders at best (and next to best) bid and ask together with stochastic spreads. As it is shown in Appendix B, our model can be used to model trading on both sides of the market. Our choice to focus on optimal liquidation is mainly justified by practitioners' needs.

In this paper, we generalize both Bayraktar and Ludkovski (2014) and Guéant et al. (2012). We indeed consider both general shapes for the intensity functions, and an investor with a CARA utility function. Moreover, we present a limiting case in which the size of the orders tends to 0 and we show that this limiting case is intrinsically linked to the usual continuous framework of Almgren and Chriss, although the latter framework only considers market orders. This limiting case helps to understand the meaning of intensity functions for quotes corresponding to marketable limit orders.

In Section 2, we present the setting of the model and the main hypotheses on execution. The third section is devoted to solving the partial differential equations arising from the control problem. Then, in Section 4, we provide illustrations of the model and we exhibit the asymptotic behavior of the quotes, generalizing therefore a result presented in Guéant et al. (2012). Section 5 is dedicated to the study of a limit regime that corresponds to orders of small size. This fifth section leads to results linked to those obtained for the *fluid limit* in Bayraktar and Ludkovski (2014), here in a risk-averse setting. This result is exploited in Section 6 that draws parallels between our model and the usual Almgren-Chriss framework.

2. OPTIMAL EXECUTION WITH LIMIT ORDERS: THE MODEL

2.1. Setup of the Model

Let us fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a filtration $(\mathcal{F}_t)_{t\geq 0}$ satisfying the usual conditions. We assume that all random variables and stochastic processes are defined on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$.

We consider a trader who has to liquidate a portfolio containing a quantity $q_0 > 0$ of a given stock. We suppose that the reference price of the stock (that can be considered

[1] See Guéant, Lehalle, and Fernandez-Tapia (2013) for the solution of the Avellaneda-Stoikov equations.

the best bid quote for example) follows a Brownian motion with a drift:

$$dS_t = \mu dt + \sigma dW_t.$$

To model limit orders and the execution process, we first introduce the set of admissible strategies:

$$\mathcal{A} = \left\{(\delta_t)_{t\in[0,T]} | (\delta_t)_t \text{ predictable process }, \delta^- \in L^\infty(\Omega \times [0,T])\right\}.$$

The trader under consideration continuously proposes an ask quote $S_t^a = S_t + \delta_t$. He will sell shares according to the rate of arrival of liquidity-taking orders at the price he quotes.

His inventory q^δ, which is the number of shares he holds, evolves according to the following dynamics:

$$dq_t^\delta = -\Delta dN_t^\delta,$$

where N^δ is a point process giving the number of executed orders, each order being of size Δ—we suppose that Δ is a fraction of q_0. The intensity process $(\lambda_t)_t$ of the point process N^δ that is the arrival rate of liquidity-taking orders depends on both the (ask) price quoted by the trader and the size of its orders:

$$\lambda_t = \Lambda_\Delta(S_t^a - S_t)1_{q_{t-}^\delta > 0} = \Lambda_\Delta(\delta_t)1_{q_{t-}^\delta > 0},$$

where $\Lambda_\Delta : \mathbb{R} \to \mathbb{R}_+$ satisfies the following assumptions:

- Λ_Δ is strictly decreasing—the cheaper the order price, the faster it will be executed,
- $\lim_{\delta\to+\infty} \Lambda_\Delta(\delta) = 0$,
- $\Lambda_\Delta \in C^2(\mathbb{R})$,
- $\Lambda_\Delta(\delta)\Lambda''_\Delta(\delta) \leq 2\Lambda'_\Delta(\delta)^2$.

As a consequence of his trades, the trader's cash account X^δ has the following dynamics:

$$dX_t^\delta = (S_t + \delta_t)\Delta dN_t^\delta.$$

Now, coming to the liquidation problem, the trader has a time horizon T to liquidate his shares and his goal is to optimize the expected utility of his P&L at time T. We focus on CARA utility functions so that the trader considers the following optimization problem:

$$\sup_{\delta\in\mathcal{A}} \mathbb{E}\left[-\exp\left(-\gamma\left(X_T^\delta + q_T^\delta\left(S_T - \ell\left(q_T^\delta\right)\right)\right)\right)\right],$$

where $\gamma > 0$ is the absolute risk aversion parameter characterizing the trader, where X_T^δ is the amount of cash at time T and where q_T^δ is the remaining quantity of shares at time T. In this setting, the trader can sell the shares remaining at time T in his portfolio at a price below the reference price, namely $S_T - \ell(q_T^\delta)$, the function ℓ being a positive and increasing penalization function, measuring execution cost.

We associate to this stochastic control problem the value function V_Δ defined by:

$$V_\Delta(t,x,q,s) = \sup_{\delta\in\mathcal{A}(t)} \mathbb{E}\left[-\exp\left(-\gamma\left(X_T^{t,x,\delta} + q_T^{t,q,\delta}\left(S_T^{t,s} - \ell\left(q_T^{t,q,\delta}\right)\right)\right)\right)\right],$$

where $\mathcal{A}(t)$ is the set of predictable processes on $[t, T]$, bounded from below and where:

$$dS_\tau^{t,s} = \mu d\tau + \sigma dW_\tau, \qquad S_t^{t,s} = s,$$

$$dX_\tau^{t,x,\delta} = (S_\tau + \delta_\tau)\Delta dN_\tau, \qquad X_t^{t,x,\delta} = x,$$

$$dq_\tau^{t,q,\delta} = -\Delta dN_\tau, \qquad q_t^{t,q,\delta} = q,$$

the point process N having stochastic intensity $(\lambda_\tau)_\tau$ with $\lambda_\tau = \Lambda_\Delta(\delta_\tau)1_{q_{\tau-}>0}$.

This setting deserves several remarks. First, orders are of constant size Δ. This is a modeling choice corresponding to the way practitioners proceed with liquidation, Δ being then a fraction of the average trade size (ATS). Also, we implicitly assume that our orders are either entirely filled or not filled at all. In other words, there is no partial fill in this model. This hypothesis is a questionable one since partial fills are common in practice. When using the model in practice, one can always consider a convex combination of optimal quotes between two multiples of Δ. Allowing for partial fills would make the model more realistic. However, it is complicated from a mathematical point of view.

A second important point regards negative δs. We indeed assume that Λ_Δ is defined on the entire real line and not only on $\mathbb{R}_+$. Our model allows to post orders at a price below the reference price. If the reference price is the first bid quote, these orders correspond to marketable limit orders. One may then wonder why there is execution uncertainty associated to these orders. An answer is linked to high-frequency traders whose capacity to rapidly cancel trades forces practitioners to use *fill and/or kill* orders or other types of marketable limit orders and not market orders. Also, considering the entire real line allows to introduce indirectly execution costs for liquidity-taking orders. We shall see in Section 6 that there is a link between the execution cost functions of Almgren–Chriss models and the intensity functions on $\{\delta < 0\}$. It is noteworthy that if one wants to avoid negative δ, adding a hard constraint $\delta \geq 0$ is also possible and does not raise any difficulty. This constrained framework is discussed later in this paper (see Section 3.3). To avoid negative δ, some authors (see, for instance, Bayraktar and Ludkovski 2014) considered an intensity function that blows up at $\delta = 0$. A natural consequence of this modeling choice is that there is unlimited liquidity available at $\delta = 0$. This is not a correct approach in our view.

The third and last point regards the structural assumption $\Lambda_\Delta(\delta)\Lambda''_\Delta(\delta) \leq 2\Lambda'_\Delta(\delta)^2$. This hypothesis, already present in Bayraktar and Ludkovski (2014), is a sufficient condition to guarantee uniqueness of the optimal trading quote. To understand the intuition, let us consider the expected PnL when posting an order at a distance δ from the reference price. This expected PnL is proportional to $\delta\Lambda_\Delta(\delta)$: δ is the premium over the reference price and $\Lambda_\Delta(\delta)$ is the instantaneous probability that a trade takes place at a distance δ from the reference price. A natural condition for this expression to have a unique maximizer is: $\Lambda_\Delta\Lambda''_\Delta < 2\Lambda'^2_\Delta$. In our case, the inequality can be binding because of risk-aversion. It is noteworthy that our framework can be used even if this hypothesis is relaxed.[2] However, the consequence is that there may be multiple optimal quotes. For the sake of exposition, we chose to present the model under this structural assumption.

[2]The same is true for the assumption $\Lambda_\Delta \in C^2$.

2.2. A System of ODEs for the Value Function

The optimization problem set up in the preceding paragraphs can be solved using classical Bellman tools. To this purpose, we introduce the Hamilton–Jacobi–Bellman (HJB) equation associated to the optimization problem, where the unknown u_Δ is going to be equal to the value function V_Δ defined above:

$$\text{(HJB)} \qquad 0 = \partial_t u_\Delta(t,x,q,s) + \mu \partial_s u_\Delta(t,x,q,s) + \frac{1}{2}\sigma^2 \partial^2_{ss} u_\Delta(t,x,q,s) + \sup_{s^a} \Lambda_\Delta(s^a - s)\left[u_\Delta(t, x + \Delta s^a, q - \Delta, s) - u_\Delta(t,x,q,s)\right],$$

with the final condition:

$$u_\Delta(T,x,q,s) = -\exp\left(-\gamma\left(x + q(s - \ell(q))\right)\right),$$

and the boundary condition:

$$u_\Delta(t,x,0,s) = -\exp(-\gamma x).$$

Since we use a CARA function, we can factor out the mark-to-market (MtM) value $x + qs$ of the portfolio. This remark leads to considering the change of variables $u_\Delta(t,x,q,s) = -\exp(-\gamma(x + qs + \theta_\Delta(t,q)))$. In that case, the above HJB equation with four variables is (formally) reduced to the following system of ordinary differential equations indexed by q:

$$\text{(HJ}_{\theta_\Delta}\text{)} \qquad 0 = \gamma \partial_t \theta_\Delta(t,q) + \gamma\mu q - \frac{1}{2}\gamma^2\sigma^2 q^2 + H_\Delta\left(\frac{\theta_\Delta(t,q) - \theta_\Delta(t,q-\Delta)}{\Delta}\right),$$

with

$$\theta_\Delta(T,q) = -\ell(q)q, \qquad \theta_\Delta(t,0) = 0,$$

where

$$H_\Delta(p) = \sup_\delta \Lambda_\Delta(\delta)\left(1 - e^{-\gamma\Delta(\delta - p)}\right).$$

3. SOLUTION OF THE OPTIMAL CONTROL PROBLEM

This section aims at solving the optimal control problem set up in the preceding section. We first concentrate on the equation (HJ$_{\theta_\Delta}$). Then, we provide a verification theorem that indeed gives a solution to the control problem and characterizes in a simple way the optimal quotes. The last subsection is dedicated to the addition of a hard constraint $\delta \geq \delta^{\min}$.

3.1. A Solution to (HJ$_{\theta_\Delta}$)

We start with a lemma about the hamiltonian function H_Δ.

LEMMA 3.1. *Let us define* $L_\Delta(p,\delta) = \Lambda_\Delta(\delta)\left(1 - e^{-\gamma\Delta(\delta - p)}\right)$.

$\forall p \in \mathbb{R}$, $\delta \mapsto L_\Delta(p,\delta)$ is strictly increasing on $(-\infty, \tilde{\delta}^*_\Delta(p)]$ and strictly decreasing on $[\tilde{\delta}^*_\Delta(p), +\infty)$, where $\tilde{\delta}^*_\Delta(p)$ is uniquely characterized by $\left(E_{\delta^*_\Delta}\right)$:

$$\tilde{\delta}^*_\Delta(p) - \frac{1}{\gamma\Delta}\log\left(1 - \gamma\Delta\frac{\Lambda_\Delta(\tilde{\delta}^*_\Delta(p))}{\Lambda'_\Delta(\tilde{\delta}^*_\Delta(p))}\right) = p, \qquad \left(E_{\delta^*_\Delta}\right).$$

Moreover, $p \mapsto \tilde{\delta}^*_\Delta(p)$ is a C^1 function.
Subsequently, H_Δ is a C^1 function with:

$$H_\Delta(p) = \gamma\Delta\frac{\Lambda_\Delta(\tilde{\delta}^*_\Delta(p))^2}{\gamma\Delta\Lambda_\Delta(\tilde{\delta}^*_\Delta(p)) - \Lambda'_\Delta(\tilde{\delta}^*_\Delta(p))}.$$

Proof. Straightforwardly, $\delta \mapsto L_\Delta(p,\delta)$ is strictly increasing on $(-\infty, p]$.

Regarding the behavior of this function on $[p, +\infty)$, let us notice that $L_\Delta(p,p) = 0$ and that $\lim_{\delta\to+\infty} L_\Delta(p,\delta) = 0$.

Now, if we differentiate, we get

$$\partial_\delta L_\Delta(p,\delta) = \Lambda'_\Delta(\delta)\left(1 - e^{-\gamma\Delta(\delta-p)}\right) + \gamma\Delta\Lambda_\Delta(\delta)e^{-\gamma\Delta(\delta-p)}.$$

Hence, $\partial_\delta L_\Delta(p,p) = \gamma\Delta\Lambda_\Delta(p) > 0$ and there is at least one $\delta^* \in (p, +\infty)$ such that $\partial_\delta L_\Delta(p,\delta^*) = 0$.

Such a δ^* must satisfy:

$$\delta^* - \frac{1}{\gamma\Delta}\log\left(1 - \gamma\Delta\frac{\Lambda_\Delta(\delta^*)}{\Lambda'_\Delta(\delta^*)}\right) = p.$$

Now, $f(x) = x - \frac{1}{\gamma\Delta}\log\left(1 - \gamma\Delta\frac{\Lambda_\Delta(x)}{\Lambda'_\Delta(x)}\right)$ defines a strictly increasing function since

$$\begin{aligned} f'(x) &= 1 + \frac{\left(\frac{\Lambda_\Delta(x)}{\Lambda'_\Delta(x)}\right)'}{1 - \gamma\Delta\frac{\Lambda_\Delta(x)}{\Lambda'_\Delta(x)}} \\ &= 1 + \frac{\Lambda'_\Delta(x)^2 - \Lambda_\Delta(x)\Lambda''_\Delta(x)}{\Lambda'_\Delta(x)^2 - \gamma\Delta\Lambda_\Delta(x)\Lambda'_\Delta(x)} \\ &= \frac{-\gamma\Delta\Lambda_\Delta(x)\Lambda'_\Delta(x)}{\Lambda'_\Delta(x)^2 - \gamma\Delta\Lambda_\Delta(x)\Lambda'_\Delta(x)} + \frac{2\Lambda'_\Delta(x)^2 - \Lambda_\Delta(x)\Lambda''_\Delta(x)}{\Lambda'_\Delta(x)^2 - \gamma\Delta\Lambda_\Delta(x)\Lambda'_\Delta(x)} \end{aligned}$$

is strictly positive because of the hypotheses onΛ_Δ.

Hence δ^*, defined by $f(\delta^*) = p$, is unique and $L_\Delta(p,\cdot)$ is strictly increasing on $(-\infty, \tilde{\delta}^*_\Delta(p)]$ and strictly decreasing on $[\tilde{\delta}^*_\Delta(p), +\infty)$, where $\tilde{\delta}^*_\Delta(p)$ is uniquely characterized by:

$$f(\tilde{\delta}^*_\Delta(p)) = \tilde{\delta}^*_\Delta(p) - \frac{1}{\gamma\Delta}\log\left(1 - \gamma\Delta\frac{\Lambda_\Delta(\tilde{\delta}^*_\Delta(p))}{\Lambda'_\Delta(\tilde{\delta}^*_\Delta(p))}\right) = p.$$

Using the implicit function theorem, this also gives that $p \mapsto \tilde{\delta}^*_\Delta(p)$ is a C^1 function.

Plugging the relation for $\tilde{\delta}^*_\Delta(p)$ in the definition of H_Δ then gives the last part of the lemma. □

Now, we are going to prove a comparison principle for the system of ODEs. This result is useful in two ways. First, it gives a priori bounds that will allow us to prove the existence of a solution to (HJ_{θ_Δ}). Second, it will provide bounds to θ_Δ independently of Δ in Section 5, when we shall consider the limiting behavior of θ_Δ as $\Delta \to 0$.

PROPOSITION 3.2. *(Comparison principle) Let $\tau \in [0, T)$, $k \in \mathbb{N}$.*
Let $\underline{\theta}_\Delta : [\tau, T] \times \{0, \Delta, \dots, k\Delta\} \to \mathbb{R}$ be a C^1 function with respect to time with

$$\forall q \in \{0, \Delta, \dots, k\Delta\}, \quad \underline{\theta}_\Delta(T, q) \leq -\ell(q)q \qquad \forall t \in [\tau, T], \quad \underline{\theta}_\Delta(t, 0) \leq 0,$$

and

$$0 \leq \gamma \partial_t \underline{\theta}_\Delta(t, q) + \gamma \mu q - \frac{1}{2}\gamma^2\sigma^2 q^2 + H_\Delta\left(\frac{\underline{\theta}_\Delta(t, q) - \underline{\theta}_\Delta(t, q - \Delta)}{\Delta}\right).$$

Let $\overline{\theta}_\Delta : [\tau, T] \times \{0, \Delta, \dots, k\Delta\} \to \mathbb{R}$ be a C^1 function with respect to time with

$$\forall q \in \{0, \Delta, \dots, k\Delta\}, \quad \overline{\theta}_\Delta(T, q) \geq -\ell(q)q \qquad \forall t \in [\tau, T], \quad \overline{\theta}_\Delta(t, 0) \geq 0,$$

and

$$0 \geq \gamma \partial_t \overline{\theta}_\Delta(t, q) + \gamma \mu q - \frac{1}{2}\gamma^2\sigma^2 q^2 + H_\Delta\left(\frac{\overline{\theta}_\Delta(t, q) - \overline{\theta}_\Delta(t, q - \Delta)}{\Delta}\right).$$

Then,

$$\overline{\theta}_\Delta \geq \underline{\theta}_\Delta.$$

Proof. Let $\alpha > 0$.
Let us consider a point (t^*_α, q^*_α) such that

$$\underline{\theta}_\Delta(t^*_\alpha, q^*_\alpha) - \overline{\theta}_\Delta(t^*_\alpha, q^*_\alpha) - \alpha(T - t^*_\alpha) = \sup_{(t,q)\in[\tau,T]\times\{0,\dots,k\Delta\}} \underline{\theta}_\Delta(t, q) - \overline{\theta}_\Delta(t, q) - \alpha(T - t).$$

If $t^*_\alpha \neq T$ and $q^*_\alpha \neq 0$, then

$$\partial_t \underline{\theta}_\Delta(t^*_\alpha, q^*_\alpha) - \partial_t \overline{\theta}_\Delta(t^*_\alpha, q^*_\alpha) \leq -\alpha.$$

By definition of (t^*_α, q^*_α), since $q^*_\alpha \neq 0$:

$$\underline{\theta}_\Delta(t^*_\alpha, q^*_\alpha) - \overline{\theta}_\Delta(t^*_\alpha, q^*_\alpha) \geq \underline{\theta}_\Delta(t^*_\alpha, q^*_\alpha - \Delta) - \overline{\theta}_\Delta(t^*_\alpha, q^*_\alpha - \Delta),$$

and

$$\underline{\theta}_\Delta(t^*_\alpha, q^*_\alpha) - \underline{\theta}_\Delta(t^*_\alpha, q^*_\alpha - \Delta) \geq \overline{\theta}_\Delta(t^*_\alpha, q^*_\alpha) - \overline{\theta}_\Delta(t^*_\alpha, q^*_\alpha - \Delta).$$

Now, by definition of the functions $\underline{\theta}_\Delta$ and $\overline{\theta}_\Delta$, we have

$$\begin{aligned} 0 \leq{}& \gamma\left[\partial_t \underline{\theta}_\Delta(t^*_\alpha, q^*_\alpha) - \partial_t \overline{\theta}_\Delta(t^*_\alpha, q^*_\alpha)\right] \\ &+ \left[H_\Delta\left(\frac{\underline{\theta}_\Delta(t^*_\alpha, q^*_\alpha) - \underline{\theta}_\Delta(t^*_\alpha, q^*_\alpha - \Delta)}{\Delta}\right) - H_\Delta\left(\frac{\overline{\theta}_\Delta(t^*_\alpha, q^*_\alpha) - \overline{\theta}_\Delta(t^*_\alpha, q^*_\alpha - \Delta)}{\Delta}\right)\right]. \end{aligned}$$

Since H_Δ is a decreasing function, we have

$$H_\Delta\left(\frac{\underline{\theta}_\Delta(t_\alpha^*, q_\alpha^*) - \underline{\theta}_\Delta(t_\alpha^*, q_\alpha^* - \Delta)}{\Delta}\right) \leq H_\Delta\left(\frac{\overline{\theta}_\Delta(t_\alpha^*, q_\alpha^*) - \overline{\theta}_\Delta(t_\alpha^*, q_\alpha^* - \Delta)}{\Delta}\right).$$

This leads to $0 \leq -\gamma\alpha$, which is not possible.
Therefore, $t_\alpha^* = T$ or $q_\alpha^* = 0$ so that

$$\sup_{(t,q)\in[\tau,T]\times\{0,\ldots,k\Delta\}} \underline{\theta}_\Delta(t,q) - \overline{\theta}_\Delta(t,q) - \alpha(T-t)$$

$$= \max\left(\sup_{q\in\{0,\ldots,k\Delta\}} \underline{\theta}_\Delta(T,q) - \overline{\theta}_\Delta(T,q),\ \sup_{t\in[\tau,T]} \underline{\theta}_\Delta(t,0) - \overline{\theta}_\Delta(t,0) - \alpha(T-t)\right) \leq 0.$$

Thus, $\forall (t,q), \underline{\theta}_\Delta(t,q) - \overline{\theta}_\Delta(t,q) \leq \alpha(T-t) \leq \alpha T$. Sending α to 0 proves our result. □

We are now ready to prove that the equation $(\mathrm{HJ}_{\theta_\Delta})$ has a unique solution.

PROPOSITION 3.3. *There exists a unique function* $\theta_\Delta : [0,T] \times \Delta\mathbb{N} \to \mathbb{R}$ *such that:*

- $t \in [0,T] \mapsto (\theta_\Delta(\cdot, q))_{q\in\Delta\mathbb{N}}$ *is continuously differentiable.*
- θ_Δ *is a solution of* $(\mathrm{HJ}_{\theta_\Delta})$.

Proof. We proceed by induction on q. For $q = 0$, we have by definition $\theta_\Delta(t,0) = 0$.
Now, for a given $q \in \Delta\mathbb{N}^*$, let us suppose that $\theta_\Delta(\cdot, q') : [0,T] \to \mathbb{R}$ is a C^1 function $\forall q' \leq q - \Delta$. Then, the ODE

$$0 = \gamma\partial_t\theta_\Delta(t,q) + \gamma\mu q - \frac{1}{2}\gamma^2\sigma^2q^2 + H_\Delta\left(\frac{\theta_\Delta(t,q) - \theta_\Delta(t,q-\Delta)}{\Delta}\right),$$

with the terminal condition

$$\theta_\Delta(T,q) = -\ell(q)q,$$

satisfies the assumptions of Cauchy–Lipschitz theorem. Consequently, there exists a unique solution $t \mapsto \theta_\Delta(t,q)$ on a maximal interval that is a subinterval of $[0,T]$ and we want to show that this subinterval is $[0,T]$ itself.

To prove this, let suppose by contradiction that $(\tilde{t}, T]$ is the maximal interval with $\tilde{t} \geq 0$.

Let us notice that because H_Δ is positive, $t \mapsto \theta_\Delta(t,q) - \mu q(T-t) + \frac{1}{2}\gamma\sigma^2q^2(T-t)$ is decreasing. Hence, the only possibility for $(\tilde{t}, T]$ to be a maximal interval in $[0,T]$ is that $\lim_{t\to\tilde{t}^+} \theta_\Delta(t,q) = +\infty$.

Now let us consider $\eta > 0$, $k = \frac{q}{\Delta}$ and $\tau = \tilde{t} + \eta$. We define on $[\tau, T] \times \{\Delta, \ldots, q\}$ the two functions $\underline{\theta}_\Delta$ and $\overline{\theta}_\Delta$ defined by

$$\underline{\theta}_\Delta = \theta_\Delta$$

and

$$\forall q' \leq q, \overline{\theta}_\Delta(t,q') = \mu^+q(T-t) + \frac{1}{\gamma}H_\Delta(0)(T-t).$$

These two functions satisfy the assumptions of the above comparison principle. We indeed have that

$$\forall q' \leq q, \overline{\theta}_\Delta(T,q') = 0 \geq -\ell(q')q' \qquad \forall t \in [\tau,T], \overline{\theta}_\Delta(t,0) = \mu^+q(T-t) + \frac{1}{\gamma}H_\Delta(0)(T-t) \geq 0,$$

and

$$\forall t \in [\tau, T],\ \ \forall q' \in \{\Delta, \dots, q\}, \gamma \partial_t \overline{\theta}_\Delta(t, q') + \gamma \mu q' - \frac{1}{2}\gamma^2\sigma^2 q'^2 + H_\Delta\left(\frac{\overline{\theta}_\Delta(t, q') - \overline{\theta}_\Delta(t, q' - \Delta)}{\Delta}\right)$$

$$= -\gamma \frac{1}{\gamma} H_\Delta(0) - \gamma \mu^+ q + \gamma \mu q' - \frac{1}{2}\gamma^2\sigma^2 q'^2 + H_\Delta(0)$$

$$= -\gamma(\mu^+ q - \mu q') - \frac{1}{2}\gamma^2\sigma^2 q'^2 \leq 0.$$

Hence, $\forall \eta > 0, \forall t \in [\tilde{t} + \eta, T], \theta_\Delta(t, q) \leq \mu^+ q(T - \tilde{t}) + \frac{1}{\gamma} H_\Delta(0)(T - \tilde{t})$, in contradiction with the fact that $\lim_{t \to \tilde{t}^+} \theta_\Delta(t, q) = +\infty$.
Hence, $t \mapsto \theta_\Delta(t, q)$ is defined on $[0, T]$ and this proves the result. □

3.2. Verification Theorem and Optimal Quotes

Now, we can solve the initial optimal control problem and find the optimal quotes at which the trader should post his limit orders.

THEOREM 3.4 (Verification theorem and optimal quotes). *Let us consider the solution θ_Δ of the system* ($\text{HJ}_{\theta_\Delta}$).
Let us define $u_\Delta(t, x, q, s) = -\exp(-\gamma(x + qs + \theta_\Delta(t, q)))$.
We have:

- *u_Δ is a solution to (HJB),*
- *u_Δ is equal to the value function V_Δ.*

Moreover, the optimal ask quote $S^a_t = S_t + \delta^*_\Delta(t)$, for $q_t > 0$, is characterized by:

$$\delta^*_\Delta(t) = \tilde{\delta}^*_\Delta\left(\frac{\theta_\Delta(t, q_t) - \theta_\Delta(t, q_t - \Delta)}{\Delta}\right),$$

where $\tilde{\delta}^*_\Delta(\cdot)$ is the function defined in Lemma 3.1.

Proof. From the very definition of θ_Δ and u_Δ, it is straightforward to see that u_Δ is a solution of (HJB).

We indeed have that the boundary condition and the terminal condition are satisfied, and for $q \geq \Delta$, we have that

$$\partial_t u_\Delta(t, x, q, s) + \mu \partial_s u_\Delta(t, x, q, s) + \frac{1}{2}\sigma^2 \partial^2_{ss} u_\Delta(t, x, q, s) + \sup_\delta \Lambda_\Delta(\delta)\left[u_\Delta(t, x + \Delta s + \Delta\delta, q - \Delta, s) - u_\Delta(t, x, q, s)\right]$$
$$= -\gamma u_\Delta(t, x, q, s)\partial_t \theta_\Delta(t, q) - \gamma \mu q u_\Delta(t, x, q, s) + \frac{1}{2}\sigma^2\gamma^2 q^2 u_\Delta(t, x, q, s)$$

$$+ \sup_\delta \Lambda_\Delta(\delta) u_\Delta(t, x, q, s)\left(\exp\left(-\gamma(\Delta\delta - (\theta_\Delta(t, q) - \theta_\Delta(t, q - \Delta)))\right) - 1\right)$$

$$= -u_\Delta(t,x,q,s)\Bigg[\gamma\partial_t\theta_\Delta(t,q) + \gamma\mu q - \frac{1}{2}\sigma^2\gamma^2q^2$$

$$+ \sup_\delta \Lambda_\Delta(\delta)\left(1 - \exp\left(-\gamma(\Delta\delta - (\theta_\Delta(t,q) - \theta_\Delta(t,q-\Delta)))\right)\right)\Bigg]$$

$$= -u_\Delta(t,x,q,s)\left[\gamma\partial_t\theta_\Delta(t,q) + \gamma\mu q - \frac{1}{2}\sigma^2\gamma^2q^2 + H_\Delta\left(\frac{\theta_\Delta(t,q) - \theta_\Delta(t,q-\Delta)}{\Delta}\right)\right] = 0.$$

Now, we need to verify that u_Δ is indeed the value function associated to the problem and to prove that our candidate $(\delta^*_\Delta)_t$ is indeed the optimal control. To that purpose, let us consider a control $\delta \in \mathcal{A}(t)$ and let us consider the following processes for $\tau \in [t, T]$:

$$dS^{t,s}_\tau = \mu d\tau + \sigma dW_\tau, \qquad S^{t,s}_t = s,$$

$$dX^{t,x,\delta}_\tau = (S_\tau + \delta_\tau)\Delta dN_\tau, \qquad X^{t,x,\delta}_t = x,$$

$$dq^{t,q,\delta}_\tau = -\Delta dN_\tau, \qquad q^{t,q,\delta}_t = q,$$

where the point process has stochastic intensity $(\lambda_\tau)_\tau$ with $\lambda_\tau = \Lambda_\Delta(\delta_\tau)1_{q_{\tau-}>0}$.[3]

Now, let us write Itô's formula for u_Δ:[4]

$$u_\Delta\left(T, X^{t,x,\delta}_{T-}, q^{t,q,\delta}_{T-}, S^{t,s}_T\right) = u_\Delta(t,x,q,s)$$

$$+ \int_t^T \left(\partial_\tau u_\Delta\left(\tau, X^{t,x,\delta}_{\tau-}, q^{t,q,\delta}_{\tau-}, S^{t,s}_\tau\right) + \mu\partial_s u_\Delta\left(\tau, X^{t,x,\delta}_{\tau-}, q^{t,q,\delta}_{\tau-}, S^{t,s}_\tau\right)\right.$$

$$\left. + \frac{\sigma^2}{2}\partial^2_{ss}u_\Delta(\tau, X^{t,x,\delta}_{\tau-}, q^{t,q,\delta}_{\tau-}, S^{t,s}_\tau)\right)d\tau$$

$$+ \int_t^T \left(u_\Delta\left(\tau, X^{t,x,\delta}_{\tau-} + \Delta S^{t,s}_\tau + \Delta\delta_\tau, q^{t,q,\delta}_{\tau-} - \Delta, S^{t,s}_\tau\right) - u_\Delta\left(\tau, X^{t,x,\delta}_{\tau-}, q^{t,q,\delta}_{\tau-}, S^{t,s}_\tau\right)\right)\lambda_\tau d\tau$$

$$+ \int_t^T \mu\partial_s u_\Delta\left(\tau, X^{t,x,\delta}_{\tau-}, q^{t,q,\delta}_{\tau-}, S^{t,s}_\tau\right)d\tau + \int_t^T \sigma\partial_s u_\Delta\left(\tau, X^{t,x,\delta}_{\tau-}, q^{t,q,\delta}_{\tau-}, S^{t,s}_\tau\right)dW_\tau$$

$$+ \int_t^T \left(u_\Delta\left(\tau, X^{t,x,\delta}_{\tau-} + \Delta S^{t,s}_\tau + \Delta\delta_\tau, q^{t,q,\delta}_{\tau-} - \Delta, S^{t,s}_\tau\right) - u_\Delta\left(\tau, X^{t,x,\delta}_{\tau-}, q^{t,q,\delta}_{\tau-}, S^{t,s}_\tau\right)\right)dM_\tau,$$

where M is the compensated process associated to N for the intensity process $(\lambda_\tau)_\tau$.

Now, we have to ensure that the last two integrals consist of martingales so that their mean is 0. To that purpose, let us notice that $\partial_s u = -\gamma q u$, and hence, since the process $q^{t,q,\delta}$ takes values between 0 and q, we just have to prove that

$$\mathbb{E}\left[\int_t^T u_\Delta\left(\tau, X^{t,x,\delta}_{\tau-}, q^{t,q,\delta}_{\tau-}, S^{t,s}_\tau\right)^2 d\tau\right] < +\infty,$$

$$\mathbb{E}\left[\int_t^T \left|u_\Delta\left(\tau, X^{t,x,\delta}_{\tau-} + \Delta S^{t,s}_\tau + \Delta\delta_\tau, q^{t,q,\delta}_{\tau-} - \Delta, S^{t,s}_\tau\right)\right|\lambda_\tau d\tau\right] < +\infty,$$

[3]This intensity being bounded since δ is bounded from below.

[4]The equality is still valid when $q_\tau = 0$ because of the boundary condition for u_Δ, and because the intensity process is then assumed to be 0.

and

$$\mathbb{E}\left[\int_t^T \left|u_\Delta\left(\tau,\, X^{t,x,\delta}_{\tau-},\, q^{t,q,\delta}_{\tau-},\, S^{t,s}_\tau\right)\right| \lambda_\tau d\tau\right] < +\infty.$$

We have

$$u_\Delta\left(\tau,\, X^{t,x,\delta}_\tau,\, q^{t,q,\delta}_\tau,\, S^{t,s}_\tau\right)^2 \leq \exp\left(2\gamma\|\theta_\Delta\|_\infty\right)\exp\left(-2\gamma\left(X^{t,x,\delta}_\tau + q^{t,q,\delta}_\tau S^{t,s}_\tau\right)\right)$$

$$\leq \exp\left(2\gamma\|\theta_\Delta\|_\infty\right)\exp\left(-2\gamma\left(x - q\|\delta^-\|_\infty + 2q\inf_{\tau\in[t,T]} S^{t,s}_\tau 1_{\inf_{\tau\in[t,T]} S^{t,s}_\tau<0}\right)\right)$$

$$\leq \exp\left(2\gamma\|\theta_\Delta\|_\infty\right)\exp\left(-2\gamma(x - q\|\delta^-\|_\infty)\right)\left(1 + \exp\left(-2\gamma q\inf_{\tau\in[t,T]} S^{t,s}_\tau\right)\right).$$

Hence,

$$\mathbb{E}\left[\int_t^T u_\Delta(\tau,\, X^{t,x,\delta}_\tau,\, q^{t,q,\delta}_\tau,\, S^{t,s}_\tau)^2 d\tau\right] = \mathbb{E}\left[\int_t^T u_\Delta(\tau,\, X^{t,x,\delta}_{\tau-},\, q^{t,q,\delta}_{\tau-},\, S^{t,s}_\tau)^2 d\tau\right]$$

$$\leq \exp\ \left(2\gamma\|\theta_\Delta\|_\infty\right)\exp\left(-2\gamma(x - q\|\delta^-\|_\infty)\right)(T-t)$$
$$\times\left(1 + \mathbb{E}\left[\exp\left(-2\gamma q\inf_{\tau\in[t,T]} S^{t,s}_\tau\right)\right]\right) < +\infty,$$

because of the law of $\inf_{\tau\in[t,T]} S^{t,s}_\tau$.

Now, the same argument works for the second and third integrals, noticing that δ is bounded from below and that λ is bounded.

Hence, since we have, by construction[5]

$$\partial_\tau u_\Delta\left(\tau,\, X^{t,x,\delta}_{\tau-},\, q^{t,q,\delta}_{\tau-},\, S^{t,s}_\tau\right) + \mu\partial_s u_\Delta\left(\tau,\, X^{t,x,\delta}_{\tau-},\, q^{t,q,\delta}_{\tau-},\, S^{t,s}_\tau\right) + \frac{\sigma^2}{2}\partial^2_{ss} u_\Delta\left(\tau,\, X^{t,x,\delta}_{\tau-},\, q^{t,q,\delta}_{\tau-},\, S^{t,s}_\tau\right)$$

$$+\left(u_\Delta\left(\tau,\, X^{t,x,\delta}_{\tau-} + \Delta S^{t,s}_\tau + \Delta\delta_t,\, q^{t,q,\delta}_{\tau-} - \Delta,\, S^{t,s}_\tau\right) - u_\Delta\left(\tau,\, X^{t,x,\delta}_{\tau-},\, q^{t,q,\delta}_{\tau-},\, S^{t,s}_\tau\right)\right)\lambda_\tau \leq 0,$$

we obtain that

$$\mathbb{E}\left[u_\Delta\left(T,\, X^{t,x,\delta}_T,\, q^{t,q,\delta}_T,\, S^{t,s}_T\right)\right] = \mathbb{E}\left[u_\Delta\left(T,\, X^{t,x,\delta}_{T-},\, q^{t,q,\delta}_{T-},\, S^{t,s}_T\right)\right] \leq u_\Delta(t,x,q,s),$$

and this is true for all $\delta \in \mathcal{A}(t)$. Since for $\delta_t = \delta^*_\Delta(t)$, we have an equality in the above inequality by construction of the function $\tilde{\delta}^*_\Delta$, we obtain that

$$\sup_{\delta\in\mathcal{A}(t)} \mathbb{E}\left[u_\Delta(T,\, X^{t,x,\delta}_T,\, q^{t,q,\delta}_T,\, S^{t,s}_T)\right] \leq u_\Delta(t,x,q,s) = \mathbb{E}\left[u_\Delta(T,\, X^{t,x,\delta^*_\Delta}_T,\, q^{t,q,\delta^*_\Delta}_T,\, S^{t,s}_T)\right],$$

i.e.,

$$\sup_{\delta\in\mathcal{A}(t)} \mathbb{E}\left[-\exp\left(-\gamma\left(X^{t,x,\delta}_T + q^{t,q,\delta}_T\left(S^{t,s}_T - \ell(q^{t,q,\delta}_T)\right)\right)\right)\right]$$

[5]This inequality is also true when the portfolio is empty because of the boundary conditions.

$$\leq u_\Delta(t,x,q,s) = \mathbb{E}\left[-\exp\left(-\gamma\left(X_T^{t,x,\delta^*_\Delta} + q_T^{t,q,\delta^*_\Delta}\left(S_T^{t,s} - \ell(q_T^{t,q,\delta^*_\Delta})\right)\right)\right)\right].$$

This proves that u_Δ is the value function and that $t \mapsto \delta^*_\Delta(t)$ is optimal. □

Theorem 3.4 proves that the optimal quotes are deterministic. This is linked to the use of a CARA utility function as in the usual Almgren–Chriss framework. Theorem 3.4 also provides a simple way to compute the optimal quotes. One has indeed to solve the triangular system of ODEs ($\mathrm{HJ}_{\theta_\Delta}$) to obtain the function θ_Δ. Numerically, this does not constitute any difficulty and one may use, for instance, a Euler scheme. Then, once θ_Δ has been computed, the optimal quotes are given by the simple expression $\tilde{\delta}^*_\Delta\left(\frac{\theta_\Delta(t,q_t)-\theta_\Delta(t,q_t-\Delta)}{\Delta}\right)$ where the function $\tilde{\delta}^*_\Delta$ is implicitly characterized by the equation $\left(E_{\delta^*_\Delta}\right)$ of Lemma 3.1, and can be easily computed using Newton's method for instance.

3.3. Introducing a Hard Constraint $\delta \geq \delta^{\min}$

In the above framework, δ was allowed to take any value on the real line. To avoid marketable limit orders, one might want to impose a constraint $\delta \geq \delta^{\min}$ where $\delta^{\min}$ would be positive. Using the same tools as above, the problem with the additional constraint $\delta \geq \delta^{\min}$ can be solved easily.

The (HJB) equation becomes

$$0 = \partial_t u^{\min}_\Delta(t,x,q,s) + \mu\partial_s u^{\min}_\Delta(t,x,q,s) + \frac{1}{2}\sigma^2\partial^2_{ss}u^{\min}_\Delta(t,x,q,s)$$

$$+ \sup_{s^a \geq s+\delta^{\min}} \Lambda_\Delta(s^a - s)\left[u^{\min}_\Delta(t,x+\Delta s^a, q-\Delta, s) - u^{\min}_\Delta(t,x,q,s)\right],$$

with the final condition:

$$u^{\min}_\Delta(T,x,q,s) = -\exp\left(-\gamma\left(x+q(s-\ell(q))\right)\right),$$

and the boundary condition:

$$u^{\min}_\Delta(t,x,0,s) = -\exp(-\gamma x).$$

We consider the change of variables $u^{\min}_\Delta(t,x,q,s) = -\exp\left(-\gamma(x+qs+\theta^{\min}_\Delta(t,q))\right)$, as above. We then obtain the following system of ODEs indexed by q

$$(\mathrm{HJ}_{\theta^{\min}_\Delta}) \qquad 0 = \gamma\partial_t\theta^{\min}_\Delta(t,q) + \gamma\mu q - \frac{1}{2}\gamma^2\sigma^2q^2 + H^{\min}_\Delta\left(\frac{\theta^{\min}_\Delta(t,q) - \theta^{\min}_\Delta(t,q-\Delta)}{\Delta}\right),$$

with

$$\theta^{\min}_\Delta(T,q) = -\ell(q)q, \qquad \theta^{\min}_\Delta(t,0) = 0,$$

where

$$H^{\min}_\Delta(p) = \sup_{\delta\geq\delta^{\min}} \Lambda_\Delta(\delta)\left(1 - e^{-\gamma\Delta(\delta-p)}\right) = \sup_{\delta\geq\delta^{\min}} L_\Delta(p,\delta).$$

The important point here is to recall that $\delta \mapsto L_\Delta(p,\delta)$ is strictly increasing on $(-\infty, \tilde{\delta}^*_\Delta(p)]$ and strictly decreasing on $[\tilde{\delta}^*_\Delta(p), +\infty)$. Hence, the unique maximizer of $\delta \mapsto L_\Delta(p,\delta)$ over $\{\delta \geq \delta^{\min}\}$ is $\max(\delta^{\min}, \tilde{\delta}^*_\Delta(p))$.

Let us define $p^{\min} = \delta^{\min} - \frac{1}{\gamma\Delta}\log\left(1 - \gamma\Delta\frac{\Lambda_\Delta(\delta^{\min})}{\Lambda'_\Delta(\delta^{\min})}\right)$. The hamiltonian function $H^{\min}_\Delta$ can be written as

$$H^{\min}_\Delta(p) = \begin{cases} H_\Delta(p) & \text{if } p \geq p^{\min} \\ \Lambda_\Delta(\delta^{\min})\left(1 - e^{-\gamma\Delta(\delta^{\min} - p)}\right) & \text{if } p \leq p^{\min}. \end{cases}$$

It is a locally Lipschitz and decreasing function. In particular, the counterpart of Proposition 3.2 and Proposition 3.3 holds: there exists a unique C^1 function $\theta^{\min}_\Delta$ solution of $(\mathrm{HJ}_{\theta^{\min}_\Delta})$. Therefore, we can enounce a verification theorem and find the optimal quotes. The proof is *mutatis mutandis* the same as for Theorem 3.4.

THEOREM 3.5 (Verification theorem and optimal quotes). *Let us consider the solution θ^{min}_Δ of the system $(\mathrm{HJ}_{\theta^{min}_\Delta})$.*

Then:

$$\begin{aligned} &-\exp\left(-\gamma\left(x + qs + \theta^{min}_\Delta(t,q)\right)\right) \\ &\quad = \sup_{\delta\in\mathcal{A}^{min}(t)} \mathbb{E}\left[-\exp\left(-\gamma\left(X^{t,x,\delta}_T + q^{t,q,\delta}_T\left(S^{t,s}_T - \ell(q^{t,q,\delta}_T)\right)\right)\right)\right], \end{aligned}$$

where $\mathcal{A}^{min}(t)$ is the set of predictable processes on $[t, T]$, bounded from below by δ^{min} and where:

$$dS^{t,s}_\tau = \mu d\tau + \sigma dW_\tau, \qquad S^{t,s}_t = s,$$

$$dX^{t,x,\delta}_\tau = (S_\tau + \delta_\tau)\Delta dN_\tau, \qquad X^{t,x,\delta}_t = x,$$

$$dq^{t,q,\delta}_\tau = -\Delta dN_\tau, \qquad q^{t,q,\delta}_t = q,$$

the point process N having stochastic intensity $(\lambda_\tau)_\tau$ with $\lambda_\tau = \Lambda_\Delta(\delta_\tau)1_{q_{\tau-}>0}$.

Moreover, the optimal ask quote $S^a_t = S_t + \delta^{min*}_\Delta(t)$, for $q_t > 0$, is characterized by

$$\delta^{min*}_\Delta(t) = \max\left(\delta^{min}, \tilde{\delta}^*_\Delta\left(\frac{\theta^{min}_\Delta(t,q_t) - \theta^{min}_\Delta(t,q_t - \Delta)}{\Delta}\right)\right),$$

where $\tilde{\delta}^*_\Delta(\cdot)$ is the function defined in Lemma 3.1.

4. EXAMPLES AND PROPERTIES

4.1. The Case of an Exponential Intensity Function

In the above section, we generalized a model already used in Guéant et al. (2012), in which the intensity functions had exponential shape: $\Lambda_\Delta(\delta) = A_\Delta e^{-k_\Delta\delta}$.

In the case of exponential intensity, we can write the results of Guéant et al. (2012) (in a slightly more general case than in the original paper) in the language of this paper. In fact, the reason why closed-form solutions can be obtained in the exponential case is that the equation $(\mathrm{HJ}_{\theta_\Delta})$ simplifies to a linear system of equations when we replace the unknown θ_Δ by $\exp\left(\frac{k_\Delta}{\Delta}\theta_\Delta\right)$:

PROPOSITION 4.1. *Assume that* $\Lambda_\Delta(\delta) = A_\Delta e^{-k_\Delta \delta}$.

Then, $H_\Delta(p) = \frac{\gamma\Delta}{k_\Delta}\left(1+\frac{\gamma\Delta}{k_\Delta}\right)^{-1-\frac{k_\Delta}{\gamma\Delta}} A_\Delta e^{-k_\Delta p}$.

Also, if we consider $\theta_\Delta : [0,T]\times\Delta\mathbb{N}\to\mathbb{R}$, *the unique* C^1 *solution of* $(\mathrm{HJ}_{\theta_\Delta})$, *then* $w_\Delta = \exp\left(\frac{k_\Delta}{\Delta}\theta_\Delta\right)$ *is the unique solution of:*

$$\partial_t w_\Delta(t,q) = -\frac{1}{\Delta}k_\Delta\mu q w_\Delta(t,q) + \frac{1}{2\Delta}\gamma k_\Delta\sigma^2 q^2 w_\Delta(t,q) - A_\Delta\left(1+\frac{\gamma\Delta}{k_\Delta}\right)^{-1-\frac{k_\Delta}{\gamma\Delta}} \times w_\Delta(t,q-\Delta),$$

with

$$w_\Delta(T,q) = e^{-\frac{k_\Delta}{\Delta}\ell(q)q}, \qquad w_\Delta(t,0) = 1,$$

and the optimal quote, for $q_t > 0$, *is given by*

$$\delta^*_\Delta(t) = \frac{1}{k_\Delta}\log\left(\frac{w_\Delta(t,q_t)}{w_\Delta(t,q_t-\Delta)}\right) + \frac{1}{\gamma\Delta}\log\left(1+\frac{\gamma\Delta}{k_\Delta}\right).$$

4.2. Numerical Examples

We now provide numerical approximations of both the function $\theta_\Delta(t,q)$ and the optimal control function $\delta^*_\Delta(t,q)$. These numerical approximations allow to compare what happens in the pure exponential case and what happens when another intensity function is considered, especially for negative δs. In this section, we consider as an alternative to the exponential form for Λ_Δ a functional form $\tilde{\Lambda}_\Delta$ (see Figure 4.1) that prevents the use of marketable limit orders. The intensity function $\tilde{\Lambda}_\Delta$ prevents the use of marketable limit orders since $\tilde{\Lambda}_\Delta$ is constant (in fact, decreasing very slowly to satisfy the hypotheses of the paper) for negative δ. Also, we included the commonly observed fact that the probability to be executed does not correspond to the exponential intensity framework for small positive δs.

Figures 4.2 and 4.3 represent, respectively, the solution θ_Δ and the optimal quotes $\delta^*_\Delta(t,q)$ as given by Theorem 3.4.[6] From Figure 4.2, we know that θ_Δ is not a monotonic function of q. It is important here to recall the economic meaning of θ_Δ. The certainty equivalent of holding q shares at time t is $qs + \theta_\Delta(t,q)$. Hence, $\theta_\Delta(t,q)$ is a risk-adjusted value of holding q shares at time t in excess of the MtM value qs. The reason why $\theta_\Delta(t,q)$ is not a monotonic function of q can then be understood easily. At the time horizon T, the function is decreasing but far from T two effects are at stake. On the one hand, when there are many shares in the portfolio, there will be many trades and hence more opportunities to make money through limit orders: this goes in the direction of an increasing function $\theta_\Delta(t,\cdot)$. On the other hand, the larger the inventory to liquidate, the more price risk. This goes in the direction of a decreasing function $\theta_\Delta(t,\cdot)$ since it is a risk-adjusted value.

Although there is almost no difference between the two cases Λ_Δ and $\tilde{\Lambda}_\Delta$ as far as θ_Δ is concerned, this is not true anymore when it comes to the optimal quotes δ^*_Δ. We

[6]One may wonder why we choose a risk aversion parameter $\gamma = 0.001$. This figure seems small but it has, in fact, an important impact since the shares are sold by groups of 50.

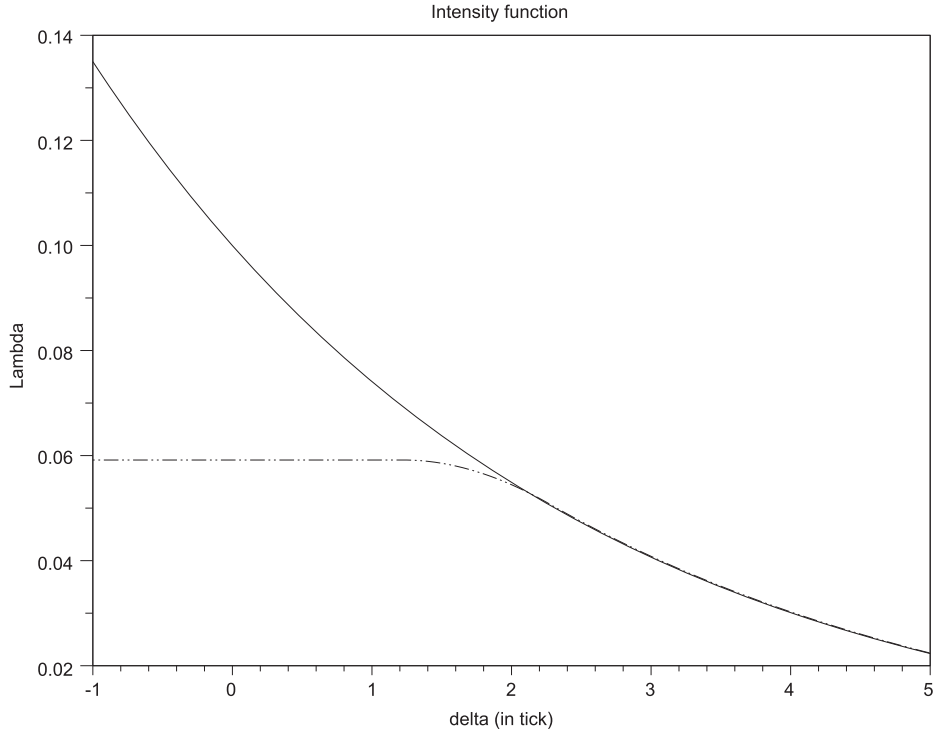

FIGURE 4.1. Intensity functions (here represented on $[-1, 5]$). Line: $\Lambda_\Delta(\delta) = Ae^{-k\delta}$, $A = 0.1$ (s^{-1}), $k = 0.3$ ($Tick^{-1}$). Dotted line: $\tilde{\Lambda}_\Delta(\delta)$, identical to $\Lambda_\Delta(\delta)$ for $\delta \geq 2$.

indeed see in Figure 4.3, as expected, that, in the case of the intensity function $\tilde{\Lambda}_\Delta$, there is no negative optimal quotes. This very conservative choice for the intensity function is a way to avoid any influence of the intensity function on the set $\{\delta < 0\}$. Also, we see that since $\tilde{\Lambda}_\Delta$ is not of exponential form for small positive δs, the lower bound for the optimal quotes is higher than expected.

Another way to prevent marketable limit orders is to impose $\delta \geq \delta^{\min} = 0$ as in Section 3.3. In that case, if we consider the intensity function Λ_Δ and the same parameters as above, we obtain the quotes given in Figure 4.4. These quotes are almost exactly the same as if we had floored the optimal quote δ^*_Δ of the unconstrained problem to 0.

4.3. Asymptotic Quote

In Guéant et al. (2012), we obtained a limiting regime when $T \to \infty$. This result generalizes to our general framework. More exactly, we obtain:

PROPOSITION 4.2 (Asymptotic behavior). *Let us suppose that:*

- $\lim_{p \to +\infty} H_\Delta(p) = 0$.[7]
- $\gamma > 0$,
- $\mu < \frac{1}{2}\gamma\sigma^2\Delta$.

[7]This is guaranteed if $\lim_{\delta \to +\infty} \delta\Lambda_\Delta(\delta) = 0$.

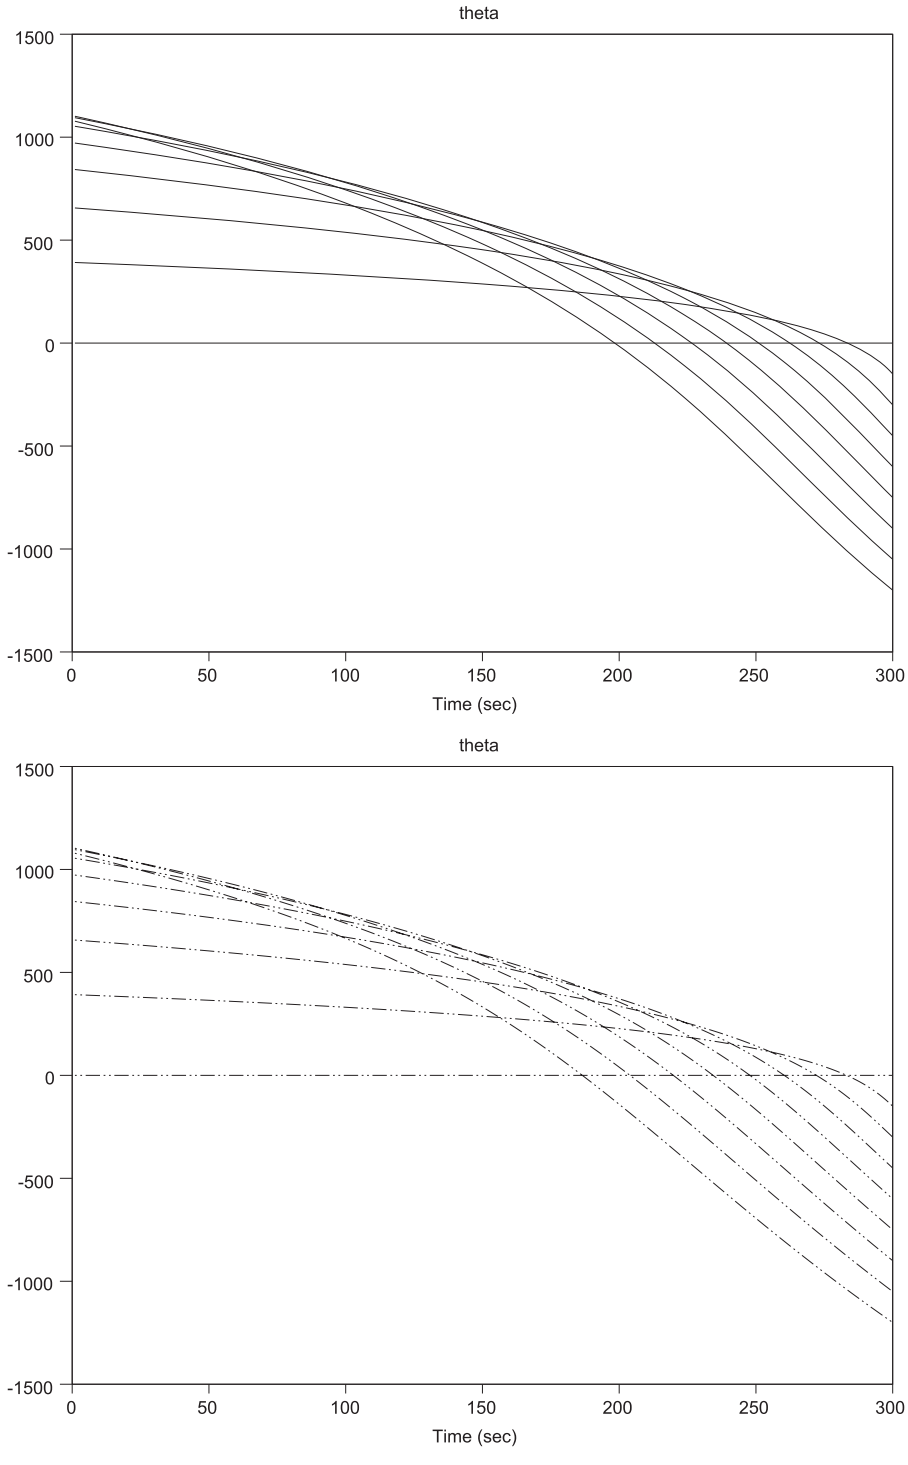

FIGURE 4.2. Solution $\theta_\Delta(t,q)$ for Λ_Δ (top) and $\tilde{\Lambda}_\Delta$ (down), $q_0 = 400$, $\Delta = 50$, $T = 300$ (s), $\mu = 0$, $\sigma = 0.3$ (Tick.s$^{-\frac{1}{2}}$), $\gamma = 0.001$ (Tick^{-1}), and $\ell(q) = \ell = 3$ (Tick). The index $q \in \{0, 50, \ldots, 400\}$ of each curve can be read from the terminal values.

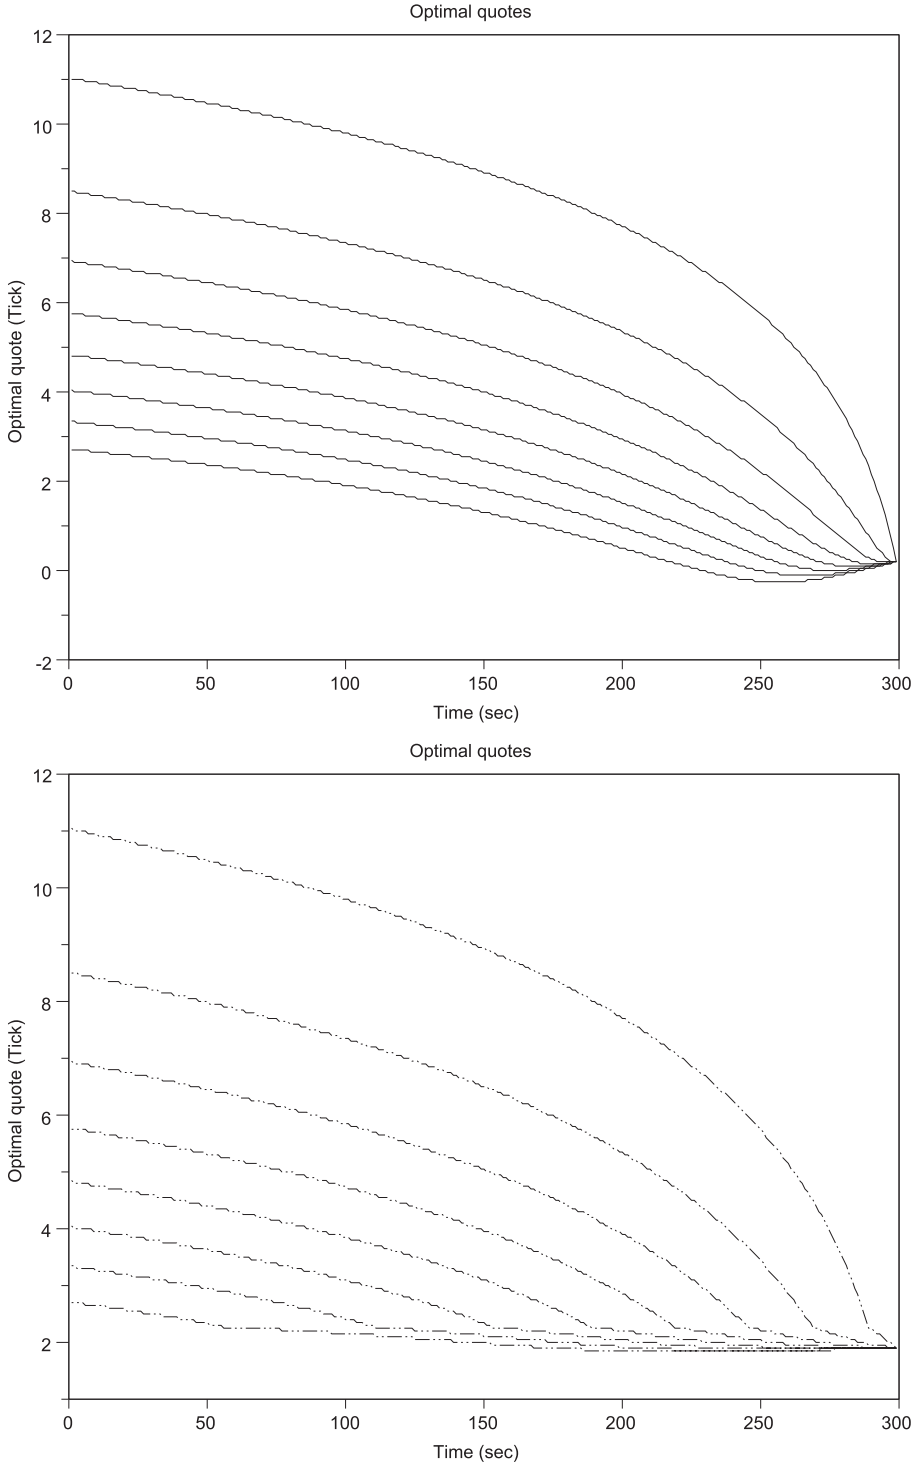

FIGURE 4.3. Solution $\delta^*_\Delta(t,q)$ for Λ_Δ (top) and $\tilde{\Lambda}_\Delta$ (down), $q_0 = 400$, $\Delta = 50$, $T = 300$ (s), $\mu = 0$, $\sigma = 0.3$ (Tick.s$^{-\frac{1}{2}}$), $\gamma = 0.001$ (Tick^{-1}), and $\ell(q) = \ell = 3$ (Tick). The lower the quotes, the higher $q \in \{50, \ldots, 400\}$.

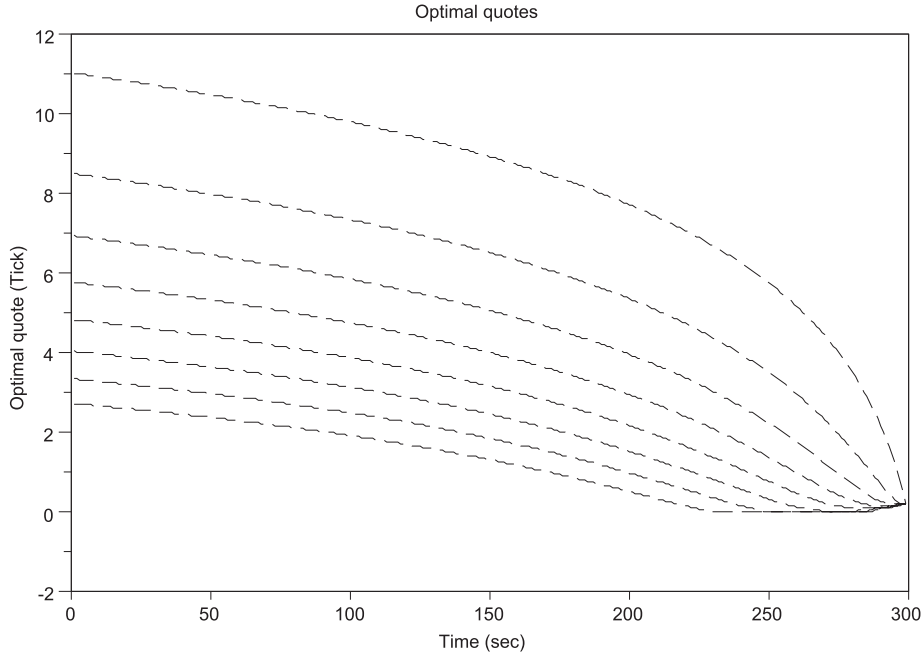

FIGURE 4.4. Optimal quote for Λ_Δ when the constraint $\delta \geq 0$ is imposed. $q_0 = 400$, $\Delta = 50$, $T = 300$ (s), $\mu = 0$, $\sigma = 0.3$ (Tick.s$^{-\frac{1}{2}}$), $\gamma = 0.001$ (Tick^{-1}), and $\ell(q) = \ell = 3$ (Tick). The lower the quotes, the higher $q \in \{50, \ldots, 400\}$.

Then, the asymptotic behavior of θ_Δ is:

$$\lim_{T\to+\infty} \theta_\Delta(0,q) = \Delta \sum_{q'\in\{\Delta,2\Delta,\ldots,q\}} H_\Delta^{-1}\left(\frac{1}{2}\gamma^2\sigma^2 q'^2 - \gamma\mu q'\right) = \theta_\Delta^\infty(q).$$

The resulting asymptotic behavior of the optimal quote is

$$\lim_{T\to+\infty} \delta_\Delta^*(t=0) = \delta_\Delta^*\left(H_\Delta^{-1}\left(\frac{1}{2}\gamma^2\sigma^2 q_0^2 - \gamma\mu q_0\right)\right) = \delta_\Delta^{*\infty}.$$

Proof. Let us define for $q \in \Delta\mathbb{N}$:

$$\theta_\Delta^\infty(q) = \Delta \sum_{q'\in\{\Delta,2\Delta,\ldots,q\}} H_\Delta^{-1}\left(\frac{1}{2}\gamma^2\sigma^2 q'^2 - \gamma\mu q'\right).^8$$

Let us define for $t \geq 0$ and $q \in \Delta\mathbb{N}$, θ_Δ^r, the unique solution[9] of:

$$0 = -\gamma\partial_t^r\theta_\Delta(t,q) + \gamma\mu q - \frac{1}{2}\gamma^2\sigma^2 q^2 + H_\Delta\left(\frac{\theta_\Delta^r(t,q) - \theta_\Delta^r(t,q-\Delta)}{\Delta}\right),$$

[8]This is well defined because of the assumptions on H_Δ and on the parameters, and because $\gamma > 0$. In the risk-neutral case, there is no upper bound for the optimal quotes when $T \to \infty$. This constitutes an important difference between our model and the model of Bayraktar and Ludkovski (2014).

[9]To prove that this function is well defined, one can use the same tools as in Proposition 3.3.

with

$$\theta^r_\Delta(0,q) = -\ell(q)q, \qquad \theta^r_\Delta(t,0) = 0.$$

Then, because we just reversed time, we want to prove that

$$\forall q \in \Delta\mathbb{N}, \quad \lim_{t\to+\infty} \theta^r_\Delta(t,q) = \theta^\infty_\Delta(q).$$

We proceed by induction. The result is true for $q = 0$. Let us suppose that the result is true for $q - \Delta$ for some $q \in \Delta\mathbb{N}^*$.
Then,

$$\forall \epsilon > 0, \exists t_{q-\Delta}, \forall t \geq t_{q-\Delta}, |\theta^r_\Delta(t, q-\Delta) - \theta^\infty_\Delta(q-\Delta)| \leq \epsilon.$$

Since H_Δ is a strictly decreasing function, we obtain that $\forall t \geq t_{q-\Delta}$:

$$\gamma\mu q - \frac{1}{2}\gamma^2\sigma^2 q^2 + H_\Delta\left(\frac{\theta^r_\Delta(t,q) - \theta^\infty_\Delta(q-\Delta) + \epsilon}{\Delta}\right)$$

$$\leq \gamma\,\partial_t\theta^r_\Delta(t,q) \leq \gamma\mu q - \frac{1}{2}\gamma^2\sigma^2 q^2 + H_\Delta\left(\frac{\theta^r_\Delta(t,q) - \theta^\infty_\Delta(q-\Delta) - \epsilon}{\Delta}\right),$$

or equivalently:

$$H_\Delta\left(\frac{\theta^r_\Delta(t,q) - \theta^\infty_\Delta(q-\Delta) + \epsilon}{\Delta}\right) - H_\Delta\left(\frac{\theta^\infty_\Delta(q) - \theta^\infty_\Delta(q-\Delta)}{\Delta}\right)$$

$$\leq \gamma\,\partial_t\theta^r_\Delta(t,q) \leq H_\Delta\left(\frac{\theta^r_\Delta(t,q) - \theta^\infty_\Delta(q-\Delta) - \epsilon}{\Delta}\right) - H_\Delta\left(\frac{\theta^\infty_\Delta(q) - \theta^\infty_\Delta(q-\Delta)}{\Delta}\right).$$

Hence, $\forall t \geq t_{q-\Delta}$:

$$\theta^r_\Delta(t,q) > \theta^\infty_\Delta(q) + \epsilon \Rightarrow \partial_t\theta^r_\Delta(t,q) < 0,$$

and

$$\theta^r_\Delta(t,q) < \theta^\infty_\Delta(q) - \epsilon \Rightarrow \partial_t\theta^r_\Delta(t,q) > 0.$$

As a consequence, if there exists $t' \geq t_{q-\Delta}$ such that $|\theta^r_\Delta(t',q) - \theta^\infty_\Delta(q)| \leq \epsilon$, then $\forall t \geq t'$, $|\theta^r_\Delta(t,q) - \theta^\infty_\Delta(q)| \leq \epsilon$.

In particular, if $|\theta^r_\Delta(t_{q-\Delta},q) - \theta^\infty_\Delta(q)| \leq \epsilon$, then $\forall t \geq t_{q-\Delta}$, $|\theta^r_\Delta(t,q) - \theta^\infty_\Delta(q)| \leq \epsilon$.

Now, if $\theta^r_\Delta(t_{q-\Delta},q) > \theta^\infty_\Delta(q) + \epsilon$, then there are two possibilities. The first one is that the function $t \geq t_{q-\Delta} \mapsto \theta^r_\Delta(t,q)$ is decreasing and in that case, it is bounded from below by $\theta^\infty_\Delta(q) - \epsilon$ and must converge. Since $\lim_{t\to+\infty}\theta^r_\Delta(t,q-\Delta) = \theta^\infty_\Delta(q-\Delta)$, the only possible limit for $\theta^r_\Delta(t,q)$ is $\theta^\infty_\Delta(q)$. The second possibility is that $t \geq t_{q-\Delta} \mapsto \theta^r_\Delta(t,q)$ is not a decreasing function and in that case, there must exist $t' \geq t_{q-\Delta}$ such that $\theta^r_\Delta(t',q) \leq \theta^\infty_\Delta(q) + \epsilon$. Since $\theta^r_\Delta(t',q) \geq \theta^\infty_\Delta(q) - \epsilon$, we now obtain that $\forall t \geq t_q$, $|\theta^r_\Delta(t,q) - \theta^\infty_\Delta(q)| \leq \epsilon$.

Finally, if $\theta^r_\Delta(t_{q-\Delta},q) < \theta^\infty_\Delta(q) - \epsilon$, then there are two possibilities. The first one is that the function $t \geq t_{q-\Delta} \mapsto \theta^r_\Delta(t,q)$ is increasing and in that case, it is bounded from above

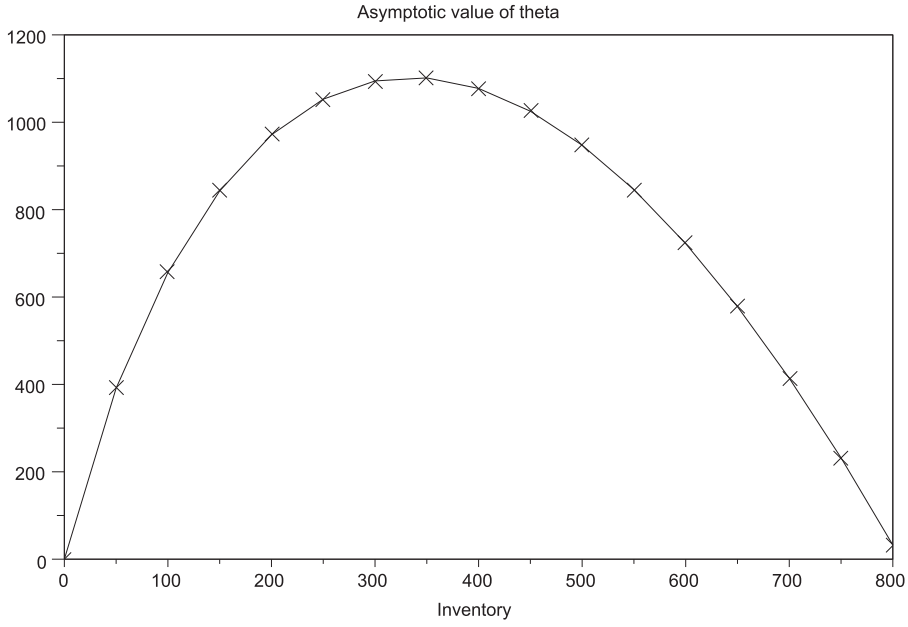

FIGURE 4.5. $\theta^{\infty}_{\Delta}(q)$ for Λ_{Δ} defined above, $q_0 = 800, \Delta = 50, \mu = 0, \sigma = 0.3$ (Tick.s$^{-\frac{1}{2}}$), and $\gamma = 0.001$ (Tick^{-1}).

by $\theta^{\infty}_{\Delta}(q) + \epsilon$ and must converge. Since $\lim_{t\to+\infty} \theta^{r}_{\Delta}(t, q-\Delta) = \theta^{\infty}_{\Delta}(q-\Delta)$, the only possible limit for $\theta^{r}_{\Delta}(t,q)$ is $\theta^{\infty}_{\Delta}(q)$. The second possibility is that $t \geq t_{q-\Delta} \mapsto \theta^{r}_{\Delta}(t,q)$ is not an increasing function and in that case, there must exist $t' \geq t_{q-\Delta}$ such that $\theta^{r}_{\Delta}(t',q) \geq \theta^{\infty}_{\Delta}(q) - \epsilon$. Since $\theta^{r}_{\Delta}(t',q) \leq \theta^{\infty}_{\Delta}(q) + \epsilon$, we now obtain that $\forall t \geq t_q$, $|\theta^{r}_{\Delta}(t,q) - \theta^{\infty}_{\Delta}(q)| \leq \epsilon$.

The conclusion is that $\limsup_{t\to+\infty} |\theta^{r}_{\Delta}(t,q) - \theta^{\infty}_{\Delta}(q)| \leq \epsilon$. Sending ϵ to 0, we get the result for θ_{Δ}.

The result for the optimal quote is then straightforward. □

These asymptotic formulae deserve some comments. First, regarding the above discussion on monotonicity, we know that H_{Δ} is a decreasing function; therefore, the asymptotic limit $\theta^{\infty}_{\Delta}(\cdot)$ is either a decreasing function (when $\frac{1}{2}\gamma^2\sigma^2\Delta^2 - \gamma\mu\Delta > H_{\Delta}(0)$) or a function that is first increasing and then decreasing (otherwise)—see Figure 4.5 in our case. Second, coming to the optimal quotes and the role of the parameters, we can analyze the way $\delta^{*\infty}_{\Delta}$ depends on μ, σ, γ, and Λ_{Δ}. The best way to proceed is to use the expression for H_{Δ} found in Lemma 3.1 and to notice that an equivalent way to define $\delta^{*\infty}_{\Delta}$ is through the following implicit characterization:

$$-\mu q_0 + \frac{1}{2}\gamma\sigma^2 q_0^2 = \Delta \frac{\Lambda_{\Delta}(\delta^{*\infty}_{\Delta})^2}{\gamma\Delta\Lambda_{\Delta}(\delta^{*\infty}_{\Delta}) - \Lambda'_{\Delta}(\delta^{*\infty}_{\Delta})}.$$

It is then straightforward to see that $\delta^{*\infty}_{\Delta}$ is an increasing function of μ. A trader expecting the stock price to go up is indeed encouraged to slow down the liquidation process. Similarly, we see that $\delta^{*\infty}_{\Delta}$ decreases as σ increases. An increase in σ corresponds to an increase in price risk and this provides the trader with an incentive to speed up

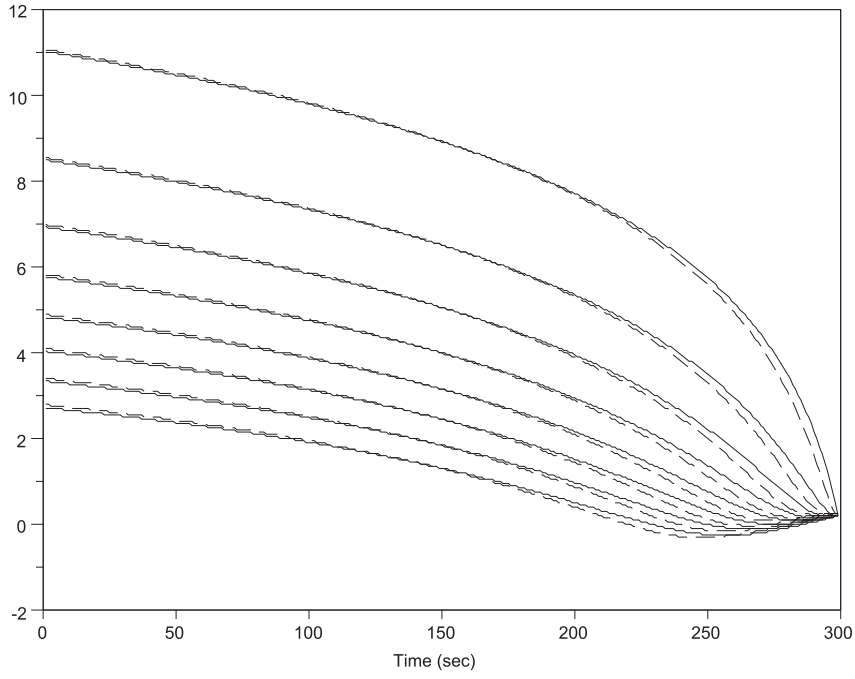

FIGURE 4.6. Optimal quotes for $q \in \{50, 100, \ldots, 400\}$, for $q_0 = 600$, $T = 1200$ (s), $\mu = 0$, $\sigma = 0.3$ (Tick.s$^{-\frac{1}{2}}$), $A = 0.1$ (s^{-1}), $k = 0.3$ (Tick^{-1}), $\gamma = 0.001$ (Tick^{-1}), and $\ell(q) = \ell = 3$ (Tick). Line: $\Lambda(\delta) = Ae^{-k\delta}$ and $\Delta = 50$. Dotted line: $\Lambda(\delta) = 2Ae^{-k\delta}$ and $\Delta = 25$.

the execution process. Therefore, it is natural that the asymptotic quote be a decreasing function of σ.

Differentiating the above expression with respect to γ, we see that the asymptotic quote decreases as the risk aversion increases. An increase in risk aversion forces indeed the trader to reduce both nonexecution risk and price risk and this leads to posting orders with lower prices.

Now, if one replaces the intensity function Λ_Δ by $\lambda\Lambda_\Delta$ where $\lambda > 1$, then it results in an increase in $\delta_\Delta^{*\infty}$. This is natural because when the rate of arrival of liquidity-taking orders increases, the trader is more likely to liquidate his shares faster and posting deeper into the book allows for larger profits.

4.4. The Influence of Δ

In addition to the asymptotic regime, we can consider different sizes Δ of orders.

We see in Figure 4.6 that there is little difference between the two cases we considered. This is linked to the existence of a limit regime as $\Delta \to 0$ and the next section is dedicated to its analysis. The limiting equation for θ_Δ will turn out to be a classical equation in optimal liquidation theory (see Section 5).

5. LIMIT REGIME $\Delta \to 0$

In the preceding sections, the size of the order posted by the trader was constant equal to Δ, a size that is supposed to be small with respect to q_0. As a consequence, the question of the limiting behavior when Δ tends to 0 is relevant.[10] To that purpose, we need to make an assumption on the behavior of the intensity function with respect to the order size Δ. The "right" scaling (already used above for the numerics underlying Figure 4.6) is to suppose that

$$\Lambda_\Delta(\delta) = \frac{\Lambda(\delta)}{\Delta}.$$

With this scaling , along with additional technical hypotheses, our goal is to prove the following theorem that is rather technical and echoes the results obtained by Bayraktar and Ludkovski (2014), here with risk aversion (i.e., $\gamma > 0$), whereas Bayraktar and Ludkovski (2014) deals with the risk-neutral case:

THEOREM 5.1 (Limit regime $\Delta \to 0$). *Let us suppose that*

- $\Lambda(\delta)\Lambda''(\delta) < 2\Lambda'(\delta)^2$,
- $\lim_{\delta \to +\infty} \delta\Lambda(\delta) = 0$,
- ℓ *is a continuous function.*

For a given $\Delta > 0$, *let us define* θ^c_Δ *on* $[0, T] \times [0, q_0]$ *by*

$$\theta^c_\Delta(t, q) = \begin{cases} \theta_\Delta(t, 0), & \text{if } q = 0, \\ \theta_\Delta\left(t, (k+1)\Delta\right), & \text{if } q \in (k\Delta, (k+1)\Delta]. \end{cases}$$

Then, θ^c_Δ *converges uniformly toward a continuous function* $\theta : [0, T] \times [0, q_0] \to \mathbb{R}$ *that is the unique viscosity solution of the equation* (HJ_{lim})*:*

$$\begin{cases} -\gamma\,\partial_t\theta(t, q) - \gamma\mu q + \frac{1}{2}\gamma^2\sigma^2 q^2 - H\left(\partial_q\theta(t, q)\right) = 0, & \text{on } [0, T) \times (0, q_0], \\ \theta(t, q) = 0, & \text{on } [0, T] \times \{0\}, \\ \theta(t, q) = -\ell(q)q, & \text{on } \{T\} \times [0, q_0], \end{cases} \qquad (HJ_{lim})$$

where $H(p) = \gamma \sup_\delta \Lambda(\delta)(\delta - p)$ *and where the terminal condition and the boundary condition are, in fact, satisfied is the classical sense.*

To prove this theorem, we first need to study H and the convergence of the Hamiltonian functions H_Δ toward H. We start with a counterpart of Lemma 3.1 that requires $\Lambda(\delta)\Lambda''(\delta) < 2\Lambda'(\delta)^2$.

LEMMA 5.2. *Let us define* $L(p, \delta) = \Lambda(\delta)(\delta - p)$.
$\forall p \in \mathbb{R}$, $\delta \mapsto L(p, \delta)$ *attains its maximum at* $\tilde{\delta}^*(p)$ *uniquely characterized by:*

$$\tilde{\delta}^*(p) + \frac{\Lambda(\tilde{\delta}^*(p))}{\Lambda'(\tilde{\delta}^*(p))} = p.$$

Moreover, $p \mapsto \tilde{\delta}^*(p)$ *is a* C^1 *function.*

[10] Although this is not recalled, it is assumed that Δ is always chosen as a fraction of q_0.

Subsequently, H is a C^1 function with

$$H(p) = \gamma \frac{\Lambda(\tilde{\delta}^*(p))^2}{-\Lambda'(\tilde{\delta}^*(p))}.$$

Proof. The proof is similar to the proof of Lemma 3.1. □

Now, we can state a result about convergence that also provides a uniform bound for the hamiltonian functions:

LEMMA 5.3. *H_Δ converges locally uniformly toward H when $\Delta \to 0$ with $\forall p \in \mathbb{R}$, $H_\Delta(p) \leq H(p)$.*

Proof. For a fixed $x \geq 0$, the function $f : \Delta \in \mathbb{R}_+ \mapsto \frac{1-e^{-\Delta x}}{\Delta}$ is a decreasing function ($f(0) = x$).

Hence,

$$\forall 0 < \Delta' < \Delta, \quad \sup_{\delta \geq p} \Lambda(\delta) \frac{1 - e^{-\gamma \Delta (\delta - p)}}{\Delta} \leq \sup_{\delta \geq p} \Lambda(\delta) \frac{1 - e^{-\gamma \Delta' (\delta - p)}}{\Delta'} \leq \gamma \sup_{\delta \geq p} \Lambda(\delta)(\delta - p).$$

This gives

$$H_\Delta(p) \leq H_{\Delta'}(p) \leq H(p).$$

Now, because H is continuous, using Dini's theorem, if we prove that convergence is pointwise, convergence will be locally uniform. We then only need to prove pointwise convergence of H_Δ toward H. Using Lemma 3.1 and Lemma 5.2, we see that it is sufficient to prove that $\tilde{\delta}^*_\Delta$ converges pointwise toward $\tilde{\delta}^*$.

For that purpose, notice that the sequence of functions $f_\Delta(x) = x - \frac{1}{\gamma \Delta} \log\left(1 - \gamma \Delta \frac{\Lambda(x)}{\Lambda'(x)}\right)$ is an increasing sequence of increasing functions. Hence, by the unique characterizations of $\tilde{\delta}^*_\Delta(p)$ and $\tilde{\delta}^*(p)$, we see that $\tilde{\delta}^*_\Delta(p)$ increases as Δ decreases to 0 and is bounded from above by $\tilde{\delta}^*(p)$ that is the only possible limit. Hence, $\tilde{\delta}^*_\Delta(p) \to \tilde{\delta}^*(p)$ as $\Delta \to 0$ and this proves the result. □

Now, we provide a uniform bound for the θ_Δ that will be important in the proof of Theorem 5.1.

PROPOSITION 5.4 (Bounds for θ_Δ). *$\forall t \in [0, T]$, $\forall q \in \{0, \Delta, \dots, q_0\}$,*

$$-\ell(q_0)q_0 - \mu^- q_0 (T - t) - \frac{1}{2}\gamma \sigma^2 q_0^2 (T - t) \leq \theta_\Delta(t, q) \leq \mu^+ q_0 (T - t) + \frac{1}{\gamma} H(0)(T - t).$$

Proof. To prove these inequalities, we use the comparison principle of Proposition 3.2. If $\overline{\theta}_\Delta(t, q) = \mu^+ q_0 (T - t) + \frac{1}{\gamma} H_\Delta(0)(T - t)$, then

$$\overline{\theta}_\Delta(T, q) = 0 \geq -\ell(q)q, \qquad \overline{\theta}_\Delta(t, 0) = \mu^+ q_0 (T - t) + \frac{1}{\gamma} H_\Delta(0)(T - t) \geq 0,$$

and

$$\gamma \partial_t \overline{\theta}_\Delta(t, q) + \gamma \mu q - \frac{1}{2}\gamma^2 \sigma^2 q^2 + H_\Delta\left(\frac{\overline{\theta}_\Delta(t, q) - \overline{\theta}_\Delta(t, q - \Delta)}{\Delta}\right)$$

$$= -\gamma \frac{1}{\gamma} H_\Delta(0) - \gamma \mu^+ q_0 + \gamma \mu q - \frac{1}{2}\gamma^2\sigma^2 q^2 + H_\Delta(0)$$
$$= -\gamma(\mu^+ q_0 - \mu q) - \frac{1}{2}\gamma^2\sigma^2 q^2 \le 0.$$

Therefore, $\theta_\Delta(t,q) \le \overline{\theta}_\Delta(t,q) = \mu^+ q_0 (T-t) + \frac{1}{\gamma} H_\Delta(0)(T-t)$.
The uniform upper bound is then obtained using Lemma 5.3.
Now, if $\underline{\theta}_\Delta(t,q) = -\ell(q_0)q_0 - \mu^- q_0 (T-t) - \frac{1}{2}\gamma\sigma^2 q_0^2 (T-t)$, then

$$\underline{\theta}_\Delta(T,q) = -\ell(q_0)q_0 \le -\ell(q)q, \qquad \underline{\theta}_\Delta(t,0) = -\ell(q_0)q_0 - \mu^- q_0 (T-t) - \frac{1}{2}\gamma^2\sigma^2 q_0^2 (T-t) \le 0,$$

and

$$\gamma \partial_t \underline{\theta}_\Delta(t,q) + \gamma \mu q - \frac{1}{2}\gamma^2\sigma^2 q^2 + H_\Delta\left(\frac{\underline{\theta}_\Delta(t,q) - \underline{\theta}_\Delta(t,q-\Delta)}{\Delta}\right)$$
$$= \gamma \mu^- q_0 + \frac{1}{2}\gamma^2\sigma^2 q_0^2 + \gamma \mu q - \frac{1}{2}\gamma^2\sigma^2 q^2 + H_\Delta(0)$$
$$\ge \gamma(\mu^- q_0 + \mu q) + \frac{1}{2}\gamma^2\sigma^2(q_0^2 - q^2) \ge 0.$$

Therefore, $\theta_\Delta(t,q) \ge \underline{\theta}_\Delta(t,q) = -\ell(q_0)q_0 - \mu^- q_0 (T-t) - \frac{1}{2}\gamma\sigma^2 q_0^2 (T-t)$. □

We are now ready to start the proof of Theorem 5.1.

Proof of Theorem 5.1. We first introduce the following half-relaxed limit functions:

$$\overline{\theta}(t,q) = \limsup_{j\to+\infty} \limsup_{\Delta\to 0} \sup\left\{\theta^c_\Delta(t',q'), \quad |t'-t| + |q'-q| \le \frac{1}{j}\right\},$$

$$\underline{\theta}(t,q) = \liminf_{j\to+\infty} \liminf_{\Delta\to 0} \inf\left\{\theta^c_\Delta(t',q'), \quad |t'-t| + |q'-q| \le \frac{1}{j}\right\}.$$

$\overline{\theta}$ and $\underline{\theta}$ are, respectively, upper semicontinuous and lower semicontinuous and the goal of the proof is to show that they are equal to one another and solution of the partial differential equation ($\mathrm{HJ}_{\mathrm{lim}}$).

Step 1: $\overline{\theta}$ and $\underline{\theta}$ are, respectively, viscosity subsolution and viscosity supersolution of the equation ($\mathrm{HJ}_{\mathrm{lim}}$).

To prove this point, let us consider (t^*, q^*) and a test function ϕ such that $\overline{\theta} - \phi$ attains a local maximum at (t^*, q^*). Without loss of generality, we can assume that $\phi(t^*, q^*) = \overline{\theta}(t^*, q^*)$ and consider $r > 0$ such that

- the maximum is global on the ball of radius r centered in (t^*, q^*),
- outside of this ball, $\phi \ge 2\sup_\Delta \|\theta_\Delta\|_\infty$—this value being finite because of the uniform bound obtained in the above proposition.

Following Barles and Souganidis (1990) methodology, we know that there exists a sequence $(\Delta_n, t_n, q_n)_n$ such that

- $\Delta_n \to 0$, $(t_n, q_n) \to (t^*, q^*)$,
- $\theta^c_{\Delta_n}(t_n, q_n) \to \overline{\theta}(t^*, q^*)$,
- $\theta^c_{\Delta_n} - \phi$ has a global maximum at (t_n, q_n).

Now, because of the definition of $\theta^c_{\Delta_n}$, if $(t^*, q^*) \in [0, T) \times (0, q_0]$, we can always suppose that $q_n \geq \Delta_n$ and $t_n \neq T$ and then by definition of θ_{Δ_n}:

$$-\gamma \partial_t \theta_{\Delta_n}(t_n, q_n) - \gamma \mu q_n + \frac{1}{2}\gamma^2\sigma^2 q_n^2 - H_{\Delta_n}\left(\frac{\theta_{\Delta_n}(t_n, q_n) - \theta_{\Delta_n}(t_n, q_n - \Delta_n)}{\Delta_n}\right) = 0.$$

Because H_{Δ_n} is decreasing we have, by definition of (t_n, q_n):

$$-H_{\Delta_n}\left(\frac{\theta_{\Delta_n}(t_n, q_n) - \theta_{\Delta_n}(t_n, q_n - \Delta_n)}{\Delta_n}\right) \geq -H_{\Delta_n}\left(\frac{\phi(t_n, q_n) - \phi(t_n, q_n - \Delta_n)}{\Delta_n}\right).$$

Similarly, since $t_n < T$, we have, for h sufficiently small:

$$\theta_{\Delta_n}(t_n + h, q_n) - \theta_{\Delta_n}(t_n, q_n) \leq \phi(t_n + h, q_n) - \phi(t_n, q_n).$$

Hence,

$$\partial_t \theta_{\Delta_n}(t_n, q_n) \leq \partial_t \phi(t_n, q_n).$$

These inequalities give

$$-\gamma \partial_t \phi(t_n, q_n) - \gamma \mu q_n + \frac{1}{2}\gamma^2\sigma^2 q_n^2 - H_{\Delta_n}\left(\frac{\phi(t_n, q_n) - \phi(t_n, q_n - \Delta_n)}{\Delta_n}\right) \leq 0.$$

Using now the convergence of (t_n, q_n) toward (t^*, q^*) and the local uniform convergence of H_{Δ_n} toward H, we eventually obtain the desired inequality:

$$-\gamma \partial_t \phi(t^*, q^*) - \gamma \mu q^* + \frac{1}{2}\gamma^2\sigma^2 q^{*2} - H\left(\partial_q \phi(t^*, q^*)\right) \leq 0.$$

We see that the boundaries corresponding to $q = q_0$ and $t = 0$ play no role. However, we need to consider the cases $t^* = T$ and $q^* = 0$.

If $t^* = T$ and $q^* \neq 0$, then there are two cases. If there are infinitely many indices n such that $t_n < T$, then the preceding proof still works. Otherwise, for all n sufficiently large, $\theta_{\Delta_n}(t_n, q_n) = -\ell(q_n)q_n$ and hence, passing to the limit, $\overline{\theta}(t^*, q^*) = -\ell(q^*)q^*$.

Eventually, we indeed have that

$$\min\left(-\gamma \partial_t \phi(T, q^*) - \gamma \mu q^* + \frac{1}{2}\gamma^2\sigma^2 q^{*2} - H\left(\partial_q \phi(T, q^*)\right), \overline{\theta}(T, q^*) + \ell(q^*)q^*\right) \leq 0.$$

If $q^* = 0$, then there are also two cases. If there are infinitely many indices n such that $q_n^* > 0$ and $t_n < T$, then the initial proof still works. Otherwise, for n sufficiently large $\theta_{\Delta_n}(t_n, q_n) = -\ell(q_n)q_n$ or $\theta_{\Delta_n}(t_n, q_n) = 0$ and hence, passing to the limit, we obtain $\overline{\theta}(t^*, q^*) = 0$.

Eventually, we indeed have that

$$\min\left(-\gamma \partial_t \phi(t^*, 0) - H\left(\partial_q \phi(t^*, 0)\right), \overline{\theta}(t^*, 0)\right) \leq 0.$$

We have proved that $\overline{\theta}$ is a subsolution of the equation in the viscosity and one can similarly prove that $\underline{\theta}$ is a supersolution.

Step 2: $\forall q \in [0, q_0], \quad \overline{\theta}(T, q) = \underline{\theta}(T, q) = -\ell(q)q.$

We consider the test function $\phi(t, q) = C_\epsilon(T - t) + \frac{1}{\epsilon}(q - q_{ref})^2$ where $q_{ref} \in [0, q_0]$ is fixed, where $\epsilon > 0$ is a constant and where C_ϵ is a constant that depends on ϵ, and that will be fixed later.

Let (t_ϵ, q_ϵ) be a maximum point of $\overline{\theta} - \phi$ on $[0, T] \times [0, q_0]$. Then,

$$\overline{\theta}(T, q_{\text{ref}}) \leq \overline{\theta}(t_\epsilon, q_\epsilon) - C_\epsilon(T - t_\epsilon) - \frac{1}{\epsilon}(q_\epsilon - q_{\text{ref}})^2.$$

This inequality gives $q_\epsilon \to q_{\text{ref}}$ as $\epsilon \to 0$.

Now,

$$-\gamma \partial_t \phi(t_\epsilon, q_\epsilon) - \gamma \mu q_\epsilon + \frac{1}{2}\gamma^2\sigma^2 {q_\epsilon}^2 - H\left(\partial_q \phi(t_\epsilon, q_\epsilon)\right) = \gamma C_\epsilon - \gamma \mu q_\epsilon + \frac{1}{2}\gamma^2\sigma^2 {q_\epsilon}^2 - H\left(\frac{2}{\epsilon}(q_\epsilon - q_{ref})\right)$$

$$\geq \gamma C_\epsilon - H\left(-\frac{2q_0}{\epsilon}\right) + \inf_{q \in [0, q_0]}\left(-\gamma \mu q + \frac{1}{2}\gamma^2\sigma^2 q^2\right).$$

Hence, if C_ϵ is positive and greater than $\frac{1}{\gamma} H\left(-\frac{2q_0}{\epsilon}\right) - \inf_{q \in [0, q_0]}\left(-\gamma \mu q + \frac{1}{2}\gamma^2\sigma^2 q^2\right) + 1$, we see that we must have either $t_\epsilon = T$ and $\overline{\theta}(t_\epsilon, q_\epsilon) \leq -\ell(q_\epsilon)q_\epsilon$ or (for sufficiently small ϵ, only in the case where $q_{\text{ref}} = 0$), $\overline{\theta}(t_\epsilon, q_\epsilon) \leq 0$.

If $q_{ref} \neq 0$, we then write, for ϵ sufficiently small:

$$\overline{\theta}(T, q_{\text{ref}}) \leq \overline{\theta}(t_\epsilon, q_\epsilon) - C_\epsilon(T - t_\epsilon) - \frac{1}{\epsilon}(q_\epsilon - q_{\text{ref}})^2 \leq \overline{\theta}(t_\epsilon, q_\epsilon) \leq -\ell(q_\epsilon)q_\epsilon.$$

Sending ϵ to 0, we obtain

$$\overline{\theta}(T, q_{\text{ref}}) \leq -\ell(q_{\text{ref}})q_{\text{ref}}.$$

If $q_{\text{ref}} = 0$, then we know from the above two inequalities that

$$\overline{\theta}(T, q_{\text{ref}}) \leq \overline{\theta}(t_\epsilon, q_\epsilon) - C_\epsilon(T - t_\epsilon) - \frac{1}{\epsilon}(q_\epsilon - q_{\text{ref}})^2 \leq \overline{\theta}(t_\epsilon, q_\epsilon) \leq \max(-\ell(q_\epsilon)q_\epsilon, 0) = 0.$$

Hence, $\forall q \in [0, q_0], \overline{\theta}(T, q) \leq -\ell(q)q$ and the same proof works for the supersolution[11] to get $\forall q \in [0, q_0], \underline{\theta}(T, q) \geq -\ell(q)q$.

As a consequence, $\forall q \in [0, q_0], \overline{\theta}(T, q) \leq -\ell(q)q \leq \underline{\theta}(T, q)$, and eventually, because $\underline{\theta}(T, q) \leq \overline{\theta}(T, q)$, we obtain that $\forall q \in [0, q_0], \overline{\theta}(T, q) = -\ell(q)q = \underline{\theta}(T, q)$.

Step 3: $\forall t \in [0, T], \quad \underline{\theta}(t, 0) = 0$

Concerning the boundary condition corresponding to $q = 0$, we can apply the same ideas but only to the supersolution $\underline{\theta}$.

[11]The only difference is that we have to pass to the limit in the case $q_{\text{ref}} = 0$ to obtain the conclusion.

Let us consider indeed $t_{ref} \in [0, T)$ and the test function $\phi(t, q) = -C_\epsilon q - \frac{1}{\epsilon}(t - t_{ref})^2$. Then, let (t_ϵ, q_ϵ) be a minimum point of $\underline{\theta} - \phi$ on $[0, T] \times [0, q_0]$. We have

$$\underline{\theta}(t_{\text{ref}}, 0) \geq \underline{\theta}(t_\epsilon, q_\epsilon) + C_\epsilon q_\epsilon + \frac{1}{\epsilon}(t_\epsilon - t_{\text{ref}})^2.$$

This inequality gives $t_\epsilon \to t_{\text{ref}}$ as $\epsilon \to 0$.

Now,

$$\begin{aligned} -\gamma \partial_t \phi(t_\epsilon, q_\epsilon) - \gamma \mu q_\epsilon + \frac{1}{2}\gamma^2\sigma^2 {q_\epsilon}^2 - H\left(\partial_q \phi(t_\epsilon, q_\epsilon)\right) &= 2\gamma \frac{t - t_{\text{ref}}}{\epsilon} - \gamma \mu q_\epsilon + \frac{1}{2}\gamma^2\sigma^2 {q_\epsilon}^2 \\ &\quad - H(-C_\epsilon) \\ &\leq 2\gamma \frac{T}{\epsilon} - \gamma \mu q_0 + \frac{1}{2}\gamma^2\sigma^2 {q_0}^2 - H(-C_\epsilon). \end{aligned}$$

Since $\lim_{p \to -\infty} H(p) = +\infty$, we can always choose $C_\epsilon \geq 0$ so that the above expression is strictly negative.

As a consequence, for ϵ sufficiently small, we must have $q_\epsilon = 0$ and $\underline{\theta}(t_\epsilon, 0) \geq 0$.

Consequently,

$$\underline{\theta}(t_{\text{ref}}, 0) \geq \underline{\theta}(t_\epsilon, q_\epsilon) + C_\epsilon q_\epsilon + \frac{1}{\epsilon}(t_\epsilon - t_{\text{ref}})^2 \geq 0.$$

This result being already true for $t_{\text{ref}} = T$, we have that $\underline{\theta}(t, 0) \geq 0$, $\forall t \in [0, T]$ and, in fact, $\underline{\theta}(t, 0) = 0$, $\forall t \in [0, T]$ because of the definition of $\underline{\theta}(t, 0)$.

Step 4: $\forall t \in [0, T]$, there exists a sequence $(t_n, q_n)_n$ such that $t_n \neq t$, $q_n \neq 0$, $(t_n, q_n) \to (t, 0)$, and $\underline{\theta}(t_n, q_n) \to 0$

To prove this claim, we prove that $g_\Delta(t) = \theta_\Delta(T - t, \Delta)$ converges uniformly (in t) toward 0 as $\Delta \to 0$.

By definition, $g_\Delta(0) = -\ell(\Delta)\Delta$ and $g'_\Delta(t) = \mu\Delta - \frac{1}{2}\gamma\sigma^2\Delta^2 + \frac{1}{\gamma} H_\Delta\left(\frac{g_\Delta(t)}{\Delta}\right)$.

We now distinguish two cases:

Case 1: $\mu \leq 0$.

The stationary state of the above ODE is $g^\infty_\Delta = \Delta H_\Delta^{-1}(\frac{1}{2}\gamma^2\sigma^2\Delta^2 - \gamma\mu\Delta)$ and g_Δ is increasing on $\{g_\Delta \leq g^\infty_\Delta\}$. Since $\Delta H_\Delta^{-1}(\frac{1}{2}\gamma^2\sigma^2\Delta^2 - \gamma\mu\Delta) \geq \Delta H_{\Delta'}^{-1}(\frac{1}{2}\gamma^2\sigma^2\Delta^2 - \gamma\mu\Delta)$ as soon as $\Delta < \Delta'$, g^∞_Δ is positive for Δ sufficiently small. As a consequence, since $g_\Delta(0) \leq 0$, g_Δ is increasing on [0,T].

Now, $g'_\Delta(t) \leq \frac{1}{\gamma} H\left(\frac{g_\Delta(t)}{\Delta}\right)$. Hence,

$$\int_0^{g_\Delta(t)} \frac{\gamma}{H\left(\frac{y}{\Delta}\right)} dy \leq \int_{-\ell(\Delta)\Delta}^{g_\Delta(t)} \frac{\gamma}{H\left(\frac{y}{\Delta}\right)} dy \leq t,$$

and this gives $g_\Delta(t) \leq \Delta G^{-1}\left(\frac{T}{\gamma\Delta}\right)$, where $G(x) = \int_0^x \frac{1}{H(y)} dy$.

Consequently,

$$\forall t \in [0, T], \quad -\ell(\Delta)\Delta \leq g_\Delta(t) \leq \Delta G^{-1}\left(\frac{T}{\gamma\Delta}\right).$$

Since $\lim_{\delta\to+\infty}\delta\Lambda(\delta)=0$, we have $\lim_{x\to+\infty}H(x)=0$, and therefore $\lim_{x\to+\infty}\frac{G(x)}{x}=+\infty$. This proves that g_Δ converges uniformly toward 0 on $[0,T]$.

Case 2: $\mu>0$.

In this second case, we know that $\exists\Delta_0$ such that $\Delta\in(0,\Delta_0)\mapsto\mu\Delta-\frac{1}{2}\gamma\sigma^2\Delta^2$ is an increasing and positive function. Now, let us introduce $\Delta_1\in(0,\Delta_0)$. $\forall\Delta\in(0,\Delta_1)$, we have

$$\int_0^{g_\Delta(T)}\frac{\gamma}{\gamma\mu\Delta_1-\frac{1}{2}\gamma^2\sigma^2\Delta_1^2+H\left(\frac{y}{\Delta}\right)}dy$$

$$\leq\int_{-\ell(\Delta)\Delta}^{g_\Delta(T)}\frac{\gamma}{\gamma\mu\Delta-\frac{1}{2}\gamma^2\sigma^2\Delta^2+H_\Delta\left(\frac{y}{\Delta}\right)}dy=T.$$

This gives $g_\Delta(T)\leq\Delta G_{\Delta_1}^{-1}\left(\frac{T}{\gamma\Delta}\right)$, where $G_{\Delta_1}(x)=\int_0^x\frac{1}{\gamma\mu\Delta_1-\frac{1}{2}\gamma^2\sigma^2\Delta_1^2+H(y)}dy$.
Now, because $\lim_{x\to+\infty}H(x)=0$, we have

$$\liminf_{x\to+\infty}\frac{G_{\Delta_1}(x)}{x}\geq\frac{1}{\gamma\mu\Delta_1-\frac{1}{2}\gamma^2\sigma^2\Delta_1^2}.$$

Hence,

$$\liminf_{\Delta\to0}\frac{T}{\gamma\Delta G_{\Delta_1}^{-1}\left(\frac{T}{\gamma\Delta}\right)}\geq\frac{1}{\gamma\mu\Delta_1-\frac{1}{2}\gamma^2\sigma^2\Delta_1^2}.$$

This gives

$$\limsup_{\Delta\to0}g_\Delta(T)\leq T(\mu\Delta_1-\frac{1}{2}\gamma\sigma^2\Delta_1^2),$$

and sending Δ_1 to 0, we get

$$\limsup_{\Delta\to0}g_\Delta(T)\leq0.$$

The result is then proved since

$$\forall t\in[0,T],\quad-\ell(\Delta)\Delta\leq g_\Delta(t)\leq g_\Delta(T).$$

Step 5: Comparison principle: $\overline{\theta}\leq\underline{\theta}$.

Let us consider $\alpha>0$ and the maximum $M=\max_{(t,q)\in[0,T]\times[0,q_0]}\overline{\theta}(t,q)-\underline{\theta}(t,q)-\alpha(T-t)$. If $m=\max_{t\in[0,T]}\overline{\theta}(t,0)-\underline{\theta}(t,0)-\alpha(T-t)<M$, then we distinguish two cases.

Case 1: $m\leq0$.

We introduce $\Phi_\epsilon(t,q,t',q')=\overline{\theta}(t,q)-\underline{\theta}(t',q')-\alpha(T-t)-\frac{(q-q')^2}{\epsilon}-\frac{(t-t')^2}{\epsilon}$.
Let us consider $(t_\epsilon,q_\epsilon,t'_\epsilon,q'_\epsilon)$ a maximum point of Φ_ϵ. We have $M\leq\Phi_\epsilon(t_\epsilon,q_\epsilon,t'_\epsilon,q'_\epsilon)$ and we are going to prove that $\liminf_{\epsilon\to0}\Phi_\epsilon(t_\epsilon,q_\epsilon,t'_\epsilon,q'_\epsilon)\leq0$.
We have $\Phi_\epsilon(t_\epsilon,q_\epsilon,t_\epsilon,q_\epsilon)\leq\Phi_\epsilon(t_\epsilon,q_\epsilon,t'_\epsilon,q'_\epsilon)$ and hence $\frac{(q_\epsilon-q'_\epsilon)^2}{\epsilon}+\frac{(t_\epsilon-t'_\epsilon)^2}{\epsilon}$ is bounded. As a consequence, $t_\epsilon-t'_\epsilon\to0$ and $q_\epsilon-q'_\epsilon\to0$.

Now, if for all ϵ sufficiently small we have $(t_\epsilon, q_\epsilon, t'_\epsilon, q'_\epsilon) \in [0, T) \times (0, q_0] \times [0, T) \times (0, q_0]$, then we have

$$-2\gamma \frac{t_\epsilon - t'_\epsilon}{\epsilon} + \gamma\alpha - \gamma\mu q_\epsilon + \frac{1}{2}\gamma^2\sigma^2 q_\epsilon^2 - H(2\frac{q_\epsilon - q'_\epsilon}{\epsilon}) \leq 0,$$

and

$$-2\gamma \frac{t_\epsilon - t'_\epsilon}{\epsilon} - \gamma\mu q'_\epsilon + \frac{1}{2}\gamma^2\sigma^2 {q'_\epsilon}^2 - H(2\frac{q_\epsilon - q'_\epsilon}{\epsilon}) \geq 0.$$

Hence, $\gamma\alpha - \gamma\mu(q_\epsilon - q'_\epsilon) + \frac{1}{2}\gamma^2\sigma^2(q_\epsilon^2 - {q'_\epsilon}^2) \leq 0$ and this is a contradiction as we send ϵ to 0.

The consequence is that there exists a sequence $\epsilon_n \to 0$ such that $(t_{\epsilon_n}, q_{\epsilon_n}, t'_{\epsilon_n}, q'_{\epsilon_n})$ verifies $t_{\epsilon_n} = T, \forall n$ or $q_{\epsilon_n} = 0, \forall n$ or $t'_{\epsilon_n}, \forall n$ or $q'_{\epsilon_n}, \forall n$.

Then,

$$\limsup_{n\to+\infty} \Phi_{\epsilon_n}(t_{\epsilon_n}, q_{\epsilon_n}, t'_{\epsilon_n}, q'_{\epsilon_n}) \leq \limsup_{n\to+\infty} \overline{\theta}(t_{\epsilon_n}, q_{\epsilon_n}) - \underline{\theta}(t'_{\epsilon_n}, q'_{\epsilon_n}) - \alpha(T - t_{\epsilon_n})$$

$$\leq \max_{(t,q)\in(\{T\}\times[0,q_0])\cup([0,T]\times\{q_0\})} \overline{\theta}(t, q) - \underline{\theta}(t, q) - \alpha(T - t) \leq \max(0, m) \leq 0.$$

Hence, in that case. $M \leq 0$.

Case 2: $m > 0$.

In that case, we replace $\underline{\theta}$ by $\underline{\theta} + m$ in the above case, and we obtain, instead of $M \leq 0$, the inequality $M \leq m$ that contradicts our hypothesis.

It remains to consider the case $m = M$. We know then that the maximum M is attained at a point $(t_{\max}, 0)$ and we suppose that $M > 0$.

From Step 4, we consider a sequence (t_n, q_n) such that $t_n \neq t_{\max}$, $q_n \neq 0$, $(t_n, q_n) \to (t_{\max}, 0)$ and $\underline{\theta}(t_n, q_n) \to 0$.

We define

$$\Psi_n(t, q, t', q') = \overline{\theta}(t, q) - \underline{\theta}(t', q') - \alpha(T - t) - \frac{(t - t')^2}{|t_n - t_{\max}|} - \left(\frac{q' - q}{q_n} - 1\right)^2.$$

This function attains its maximum at a point $(t_n^*, q_n^*, {t'}_n^*, {q'}_n^*)$. We first consider the inequality $\Psi_n(t_{\max}, 0, t_n, q_{\max}) \leq \Psi_n(t_n^*, q_n^*, {t'}_n^*, {q'}_n^*)$:

$$\overline{\theta}(t_{\max}, 0) - \underline{\theta}(t_n, q_n) - \alpha(T - t_{\max}) - |t_n - t_{\max}|$$

$$\leq \overline{\theta}(t_n^*, q_n^*) - \underline{\theta}({t'}_n^*, {q'}_n^*) - \alpha(T - t_n^*) - \frac{(t_n^* - {t'}_n^*)^2}{|t_n - t_{\max}|} - \left(\frac{{q'}_n^* - q_n^*}{q_n} - 1\right)^2.$$

We then have $t_n^* - {t'}_n^* \to 0$ and ${q'}_n^* - q_n^* \to 0$.

Hence, $\limsup_n \overline{\theta}(t_n^*, q_n^*) - \underline{\theta}({t'}_n^*, {q'}_n^*) - \alpha(T - t_n^*) \leq M$. Now, the maximum M is also given by $\lim_n \overline{\theta}(t_{\max}, 0) - \underline{\theta}(t_n, q_n) - \alpha(T - t_{\max})$, and we obtain that the penalization term $\left(\frac{{q'}_n^* - q_n^*}{q_n} - 1\right)^2$ converges to 0.

This gives ${q'}_n^* = q_n^* + q_n + o(q_n) > 0$.

Now, if we have infinitely many n such that $t'^*_n = T$, then

$$M = \lim_n \overline{\theta}(t_{\max}, 0) - \underline{\theta}(t_n, q_n) - \alpha(T - t_{\max}) \leq \sup_{q \in [0, q_0]} \overline{\theta}(T, q) - \underline{\theta}(T, q) = 0,$$

and this contradicts our hypothesis.

Otherwise, for all n sufficiently large:

$$-2\gamma \frac{t^*_n - t'^*_n}{|t_n - t_{\max}|} - \gamma \mu q'^*_n + \frac{1}{2}\gamma^2\sigma^2 q'^{*2}_n - H\left(-\frac{2}{q_n}\left(\frac{q'^*_n - q^*_n}{q_n} - 1\right)\right) \geq 0.$$

Now, going to the subsolution, if there are infinitely many n such that $t^*_n = T$, then

$$M = \lim_n \overline{\theta}(t_{\max}, 0) - \underline{\theta}(t_n, q_n) - \alpha(T - t_{\max}) \leq \sup_{q \in [0, q_0]} \overline{\theta}(T, q) - \underline{\theta}(T, q) = 0,$$

in contradiction with our hypothesis.

Else, if $q^*_n = 0$ and $\overline{\theta}(t^*_n, q^*_n) \leq 0$ for infinitely many n, then we obtain $M = \lim_n \overline{\theta}(t_{\max}, 0) - \underline{\theta}(t_n, q_n) - \alpha(T - t_{\max}) - |t_n - t_{\max}| \leq 0$ straightforwardly, and this contradicts our hypothesis. Hence, the viscosity inequality must be satisfied and we get

$$-2\gamma \frac{t^*_n - t'^*_n}{|t_n - t_{\max}|} + \gamma\alpha - \gamma\mu q^*_n + \frac{1}{2}\gamma^2\sigma^2 q^{*2}_n - H\left(-\frac{2}{q_n}\left(\frac{q'^*_n - q^*_n}{q_n} - 1\right)\right) \leq 0.$$

Combining the two inequalities eventually leads to $\alpha\gamma \leq 0$ as $n \to \infty$ and this is a contradiction.

We have obtained that $M \leq 0$ and hence $\overline{\theta} - \underline{\theta} \leq \alpha T$. Sending α to 0, we get $\overline{\theta} \leq \underline{\theta}$.

This proves that $\overline{\theta} = \underline{\theta}$ is, in fact, a continuous function that we call θ, solution of the PDE ($\mathrm{HJ}_{\lim}$). Using the same techniques as above, it is clear that θ is the unique viscosity solution of ($\mathrm{HJ}_{\lim}$).

We have

$$\theta(t, q) = \underline{\theta}(t, q) \leq \liminf_{\Delta \to 0} \theta^c_\Delta(t, q) \leq \limsup_{\Delta \to 0} \theta^c_\Delta(t, q) \leq \overline{\theta}(t, q) = \theta(t, q).$$

Hence, $\theta(t, q) = \lim_{\Delta \to 0} \theta^c_\Delta(t, q)$ and by the same token, $\lim_{\Delta \to 0, (t', q') \to (t, q)} \theta^c_\Delta(t', q') = \theta(t, q)$ so that the convergence is locally uniform and then uniform on the compact set $[0, T] \times [0, q_0]$. □

6. LINK WITH ALMGREN–CHRISS

6.1. Interpretation of the PDE in the Limit Regime

In the above section, we proved that θ_Δ converged to θ, which is the unique continuous viscosity solution of ($\mathrm{HJ}_{\lim}$), the limit condition and boundary condition being satisfied in the classical sense.

Now, we are going to link this equation to a classical equation of Almgren–Chriss–like models. The intuition behind the link between our framework in the limit regime $\Delta \to 0$ and the Almgren–Chriss framework is that nonexecution risk vanishes as Δ tends to 0. Hence, the only remaining risk is price risk, corresponding to the term $\frac{1}{2}\gamma^2\sigma^2 q^2$ in the above equation. To see more precisely the correspondence between the two approaches,

let us write the hamiltonian function as

$$\begin{aligned} H(p) &= \gamma \sup_{\delta} \Lambda(\delta)(\delta - p) \\ &= \gamma \sup_{v>0} v(\Lambda^{-1}(v) - p) \\ &= \gamma \sup_{v\geq 0} v\Lambda^{-1}(v) - pv, \end{aligned}$$

where the last equality holds since $\lim_{\delta\to+\infty} \delta\Lambda(\delta) = 0$.

Hence, if we define for $v > 0$, $f(v) = -\Lambda^{-1}(v)$, then we can define the function $\tilde{H}(p) = \sup_{v\geq 0} -f(v)v - pv$ and write the partial differential equation for θ as

$$\begin{cases} -\partial_t\theta(t,q) - \mu q + \frac{1}{2}\gamma\sigma^2 q^2 - \tilde{H}\left(\partial_q\theta(t,q)\right) = 0, & \text{on } [0,\, T) \times (0,\, q_0], \\ \theta(t,q) = 0, & \text{on } [0,\, T] \times \{0\}, \\ \theta(t,q) = -\ell(q)q, & \text{on } \{T\} \times [0,\, q_0]. \end{cases}$$

This equation is the Hamilton–Jacobi equation associated to the Almgren–Chriss optimal liquidation problem with an instantaneous market impact function (per share) f and with a final discount $\ell(q_T)$ per share. More precisely, the above Hamilton–Jacobi equation is the Hamilton–Jacobi equation associated to the optimization problem:[12]

$$\inf_{q\in AC(0,T), q(0)=q_0} \int_0^T \left(-q'(t)f(-q'(t)) - \mu q(t) + \frac{1}{2}\gamma\sigma^2 q(t)^2\right) dt + q(T)\ell(q(T)),$$

where $AC(0,\, T)$ is the set of absolutely continuous functions on $[0,\, T]$.

However, the instantaneous market impact function f has here a rather unusual form since $f(v)$ is negative for $v < \Lambda(0)$ and positive for $v > \Lambda(0)$, whereas it is usually a positive function. This must be interpreted in a very simple way: if one needs to obtain an instantaneous volume lesser than $\Lambda(0)$, then one will choose a positive δ. In other words, he will post a classical limit order, since we assumed that the reference price is the first bid quote—this makes sense since nonexecution risk disappears in the limit regime $\Delta \to 0$. On the contrary, if one needs an instantaneous volume greater than $\Lambda(0)$, then one will rely on a marketable limit order ($\delta < 0$).

The above discussion only makes sense at the limit, when nonexecution risk does not exist anymore. However, it clarifies the meaning of negative δs. In particular, the above correspondence between our model and a model *à la* Almgren–Chriss provides a possible way to solve one of the main practical problems of the model discussed in Guéant et al. (2012): the interpretation of the intensity functions for negative values of δ.

6.2. Discussion on the Choice of Λ

The model we discuss in this paper does not consider explicitly market orders or limit orders but rather considers that there is, for each price $s^a = s + \delta$, an instantaneous probability to obtain a trade at that price. In practice, this interpretation is perfectly suited to classical limit orders, but we need to provide an interpretation for the intensity function Λ on the entire real line. In practice, since the model has been designed to liquidate a

[12] In the usual Almgren–Chriss framework, liquidation is mandatory but the theory can easily be adapted to allow for a penalization term.

position with limit orders, it should not be used if the liquidation evidently requires liquidity-taking orders from the very beginning. However, it may happen, because of a slow execution, that the optimal quote in the model turns out to be negative[13] after some time.

This issue of negative δ was present in the model with exponential intensity functions introduced in Guéant et al. (2012). Although the exponential form was justified for many stocks for positive δs, the intensity function was also of exponential form for $\delta < 0$, and this choice was dictated by mathematical needs rather than by empirical rationale. Since Λ (or in practice Λ_Δ) can be chosen in our setting, we can improve the initial model.

First, using statistics on execution, we can estimate the probability to be executed at any positive distance from the first bid limit. In practice, the profile of the empirical intensity for positive δ is decreasing and may not be convex, especially when the bid-ask spread is large (this is the rationale underlying our choice – for $\delta \geq 0$—in Figure 4.1). Then, once the function Λ_Δ has been calibrated for positive δ, several natural choices are possible for $\Lambda_\Delta(\delta)$ when $\delta \leq 0$. Instead of extending the function for negative δ using a specific functional form as in Guéant et al. (2012), we can assign to $\Lambda_\Delta(\delta)$ a constant value when $\delta \leq 0$ as above for $\tilde{\Lambda}_\Delta$. This corresponds to a very conservative choice that basically prevents the use of marketable limit orders since, intuitively, the optimal quote will then always be positive. Another choice consists in using the parallel made between the usual literature and our framework in the limit regime $\Delta \to 0$. If we indeed omit nonexecution risk, we can consider, for $\delta \leq 0$, that $\Lambda_\Delta(\delta) = \frac{1}{\Delta} f^{-1}(-\delta) = \frac{1}{\Delta} \sup\{v \geq 0,\, f(v) \leq -\delta\}$ where $v \mapsto f(v)$ is an instantaneous market impact function (average execution cost per share) that is typically equal to nought for small values of v and increasing after a certain threshold.

This choice for Λ_Δ is, however, subject to several comments. First, the function Λ_Δ must satisfy the hypotheses of the model. In particular, it must be decreasing. However, it may happen that although the specifications for positive δ and negative δ are both decreasing, the function is not decreasing on the entire real line. Since the function $\delta \leq 0 \mapsto \frac{1}{\Delta} f^{-1}(-\delta)$ can be considered a lower bound to $\Lambda_\Delta(\delta)$ because we have to take into account the risk of nonexecution, we can always scale the function $\delta \leq 0 \mapsto \Lambda_\Delta(\delta)$ so that the resulting function Λ_Δ is decreasing (and strictly decreasing if we smooth the function). In general, the inequality $\Lambda'' \Lambda \leq 2\Lambda'^2$ may not be satisfied but the model can still be used although the optimal quotes δ^* may not be unique. Second, a question remains regarding the interpretation of the model when an optimal quote δ^* turns out to be negative. A possible answer, in line with the parallel made with Almgren–Chriss–like models, is to send a market order of size $\Lambda(\delta^*) = \Delta \Lambda_\Delta(\delta^*)$ (or, in practice, to use marketable limit orders to obtain this size).

7. CONCLUSION

In this paper, we analyze optimal liquidation using limit orders. The classical literature on optimal liquidation, following Almgren–Chriss, only considers optimal scheduling and a new strand of research has recently emerged that uses either dark pools or limit orders to tackle the issue of the actual optimal way to liquidate. Our paper provides a general model for optimal liquidation with limit orders and extends both Bayraktar and

[13] Although quotes evolve continuously between execution times, using the model, in practice, requires to post orders and to keep them in the order book for some time. Hence, the optimal quote at some point may, in practice, be strictly negative.

Ludkovski (2014) that only considers a risk-neutral framework and Guéant et al. (2012) that was restricted to exponential intensity functions. Our general framework also sheds new light on the important topic of negative quotes. An important improvement of our model would consist in linking the Brownian motion that drives the price and the point process modeling execution. Research in this direction has recently been made by Cartea et al. (2014) to model market making and it may adapt to our case.

APPENDIX A: THE MULTIASSET CASE

This appendix is devoted to the generalization of our results to the case of a portfolio with multiple stocks. Our main result (Theorem 3.4) generalizes to the multiasset case.

We consider a trader who has to liquidate a portfolio made of d different stocks with a quantity q_0^i of stock i ($i \in \{1, \dots, d\}$).

We suppose that the reference price of stock i evolves as

$$dS_t^i = \mu^i dt + \sigma^i dW_t^i,$$

with $\mathbb{V}(W^1, \dots, W^d) = (\rho^{i,j})_{1 \leq i,j \leq d}$ definite positive.

The trader under consideration continuously proposes an ask quote for each stock ($S_t^{i,a} = S_t^i + \delta_t^i$ for stock i), and will hence sell shares according to the rate of arrival of liquidity-taking orders at the prices he quotes.

The state of the portfolio is $\left(q^{1,\delta^1}, \dots, q^{d,\delta^d}\right)$. It evolves according to the following dynamics:

$$\forall i, dq_t^{i,\delta^i} = -\Delta^i N_t^{i,\delta^i},$$

where N^{i,δ^i} is a point process giving the number of executed orders for stock i, each order on stock i being of size Δ^i—we suppose as above that Δ^i is a fraction of q_0^i. The intensity process $(\lambda_t^i)_t$ of the point process N^{i,δ^i} is given by

$$\lambda_t^i = \Lambda_{\Delta^i}^i(S_t^{i,a} - S_t^i) 1_{q_{t-}^{i,\delta^i} > 0} = \Lambda_{\Delta^i}^i(\delta_t^i) 1_{q_{t-}^{i,\delta^i} > 0},$$

where the function $\Lambda_{\Delta^i}^i : \mathbb{R} \to \mathbb{R}_+$ satisfies the same assumptions as in the single-stock case.

We suppose that the point processes $N^{1,\delta^1}, \dots, N^{d,\delta^d}$ are independent.

The cash account $X^{\delta^1,\dots,\delta^d}$ of the trader has then the following dynamics:

$$dX_t^{\delta^1,\dots,\delta^d} = \sum_{i=1}^d (S_t^i + \delta_t^i)\Delta^i dN_t^{i,\delta^i}.$$

Eventually, the optimization problem is

$$\sup_{(\delta^1,\dots,\delta^d) \in \mathcal{A}^d} \mathbb{E}\left[-\exp\left(-\gamma\left(X_T^{\delta^1,\dots,\delta^d} + \sum_{i=1}^d q_T^{i,\delta^i}\left(S_T^i - \ell^i\left(q_T^{i,\delta^i}\right)\right)\right)\right)\right],$$

where the functions $\ell^1, \dots, \ell^d$ satisfy the same assumptions as in the single-asset case.

This multiasset setting deserves two remarks. First, the execution processes associated to different stocks are independent. Hence, the only difference between the multiasset case and the single-asset cases has to do with price risk: in this multiasset framework,

optimal quotes will depend on the correlation structure between stocks. Second, the trader can only sell shares, although buying shares might sometimes reduce price risk in practice.

The counterpart of Proposition 3.3 and Theorem 3.4 in this multiasset framework is the theorem below:

Theorem A.1. (Verification theorem and optimal quotes)
There exists a unique solution $\theta_{\Delta^1,\dots,\Delta^d}$, C^1 in time, of the system:

$$\forall t \in [0, T), (q^1, \dots, q^d) \neq (0, \dots, 0),$$

$$0 = \gamma \partial_t \theta_{\Delta^1,\dots,\Delta^d}(t, q^1, \dots, q^d) + \gamma \sum_{i=1}^{d} \mu^i q^i - \frac{1}{2}\gamma^2 \sum_{1 \le i,j \le d} \rho^{i,j} \sigma^i \sigma^j q^i q^j$$

$$+ \sum_{i=1}^{d} 1_{q^i > 0} H^i_{\Delta^i} \left(\frac{\theta_{\Delta^1,\dots,\Delta^d}(t, q^1, \dots, q^d) - \theta_{\Delta^1,\dots,\Delta^d}(t, q^1, \dots, q^i - \Delta^i, \dots, q^d)}{\Delta^i} \right),$$

and the conditions:

$$\theta_{\Delta^1,\dots,\Delta^d}(T, q^1, \dots, q^d) = -\sum_{i=1}^{d} \ell(q^i) q^i, \qquad \theta_{\Delta^1,\dots,\Delta^d}(t, 0, \dots, 0) = 0,$$

where

$$\forall i,\ H^i_{\Delta^i}(p) = \sup_{\delta^i} \Lambda^i_{\Delta^i}(\delta^i) \left(1 - e^{-\gamma \Delta^i (\delta^i - p)}\right).$$

If $\theta_{\Delta^1,\dots,\Delta^d}$ is this function, then

$$u_{\Delta^1,\ \dots,\ \Delta^d}(t, x, q^1, \dots, q^d, s^1, \dots, s^d) = - \exp\left(-\gamma \left(x + \sum_{i=1}^{d} q^i s^i + \theta_{\Delta^1,\dots,\Delta^d}(t, q^1, \dots, q^d)\right)\right)$$

is the value function of the optimal control problem, and the optimal ask quotes are characterized by

$$\forall i, (\delta^{i*}_{\Delta^i})_t = \tilde{\delta}^{i*}_{\Delta^i} \left(\frac{\theta_{\Delta^1,\dots,\Delta^d}(t, q^1_t, \dots, q^d_t) - \theta_{\Delta^1,\dots,\Delta^d}(t, q^1_t, \dots, q^i_t - \Delta^i, \dots, q^d_t)}{\Delta^i} \right),$$

where $\tilde{\delta}^{i*}_{\Delta^i}(p)$ is uniquely characterized by

$$\tilde{\delta}^{i*}_{\Delta^i}(p) - \frac{1}{\gamma \Delta^i} \log\left(1 - \gamma \Delta^i \frac{\Lambda^i_{\Delta^i}(\tilde{\delta}^{i*}_{\Delta^i}(p))}{{\Lambda^i}'_{\Delta^i}(\tilde{\delta}^{i*}_{\Delta^i}(p))}\right) = p.$$

Proof. The proof is *mutatis mutandis* the same as in the single-stock case. Existence of a local solution $t \mapsto \theta_{\Delta^1,\dots,\Delta^d}(t, q^1, \dots, q^d)$ comes from Cauchy–Lipschitz. To obtain existence on $[0, T]$, first notice that $\theta_{\Delta^1,\dots,\Delta^d}(t, q^1, \dots, q^d) + \sum_{i=1}^{d} \mu^i q^i (T - t) -$

$\frac{1}{2}\gamma \sum_{1\leq i,j\leq d} \rho^{i,j}\sigma^i\sigma^j q^i q^j (T-t)$ is a decreasing function of t. Now, using a comparison principle similar to Proposition 3.2, we have

$$\theta_{\Delta^1,\dots,\Delta^d}(t,q^1,\dots,q^d) \leq \sum_{i=1}^{d} \mu^{i+} q_0^i (T-t) + \frac{1}{\gamma}\sum_{i=1}^{d} H_{\Delta^i}(0)(T-t),$$

and this bound guarantees that there is no blow up. This led to global existence. Uniqueness follows from a comparison principle.

As far as verification is concerned, the proof is exactly the same as in the single-asset case. □

APPENDIX B: TRADING ON BOTH SIDES

Along with optimal liquidation, an important strand of research on optimal trading is high-frequency market making. We claimed in the introduction that many models can be used to deal with both optimal liquidation and market making. We illustrate this claim and show that our framework can easily be adapted to two-sided trading. We present the model in the single-stock case. The multistock case works the same but notations make the exposition cumbersome.

We consider a stock with a reference price following a Brownian motion with a drift:

$$dS_t = \mu dt + \sigma dW_t.$$

The trader under consideration continuously proposes a bid quote $S_t^b = S_t - \delta_t^b$ and an ask quote $S_t^a = S_t + \delta_t^a$. His inventory q^{δ^b,δ^a} evolves according to the rate of arrival of liquidity-taking orders at the prices he quotes:

$$dq_t^{\delta^b,\delta^a} = \Delta dN_t^{\delta^b} - \Delta dN_t^{\delta^a},$$

where N^{δ^b} and N^{δ^a} are the point processes giving the number of executed orders, respectively, on the bid side and on the ask side, each order being of size Δ (on both sides). The intensity processes $(\lambda_t^b)_t$ and $(\lambda_t^a)_t$ of the point processes N^{δ^b} and N^{δ^a} are given by

$$\lambda_t^b = \Lambda_\Delta(S_t - S_t^b)1_{q_{t-}^{\delta^b,\delta^a} < Q} = \Lambda_\Delta(\delta_t^b)1_{q_{t-}^{\delta^b,\delta^a} < Q},$$

and

$$\lambda_t^a = \Lambda_\Delta(S_t^a - S_t)1_{q_{t-}^{\delta^b,\delta^a} > -Q} = \Lambda_\Delta(\delta_t^a)1_{q_{t-}^{\delta^b,\delta^a} > -Q},$$

where $\Lambda_\Delta : \mathbb{R} \to \mathbb{R}_+$ satisfies the same assumptions as in the case of the optimal liquidation model, and where Q is a bound on the inventory. The bounds on the inventory have two roles: (i) they stand for the risk limits the traders have in practice, and (ii) they allow to write a verification theorem as in Guéant et al. (2013).

As a consequence of his trades, the cash account X^{δ^b,δ^a} of the trader has the following dynamics:

$$dX_t^\delta = (S_t + \delta_t^a)\Delta dN_t^{\delta^a} - (S_t - \delta_t^b)\Delta dN_t^{\delta^b}.$$

Finally, the optimization problem of the high-frequency market maker is

$$\sup_{\delta^b,\delta^a\in\mathcal{A}} \mathbb{E}\left[-\exp\left(-\gamma\left(X_T^{\delta^b,\delta^a}+q_T^{\delta^b,\delta^a}S_T-|q_T^{\delta^b,\delta^a}|\ell(|q_T^{\delta^b,\delta^a}|)\right)\right)\right],$$

where ℓ satisfies the same assumptions as in the case of the optimal liquidation model.

The counterpart of Proposition 3.3 and Theorem 3.4 in this market model framework is the theorem below:

Theorem B.1. (Verification theorem and optimal quotes)
There exists a unique solution θ_Δ, C^1 in time, of the system:

$$\forall t\in[0,T),\forall q\in\{-Q,\dots,-\Delta,0,\Delta,\dots,Q\},\quad 0=\gamma\partial_t\theta_\Delta(t,q)+\gamma\mu q-\frac{1}{2}\gamma^2\sigma^2q^2$$

$$+1_{q<Q}H_\Delta\left(\frac{\theta_\Delta(t,q)-\theta_\Delta(t,q+\Delta)}{\Delta}\right)+1_{q>-Q}H_\Delta\left(\frac{\theta_\Delta(t,q)-\theta_\Delta(t,q-\Delta)}{\Delta}\right)$$

and the terminal condition:

$$\theta_\Delta(T,q)=-\ell(|q|)|q|,$$

where $H_\Delta(p)=\sup_\delta\Lambda_\Delta(\delta)\left(1-e^{-\gamma\Delta(\delta-p)}\right)$.
If θ_Δ is this function, then

$$u_\Delta(t,x,q,s)=-\exp\left(-\gamma\left(x+qs+\theta_\Delta(t,q)\right)\right)$$

is the value function of the optimal control problem, and the optimal bid and ask quotes are characterized by

$$(\delta_\Delta^{b*})_t=\tilde{\delta}_\Delta^*\left(\frac{\theta_\Delta(t,q_t)-\theta_\Delta(t,q_t+\Delta)}{\Delta}\right),$$

$$(\delta_\Delta^{a*})_t=\tilde{\delta}_\Delta^*\left(\frac{\theta_\Delta(t,q_t)-\theta_\Delta(t,q_t-\Delta)}{\Delta}\right),$$

where $\tilde{\delta}_\Delta^*(p)$ is uniquely characterized by the equation $\left(E_{\delta_\Delta^*}\right)$ of Lemma 3.1.

Proof. The proof is close to the proof in the one-sided case. Existence of a local solution $t\mapsto(\theta_\Delta(t,q))_{q\in\{-Q,\dots,-\Delta,0,\Delta,\dots,Q\}}$ comes from Cauchy–Lipschitz. To obtain existence on $[0,T]$, first notice that $\theta_\Delta(t,q)+\mu q(T-t)-\frac{1}{2}\gamma\sigma^2q^2(T-t)$ is a decreasing function of t. Then, using a comparison principle, we have

$$\theta_\Delta(t,q)\le|\mu|Q(T-t)+\frac{2}{\gamma}H_\Delta(0)(T-t),$$

and this bound guarantees that there is no blow up, hence global existence. Uniqueness follows from a comparison principle.

As far as verification is concerned, the proof similar to the proof of Theorem 3.4. □

REFERENCES

ALFONSI, A., A. FRUTH, and A. SCHIED (2008): Constrained Portfolio Liquidation in a Limit Order Book Model, *Banach Center Publ.* 83, 9–25.

ALFONSI, A., A. FRUTH, and A. SCHIED (2010): Optimal Execution Strategies in Limit Order Books with General Shape Functions, *Quant. Finance* 10(2), 143–157.

ALFONSI, A. and A. SCHIED (2010): Optimal Trade Execution and Absence of Price Manipulations in Limit Order Book Models, *SIAM J. Financ. Math.* 1, 490–522.

ALFONSI, A., A. SCHIED, and A. SLYNKO (2012): Order Book Resilience, Price Manipulation, and the Positive Portfolio Problem, *SIAM J. Financ. Math.* 3(1), 511–533.

ALMGREN, R. (2009): Optimal Trading in a Dynamic Market. Preprint.

ALMGREN, R. (2012): Optimal Trading with Stochastic Liquidity and Volatility, *SIAM J. Financ. Math.* 3(1), 163–181.

ALMGREN, R., and N. CHRISS (1999): Value under Liquidation, *Risk* 12(12), 61–63.

ALMGREN, R., and N. CHRISS (2001): Optimal Execution of Portfolio Transactions, *J. Risk* 3, 5–40.

ALMGREN, R., and J. LORENZ (2007): Adaptive Arrival Price, *J. Trading* 2007(1), 59–66.

ALMGREN, R. F. (2003): Optimal Execution with Nonlinear Impact Functions and Trading-Enhanced Risk, *Appl. Math. Finance* 10(1), 1–18.

AVELLANEDA, M., and S. STOIKOV (2008): High-Frequency Trading in a Limit Order Book, *Quant. Finance* 8(3), 217–224.

BARLES, G., and P. E. SOUGANIDIS (1990): Convergence of Approximation Schemes for Fully Nonlinear Second Order Equations, in *Proceedings of the 29th IEEE Conference on Decision and Control*, pp. 2347–2349.

BAYRAKTAR, E., and M. LUDKOVSKI (2014): Liquidation in Limit Order Books with Controlled Intensity, *Math. Finance* 24(4), 627–650.

BERTSIMAS, D., and A. LO (1998): Optimal Control of Execution Costs, *J. Financ. Markets* 1(1), 1–50.

BOUCHARD, B., N. M. DANG, and C.-A. LEHALLE (2011): Optimal Control of Trading Algorithms: A General Impulse Control Approach, *SIAM J. Financ. Math.* 2(1), 404–438.

CARTEA, Á., and S. JAIMUNGAL (2013): Modeling Asset Prices for Algorithmic and High Frequency Trading. *Appl. Math. Finance* 20(6), 512–547.

CARTEA, Á., and S. JAIMUNGAL (2015): Risk Metrics and Fine Tuning of High Frequency Trading Strategies. *Math. Finance* 25(3), 576–611.

CARTEA, Á., S. JAIMUNGAL, and J. RICCI (2014): Buy Low Sell High: A High Frequency Trading Perspective. *SIAM J. Financ. Math*. 5(1), 415–444.

FORSYTH, P. A., J. S. KENNEDY, S. T. TSE, and H. WINDCLIFF (2012): Optimal Trade Execution: A Mean Quadratic Variation Approach, *J. Econ. Dyn. Control*. 36(12), 1971–1991.

GATHERAL, J. (2010): No-Dynamic-Arbitrage and Market Impact, *Quant. Finance* 10(7), 749–759.

GATHERAL, J., A. SCHIED, and A. SLYNKO (2012): Transient Linear Price Impact and Fredholm Integral Equations, *Math. Finance* 22(3), 445–474.

GUÉANT, O. (2012): Optimal Execution and Block Trade Pricing: A General Framework. Working Paper.

GUÉANT, O., C.-A. LEHALLE, and J. FERNANDEZ-TAPIA (2012): Optimal Portfolio Liquidation with Limit Orders, *SIAM J. Financ. Math.* 3(1), 740–764.

GUÉANT, O., C.-A. LEHALLE, and J. FERNANDEZ-TAPIA (2013): Dealing with the Inventory Risk: A Solution to the Market Making Problem, *Math. Financ. Econ*. 7(4), 477–507.

GUILBAUD, F., and H. PHAM (2013): Optimal High Frequency Trading with Limit and Market Orders, *Quant. Finance* 13(1), 79–94.

GUILBAUD, F., and H. PHAM (2015): Optimal High Frequency Trading in a Pro-Rata Microstructure with Predictive Information. *Math. Finance* 25(3), 545–575.

HE, H., and H. MAMAYSKY (2005): Dynamic Trading Policies with Price Impact, *J. Econ. Dyn. Control* 29(5), 891–930.

HO, T., and H. R. STOLL (1981): Optimal Dealer Pricing under Transactions and Return Uncertainty, *J. Financ. Econ.* 9(1), 47–73.

HUBERMAN, G., and W. STANZL (2005): Optimal Liquidity Trading, *Rev. Finance* 9(2), 165–200.

HUITEMA, R. (2012): Optimal Portfolio Execution Using Market and Limit Orders. Preprint.

KHARROUBI, I., and H. PHAM (2010): Optimal Portfolio Liquidation with Execution Cost and Risk. *SIAM J. Financ. Math.* 1, 897–931.

KRATZ, P., and T. SCHÖNEBORN (2014): Optimal Liquidation in Dark Pools, *Quant. Finance* 14(9), 1519–1539.

LARUELLE, S., C.-A. LEHALLE, and G. PAGÈS (2011a): Optimal Split of Orders across Liquidity Pools: A Stochastic Algorithm Approach, *SIAM J. Financ. Math.* 2(1), 1042–1076.

LARUELLE, S., C.-A. LEHALLE, and G. PAGÈS (2011b): Optimal Posting Distance of Limit Orders: A Stochastic Algorithm Approach. Preprint arXiv:1112.2397.

LORENZ, J., and R. ALMGREN (2011): Mean-Variance Optimal Adaptive Execution, *Appl. Math. Finance*, 18(5), 395–422.

OBIZHAEVA, A., and J. WANG (2005): Optimal Trading Strategy and Supply/Demand Dynamics. Technical report, National Bureau of Economic Research.

PREDOIU, S., G. SHAIKHET, and S. SHREVE (2011): Optimal Execution in a General One-Sided Limit-Order Book. *SIAM J. Financ. Math.* 2(1), 183–212.

SCHIED, A., and T. SCHÖNEBORN (2007): Optimal Portfolio Liquidation for CARA Investors. Available at SSRN.

SCHIED, A., and T. SCHÖNEBORN (2009): Risk Aversion and the Dynamics of Optimal Liquidation Strategies in Illiquid Markets, *Finance Stoch.* 13(2), 181–204.

SCHIED, A., T. SCHÖNEBORN, and M. TEHRANCHI (2010): Optimal Basket Liquidation for Cara Investors Is Deterministic, *Appl. Math. Finance* 17(6), 471–489.

MATHEMATICAL FINANCE
An International Journal of Mathematics, Statistics and Financial Economics

Mathematical Finance, Vol. 25, No. 3 (July 2015), 496–544

PORTFOLIO LIQUIDATION IN DARK POOLS IN CONTINUOUS TIME

PETER KRATZ

Aix-Marseille Université

TORSTEN SCHÖNEBORN

Deutsche Bank AG

We consider an illiquid financial market where a risk averse investor has to liquidate a portfolio within a finite time horizon $[0, T]$ and can trade continuously at a traditional exchange (the "primary venue") and in a dark pool. At the primary venue, trading yields a linear price impact. In the dark pool, no price impact costs arise but order execution is uncertain, modeled by a multidimensional Poisson process. We characterize the costs of trading by a linear-quadratic functional which incorporates both the price impact costs of trading at the primary exchange and the market risk of the position. The solution of the cost minimization problem is characterized by a matrix differential equation with singular boundary condition; by means of stochastic control theory, we provide a verification argument. If a single-asset position is to be liquidated, the investor slowly trades out of her position at the primary venue, with the remainder being placed in the dark pool at any point in time. For multi-asset liquidations this is generally not the case; for example, it can be optimal to oversize orders in the dark pool in order to turn a poorly balanced portfolio into a portfolio bearing less risk.

KEY WORDS: stochastic control, optimal liquidation, dark pools, singular boundary condition, illiquid markets, market microstructure.

1. INTRODUCTION

In the last years, equity trading has been transformed by the advent of so-called dark pools. These alternative trading venues differ significantly from classical exchanges and have gained considerable market share, especially in the United States. Dark pools vary in a number of properties such as crossing procedure, ownership, and accessibility (see Mittal 2008 and Degryse et al. 2009b for further details and a typology of dark pools). However, they generally share the following two stylized facts. First, the liquidity available in dark pools is not quoted, hence making trade execution uncertain and unpredictable. Second, dark pools do not determine prices. Instead, they monitor the prices determined by the classical exchanges and settle trades in the dark pool only if possible at these prices. Thus, trades in the dark pool have no or less price impact.[1]

We wish to thank Ulrich Horst, Werner Kratz, Nicholas Westray, and the anonymous referees for useful suggestions and comments. We are also grateful to seminar participants at the University of Bonn, Humboldt University Berlin, and the Bachelier seminar Paris. This research was supported by Deutsche Bank through the Quantitative Products Laboratory.

Manuscript received January 2012; final revision received December 2012.

Address correspondence to Peter Kratz, Laboratoire d'Analyse, Topologie, Probabilités, Aix-Marseille Université, 13453 Marseille Cedex 13, France; e-mail: kratz@math.hu-berlin.de.

[1] For empirical evidence of lower transaction costs or price impact of dark pools compared to classical exchanges see, e.g., Conrad et al. (2003) and Fong et al. (2013).

DOI: 10.1111/mafi.12037

In this paper, we consider an investor who has to liquidate a portfolio within a finite time horizon $[0, T]$ and has access both to a classical exchange (also called "primary venue" or "primary exchange") and to a dark pool.[2] The investor hence faces the trade-off of execution costs of the orders at the primary venue, market risk of the portfolio position, and execution risk of the dark pool orders. For liquidation of a position in a single asset, it should intuitively be optimal to slow down trading at the primary exchange initially in the hope of finding liquidity in the dark pool; no execution costs have to be paid for dark pool orders and hence the full remainder of the position should be placed in the dark pool at any point in time. For multi-asset liquidation, the correlation between the assets comes into play. Depending on whether the portfolio is well (poorly) balanced, it should be optimal to place orders in the dark pool which are smaller (larger) than the remaining portfolio position. In the first case, the portfolio bears little market risk and therefore the trader will not risk to lose her balanced position by the full execution of her dark pool order for only one of the assets. In the second case, it might be optimal to oversize the dark pool orders for risk mitigation reasons.

To our best knowledge, the mathematical framework established here and in the discrete-time model of Kratz and Schöneborn (2014) (cf. the discussion later) is the first within which optimal trade execution is analyzed for a multi-asset portfolio using a classical exchange and a dark pool simultaneously. Within our framework, we are able to derive an analytical representation of the optimal trading strategy allowing a thorough analysis of its properties which confirms the intuitions above. When a position in a single asset is to be liquidated, the current asset position is at all times being offered in the dark pool, while it is liquidated in parallel at the primary exchange. The opportunity to trade in the dark pool leads to a slower liquidation at the primary exchange compared to a market without a dark pool. Traders hence need to fundamentally adjust their trade execution algorithms when a dark pool is introduced. It is not sufficient to use an algorithm that was optimal for execution at the exchange and to add a component that also places trades in the dark pool; instead, trading at the exchange needs to be adjusted during the entire trade execution time interval. While in a single-asset setting the entire asset position is placed in the dark pool, this is not true if a multi-asset portfolio is to be liquidated. If, for example, the portfolio is balanced and thus only exposed to little market risk, then a complete liquidation of the position in one of the assets is unfavorable and thus only a fraction of the entire portfolio should be placed in the dark pool. This highlights again that overly simple adjustments to existing trade execution algorithms are exposed to potential pitfalls. For dark pools, the reluctance of traders to place balanced portfolios in a dark pool is an incentive to offer balanced executions in order to attract more liquidity.

Our model for trading and price formation at the classical exchange is a linear (infinitely resilient) price impact model. Trade execution can be enforced by selling aggressively, which however results in quadratic execution costs due to a stronger price impact. We model order execution in the dark pool by a multidimensional Poisson process. Orders submitted to the dark pool are executed at the jump times of the respective components of the process. The split of orders between dark pool and exchange is thus driven by the trade-off between execution uncertainty and price impact costs.

[2]The overall liquidity traded in dark pools in the United States is strongly fragmented among approximately 40 different venues, see, e.g., Carrie (2008). Therefore, liquidity aggregation is a major issue. Ganchev et al. (2010) and Laruelle et al. (2011) establish learning algorithms to achieve optimal order split between dark pools. Instead of analyzing the simultaneous use of several dark pools, we consider such an "aggregated" dark pool in our model which we call "the dark pool."

The investor aims to maximize her expected proceeds minus a quadratic risk term reflecting the market risk of the position. For appropriate assumptions on the price dynamics, this yields a linear-quadratic cost functional. However, the liquidation constraint implies a singularity of the value function at the terminal time T, which renders the solution of the optimization problem difficult. We approximate the constraint by a sequence of modified unconstrained optimization problems with increasing finite end-costs. Via a quadratic ansatz, the corresponding Hamilton-Jacobi-Bellman (HJB) equation suggests that the value functions of the modified optimization problems are quadratic forms for matrix-valued functions which are the solutions of initial value problems for a specific matrix differential equation. Explicit solutions for these initial value problems are not known. We establish a matrix inequality which allows us to apply a known comparison result for Riccati matrix differential equations in order to obtain closed form upper and lower bounds for the solutions of the initial value problems. The bounds are important for several reasons. First, they enable us to prove a verification theorem for the modified unconstrained optimization problems with finite end-costs. Second, they imply the existence of the limit of the solutions of these optimization problems and thus yield a well-defined candidate for the value function of the original optimization problem with liquidation constraint. Third, the limit of the bounds transfers to bounds for this candidate value function and we deduce a verification theorem for the solution of the optimal liquidation problem.

Hence, the value function of the problem is a quadratic form for a matrix-valued function which is the "principal solution" of the matrix differential equation. The optimal strategy is linear in the portfolio position; it is determined by the value function and can be computed easily from a numerical solution of the "principal solution" of the matrix differential equation. This makes the model tractable for practicable applications and allows us to investigate the properties of the solution (cf. the discussion above).

Our model complements the model by Kratz and Schöneborn (2014), where (in particular) multi-asset liquidation in dark pools is studied in discrete time; for single-asset liquidation, the two models are connected via a convergence result. The mathematical analysis of the continuous-time model is substantially different from the discrete-time model. In discrete time, the optimal liquidation strategy and its costs are given by a backward recursion obtained by standard dynamic programming; this recursion is rather unhandy and it is difficult to deduce properties of the optimal strategy from it. In continuous time, the optimization problem cannot be solved by standard methods of stochastic control because of the singular boundary condition for the value function. Instead, these methods had to be modified and extended. We believe that the analysis of the matrix differential equation characterizing the optimal strategy and the value function as well as the resulting verification theorem for the optimization problem are mathematically interesting in their own right (cf. the discussion above). In addition, the differential equation is easier to analyze than the (discrete-time) backward recursion and we can deduce additional properties of the value function and the optimal strategy from it; in particular, we are able to show that for a two-asset portfolio, the optimal strategy is monotone in the correlation of the assets.

In addition, our paper is connected to several lines of research. First, it builds on research on optimal trade execution strategies for a single trader in models with exogenously given liquidity effects. Several such models have been proposed for classical trading venues, for example, Bertsimas and Lo (1998), Almgren and Chriss (2001), Almgren (2003), Schied and Schöneborn (2009), Alfonsi et al. (2010), Schied et al. (2010), and Obizhaeva and Wang (2013). We follow Almgren and Chriss (2001) and assume a linear

temporary price impact model for the primary venue. This choice yields a tractable model which allows us to derive closed form solutions for the optimal liquidation strategy while capturing price impact effects. Furthermore, linear price impact models have become the basis of several theoretical studies, for example, Almgren and Lorenz (2007), Carlin et al. (2007), Schöneborn and Schied (2013), and Rogers and Singh (2010). The models above are limited to trading on one venue; they are concerned with the trade-off of execution costs versus market risk. By using a dark pool, the investor additionally faces order execution risk in our model. The trade-off of execution costs, market risk, and execution risk is also apparent in the models of Bayraktar and Ludkovski (2014), Guéant et al. (2012), Guéant and Lehalle (2015), Guilbaud and Pham (2015), and Huitema (2013) who consider optimal liquidation with limit orders.[3] In most of these models, the investor chooses a limit price for her orders;[4] the "probability" of execution is then dependent on this price and the left-over position can be liquidated at fixed costs at the exchange. In contrast to this, the dark pool orders in our model do not involve a limit price; they are executed at the price of the primary venue. Klöck et al. (2013) also work in a market model with exogenously specified liquidity characteristics for both a public exchange as well as a dark pool. Their focus is on the circumstances that can lead to price manipulation, while we focus on the quantitative and qualitative features of optimal trade execution strategies in markets without price manipulation. Therefore they investigate a more general class of price impact relationships, while our framework incorporates several aspects of market and investor behavior that shape optimal trade execution, such as risk aversion, multiple correlated assets, and a dynamic usage of the dark pool.

A second line of research connected to our paper focuses on the underlying mechanisms for illiquidity effects. Early examples include Kyle (1985), Glosten and Milgrom (1985), and Grossman and Miller (1988). In these models, price impact arises endogenously through the interplay of market participants. More recently, such models have been proposed to analyze the competition between classical trading venues and dark pools. Hendershott and Mendelson (2000) analyze the interaction of dealer markets and a crossing network[5] in a static one-period framework where each investor buys or sells a *single* share. Their findings include that trading in a crossing network is cheaper than in dealer markets since the trader saves half the spread, but trade execution is uncertain in the crossing network. In a similar setting, Dönges and Heinemann (2013) focus on game-theoretic refinements in order to remove the multiplicity of equilibria in Hendershott and Mendelson (2000). Degryse et al. (2009a) introduce a dynamic multiperiod framework and compare the effect of different levels of transparency of the dark pool. In order to focus on the optimal order execution of an individual trader, we exogenously specify the liquidity properties of the market. Hence, all the models cited above can shed more light on how the liquidity properties of the dark pool come about than our model can. On the other hand, by defining the model parameters exogenously, we can allow the investor to split her orders over time as well as between the two trading venues. While executing her order over time, she can dynamically react to the existence or absence of liquidity in the dark pool and adjust her trading strategy accordingly. Furthermore, we can take a multi-asset perspective and investigate how the composition of basket portfolios influences optimal trade execution strategies.

[3]In our model, all orders in the primary venue are market orders whose execution is guaranteed; hence, the execution risk only applies to the dark pool orders.

[4]Except for Guilbaud and Pham (2015), where all limit orders are half a tick away from the midprice.

[5]Crossing networks are specific types of dark pools that offer periodic matching of buy and sell orders.

Methods of stochastic control is a third line of research that this paper is connected to. A standard reference for stochastic control with jumps is the book by Øksendal and Sulem (2007). The liquidation constraint yields a singularity of the value function at the terminal time T, and thus, the resulting stochastic control problem requires non-standard considerations. For single-asset liquidation and finite end-costs (hence no singularity of the value function at time T), the solution of the linear-quadratic control problem is wellknown, see, e.g., Øksendal and Hu (2008). Multidimensional linear-quadratic control with jumps is treated in the book by Hanson (2007). However, our setting (even without the singularity of the value function) is not covered therein. The difficulty in our setting stems from the *combination* of the singularity of the value function and the fact that we consider multidimensional portfolios; thus the solution of the optimization problem involves the detailed analysis of a specific non-Riccati-type matrix differential equation, for which we establish existence results and upper and lower bounds of the solution by means of a novel matrix inequality. Several other texts are dealing with singular boundary constraints in liquidation problems. For instance, Schied et al. (2010) study optimal liquidation without dark pools for CARA investors. In this case the optimal control problem does not include jumps. They carry out a verification argument for a candidate value function obtained by considering only deterministic strategies. Naujokat and Westray 2011 and Höschler (2011) treat similar control problems with jumps. The focus of both texts is on trading with limit orders rather than with dark pools; they only treat single-asset trading and obtain the single-asset case of this paper as special cases of their respective settings. Kratz (2013) generalizes the single-asset case of this paper by allowing for adverse selection; this renders the liquidation problem nonlinear quadratic; he shows that in this case the value function is a "quasi-polynomial" of degree two instead of being quadratic; the solution for the multidimensional case is not known yet.

The remainder of this paper is structured as follows. We specify the model set-up and the optimization problem in Section 2 and obtain a candidate for the value function of the unconstrained optimization problem with finite end-costs via an initial value problem for a matrix differential equation. In Section 3, we state the main theoretical results of the paper: the solution of the initial value problem (Section 3.1), the solution of the unconstrained optimization problem (Section 3.2), and finally the solution of the constrained optimization problem (Section 3.3). We discuss the properties of the value function of the optimal strategy extensively for a single-asset position and a portfolio consisting of two assets in Section 4. The main results of Section 3 are proven in Section 5 and the results of Section 4 are proven in Section 6.

2. MODEL DESCRIPTION

For a fixed time interval $[0, T]$, we consider the stochastic basis $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F} = (\mathcal{F}_t)_{t\in[0,T]})$.[6] We investigate a market model where a risk averse trader with a personal risk aversion parameter $\alpha \geq 0$[7] has to liquidate a portfolio $x \in \mathbb{R}^n$ of n assets within a finite trading horizon $[0, T]$. The investor has the possibility to trade simultaneously at a traditional

[6]The filtration is generated by the involved random processes and is specified at the end of Section 2.2

[7]Although we use the term "risk aversion parameter," the investor is only risk averse to market but not to execution risk in our model as pointed out in Section 2.4 below (cf. also the discussion in the Appendix A).

exchange and in a dark pool, facing the trade-off of paying market impact costs in the traditional exchange against uncertain execution in the dark pool.

We specify the price dynamics at the primary exchange in Section 2.1 and the dynamics of order execution in the dark pool in Section 2.2. In Section 2.3, we define admissibility of trading strategies. In Section 2.4, we specify the trading objective and show that the resulting optimization problem is equivalent to a linear-quadratic optimization problem. Heuristic arguments suggest that the value function of the optimization problem is singular at terminal time T because of the liquidation constraint. Hence, we introduce a modified optimization problem where we drop the liquidation constraint and approximate it by finite end-costs for a portfolio not liquidated by time T as an intermediate step. In Section 2.5, we derive a candidate for the value function of the modified optimization problem via a quadratic ansatz and the corresponding HJB equation.

2.1. The Primary Exchange

In absence of transactions of the investor, the fundamental asset price at the primary exchange is given by an n-dimensional stochastic process $\tilde{P}$.

ASSUMPTION 2.1. *We assume that $\tilde{P}$ satisfies the following properties.*

(i) *$\tilde{P}$ is a square-integrable càdlàg martingale.*

(ii) *The covariance matrix of $\tilde{P}$ is constant in time, that is, for all $t \in [0, T]$, $i, j = 1, \ldots, n$,*

$$\mathrm{Cov}(\tilde{P}_i(t), \tilde{P}_j(t)) = t\Sigma_{i,j}$$

and $\Sigma = (\Sigma_{i,j})_{i,j=1,\ldots,n} \in \mathbb{R}^{n\times n}$.

The martingale property in Assumption 2.1 (i) makes the model tractable for closed form solutions. Furthermore, as the time horizon for portfolio liquidation is usually short, that is, several hours or a few days, it is a sufficient approximation of reality for our purposes.[8]

Once the trader becomes active on the primary exchange, she influences the market price P. We assume that the trader at the primary exchange can only execute trades continuously, that is, that her trading activity at the primary exchange is absolutely continuous and can hence be described by her trading intensity $\xi(t)$ with $t \in [0, T)$.[9]

[8]We do not model the period in which the investor has held the assets before time zero. This period may be long, whereas the liquidation time period is usually short. Any risk premium influences asset prices over any time horizon, however its magnitude over short time horizons is usually outweighed by the magnitude of market risk and price impact costs (see, e.g., Almgren and Chriss 2001). Furthermore, models with non-zero drift can result in profits from trading even if no initial position is to be liquidated, thus making it harder to differentiate between effects of portfolio liquidation and effects of optimal investment. We want to remark here that also the optimal dark pool orders are sensitive to the martingale assumption (cf., e.g., Footnote 16).

[9]In the single-asset setting without dark pools of Obizhaeva and Wang (2013), impulse trades are allowed in addition to absolute continuous trading. In order to include impulse trades in our model, it would first be necessary to define plausible costs for such trades; in Obizhaeva and Wang (2013) these costs arise naturally from their model of the order book. Unfortunately, this renders the model intractable for multi-asset trading using dark pools. Furthermore, the introduction of impulse trades would render the model less applicable for practitioners. We want to remark that the results of our paper are sensitive to the exclusion of impulse trading. If, for example, a discrete trade at time T is allowed (with quadratic costs), the investor will liquidate

Trading in the traditional exchange generates price impact, which we assume to be temporary and linear in the trading rate $\xi(t)$. Given a strategy $(\xi(t))_{t\in[0,T]}$, the transaction price at time $t \in [0, T]$ is given by

$$P(t) = \tilde{P}(t) - \Lambda\xi(t),$$

where $\Lambda \in \mathbb{R}^{n\times n}$ is a positive definite matrix constant in time. Execution of the trades at the primary exchange is *certain*; we hence consider only market orders and no limit orders.

By assuming linear price impact for the primary venue, we follow Almgren and Chriss (2001). This choice yields a tractable model which allows us to derive closed form solutions for the optimal liquidation strategy while capturing price impact effects. Furthermore, linear price impact models have become the basis of several theoretical studies, for example, Almgren and Lorenz (2007), Carlin et al. (2007), Schöneborn and Schied (2013), and Rogers and Singh (2010). In contrast to Almgren and Chriss (2001), the price impact in our model is purely temporary; Klöck et al. (2013) analyze the influence of *permanent* price impact on the existence of market manipulation strategies. Such an analysis is not in the scope of our paper; hence we allow only for temporary price impact.[10]

2.2. Order Transaction in the Dark Pool

In addition to the primary exchange, the trader can also use a dark pool. Dark pools often have rather complex order allocation mechanisms; most of them use some sort of a pro-rata or time-priority rule for matching orders from opposite sides of the market. Here, we consider a dark pool with a time-priority matching rule: the investor's order η enters a queue and is matched with liquidity from the opposite side of the market (if there is any) once it has reached the front of the queue.

We allow for continuous updating of the orders $\eta(t)$ in the dark pool at any time $t \in [0, T]$. Orders for the ith asset in the dark pool are executed fully at the jump times of the ith component of an n-dimensional Poisson process

$$\pi = (\pi_1, \dots, \pi_n) \text{ with intensities } \theta_1, \dots, \theta_n \geq 0, \text{ respectively.}$$

Else, the orders are not executed at all. This mechanism implies two main simplifications of reality which allow a thorough mathematical analysis of the model. First, we exclude partial execution; the probability of execution does not depend on the size of the order. Second, we assume independence of the increments of the dark pool liquidity. On the other hand, the resulting model captures the stylized facts outlined above and in the introduction; we believe that it constitutes a sufficiently good approximation of reality for our purposes.

at least a small position at time T (unless she finds liquidity in the dark pool before; cf. Section 4.1), that is, $X^u(T-) \neq 0$ in contradiction to Assumption 2.3 (v) below.

[10] As the price is not influenced permanently by the investor, the term *price impact* might be misleading. Alternatively, we could rename the quadratic costs caused by trading in the primary venue as the *execution costs* of the investor; indeed, the costs of trading can be interpreted to include such different effects as price impact and transaction costs. However, the term *execution costs* does not capture the possible impact of trading in one asset on the price of another asset which we include by allowing the off-diagonal elements of Λ to be non-zero (cf. the discussion of cross price impact in Section 4.2.2).

ASSUMPTION 2.2. *We assume that π satisfies the following conditions.*

(i) $\pi_1, \dots, \pi_n$ *are independent.*
(ii) π *and* $\tilde{P}$ *are independent.*

By Assumption 2.2 (ii) we exclude correlations between dark pool liquidity and the fundamental asset price in the market. In particular, this rules out adverse selection.[11] For single-asset liquidation, adverse selection was incorporated by Kratz (2013). This renders the liquidation problem nonlinear quadratic; he shows that in this case the value function is a "quasi-polynomial" of degree two instead of being quadratic. A solution for the multidimensional case is not known yet.

While the dark pool has no impact on prices at the primary venue, it is less clear to which extent the price impact of the primary venue $\Lambda\xi(t)$ is reflected in the trade price of the dark pool. If, for example, the price impact is realized predominantly in the form of a widening spread, then the impact on dark pools that monitor the mid-quote can be much smaller than $\Lambda\xi(t)$. We will make the simplifying assumption that trades in the dark pool are not influenced by the price impact at all, that is, that they are executed at the fundamental price $\tilde{P}$. If alternatively the transaction price in the dark pool is the price P at the primary exchange including the trader's price impact, market manipulation strategies can become profitable unless the parameters are chosen with great care, as has been shown by Kratz and Schöneborn (2014) for the discrete-time case. For a detailed discussion see also Klöck et al. (2013) who analyze the circumstances that can lead to price manipulation in dark pools.

We are now ready to specify the filtration $(\mathcal{F}_t)_t$ as the completion of $(\sigma(\tilde{P}(s), \pi(r) | 0 \leq s \leq t, 0 \leq r < t))_t$.

2.3. Admissible Trading Strategies

Let $t \in [0, T)$ be a given point in time and $x \in \mathbb{R}^n$ be the portfolio position of the trader at time t. The trader has the possibility to trade asset k in the traditional exchange with trading intensity $\xi_k(s)$ at time $s \in [t, T)$ and to place orders $\eta_k(s)$ in the dark pool at time s.

We call a $2n$-dimensional stochastic process

$$(u(s))_{s\in[t,T)} = (\xi(s), \eta(s))_{s\in[t,T)}$$

a *trading strategy* if ξ is progressively measurable and η is predictable. Given a trading strategy u, the portfolio position at time $s \in [t, T)$ is given by the following controlled stochastic differential equation:

$$\begin{aligned} dX^u(s) &= -\xi(s)ds - \eta(s)d\pi(s) \\ X^u(t) &= x \end{aligned} \tag{2.1}$$

such that the left-hand side in (2.1) is well defined.

For technical reasons we require all trading strategies to fulfill the following conditions.

[11] Here, adverse selection refers to the phenomenon that liquidity seeking traders find that their trades in the dark pool are usually executed just before a favorable price move, i.e., exactly when they *do not* want them to be executed since they miss out on the price improvement.

DEFINITION 2.3. Let $t \in [0, T)$ and $x \in \mathbb{R}^n$ be fixed. Let $u = (u(s))_{s\in[t,T)} = ((\xi(s), \eta(s)))_{s\in[t,T)}$ be a trading strategy, that is, ξ is progressively measurable and η is predictable.

(a) We call u an *admissible trading strategy* if it fulfills the following conditions.

(i) The Stochastic Differential Equation (2.1) possesses a unique solution on $[t, T)$.

(ii) $$\mathbb{E}\left[\int_t^T \|\xi(s)\|_2^4 ds\right] < +\infty, \quad \mathbb{E}\left[\int_t^T \|\eta(s)\|_2^8 ds\right] < +\infty.$$

(iii) If $\theta_i = 0$, then $\eta_i(s) = 0$ for all $s \in [t, T)$.

We denote the set of admissible trading strategies by $\tilde{\mathbb{A}}(t)$.

(b) We call $u \in \tilde{\mathbb{A}}(t)$ an *admissible liquidation strategy* or just *liquidation strategy* if additionally

(iv) $\lim_{s\to T-} X^u(s) = 0$ a.s.

and denote the set of admissible liquidation strategies by $\mathbb{A}(t, x)$.

Let us shortly comment on Definition 2.3. Condition (ii) is required for the moment bound in Lemma 5.5 which is in turn needed for the verification later. Condition (iii) is needed in order to ensure uniqueness of optimal trading strategies: if $\theta_i = 0$, no additional gain can be achieved by nonzero dark pool orders. If the portfolio is liquidated at constant speed at the primary exchange only, that is, $\xi(s) \equiv \frac{x}{T-t}$, $\eta(s) = 0$ for $s \in [t, T)$, Definition 2.3 (in particular the liquidation constraint (iv)) is satisfied and hence $\mathbb{A}(t, x) \neq \emptyset$.

REMARK 2.4. We expect that the stochastic control problems we solve are such that the optimal control is of *Markovian form* (see, e.g., the book by Øksendal 2007, theorem 11.2.3):

$$u(s) = (\xi(s), \eta(s)) = (\xi(s, X(s)), \eta(s, X(s-))$$

for deterministic functions $\xi, \eta : [t, T) \times \mathbb{R}^n \to \mathbb{R}^n$. The deterministic initial value problem

$$X' = -\xi(\cdot, X), \quad X(t) = x \tag{2.2}$$

possesses a unique solution on $[t, T)$ if $\|\xi(s, y)\| \leq f(s)\|y\| + g(s)$ on $[t, T) \times \mathbb{R}^n$ for $f, g \in C([t, T))$ and $\xi(s, \cdot)$ is locally Lipschitz (e.g., C^1).[12] Let $\xi : [t, T) \times \mathbb{R}^n \to \mathbb{R}^n$ fulfill these conditions and let $\eta : [t, T) \times \mathbb{R}^n \to \mathbb{R}^n$. We can pathwise construct the solution of the Stochastic Differential Equation (2.1) inductively by interlacing the jumps (see, e.g., Applebaum 2004, example 1.3.13): as the n Poisson processes are independent, they jump at distinct times almost surely. Let $(\tau_i)_{i\geq 1}$ be the jump times of π such that $t =: \tau_0 < \tau_1 < \ldots$ almost surely, and let X be the solution of the Initial Value Problem (2.2) on $[\tau_i, \tau_{i+1} \wedge T)$ with initial value $x = X(\tau_i)$ for $i \in \mathbb{N}$ such that $\tau_i \leq T$. For $\tau_{i+1} \leq T$ and $\Delta\pi_k(\tau_{i+1}) > 0$, we set

$$X(\tau_{i+1}) := X(\tau_{i+1}-) - \eta_k(\tau_{i+1}, X(\tau_{i+1}-))e_k,$$

where e_k is the kth unit vector.

[12] This follows, for example, by Peano's existence theorem and Gronwall's inequality.

2.4. Cost Functional

The proceeds of selling the portfolio $x \in \mathbb{R}^n$ during $[t, T]$ according to the strategy $(u(s))_s = (\xi(s), \eta(s))_s \in \mathbb{A}(t, x)$ are given by

$$\phi(t, x, u) := \int_t^T \xi(s)^\top (\tilde{P}(s) - \Lambda\xi(s))ds + \int_t^T \eta(s)^\top \tilde{P}(s)d\pi(s).$$

The first term in the above equation represents the proceeds of selling at the primary exchange at a price of $P(s) = \tilde{P}(s) - \Lambda\xi(s)$, while the second term accounts for the proceeds of selling in the dark pool at the unaffected price $\tilde{P}(s)$. Applying integration by parts and using $X^u(s) = x - \int_t^s \xi(r)dr - \int_t^s \eta(r)d\pi(r)$, Assumption 2.1 (i), the fact that $\tilde{P}$ and π are independent (Assumption 2.2 (ii)) and the liquidation constraint (Definition 2.3 (iv)), we obtain

$$\phi(t, x, u) = -\int_t^T \xi(s)^\top \Lambda\xi(s)ds + x^\top \tilde{P}(t) + \int_t^T X^u(s-)d\tilde{P}(s).$$

This yields (cf. Assumption 2.1)

$$\mathbb{E}[\phi(t, x, u)] = x^\top \tilde{P}(t) - \mathbb{E}\left[\int_t^T \xi(s)^\top \Lambda\xi(s)ds\right].$$

Instead of maximizing expected proceeds, we can thus equivalently minimize expected price impact costs. We assume that the trader is not only interested in expected liquidation proceeds, but in addition also wants to minimize risk during liquidation. We incorporate both aspects in the following cost functional:[13]

$$\begin{aligned} J(t, x, u) &:= x^\top \tilde{P}(t) - \mathbb{E}[\phi(t, x, u)] + \mathbb{E}\left[\alpha \int_t^T X^u(s)^\top \Sigma X^u(s)ds\right] \\ &= \mathbb{E}\left[\int_t^T f(\xi(s), X^u(s))ds\right], \end{aligned}$$

where $f : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ is given by $f(\xi, x) := \xi^\top \Lambda\xi + \alpha x^\top \Sigma x$. The first two terms in the cost functional capture the expected liquidation shortfall, while the last term is an additive penalty function $\alpha \int_t^T X^u(s)^\top \Sigma X^u(s)\, ds$, which reflects the market risk of the portfolio (recall that $\alpha \geq 0$ is the risk aversion parameter of the investor); it penalizes slow liquidation and poorly balanced portfolios. It does not incorporate execution risk; hence the investor is only risk averse with respect to market risk but not with respect to execution risk. This simplification allows us to obtain analytical solutions for the cost minimization problem below. We investigate the impact of this model choice in Appendix A. As outlined there, we believe that for realistic parameters, market risk outweighs execution risk; hence we expect that the inclusion of risk aversion with respect to execution risk would not change the optimal strategy substantially. For deterministic liquidation strategies without dark pools, the risk term reflects

[13] Both mean and variance of execution costs are often used as measures of execution quality. The cost functional J is inspired by such mean variance measures. An alternative approach is the maximization of expected utility: $\tilde{J}(t, x, u) = E[U(\phi(t, x, u))]$ for some utility function U. In this alternative set-up analytical solutions are unfortunately not directly available through the methods presented in this paper and are hence left for future research.

the variance of the liquidation costs (see Almgren and Chriss 2001). In this case, minimizing a mean-variance functional of the liquidation costs over all deterministic strategies is equivalent to maximizing the expected utility of the proceeds of an investor with CARA preferences over all strategies (see Schied et al. 2010).

We assume that the trader aims to minimize the cost functional and hence considers the following optimization problem:[14]

$$\text{(OPT)} \qquad v(t,x) := \inf_{u \in \mathbb{A}(t,x)} J(t,x,u).$$

Note that the optimization problem is well defined and the value function satisfies $v(t,x) < +\infty$ for $t < T$ (consider, e.g., constant liquidation exclusively in the primary exchange). Because of the liquidation constraint (cf. Definition 2.3 (iv)), we expect the value function to fulfill

$$\lim_{s \to T-} v(s,x) = \begin{cases} 0 & \text{if } x = 0 \\ +\infty & \text{else,} \end{cases}$$

that is, v has a singularity at the terminal time T. Because of this singularity, nonstandard considerations are necessary for solving the Optimization Problem (OPT) via a verification argument using the HJB equation.

As an intermediate step, we hence weaken the liquidation constraint by allowing for all strategies $u \in \tilde{\mathbb{A}}(t)$ and by penalizing nonliquidation by finite end-costs. More precisely, for $l > 0$ and $u = (\xi, \eta) \in \tilde{\mathbb{A}}(t)$, we define the following cost functional

$$\tilde{J}(l,t,x,u) := \mathbb{E}\left[\int_t^T f(\xi(s), X^u(s))ds + l \cdot X^u(T)^\top I X^u(T)\right].$$

The resulting optimization problem is

$$\widetilde{\text{(OPT)}} \qquad \tilde{v}(l,t,x) := \inf_{u \in \tilde{\mathbb{A}}(t)} \tilde{J}(l,t,x,u).$$

The Optimization Problem $\widetilde{\text{(OPT)}}$ mainly serves as an approximation of the Optimization Problem (OPT). However, it is also of interest itself: the penalization term $l \cdot X^u(T)^\top I X^u(T)$ can be considered as the liquidation cost of the left-over position. Note that the identity matrix in the term can be replaced by any positive definite matrix reflecting this interpretation (e.g., Λ) without changing any of the proofs significantly.

In the following, we solve the unconstrained Optimization Problem $\widetilde{\text{(OPT)}}$ first (Section 3.2). Then, we show that the solution of the Optimization Problem $\widetilde{\text{(OPT)}}$ converges to the solution of the original constrained Optimization Problem (OPT) as $l \to \infty$ (Section 3.3).

2.5. Hamilton-Jacobi-Bellman Equation

In this section, we derive a candidate for the value function of the Optimization Problem $\widetilde{\text{(OPT)}}$. Heuristic considerations suggest that it should satisfy the following HJB

[14] An alternative interesting set-up is to consider trade execution under minimum proceeds constraints. Such problems have recently been addressed using theory about stochastic target problems (see, e.g., Bouchard et al. 2009 and Bouchard and Dang 2013).

equation (see, e.g., Øksendal and Sulem 2007):

($\widetilde{\text{HJB}}$)

$$\frac{\partial w}{\partial t}(t,x) = \sup_{u=(\xi,\eta)\in\mathbb{R}^n\times\mathbb{R}^n}\left[\sum_{i=1}^{n}\theta_i\big(w(t,x)-w(t,x-\eta^\top e_i)\big)+\nabla_x w(t,x)\xi - f(\xi,x)\right]$$

$$w(T,x) = lx^\top x. \tag{2.3}$$

The linear-quadratic form of the cost functional suggests that the value function is quadratic. Assuming that the above guesses are correct, the following proposition provides candidates both for the value function and for the optimal strategy.

PROPOSITION 2.5. *Let $l > 0$ and assume that the initial value problem for a matrix differential equation*

$$C' = C^\top\Lambda^{-1}C + C^\top\tilde{C}C - \alpha\Sigma, \quad C(T) = lI,$$

where $\tilde{C}(l,t) := \operatorname{diag}(\frac{\theta_i}{c_{i,i}(l,t)})$ possesses a positive definite solution $C(l,t) = (c_{i,j}(l,t))_{i,j=1,\dots,n}$ on $[0,T]$. Then

$$w(l,t,x) := x^\top C(l,t)x$$

satisfies the HJB equation ($\widetilde{HJB}$) with maximizer $u^* = (\xi^*,\eta^*)$ for

$$\xi^* := \xi^*(l,t,x) := \Lambda^{-1}C(l,t)x, \quad \eta^* := \eta^*(l,t,x) := \bar{C}(l,t)C(l,t)x,$$

where $\bar{C}(l,t) := \operatorname{diag}(\frac{1}{c_{i,i}(l,t)})$.

If there exist $i_1,\dots,i_k \in \{1,\dots,n\}$ such that $\theta_{i_j} = 0$ $(j = 1,\dots,k)$, then η_{i_j} can be chosen arbitrarily. Up to arbitrary choices of η_{i_j}, the maximizer is unique.

Proof. The assertion follows directly from plugging the quadratic ansatz $w(l,t,x) = x^\top C(l,t)x$ into the HJB equation ($\widetilde{\text{HJB}}$); the resulting function can be maximized by standard calculus. □

3. MAIN RESULTS

Proposition 2.5 suggests that the solution of the Optimization Problem ($\widetilde{\text{OPT}}$) solves the initial value problem for the matrix differential equation

$$\begin{aligned} C' &= C^\top\Lambda^{-1}C + C^\top\tilde{C}C - \alpha\Sigma \\ C(T) &= lI, \end{aligned} \tag{3.1}$$

where

$$\tilde{C} := \operatorname{diag}\left(\frac{\theta_i}{c_{i,i}}\right). \tag{3.2}$$

In the remainder of the section, we state the main results of the paper. In Section 3.1, we show that (3.1) admits a unique solution C on $[0,T]$ and establish appropriate upper and lower bounds for C. Subsequently, we deduce the solution of the Optimization Problem

($\widetilde{\text{OPT}}$) in Section 3.2 and as a limit of this (as $l \to \infty$) the solution of the Optimization Problem (OPT) in Section 3.3. Proofs of these results are presented in Section 5.

Before we proceed, we introduce the following notations.

NOTATION 3.1.

(i) For symmetric matrices A and B we say $A > B$ ($A \geq B$) if $A - B$ is positive (nonnegative) definite.

(ii) We denote the smallest and the largest eigenvalues of a real-symmetric matrix A by $a_{\min}$ and $a_{\max}$, respectively.

(iii) We define the following nonnegative definite matrix: $D := \sqrt{\Lambda^{-1}}\Sigma\sqrt{\Lambda^{-1}}$.

3.1. Solution of the Initial Value Problem (3.1)

It is not immediately clear that the Initial Value Problem (3.1) possesses a positive definite solution on the whole interval $[0, T]$ for $n \geq 2$. For $n = 1$, it reduces to $C' = \frac{C^2}{\Lambda} + \theta_1 C - \alpha\Sigma$, $C(T) = l$. This is an initial value problem for a scalar Riccati differential equation with constant coefficients, whose unique solution is explicitly known and exists on the whole interval $[0, T]$ (cf. Section 4.1). For $n \geq 2$, the following theorem establishes the existence and uniqueness of the solution of (3.1).

THEOREM 3.2. *Let $\theta_i \geq 0$ for $i = 1, \ldots, n$, $\theta = \sum_{i=1}^n \theta_i$ and $l > l_0$, where*

$$l_0 := \max\left\{\lambda_{\max}\left(\sqrt{\frac{\theta^2}{4} + \alpha d_{\min}} - \frac{\theta}{2}\right), \lambda_{\min}\left(\sqrt{\alpha d_{\max}}\right)\right\} \tag{3.3}$$

(cf. Notation 3.1). Then the Initial Value Problem (3.1) possesses a unique solution $C(l, \cdot)$ on $(-\infty, T]$. The solution is symmetric for all $t \in (-\infty, T]$ and

$$0 < P(l, t) \leq \sqrt{\Lambda^{-1}}C(l, t)\sqrt{\Lambda^{-1}} \leq Q(l, t),$$

where P and Q are the solutions of the initial value problems

$$\begin{aligned} P' &= P^2 + \theta P - \alpha d_{\min} I, \quad P(T) = \tfrac{l}{\lambda_{\max}} I, \quad \text{and} \quad Q' \\ &= Q^2 - \alpha d_{\max} I, \quad Q(T) = \tfrac{l}{\lambda_{\min}} I, \end{aligned} \tag{3.4}$$

respectively.

REMARK 3.3.

(i) The solutions of the initial value problems for Riccati matrix differential equations in (3.4) exist on the whole interval $[0, T]$ and can be computed in closed form (cf. equations (5.6)–(5.9)). For technical reasons, we prefer to establish bounds for $\sqrt{\Lambda^{-1}}C\sqrt{\Lambda^{-1}}$ instead of bounds for C. P and Q are constructed in terms of multiples of the identity matrix and hence commute with all matrices. Therefore, they transfer to bounds of C directly by multiplying them with Λ.

(ii) The bounds of C are an essential component for the proof of Theorem 3.2. Additionally, they are required for all key steps of the solution of the Optimization Problem ($\widetilde{\text{OPT}}$) (Proposition 3.4 and Theorem 3.5) and of the solution of the Optimization Problem (OPT) (Theorem 3.6, Theorem 3.8, and Theorem 3.9).

3.2. Solution of the Optimization Problem ($\widetilde{\text{OPT}}$)

Combining Proposition 2.5 and Theorem 3.2, we obtain well-defined candidates both for the value function $(x^\top C(l,t)x)$ and for the optimal strategy $u^* = (\xi^*, \eta^*)$ of the Optimization Problem ($\widetilde{\text{OPT}}$). The latter is given by

$$\xi^*(l) := \xi^*(l,t,x) := \Lambda^{-1} C(l,t)x, \tag{3.5}$$

$$\eta^*(l) := \eta^*(l,t,x) := \tilde{I}\bar{C}(l,t)C(l,t)x, \tag{3.6}$$

where $\tilde{I} = (e_{i,j})_{i,j=1,\dots,n}$ is the diagonal matrix with

$$e_{i,i} = \begin{cases} 1 & \text{if } \theta_i > 0 \\ 0 & \text{else} \end{cases} \qquad \text{and} \qquad \bar{C}(l,t) := \operatorname{diag}\left(\frac{1}{c_{i,i}(l,t)}\right).$$

The following proposition confirms that u^* is admissible.

PROPOSITION 3.4. *Let $l > l_0$ for l_0 as in equation (3.3) and $(t,x) \in [0,T) \times \mathbb{R}^n$. Then $u^*(l) = (\xi^*(l), \eta^*(l)) \in \tilde{\mathbb{A}}(t)$, where $\xi^*(l)$ and $\eta^*(l)$ are as in equations (3.5) and (3.6), respectively.*

Finally, we obtain the solution of the Optimization Problem ($\widetilde{\text{OPT}}$).

THEOREM 3.5. *Let $l \geq l_0$ for l_0 as in equation (3.3) and let $C(l,t)$ be the unique solution of the Initial Value Problem (3.1). Then the value function of the Optimization Problem ($\widetilde{\text{OPT}}$) is given by*

$$\tilde{v}(l,t,x) = x^\top C(l,t)x$$

and the $\mathbb{P} \otimes \lambda$ - almost surely unique optimal strategy is given by $u^*(l)$ as in equations (3.5) and (3.6).

3.3. Solution of the Optimization Problem (OPT)

Intuitively, infinite end-costs should force the controlled process $X^u(s)$ to approach zero as $s \to T-$; furthermore the solution of the Optimization Problem (OPT) should be the limit of the solution of the Optimization Problem ($\widetilde{\text{OPT}}$) as $l \to \infty$. The following theorem confirms (in particular) that this limit is well defined.

THEOREM 3.6. *Let $t \in [0,T)$.*

(i) The element-wise limit of the value function matrix

$$C(t) := \lim_{l \to \infty} C(l,t)$$

exists on $[0,T)$, and $C(l,\cdot)$ converges compactly to C on $[0,T)$. Furthermore, $\lim_{l\to\infty} c_{\min}(l,T) = +\infty$.

(ii) C solves the matrix differential equation

$$C' = C^\top \Lambda^{-1} C + C^\top \tilde{C} C - \alpha\Sigma \tag{3.7}$$

on $[0, T)$ *with boundary condition* $\lim_{s\to T-} c_{\min}(s) = +\infty$. *Moreover, the following inequalities hold:*

$$0 < P(t) \leq \sqrt{\Lambda^{-1}}C(t)\sqrt{\Lambda^{-1}} \leq Q(t), \tag{3.8}$$

where $P(t) := \lim_{l\to\infty} P(l, t)$ *and* $Q(t) := \lim_{l\to\infty} Q(l, t)$.

REMARK 3.7. For Riccati matrix differential equations, there exists a unique solution F with $\lim_{s\to T-} f_{\min}(s) = +\infty$. This solution is called the principal solution (see, e.g., Coppel 1971). In this spirit, C is the principal solution of the Matrix Differential Equation (3.7). Note however that it is not entirely clear that C is the *only* solution of (3.7) satisfying $\lim_{s\to T-} c_{\min}(s) = +\infty$ since (3.7) is *not* a Riccati matrix differential equation.

By Theorem 3.6, we also obtain the existence of the limits of the optimal strategy:

$$\begin{aligned} \xi^* :&= \xi^*(t, x) := \lim_{l\to\infty} \xi^*(l, t, x) = \Lambda^{-1}C(t)x, \\ \eta^* :&= \eta^*(t, x) := \lim_{l\to\infty} \eta^*(l, t, x) = \tilde{I}\bar{C}(t)C(t)x. \end{aligned} \tag{3.9}$$

It turns out that $u^* := (\xi^*, \eta^*)$ is an admissible liquidation strategy, in particular that it satisfies the liquidation constraint

$$\lim_{s\to T-} X^*(s) = 0 \quad \text{for} \quad X^*(s) := X^{u^*}(s).$$

THEOREM 3.8. *Let* $t \in [0, T)$, $x \in \mathbb{R}^n$ *and* $u^* = (\xi^*, \eta^*)$ *for* ξ^* *and* η^* *as in (3.9). Then* $u^* \in \mathbb{A}(t, x)$.

We are now ready to present the main result of this paper: the solution of the Optimization Problem (OPT).

THEOREM 3.9. *The value function of the Optimization Problem (OPT) is given by*

$$v(t, x) = x^\top C(t)x$$

for all $t \in [0, T)$, $x \in \mathbb{R}^n$ *and*

$$\lim_{s\to T-} v(s, x) = \begin{cases} 0 & \text{if } x = 0 \\ +\infty & \text{else.} \end{cases}$$

The $\mathbb{P} \otimes \lambda$ *- almost surely unique optimal strategy is given by* $u^* = (\xi^*, \eta^*)$ *as in (3.9) .*

4. PROPERTIES OF THE VALUE FUNCTION AND THE OPTIMAL STRATEGY

The characterization of the solution of the Optimization Problem (OPT) enables us to analyze the properties of the optimal strategy and the value function in detail.[15] For single-asset liquidation ($n = 1$) the Initial Value Problem (3.1) can be solved in closed form as the differential equation is a scalar Riccati equation with constant coefficients.

[15]We limit the analysis to the Optimization Problem (OPT). Most of the properties transfer directly to the Optimization Problem ($\widetilde{\text{OPT}}$) with the same or similar proofs.

This allows us to prove monotonicity properties of the value function and the optimal strategy in Section 4.1. In Section 4.2, we discuss the multi-asset case by analyzing a portfolio of two assets. Although a closed form solution of the Initial Value Problem (3.1) is not known in general for $n \geq 2$, it is possible to derive analytical results about the dependence of the value function and the optimal strategy on the model parameters, in particular on the correlation of the assets. We illustrate that overly simple adjustments of existing trading algorithms for optimal liquidation without dark pools can have undesirable properties. The proofs of the results of this section are presented in Section 6.

4.1. Single-Asset Liquidation

We let $n = 1$ and set $\theta = \theta_1$. The solution of the Initial Value Problem (3.1) is given by

$$C(l, t) = \frac{\Lambda\tilde{\theta}}{2} \coth\left(\frac{\tilde{\theta}}{2}(T-t) + \kappa(l)\right) - \frac{\Lambda\theta}{2},$$

where

$$\kappa(l) := \operatorname{arcoth}\left(\frac{\frac{2l}{\Lambda} + \theta}{\tilde{\theta}}\right), \quad \tilde{\theta} := \sqrt{\theta^2 + \frac{4\alpha\Sigma}{\Lambda}}$$

for $\theta > 0$ or $\alpha\Sigma > 0$ and $C(l, t) = \frac{\Lambda}{T - t + \frac{\Lambda}{l}}$ for $\theta = \alpha\Sigma = 0$. In order to highlight the dependence of the value function on the parameters θ, Λ, and $\alpha\Sigma$, we define, for example,

$$C(t;\theta) := C(t) = \lim_{l\to\infty} C(l, t) = \begin{cases} \dfrac{\Lambda\tilde{\theta}}{2} \coth\left(\dfrac{\tilde{\theta}}{2}(T-t)\right) - \dfrac{\Lambda\theta}{2} & \text{if } \theta > 0 \text{ or } \alpha\Sigma > 0 \\ \dfrac{\Lambda}{T-t} & \text{if } \theta = \alpha\Sigma = 0, \end{cases}$$

in particular

$$C(t;0) = \begin{cases} \sqrt{\alpha\Sigma\Lambda} \coth\left(\sqrt{\dfrac{\alpha\Sigma}{\Lambda}}(T-t)\right) & \text{if } \alpha\Sigma > 0 \\ \dfrac{\Lambda}{T-t} & \text{if } \alpha\Sigma = 0. \end{cases}$$

We will apply similar notations throughout Section 4 to make the dependence of other model components (such as the optimal strategy) on the respective parameters explicit whenever this clarifies the exposition. For $\theta = 0$, we obtain the special case of optimal liquidation without dark pool, see Almgren and Chriss (2001) for the discrete-time case and Schied et al. (2010) for the continuous-time version.

No transaction costs must be paid in the dark pool; intuitively, the investor should hence try to liquidate as much as possible in the dark pool. Indeed, we have

$$\eta^*(t, x;\theta) = \bar{C}(t;\theta)C(t;\theta)x = x$$

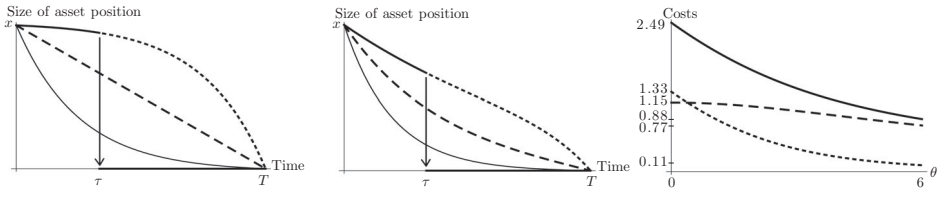

FIGURE 4.1. The left and the middle picture show optimal trading trajectories for a risk neutral, respectively, a risk averse trader. The thick solid lines denote the optimal trading trajectory with dark pool in the scenario where the dark pool order is executed at time τ. The dotted lines denote the scenario where the dark pool order is not executed during the entire trading horizon and the thin solid lines refer to the expected asset positions. The dashed lines denote optimal liquidation without dark pools. The right picture shows the dependence of the value function (solid line), the risk costs (dashed line), and the impact costs (dotted line) on the dark pool execution intensity θ. $x = 1$, $T = 1$, $\Lambda = 1$, $\theta = 4$ (left and middle picture), respectively, $\theta \in (0, 6)$ (right picture) and $\alpha = 6$, $\Sigma = 1$ (middle and right picture).

for $\theta > 0$ by Theorem 3.9, that is, it is optimal to always place the full remaining asset position in the dark pool;[16] note that the execution of the dark pool order immediately stops the trading activity by the linearity of the optimal strategy in the position.

In Section 4.1.1, we discuss the dependence of the optimal strategy and the value function on θ. Subsequently, we analyze the dependence on the price impact Λ and the risk parameter $\alpha\Sigma$ in Section 4.1.2.

4.1.1. Dependence on θ. We expect it to be optimal to slow down trading in the primary venue initially as the trader hopes to trade cheaper in the dark pool. This intuition is confirmed by Proposition 4.1 (ii) and (iii) below. In order to state the property rigorously, we first denote the optimal trading trajectory until execution in the dark pool by $\tilde{X}(\cdot;\theta)$, that is, $\tilde{X}$ is the solution of the linear initial value problem $X' = -\xi^*(\cdot, X;\theta)$, $X(0) = x$, where x is the initial asset position at time zero. Then for $t \in [0, T)$,

$$\tilde{X}(t;\theta) = x \exp\left(-\int_0^t \frac{C(s;\theta)}{\Lambda} ds\right) = \frac{\sinh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\exp\left(\frac{\theta}{2}t\right)}{\sinh\left(\frac{\tilde{\theta}}{2}T\right)} x. \tag{4.1}$$

We obtain the following monotonicity properties. For simplicity of exposition we assume $\alpha\Sigma > 0$. Similar results hold for the simpler case $\alpha\Sigma = 0$ (cf. also the left picture of Figure 4.1).

PROPOSITION 4.1.

(i) For $t \in [0, T)$, $C(t;\theta)$ is strictly decreasing in θ.
(ii) For $x > 0$ and $t \in (0, T]$, $\xi^*(t, x;\theta)$ is strictly decreasing in θ.

[16] We want to remark that this property is sensitive to the assumption that $\tilde{P}$ is a martingale (Assumption 2.1 (i)). If, for example, the investor holds a long position in the asset and $\tilde{P}$ has a positive drift, she should be reluctant to sell her entire position too early. Therefore, we expect that a drift changes this property. Similarly, if adverse selection is included into the model, the property $\eta^*(t, x) = x$ does not hold; this was shown by Kratz (2013).

(iii) *For $x > 0$ and $t \in (0, T)$, $\tilde{X}(t;\theta)$ is strictly increasing in θ.*
(iv) *For $x > 0$ and $t \in (0, T)$, the expected asset position if the optimal strategy is applied, $\mathbb{E}[X^*(t;\theta)]$, is strictly decreasing in θ.*
(v) *For $\alpha\Sigma > 0$, the risk costs $\alpha\Sigma \cdot \mathbb{E}[\int_0^T X^*(t;\theta)^2 dt]$ are strictly decreasing in θ.*
(vi) *The impact costs $\Lambda \cdot \mathbb{E}[\int_0^T \xi^*(t, X^*(t;\theta);\theta)^2 dt]$ are strictly decreasing in θ.*

Let us shortly comment on Proposition 4.1. The fact that the overall costs are decreasing in θ (part (i)) is quite intuitive and can be deduced directly from the definition of the cost functional J. Proposition 4.1 (iv) states that the introduction of a dark pool decreases the expected asset position despite slower initial trading in the primary exchange (parts (ii) and (iii)). Parts (v) and (vi) of Proposition 4.1 confirm that the introduction of the dark pool decreases both the impact costs component and the risk costs component of the value function.[17]

We illustrate these properties in Figure 4.1. In the left picture, we consider a risk neutral investor and in the middle picture a risk averse investor. The optimal trading trajectories for trading with dark pool are displayed by the thick solid lines in these pictures. In the displayed scenario the dark pool order is executed at time τ. After execution in the dark pool, the liquidation task is finished. The dotted lines denote the scenario where the dark pool order is not executed during the entire trading horizon and the thin solid lines refer to the expected asset positions. We contrast the optimal strategy with dark pool by the optimal strategy without dark pool (dashed lines). As shown in Proposition 4.1, the dark pool slows down trading in the primary venue initially. Nevertheless, the expected position is smaller than the trading trajectory without dark pool. The right picture illustrates the dependence of the value function (solid line) and the two cost components (risk costs: dashed line; impact costs: dotted line) on the dark pool execution intensity θ. As proven in Proposition 4.1 (i) respectively (v) and (iv), both the overall costs and the cost components are decreasing in θ.

4.1.2. Dependence on Λ and $\alpha\Sigma$. It follows directly from the definition of the cost functional J that the value function is strictly increasing both in the impact costs parameter Λ and in the risk costs parameter $\alpha\Sigma$. For the optimal strategy, impact costs and risk costs have conflicting influences: while larger impact costs yield a reduction of the trading intensity, larger risk costs yield faster trading (cf. also the difference of the left and the middle picture of Figure 4.1). We summarize these findings in the following proposition.

PROPOSITION 4.2.

(i) *For $t \in [0, T)$, $C(t; \Lambda, \alpha\Sigma)$ is strictly increasing in Λ and in $\alpha\Sigma$.*
(ii) *Let $t \in [0, T)$ and $x > 0$ be fixed. Then $\xi^*(t, x; \Lambda)$ is strictly decreasing in Λ. Consequently, $\tilde{X}(t; \Lambda)$ is strictly increasing in Λ for $t \in (0, T)$.*
(iii) *Let $t \in [0, T)$ and $x > 0$ be fixed. Then $\xi^*(t, x; \alpha\Sigma)$ is strictly increasing in $\alpha\Sigma$. Consequently, $\tilde{X}(t; \alpha\Sigma)$ is strictly decreasing in $\alpha\Sigma$ for $t \in (0, T)$.*

[17] In the discrete-time setting of Kratz and Schöneborn (2014) the risk costs (but not the impact costs) are increasing for small θ. Also in discrete time, the overall costs are always decreasing in θ; hence for small θ, the increase in risk costs is outweighed by the decrease in impact costs. If the lengths of the trading periods tend to zero, the effect disappears (see Kratz 2011, remark 2.6.4); the reason for the increasing risk costs in discrete time is hence the limited flexibility of the trader due to long trading intervals where she cannot trade.

4.2. Portfolio Liquidation

If a risk averse investor has to liquidate a portfolio of multiple assets ($n \geq 2$), then correlation between the assets comes into play. Depending on whether the portfolio is well (poorly) balanced, it is intuitively optimal to place orders in the dark pool which are smaller (larger) than the remaining portfolio position. In the first case, the risk costs of the portfolio are small and therefore the trader will not risk to lose her balanced position by the full execution of her dark pool order for only one of the assets; hence her orders are smaller than the remainder of the position. In the second case, it might be optimal to place orders in the dark pool which are larger than the remainder of the position for risk mitigation reasons. This illustration suggests that overly simple adjustments of existing trading algorithms for optimal liquidation without dark pools can have undesirable results.

In Section 4.2.1, we verify the above intuition by analytical results about the dependence of the value function and the optimal strategy on the correlation of a portfolio of two assets; we then deduce the general structure of the optimal strategy dependent on whether the portfolio is well or poorly balanced. We also discuss the dependence of the value function and the optimal strategy on the price impact parameter and the execution intensities in Section 4.2.2 and 4.2.3, respectively.

As a prerequisite, we introduce a characterization of the optimal dark pool order which exploits that the jump times of the Poisson processes are almost surely distinct. Intuitively, an execution of the optimal dark pool order for asset i should bring the position in asset i to its optimal value given unchanged positions in all other assets $j \neq i$. The following proposition confirms this intuition (see also Naujokat and Westray 2011 for a similar result).

PROPOSITION 4.3. *Let $t \in [0, T]$, $x \in \mathbb{R}^n$ be the portfolio position at time t and $i = 1, \ldots, n$. Then,*

$$v(t, x - \eta_i^*(t, x)e_i) = \min_{\eta \in \mathbb{R}} v(t, x - \eta e_i).$$

4.2.1. Dependence on Correlation. We will see that dark pool trading is sensitive to the correlation of price increments. In the following, we discuss the case $n = 2$.[18] In order to simplify the exposition, we assume that there is no cross asset price impact:[19]

$$\Lambda = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}, \quad \Sigma = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix}. \tag{4.2}$$

For the purposes of this section, we assume that the variances σ_1 and σ_2 of the two assets as well as the risk aversion parameter α are strictly positive.

If the correlation of the two assets is positive ($\rho > 0$), a portfolio consisting of a long position in one asset and a short position in the other asset is more desirable than long positions (or short positions) in both assets; in the former case, a part of the risk of each

[18] If $n > 2$, the situation is more complicated: positive correlation is not transitive in general. Hence, we cannot use Definition 4.4 below for the characterization of well versus poorly diversified portfolios. If the correlations of the price processes satisfy transitivity, some of the results can be generalized.

[19] Cross price impact and correlation can have conflicting influences on the value function and the optimal strategy. We discuss cross asset impact at the end of Section 4.2.2.

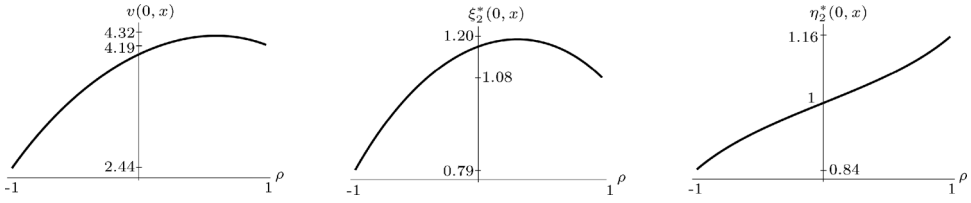

FIGURE 4.2. Dependence of the value function (left picture), the optimal trading intensity for the second asset (middle picture), and the optimal dark pool order for the second asset (right picture) on the correlation ρ. $x_1 = x_2 = 1$, $T = 1$, $\lambda_1 = 3$, $\lambda_2 = 0.2$, $\theta_1 = 0.5$, $\theta_2 = 5$, $\alpha = 4$, and $\sigma_1 = \sigma_2 = 1$.

asset is hedged by the other asset. Conversely, if $\rho < 0$, it is more desirable to have long (or short) positions in both assets.

DEFINITION 4.4. A portfolio $x = (x_1, x_2)^\top$ ($x_1, x_2 \neq 0$) is *well diversified* if either the signs of the positions are equal ($\operatorname{sgn}(x_1) = \operatorname{sgn}(x_2)$) and $\rho < 0$ or if the signs of the positions are different and $\rho > 0$. Otherwise, the portfolio is *poorly diversified*.

PROPOSITION 4.5. *Let $t \in [0, T)$ and $x_1, x_2 \neq 0$. Then $v(t, (x_1, x_2)^\top) < v(t, (x_1, -x_2)^\top)$ if and only if the portfolio x is well diversified.*

We can further specify the dependence of the value function on the correlation if the portfolio is well diversified.

PROPOSITION 4.6. *Let $t \in [0, T)$ and x be well diversified. Then $v(t, x)$ is strictly decreasing in $|\rho|$.*

The left picture of Figure 4.2 illustrates the dependence of the value function on the correlation ρ for a portfolio, that is, long in both assets. For $\rho < 0$, this portfolio is well diversified and the value function is increasing in ρ (i.e., decreasing in $|\rho|$) in line with Proposition 4.6. For $\rho > 0$, the portfolio is poorly diversified. This leads to elevated liquidation costs for small positive ρ. For large positive ρ, the increased risk costs of the current portfolio are outweighed by the (projected) smaller risk costs of a future well diversified portfolio (e.g., after the execution of an order in the dark pool; cf. also Proposition 4.8 (ii) below). The opportunity of risk reduction results in a decrease of the value function for large values of ρ in the displayed example.

We have the following symmetry and monotonicity properties for the entries of the value function matrix C.

PROPOSITION 4.7. *Let $t \in [0, T)$ and denote the entries of the value function matrix by $c_{i,j}(t; \rho)$, $i, j = 1, 2$.*

(i) $c_{1,1}(t; \rho), c_{2,2}(t; \rho) > 0$ *and* $c_{1,1}(t; \rho) = c_{1,1}(t; -\rho)$, $c_{2,2}(t; \rho) = c_{2,2}(t; -\rho)$, $\operatorname{sgn}(c_{1,2}(t; \rho)) = \operatorname{sgn}(\rho)$ *and* $c_{1,2}(t; \rho) = -c_{1,2}(t; -\rho)$.

(ii) $c_{1,1}(t; \rho)$ *and* $c_{2,2}(t; \rho)$ *are strictly increasing in ρ on* $[-1, 0)$ *and strictly decreasing in ρ on* $(0, 1]$. $c_{1,2}(t; \rho)$ *is increasing in ρ on* $[-1, 1]$.

The risk mitigation opportunity created by a strong correlation of the price increments becomes apparent again in Proposition 4.7. Liquidating a single-asset position $x = (x_1, 0)^\top$ results in the cost $v(t, (x, 0)) = c_{1,1}(t; \rho)x_1^2$, which exhibits a strict local maximum at $\rho = 0$ and decreases as the correlation between the two assets becomes stronger

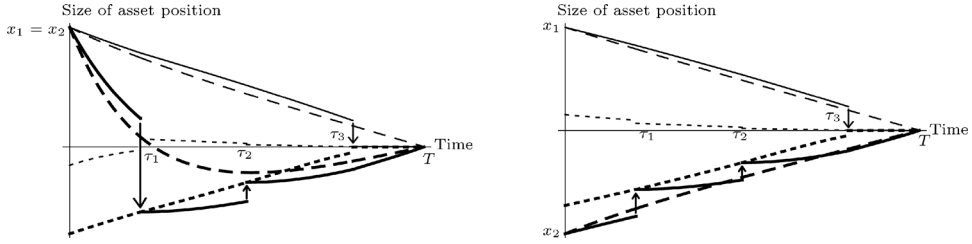

FIGURE 4.3. Evolution of a portfolio consisting of two highly correlated stocks over time. The left picture illustrates a poorly diversified portfolio, the right picture a well-diversified portfolio. In both pictures, thin lines are used for the less liquid first stock and thick lines for the more liquid second stock. Dashed lines correspond to trading without the dark pool and solid lines correspond to a realization of the liquidation process using the dark pool, where dark pool orders for the second stock are executed at times τ_1 and τ_2 and for the first stock at time τ_3. Dotted lines correspond to the position which the investor aims to reach by her dark pool order for the respective stock. $x_1 = 1$, $x_2 = 1$ (left picture), $x_2 = -1$ (right picture), $T = 1$, $\theta_1 = 0.5$, $\theta_2 = 3$, $\lambda_1 = 3$, $\lambda_2 = 0.2$, $\alpha = 4$, $\sigma_1 = \sigma_2 = 1$, and $\rho = 0.9$.

(irrespective of the sign of the correlation). This implies in particular that it is optimal for the investor to trade in both assets (unless $\rho = 0$) even if the current position in one asset is zero.

Proposition 4.5 suggests that an optimal liquidation strategy never changes the sign of the asset positions of a well-diversified portfolio and always seeks to turn a poorly diversified portfolio into a well-diversified portfolio. The following proposition confirms this conjecture.

PROPOSITION 4.8. *Let $t \in [0, T)$ and $x \in \mathbb{R}^2$ be the portfolio position at time t.*

(i) If x is well diversified, then $X^*(s)$ is well diversified for all $s \in [t, T)$ with $\operatorname{sgn}(X_i^*(s)) = \operatorname{sgn}(x_i)$ for $i = 1, 2$.

(ii) If x is poorly diversified, then $\operatorname{sgn}(X_i^*(s-) - \eta_i^*(s, X^*(s-))) \neq \operatorname{sgn}(x_i)$ for $s \in [t, \tau)$, where

$$\tau = \inf\{s \geq t | \operatorname{sgn}(X_i^*(s)) \neq \operatorname{sgn}(x_i) \text{ or } X_i^*(s) = 0 \text{ for some } i = 1, 2\} \wedge T > t \quad a.s.$$

By Proposition 4.8, the investor trades in both assets during the entire trading horizon $[0, T]$ if the portfolio is well diversified. If the portfolio is poorly diversified the execution of the dark pool order in one of the assets *always* changes the sign of the position. If dark pool orders are never executed, it can be optimal to decrease the risk costs by changing the sign of the position in one of the assets by only trading at the exchange (as it is the case in the numerical example underlying the left picture of Figure 4.3); in general, this is not the case.[20]

[20]Consider, e.g., a portfolio where $x_1 = x_2$, $\lambda_1 = \lambda_2$, $\sigma_1 = \sigma_2$, and $\theta_1 = \theta_2$. Then the optimal trading intensities for the two assets must be equal until a dark pool order is executed. In particular, if the orders in the dark pool are never executed, both positions must become zero at the same time after which further trading is not optimal.

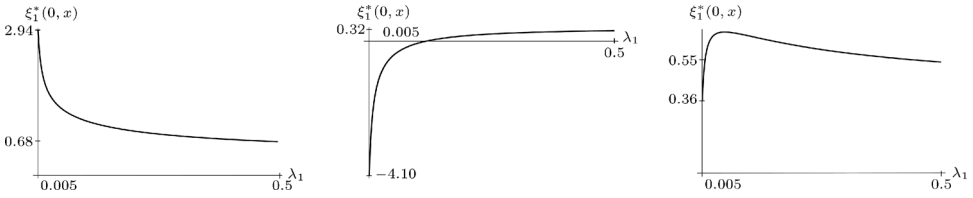

FIGURE 4.4. Dependence of the optimal trading intensity for the first asset on $\lambda_1 \in (0.005, 0.5)$ for a well-diversified portfolio x with $x_1 = 0.7$ and different positions in the second asset: $x_2 = -0.7$ (left picture), $x_2 = -1$ (middle picture), and $x_2 = -0.81$ (right picture). $T = 1$, $\lambda_2 = 1$, $\theta_1 = \theta_2 = 1$, $\alpha = 4$, $\sigma_1 = \sigma_2 = 1$, and $\rho = 0.9$.

In the following, we specify the dependence of the optimal strategy on ρ. By Theorem 3.9, we have

$$\begin{aligned} \xi_1^*(t,x) &= \tfrac{1}{\lambda_1}(c_{1,1}(t;\rho)x_1 + c_{1,2}(t;\rho)x_2), \\ \xi_2^*(t,x) &= \tfrac{1}{\lambda_2}(c_{2,2}(t;\rho)x_2 + c_{1,2}(t;\rho)x_2), \end{aligned} \tag{4.3}$$

$$\eta_1^*(t,x) = x_1 + \frac{c_{1,2}(t;\rho)}{c_{1,1}(t;\rho)}x_2, \quad \eta_2^*(t,x) = x_2 + \frac{c_{1,2}(t;\rho)}{c_{2,2}(t;\rho)}x_1. \tag{4.4}$$

PROPOSITION 4.9. *Let $t \in [0,T)$, $x = (x_1, x_2)^\top$, $x_1, x_2 > 0$, and $i = 1, 2$.*

(i) $\eta_i(t,x)$ is strictly increasing in ρ for $\rho \in [-1,1]$.
(ii) $\xi_i(t,x)$ is strictly increasing in ρ for $\rho \in [-1,0)$.

Analogue results hold for $x_1, x_2 < 0$ and $\operatorname{sgn}(x_1) \neq \operatorname{sgn}(x_2)$.

For a well-diversified portfolio, the profit from diversification is increasing in $|\rho|$ (cf. Proposition 4.6). Therefore, the investor decreases her trading activity both at the exchange and in the dark pool for larger ρ. For a poorly diversified portfolio, this is not necessarily the case (cf. the left picture of Figure 4.2 and the corresponding discussion); in contrast to the optimal dark pool orders, the trader might decrease her trading intensity for large positive ρ in order to save price impact costs while waiting for the execution of an order in the dark pool (which yields a well-diversified position by Proposition 4.8 (ii)). We illustrate the dependence of u^* on ρ in the middle and the right picture of Figure 4.2. In the displayed case, the trading intensity of the second asset is not increasing in ρ if x is poorly diversified (i.e., $\rho > 0$). On the other hand, the optimal dark pool order is strictly increasing for $\rho \in [-1,1]$. The symmetry of the graph in the right picture follows directly from the symmetries of $c_{1,2}$ and $c_{2,2}$. If $\rho = 0$, the optimal strategies for the two assets are independent and follow from the formulae of Section 4.1; in particular, the optimal dark pool order equals $x_2 = 1$.

We deduce the general structure of the optimal strategy for initial positions $x_1, x_2 > 0$ from the above results. If x is poorly diversified ($\rho > 0$), we have $\xi_i(t,x), \eta_i(t,x) > 0$ by equations (4.3) and (4.4); hence it is optimal to decrease the position. This is not necessarily the case for a well diversified portfolio ($\rho < 0$); in this case both $\xi_i(t,x)$ and $\eta_i(t,x)$ can be negative as $c_{1,2}(t;\rho) < 0$ (cf. also the middle picture of Figure 4.4 below). For the first asset, this is the case if and only if $x_1 < -\frac{c_{1,2}(t;\rho)}{c_{1,1}(t;\rho)}x_2$. It can hence be optimal to *increase* the position in order to further reduce the risk costs of the portfolio. Note that as $\operatorname{sgn}(\xi_i(t,x)) = \operatorname{sgn}(\eta_i(t,x))$ the direction of trading in the dark pool and at the

primary exchange is always the same. Furthermore, it can be optimal to neither trade at the exchange nor in the dark pool (in the first asset) at time t if $x_1 = -\frac{c_{1,2}(t;\rho)}{c_{1,1}(t;\rho)} x_2$.

As the execution of an order in the dark pool balances the trade-off of price impact costs against risk costs, it should intuitively never be optimal to increase a positive position or to decrease a negative position once the dark pool order of one of the assets has been executed. The following proposition confirms this conjecture.

PROPOSITION 4.10. *Let $t \in [0, T)$, $x \in \mathbb{R}^2$ be the portfolio position at time t and τ_1 be the first jump time of π. Then for all $s \geq \tau_1$, $i = 1, 2$,*

$$\operatorname{sgn}(\xi_i(s, X^*(s-))) = \operatorname{sgn}(\eta_i(s, X^*(s-))) = \operatorname{sgn}(X_i^*(s-)) \quad \text{or}$$
$$\xi_i(s, X^*(s-)) = \eta_i(s, X^*(s-)) = 0.$$

We close the section by illustrating the structure of the optimal strategy and its dependence on ρ by a numerical example. To this end, we consider two strongly positively correlated assets with $\lambda_1 = 3$, $\lambda_2 = 0.2$, that is, the second asset is more liquid than the first asset. We also model the dark pool in such a way that the execution of orders for the second asset is more probable than for the first asset: $\theta_1 = 0.5$, $\theta_2 = 3$.[21] We consider a poorly diversified portfolio $x = (1, 1)^\top$ and a well-diversified portfolio $x = (1, -1)^\top$ of the two stocks.

Figure 4.3 shows the evolution of the two portfolios if a risk averse investor applies the optimal strategy. The left picture corresponds to the poorly diversified portfolio, the right one to the well-diversified portfolio. In both cases, thin lines are used for the first stock and thick lines for the second. Dashed lines correspond to trading without the dark pool and the solid lines correspond to a realization of the liquidation process using the dark pool, where the dark pool orders for the second stock are executed at times τ_1 and τ_2 and for the first stock only at time τ_3, that is, dark pool orders for the more liquid stock are executed twice before any execution in the less liquid stock takes place. Dotted lines correspond to the position which the investor aims to reach by her dark pool order for the respective stock (cf. Proposition 4.3).

For the poorly diversified portfolio, the trader tries to improve her risky position by trading out of the second stock (cf. Proposition 4.8 (ii) and the subsequent discussion). For this stock, trading in the primary venue is less expensive and being executed in the dark pool is more probable. For the well-diversified portfolio, the portfolio position is decreasing almost linearly in time in all cases. In addition, orders in the dark pool are very large for the poorly diversified portfolio and comparatively small for the well diversified portfolio, in line with Proposition 4.8. The reason is that dark pool orders are such that either the risk costs are decreased significantly by an execution (in the poorly diversified case) or they are only slightly increased (in the well-diversified case). Note that both for the poorly and for the well-diversified portfolio, these effects are stronger for the liquid stock; for the illiquid stock, savings in price impact costs outweigh savings in risk costs.

4.2.2. Dependence on Price Impact. In this section, we discuss the dependence of the value function and the optimal strategy on the price impact and on the cross price impact. It follows directly from the definition of the cost functional J that the value function is

[21] Our choice of the parameters reflects the intuition that the asset which is more liquid at the exchange (smaller λ) is also more liquid in the dark pool (larger θ). Theoretical findings of Ye (2013) support this choice. However, we are not aware of any empirical evidence for this; in some cases, the opposite parameter choice can also be plausible.

increasing in the price impact matrix Λ in the sense of Notation 3.1 (i); in particular, it is increasing in the diagonal elements of Λ. On the other hand, the increase of v is in some sense bounded by the increase of Λ.

PROPOSITION 4.11. *Let $t \in [0, T)$ and $x \in \mathbb{R}^n$.*

(i) *$v(t, x)$ is increasing in Λ.*
(ii) *Let $\Lambda = \operatorname{diag}(\lambda_j, ; j = 1, \ldots, n)$. For $i = 1, \ldots, n$, $\frac{v(t,x;\lambda_i)}{\lambda_i}$ is decreasing in λ_i.*

The optimal trading strategy does not need to be monotone in the price impact parameter since risk mitigation and liquidation can be conflicting desires. We illustrate this situation in Figure 4.4 by considering three different well diversified portfolios. In the first case ($x_1 = 0.7$, $x_2 = -0.7$; left picture) further risk mitigation is not optimal; the trading intensity is decreasing in the price impact. In the second case ($x_1 = 0.7$, $x_2 = -1.1$; middle picture) it is profitable to increase the position for small price impact in order to reach a position with even less risk costs; the optimal trading intensity is increasing and further risk mitigation is only profitable for small enough λ_1. In the third case ($x_1 = 0.7, x_2 = -0.81$; right picture) the optimal intensity is increasing for very small λ_1 and then decreasing; the conflict apparent in the left and the middle picture destroys the monotonicity in this case.

We close the section by analyzing the cross asset price impact $\lambda_{1,2}$ in the price impact matrix

$$\Lambda = \begin{pmatrix} \lambda_1 & \lambda_{1,2} \\ \lambda_{1,2} & \lambda_2 \end{pmatrix} \quad \text{for} \quad 0 < |\lambda_{1,2}| < \sqrt{\lambda_1 \lambda_2}.$$

We first consider the case $\operatorname{sgn}(\lambda_{1,2}) = \operatorname{sgn}(\rho)$.[22] We obtain the following analogs of Propositions 4.5 and 4.8 (ii).

PROPOSITION 4.12. *Let $t \in [0, T)$, $\operatorname{sgn}(\lambda_{1,2}) = \operatorname{sgn}(\rho)$ and $x_1, x_2 \neq 0$.*

(i) *$v(t, (x_1, x_2)^\top) < v(t, (x_1, -x_2)^\top)$ if and only if the portfolio x is well diversified.*[23]
(ii) *If x is poorly diversified, then $\operatorname{sgn}(X_i^*(s-) - \eta_i^*(s, X^*(s-))) \neq \operatorname{sgn}(x_i)$ for $s \in [t, \tau)$, where*

$$\tau = \inf\{s \geq t | \operatorname{sgn}(X_i^*(s)) \neq \operatorname{sgn}(x_i) \text{ or } X_i^*(s) = 0 \text{ for some } i = 1, 2\} \wedge T > t \quad a.s.$$

(iii) *If x is well diversified, then $\operatorname{sgn}(X_i^*(s-) - \eta_i^*(s, X^*(s-))) = \operatorname{sgn}(x_i)$ for $s \in [t, \tau)$, where τ is as above.*

For well-diversified portfolios, we can recover only a part of Proposition 4.8 (i): by (iii), it is not optimal to change the sign of the position by placing oversized orders in the dark pool. However, it can be optimal to turn a well-diversified portfolio into a poorly diversified portfolio by trading at the primary exchange. We illustrate this in the left

[22] This case is more intuitive than the converse: if the assets are positively correlated, buying in the first asset should rather increase than decrease the price of the second asset.

[23] Similarly as before (cf. Definition 4.4), we call a portfolio $x = (x_1, x_2)^\top$ ($x_1, x_2 \neq 0$) *well diversified* if either the signs of the positions are equal ($\operatorname{sgn}(x_1) = \operatorname{sgn}(x_2)$) and $\rho, \lambda_{1,2} < 0$ or if the signs of the positions are different and $\rho, \lambda_{1,2} > 0$. Otherwise, the portfolio is *poorly diversified*.

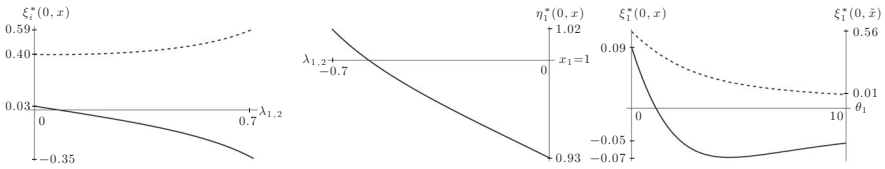

FIGURE 4.5. The left picture illustrates the dependence of the optimal trading intensity on $\lambda_{1,2} \in (0, 0.7)$ for a portfolio with $x_1 = 0$, $x_2 = 1$. The solid line refers to the first asset, the dashed line to the second asset. The middle picture illustrates the dependence of the optimal dark pool order on $\lambda_{1,2} \in (-0.7, 0)$ for a portfolio $x_1 = 1$, $x_2 = -1$. In both pictures, $T = 1$, $\lambda_1 = \lambda_2 = \alpha = \sigma_1 = \sigma_2 = 1$, $\theta_1 = \theta_2 = 3$ and $\rho = 0.2$; in particular, $\operatorname{sgn}(\rho) = \operatorname{sgn}(\lambda_{1,2})$ for the left picture and $\operatorname{sgn}(\rho) \neq \operatorname{sgn}(\lambda_{1,2})$ for the middle picture. The right picture illustrates the dependence of the optimal trading intensity on $\theta_1 \in (0, 10)$ for a well diversified portfolio x ($x_1 = 0.25$, $x_2 = -1$, $\rho = 0.9$; solid line) and a poorly diversified portfolio $\tilde{x}$ ($\tilde{x}_1 = x_1$, $\tilde{x}_2 = 1$; dashed line). $\lambda_{1,2} = 0$; the remaining parameters are as above.

picture of Figure 4.5. In the displayed scenario, the trader holds a positive position in the second asset only. The trading intensity of the first asset is denoted by the solid line. If $\lambda_{1,2} = 0$, it is optimal to sell stocks in the first asset which is positively correlated to the second asset (cf. the results of Section 4.2.1, in particular the discussion following Proposition 4.7). If $\lambda_{1,2}$ increases, it becomes profitable to buy stocks in the first asset: as the cross price impact is positive, this allows the trader to increase her trading intensity in the second asset (dashed line) without increasing the overall impact costs too much. In parallel, the trader places a sell order for the first asset in the dark pool as she does not want to sell the stocks at the exchange again (which would result in additional price impact costs due to the sign of $\lambda_{1,2}$).

The case $\operatorname{sgn}(\lambda_{1,2}) \neq \operatorname{sgn}(\rho)$ (cf. Footnote 22) is more complicated. In this case, it can be optimal to change the sign of a well diversified position (in the sense of Definition 4.4) by placing oversized orders in the dark pool; for large cross price impact, the savings in impact costs resulting from the change of the position can outweigh the increase risk costs. We illustrate this by a numerical example in the middle picture of Figure 4.5.

4.2.3. Dependence on the Execution Intensities. We conclude this section by analyzing the dependence of the value function and the optimal strategy on dark pool liquidity (more precisely on the intensities of the Poisson process π). Similarly as in the single-asset case (cf. Proposition 4.1 (i)), the costs are decreasing in the intensities θ_i.

PROPOSITION 4.13. *Let $t \in [0, T)$, $x \in \mathbb{R}^n$ and $i = 1, \ldots, n$. Then $v(t, x; \theta_i)$ is decreasing in θ_i.*

The same trade-offs that can cause a nonmonotone dependence of the optimal strategy on the price impact parameters can also give rise to a nonmonotone dependence on the dark pool liquidity parameters θ_i. We illustrate the dependence of the optimal trading intensity on θ_1 for a poorly (dashed line) and a well diversified portfolio (solid line) in the right picture of Figure 4.5.

5. PROOFS OF THE MAIN RESULTS

5.1. Proof of Theorem 3.2

For $n \geq 2$, the second summand in the matrix differential equation $C^\top \tilde{C} C$ is in general not linear (or quadratic), and (3.1) is not a Riccati matrix differential equation. Furthermore, a closed form solution for the corresponding initial value problem is not known, and the existing theory about Riccati matrix differential equations is not applicable directly.

It turns out that appropriate upper and lower bounds for the nonlinear term $C^\top \tilde{C} C$ transform to lower and upper bounds (P respectively Q) for the solution of the Matrix Initial Value Problem (3.1) and yield existence and positive definiteness of the solution on the whole interval $(-\infty, T]$ (Theorem 3.2). To this end, we require a version of a well-known comparison result for matrix Riccati differential equations, which we state in Appendix A. The main step is thus to obtain adequate matrix inequalities which enable us to transfer these results to the Initial Value Problem (3.1).

For $C > 0$, we have $0 \leq C\tilde{C}C$. The desired upper bound of $C\tilde{C}C$ is a direct consequence of the matrix inequality stated in the following result.

PROPOSITION 5.1. *Let $C = (c_{i,j})_{i,j=1,\dots,n} \in \mathbb{R}^{n\times n}$ be a positive definite matrix and $\theta_i > 0$, $i = 1, \dots, n$. Then*

$$C \leq \theta \operatorname{diag}\left(\frac{c_{i,i}}{\theta_i}\right) = \theta \tilde{C}^{-1}, \tag{5.1}$$

where $\theta := \theta(n) := \sum_{i=1}^n \theta_i$.

Proof. We prove the inequality by induction on n. It is clear for $n = 1$ with equality in (5.1). Let now $n \geq 1$ and $C = (c_{i,j})_{i,j=1,\dots,n+1} \in \mathbb{R}^{(n+1)\times(n+1)}$ be positive definite. Define $C_n \in \mathbb{R}^{n\times n}$ and $c \in \mathbb{R}^n$ such that

$$C = \begin{pmatrix} C_n & c \\ c^\top & c_{n+1,n+1} \end{pmatrix}.$$

For $z = (x, y)^\top \in \mathbb{R}^n \times \mathbb{R}$, $z \neq 0$, we have

$$\begin{aligned} z^\top C z &= x^\top C_n x + 2x^\top c y + c_{n+1,n+1} y^2, \\ z^\top \operatorname{diag}\left(\frac{c_{i,i}}{\theta_i}\right) z &= x^\top \operatorname{diag}\left(\frac{c_{i,i}}{\theta_i}\right) x + \frac{c_{n+1,n+1}}{\theta_{n+1}} y^2. \end{aligned}$$

By abuse of notation, $\operatorname{diag}(\frac{c_{i,i}}{\theta_i})$ is used both for the diagonal $n \times n$-matrix with $\frac{c_{1,1}}{\theta_1}, \dots, \frac{c_{n,n}}{\theta_n}$ in the diagonal and for the respective diagonal $(n+1) \times (n+1)$ matrix. Which one we refer to is always clear from the context. We compute

$$\begin{aligned} z^\top \left(\theta(n+1)\operatorname{diag}\left(\frac{c_{i,i}}{\theta_i}\right) - C\right) z &= x^\top \left(\theta(n)\operatorname{diag}\left(\frac{c_{i,i}}{\theta_i}\right) - C_n\right) x^\top \\ &\quad + \theta_{n+1} x^\top \operatorname{diag}\left(\frac{c_{i,i}}{\theta_i}\right) x - 2x^\top c y + \frac{\theta(n)}{\theta_{n+1}} c_{n+1,n+1} y^2. \end{aligned} \tag{5.2}$$

The first summand in equation (5.2) is nonnegative by the induction hypothesis. The remainder equals

$$\left(x - \frac{1}{\theta_{n+1}}\operatorname{diag}\left(\frac{\theta_i}{c_{i,i}}\right)cy\right)^\top \theta_{n+1}\operatorname{diag}\left(\frac{c_{i,i}}{\theta_i}\right)\left(x - \frac{1}{\theta_{n+1}}\operatorname{diag}\left(\frac{\theta_i}{c_{i,i}}\right)cy\right) + \left(c_{n+1,n+1}\frac{\theta(n)}{\theta_{n+1}} - c^\top \frac{\operatorname{diag}\left(\frac{\theta_i}{c_{i,i}}\right)}{\theta_{n+1}}c\right)y^2,$$

where the first summand is nonnegative as C_n (and therefore $\operatorname{diag}(\frac{c_{i,i}}{\theta_i})$) is positive definite and $\theta_{n+1} > 0$. We have (see, e.g., the book by Horn and Johnson 1985, corollary 7.7.4[24])

$$0 < \frac{1}{\theta(n)}\operatorname{diag}\left(\frac{\theta_i}{c_{i,i}}\right) \le C_n^{-1} \tag{5.3}$$

by the induction hypothesis. Moreover

$$C\begin{pmatrix} I_{n\times n} & -C_n^{-1}c \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} C_n & 0 \\ c^\top & c_{n+1,n+1} - c^\top C_n^{-1}c \end{pmatrix}$$

and hence

$$c_{n+1,n+1} - c^\top C_n^{-1}c = \frac{\det C}{\det C_n} > 0. \tag{5.4}$$

Finally,

$$c_{n+1,n+1} - \frac{1}{\theta(n)}c^\top \operatorname{diag}\left(\frac{\theta_i}{c_{i,i}}\right)c \overset{(18)}{\ge} c_{n+1,n+1} - c^\top C_n^{-1}c \overset{(19)}{>} 0,$$

finishing the proof. □

Applications of Horn and Johnson (1985), corollary 7.7.4 (cf. Footnote 24 again), imply the desired bound for $C\tilde{C}C$. In addition, we obtain two elementary matrix inequalities.

COROLLARY 5.2. *Let $C > 0$ and $\theta_i \ge 0$, $i = 1, \dots, n$. Then*

$$C\tilde{C}C \le \theta C, \tag{5.5}$$

$$C \le n \operatorname{diag}(c_{i,i}) \quad \textit{and} \quad C \le \operatorname{tr}(C)I.$$

The Matrix Inequality (5.5) enables us to apply Theorem B.1 to the Matrix Initial Value Problem (3.1) such that we can prove existence of a solution C of (3.1) on the whole interval $(-\infty, T]$ and at the same time construct upper and lower bounds for C via the solutions of the initial value problems in (3.4); these are given explicitly by

$$P(l,t) = p(l,t)I, \quad Q(l,t) = q(l,t)I, \tag{5.6}$$

[24]For matrices A, B with $0 < A < B$, we have $0 < B^{-1} < A^{-1}$.

where

$$p(l,t)=\sqrt{\frac{\theta^2}{4}+\alpha d_{\min}}\coth\left(\sqrt{\frac{\theta^2}{4}+\alpha d_{\min}}(T-t)+\kappa_1(l)\right)-\frac{\theta}{2}, \tag{5.7}$$

$$q(l,t):=\sqrt{\alpha d_{\max}}\coth\left(\sqrt{\alpha d_{\max}}(T-t)+\kappa_2(l)\right) \tag{5.8}$$

for $\theta+\alpha d_{\min}>0$, respectively, $\alpha d_{\max}>0$ with

$$\kappa_1(l):=\operatorname{arcoth}\left(\frac{\frac{l}{\lambda_{\max}}+\frac{\theta}{2}}{\sqrt{\frac{\theta^2}{4}+\alpha d_{\min}}}\right)>0,\quad \kappa_2(l):=\operatorname{arcoth}\left(\frac{\frac{l}{\lambda_{\min}}}{\sqrt{\alpha d_{\max}}}\right)>0$$

and

$$p(l,t):=\frac{1}{T-t+\frac{\lambda_{\max}}{l}},\quad q(l,t):=\frac{1}{T-t+\frac{\lambda_{\min}}{l}} \tag{5.9}$$

for $\theta=\alpha d_{\min}=0$, respectively, $\alpha d_{\max}=0$. Note also that $0<p(l,t),q(l,t)<+\infty$ for all $t\in(-\infty,T]$.

Proof of Theorem 3.2. Let $C(l,t)$ be a solution of (3.1) on some interval $(t_1,T]$; note that there exists a local solution by the Picard-Lindelöf theorem. The symmetry of Λ, Σ and the initial value $C(l,T)=lI$ imply that $C(l,t)$ is symmetric on $(t_1,T]$.

Let now $\hat{P}:=\Lambda P$ for P as in (5.6). Then $\hat{P}(l,t)$ solves $\hat{P}'=\hat{P}\Lambda^{-1}\hat{P}+\theta\hat{P}-\alpha d_{\min}\Lambda$, $\hat{P}(T)=\frac{l}{\lambda_{\max}}\Lambda$ on $(-\infty,T]$. As $P>0$ and P commutes with Λ, we have $\hat{P}(l,t)>0$. Assume that

$$\{t\in(t_1,T]\,|\,C(l,t)\text{ is not positive definite}\}\neq\emptyset \tag{5.10}$$

and define $\tau:=\sup\{t\in(t_1,T]\,|\,C(l,t)\text{ is not positive definite}\}$. As $C(l,T)=lI>0$ and $C(l,\cdot)$ is continuous, there exists an $\epsilon>0$ such that $C(l,t)>0$ for $t\in(T-\epsilon,T]$ and thus $\tau<T$. We apply Theorem B.1 to $\bar{P}:=-\hat{P}$ and $\bar{C}:=-C$ on $[\tau,T]$. We have

$$\bar{P}(l,T)=-\frac{l}{\lambda_{\max}}\Lambda\geq -lI=\bar{C}(l,T)$$

and

$$\begin{aligned}\bar{P}'&=-\bar{P}\Lambda^{-1}\bar{P}+\theta\bar{P}+\alpha d_{\min}\Lambda,\\ \bar{C}'&=-\bar{C}\Lambda^{-1}\bar{C}+\bar{C}\tilde{\bar{C}}\bar{C}+\alpha\Sigma=-\bar{C}\Lambda^{-1}\bar{C}+\theta\bar{C}+(\alpha\sqrt{\Lambda}D\sqrt{\Lambda}+\bar{C}\tilde{\bar{C}}\bar{C}-\theta\bar{C}).\end{aligned}$$

Let now $x\in\mathbb{R}^n$. Applying Corollary 5.2 to $-\bar{C}$, we obtain

$$\begin{aligned}&x^\top(\alpha\sqrt{\Lambda}D\sqrt{\Lambda}+\bar{C}\tilde{\bar{C}}\bar{C}-\theta\bar{C}-\alpha d_{\min}\Lambda)x=\\ &\quad\alpha x^\top(\sqrt{\Lambda}(D-d_{\min}I)\sqrt{\Lambda})x+x^\top(\bar{C}\tilde{\bar{C}}\bar{C}-\theta\bar{C})(t)x\geq 0.\end{aligned}$$

As $\Lambda>0$, Theorem B.1 implies $\bar{C}(l,t)\leq\bar{P}(l,t)$ and therefore $0<\hat{P}(l,t)\leq C(l,t)$ on $(\tau,T]$. By continuity of $C(l,\cdot)$, we have $0<\hat{P}(l,\tau)\leq C(l,\tau)$ and thus $C(l,t)>0$ in some neighborhood of τ, a contradiction to Assumption (5.10). Hence, $C(l,t)$ is positive definite on the whole interval $(t_1,T]$. Applying Theorem B.1 in the same way as above again, yields that we may choose $t_1=-\infty$ and that $0<\hat{P}(l,t)\leq C(l,t)$ on $(-\infty,T]$.

A similar argument establishes Q as an upper bound by using $0 \leq C\tilde{C}C$ instead of Inequality (5.5). □

5.2. Proof of Proposition 3.4

Before we begin, we introduce the following notation.

NOTATION 5.3.

(i) We denote the jump times of π by $(\tau_j)_{j\in\mathbb{N}}$, where $\tau_j < \tau_{j+1}$ $(j \in \mathbb{N})$ almost surely (with the convention $\tau_0 = t$).

(ii) Given the Markovian control $u^*(l)$, the Stochastic Differential Equation (2.1) possesses a unique solution. We denote the process controlled by $u^*(l)$ by

$$X^*(l, s) := X^{u^*(l)}(s).$$

In order to prove admissibility of $u^*(l)$, we show that $\|X^*(l, \cdot)\|_2$ is bounded by using Gronwall's inequality pathwise inductively on the time-intervals $[\tau_i \wedge T, \tau_{i+1} \wedge T)$ and interlacing the jumps (cf. Remark 2.4). This can be achieved by applying the upper and lower bounds of $C(l, s)$ from Theorem 3.2.

LEMMA 5.4. *Let $l > l_0$ for l_0 as in equation (3.3), $t \in [0, T)$, $x \in \mathbb{R}^n$ be the portfolio position at time t and $\theta = \sum_i \theta_i$ as before. Then the following hold.*

(i) For $s \in [t, T)$,

$$X^*(l, s)^\top C(l, s) X^*(l, s) \leq \exp(\theta(s - t)) x^\top C(l, t) x \quad a.s.$$

(ii) There exists a constant K independent of l such that for all $s \in [t, T)$, $\|X^*(l, s)\|_2 \leq K$ a.s.

Proof.

(i) Let $i \in \mathbb{N}$. On $\{\tau_i < T\}$, $X^*(l, \cdot)$ solves the initial value problem

$$X' = -\Lambda^{-1} C(l) X, \quad X(\tau_i) = X^*(l, \tau_i)$$

for $s \in [\tau_i, \tau_{i+1} \wedge T)$. Hence, as $C(l)$ solves the Initial Value Problem (3.1),

$$\begin{aligned}
&\frac{\partial}{\partial s}(X^*(l, s)^\top C(l, s) X^*(l, s)) \\
&\quad = \frac{\partial}{\partial s} X^*(l, s)^\top C(l, s) X^*(l, s) + X^*(l, s)^\top \frac{\partial}{\partial s} C(l, s) X^*(l, s) \\
&\qquad + X^*(l, s)^\top C(l, s) \frac{\partial}{\partial s} X^*(l, s) \\
&\quad = -X^*(l, s)^\top C(l, s) \Lambda^{-1} C(l, s) X^*(l, s) + X^*(l, s)^\top C(l, s) \Lambda^{-1} C(l, s) X^*(l, s) \\
&\qquad + X^*(l, s)^\top C(l, s) \tilde{C}(l, s) C(l, s) X^*(l, s) - \alpha X^*(l, s)^\top \Sigma X^*(l, s) \\
&\qquad - X^*(l, s)^\top C(l, s) \Lambda^{-1} C(l, s) X^*(l, s) \\
&\quad \leq \theta X^*(l, s)^\top C(l, s) X^*(l, s)
\end{aligned}$$

by Corollary 5.2 and the fact that $C, \Lambda^{-1}, \Sigma \geq 0$. By Gronwall's inequality, this implies

$$X^*(l,s)^\top C(l,s) X^*(l,s) \leq \exp(\theta(s-\tau_i)) x^\top C(l,t) x. \tag{5.11}$$

Now, on $\{\tau_{i+1} < T\}$, there exits an (almost surely) unique $j = 1, \ldots, n$ (cf. Assumption 2.2 (ii)) such that

$$X^*(l,\tau_{i+1}) = X^*(l,\tau_{i+1}-) - \eta_j^*(l,\tau_{i+1}, X^*(l,\tau_{i+1}-))e_j.$$

Let $\eta \in \mathbb{R}$. Then

$$\begin{aligned}
&(X^*(l,\tau_{i+1}-) - \eta e_j)^\top C(l,\tau_{i+1})(X^*(l,\tau_{i+1}-) - \eta e_j) \\
&\quad = X^*(l,\tau_{i+1}-)^\top C(l,\tau_{i+1}) X^*(l,\tau_{i+1}-) + \eta^2 e_j^\top C(l,\tau_{i+1}) e_j \\
&\qquad - 2\eta e_j^\top C(l,\tau_{i+1}) X^*(l,\tau_{i+1}-) \\
&\quad = X^*(l,\tau_{i+1}-)^\top C(l,\tau_{i+1}) X^*(l,\tau_{i+1}-) \\
&\qquad + \left(\eta\sqrt{c_{j,j}(l,\tau_{i+1})} - \frac{1}{\sqrt{c_{j,j}(l,\tau_{i+1})}} e_j^\top C(l,\tau_{i+1}) X^*(l,\tau_{i+1}-)\right)^2 \\
&\qquad - \frac{1}{c_{j,j}(l,\tau_{i+1})}(e_j^\top C(l,\tau_{i+1}) X^*(l,\tau_{i+1}-))^2,
\end{aligned}$$

which attains its (unique) minimum in η for $\eta = \eta_j^*(l,\tau_{i+1}, X^*(l,\tau_{i+1}-))$. Hence,

$$\begin{aligned}
&(X^*(l,\tau_{i+1}-) - \eta_j^*(l,\tau_{i+1}, X^*(l,\tau_{i+1}-))e_j)^\top \\
&\quad C(l,\tau_{i+1})(X^*(l,\tau_{i+1}-) - \eta_j^*(l,\tau_{i+1}, X^*(l,\tau_{i+1}-))e_j) \\
&\qquad = \min_{\eta\in\mathbb{R}} (X^*(l,\tau_{i+1}-) - \eta e_j)^\top C(l,\tau_{i+1})(X^*(l,\tau_{i+1}-) - \eta e_j)
\end{aligned}$$

and therefore

$$X^*(l,\tau_{i+1})^\top C(l,\tau_{i+1}) X^*(l,\tau_{i+1}) \leq X^*(l,\tau_{i+1}-)^\top C(l,\tau_{i+1}) X^*(l,\tau_{i+1}-). \tag{5.12}$$

Using Inequalities (5.11) and (5.12) inductively, we obtain the assertion.

(ii) By Theorem 3.2, we have $\Lambda \leq 1/p(l,s)C(l,s)$ and hence,

$$\begin{aligned}
X^*(l,s)^\top \Lambda X^*(l,s) &\leq \frac{1}{p(l,s)} X^*(l,s)^\top C(l,s) X^*(l,s) \\
&\leq \frac{1}{p(l_0,s)} X^*(l,s)^\top C(l,s) X^*(l,s) \\
&\leq \tilde{K} x^\top C(l,t) x
\end{aligned} \tag{5.13}$$

for a constant $\tilde{K}$ independent of s by (i) and the fact that $p(l_0)$ attains its minimum in $[t,T)$. The assertion follows as $q(t) := \lim_{l\to\infty} q(l,t) < +\infty$ exists (cf. Equations (5.8) and (5.9); cf. also Lemma 5.6 below) and $C(l,t) \leq \lambda_{\max} q(t)$ (cf. Theorem 3.2).

□

The bound obtained in Lemma 5.4 enables us to prove that $u^*(l)$ fulfills the moment conditions of Definition 2.3 (ii) and hence Proposition 3.4.

Proof of Proposition 3.4. Definition 2.3 (i) and (iii) are clearly satisfied.

Let $\|\cdot\|_{2,2}$ denote the matrix norm induced by the space $(\mathbb{R}^n, \|\cdot\|_2)$. Note that $\|\cdot\|_{2,2}$ is the spectral norm on $\mathbb{R}^{n\times n}$ and therefore (see, e.g., Bernstein 2005, theorem 8.4.9) $\|A\|_{2,2} \le \|B\|_{2,2}$ for $0 \le A \le B$. Using Theorem 3.2 and 5.4 (ii), we deduce

$$\begin{aligned}\mathbb{E}\left[\int_t^T \|\xi^*(l,s,X^*(s))\|_2^4 ds\right] &\le \mathbb{E}\left[\int_t^T \|\Lambda^{-1}\|_{2,2}^4 \|C(l,s)\|_{2,2}^4 \|X^*(l,s)\|_2^4 ds\right]\\ &\le \mathbb{E}\left[\int_t^T q(l,s)^4 K^4 ds\right] < +\infty\end{aligned}$$

for K as in Lemma 5.4 (ii). Similarly,

$$\begin{aligned}\mathbb{E}\left[\int_t^T \|\eta^*(l,s,X^*(s-))\|_2^8 ds\right] &\le \mathbb{E}\left[\int_t^T \|\underbrace{\bar{C}(l,s)}_{\le nC(l,s)^{-1} \text{ by Corollary 5.2}}\|_{2,2}^8 \|C(l,s)\|_{2,2}^8 \|X^*(l,s-)\|_2^8 ds\right]\\ &\le \mathbb{E}\left[\int_t^T \frac{n^8\lambda_{\max}^8 q(l,s)^8}{\lambda_{\min}^8 p(l,s)^8} K^8 ds\right] < +\infty.\end{aligned}$$

□

5.3. Proof of Theorem 3.5

We first require the following moment estimate for the controlled process.

LEMMA 5.5. *Let $t \in [0,T)$, $x \in \mathbb{R}^n$ and $u \in \tilde{\mathbb{A}}(t)$. Then*

$$\mathbb{E}\left[\sup_{t\le s\le T} \|X^u(s)\|_2^4\right] < +\infty, \quad \textit{in particular} \quad \mathbb{E}\left[\int_t^T \|X^u(s)\|_2^4 ds\right] < +\infty.$$

Proof. Let $s \in [t,T]$. Then by Hölder's inequality and a multidimensional version of Jensen's inequality (see, e.g., Kallenberg 2002, lemma 3.5) we obtain

$$\begin{aligned}\|X^u(s)\|_2^4 &\le \left(\|x\|_2 + \left\|\int_t^s \xi(r)dr\right\|_2 + \left\|\int_t^s \eta(r)d\pi(r)\right\|_2\right)^4\\ &\le 27\left(\|x\|_2^4 + (s-t)^3\int_t^s \|\xi(r)\|_2^4 dr + \left\|\int_t^s \eta(r)d\pi(r)\right\|_2^4\right). \qquad (5.14)\end{aligned}$$

By Definition 2.3 (ii), it is hence sufficient to consider the last summand of Equation (5.14). To this end, we define the compensated Poisson processes $M_i(s) := \pi_i(s) - \theta_i s$, $i = 1,\dots,n$. We note that Poisson distributed random variables have finite moments and apply Itô's isometry (note that $\langle M_i\rangle(s) = \theta_i s$) and Hölder's inequality to obtain

$$\begin{aligned}\mathbb{E}\left[\left|\int_t^s \eta_i(r)d\pi_i(r)\right|^4\right] &\le \mathbb{E}\left[(\pi_i(s)-\pi_i(t))^3\left(\int_t^s |\eta_i(r)|^4 dM_i(r) + \theta_i\int_t^s |\eta_i(r)|^4 dr\right)\right]\\ &\le \mathbb{E}\left[(\pi_i(s)-\pi_i(t))^6\right]^{\frac{1}{2}}\left(\mathbb{E}\left[\left(\int_t^s |\eta_i(r)|^4 dM_i(r)\right)^2\right]^{\frac{1}{2}}\right.\\ &\qquad\left. + \mathbb{E}\left[\left(\theta_i\int_t^s |\eta_i(r)|^4 dr\right)^2\right]^{\frac{1}{2}}\right)\\ &< K_i\end{aligned}$$

for a constant K_i independent of s by Definition 2.3 (ii), which finishes the proof. □

Proof of Theorem 3.5. Let $l > l_0$, $(t, x) \in [0, T) \times \mathbb{R}^n$ and $u = (\xi, \eta) \in \tilde{\mathbb{A}}(t)$. We apply Itô's formula (see, e.g., the book by Øksendal and Sulem 2007) to the function $w(l, t, X^u(t)) = X^u(t)^\top C(l, t) X^u(t)$:

$$\begin{aligned}
w(l, t, x) &= w(l, T, X^u(T)) + \int_t^T \nabla_x w(l, s, X^u(s))\xi(s) - \frac{\partial w}{\partial s}(l, s, X^u(s))ds \\
&\quad + \int_t^T \left(\sum_{i=1}^n w(l, s, X^u(s-)) - w(l, s, X^u(s-) - \eta_i(s)e_i) \right) \pi_i(ds) \\
&\leq w(l, T, X^u(T)) + \int_t^T f(X^u(s), \xi(s))ds \\
&\quad + \int_t^T \left(\sum_{i=1}^n w(l, s, X^u(s-)) - w(l, s, X^u(s-) - \eta_i(s)e_i) \right) \pi_i(ds) \\
&\quad - \int_t^T \left(\sum_{i=1}^n \theta_i (w(l, s, X^u(s-)) - w(l, s, X^u(s-) - \eta_i(s)e_i) \right) ds \qquad (5.15) \\
&= w(l, T, X^u(T)) + \int_t^T f(X^u(s), \xi(s))ds \\
&\quad + \sum_{i=1}^n \int_t^T (w(l, s, X^u(s-)) - w(l, s, X^u(s-) - \eta_i(s)e_i)) M_i(ds),
\end{aligned}$$

where M_i is the compensated Poisson process $M_i(s) := \pi_i(s) - \theta_i s$ and Inequality (5.15) follows from Proposition 2.5. Furthermore, we have (pathwise) equality in (5.15) if and only if $u = u^*$ λ - a.s.

Taking expectations on both sides, we obtain

$$\begin{aligned}
w(l, t, x) &\leq \tilde{J}(l, t, x, u) \\
&\quad + \sum_{i=1}^n \mathbb{E}\left[\int_t^T \left(w(l, s, X^u(s-)) - w(l, s, X^u(s-) - \eta_i(s)e_i) \right) M_i(ds) \right], \qquad (5.16)
\end{aligned}$$

with equality if and only if $u = u^*$ $\mathbb{P} \otimes \lambda$ - a.s.

It remains to show that the stochastic integrals in Inequality (5.16) are martingales. To this end, we compute

$$\begin{aligned}
&\mathbb{E}\left[\int_t^T |w(l, s, X^u(s-)) - w(l, s, X^u(s-) - \eta_i(s)e_i)|^2 ds \right] \\
&\quad = \mathbb{E}\left[\int_t^T |2X^u(s-)^\top C(l, s)e_i\eta_i(s) - \eta_i(s)^2 c_{i,i}(s, l)|^2 ds \right] \\
&\quad \leq \mathbb{E}\left[\int_t^T 2|2X^u(s-)^\top C(l, s)e_i\eta_i(s)|^2 ds \right] + \mathbb{E}\left[\int_t^T 2|\eta_i(s)^2 c_{i,i}(l, s)|^2 ds \right] \\
&\quad \leq 8\mathbb{E}\left[\int_t^T \| X^u(s-)^\top \|_2^2 \| C(l, s) \|_{2,2}^2 |\eta_i(s)|^2 ds \right] + 2\mathbb{E}\left[\int_t^T |\eta_i(s)|^4 |c_{i,i}(l, s)|^2 ds \right] \qquad (5.17) \\
&\quad \leq 8\lambda_{\max}^2 \left(\max_{s \in [t, T]} q(l, s)^2 \right) \mathbb{E}\left[\int_t^T \| X^u(s-) \|_2^4 ds \right]^{\frac{1}{2}} \mathbb{E}\left[\int_t^T \| \eta(s) \|_2^4 ds \right]^{\frac{1}{2}}
\end{aligned}$$

$$+ 2\lambda_{\max}^2 \left(\max_{s\in[t,T]} q(l,s)^2 \right) \mathbb{E}\left[\int_t^T \|\eta(s)\|_2^4 ds \right]$$
$$< +\infty$$

by Definition 2.3 (ii) and Lemma 5.5, where Inequality (5.17) follows from Hölder's inequality. As $\langle M_i \rangle(s) = \theta_i s$, this finishes the proof. □

5.4. Proof of Theorem 3.6

We start by computing the limits of the functions $p(l)$ and $q(l)$ given by Equations (5.7) and (5.8), respectively, (5.9).

LEMMA 5.6. *Let $t \in [0, T)$ and $p(l)$ and $q(l)$ as in Equations (5.7) and (5.8), respectively, (5.9). Then*

$$\lim_{l\to\infty} p(l,\cdot) = p(\cdot), \quad \lim_{l\to\infty} q(l,\cdot) = q(\cdot),$$

where p and q are given by

$$p(t) := \sqrt{\frac{\theta^2}{4} + \alpha d_{\min}} \coth\left(\sqrt{\frac{\theta^2}{4} + \alpha d_{\min}} (T-t) \right) - \frac{\theta}{2}, \tag{5.18}$$

$$q(t) := \sqrt{\alpha d_{\max}} \coth\left(\sqrt{\alpha d_{\max}} (T-t) \right) \tag{5.19}$$

for $\theta + \alpha d_{\min} > 0$, respectively, $\alpha d_{\max} > 0$ and

$$p(t) := q(t) := \frac{1}{T-t}, \tag{5.20}$$

for $\theta = \alpha d_{\min} = 0$, respectively, $\alpha d_{\max} = 0$. The convergence is compact and strictly increasing on $[0, T)$. Furthermore, $p(l, T), q(l, T) \nearrow \infty$ as $l \to \infty$.

Proof. Point-wise convergence and the formulae for the limits are straightforward by equations (5.7) and (5.8), respectively, by equation (5.9). Strict monotonicity follows from the fact that the initial values are strictly increasing in l. Finally, compact convergence follows from these observations by Dini's theorem. □

Proof of Theorem 3.6.

(i) Note first that Theorem 3.5 implies that for fixed $t \in (-\infty, T]$, $C(l, t)$ is increasing in l on $(l_0, +\infty)$ for l_0 as in equation (3.3) in the sense of Notation 3.1 (i). The existence of the element-wise limit of $(C(l, t))_{l>0}$ follows directly from this monotonicity and the boundedness by $\Lambda q(t) I$ for $q(t)$ as in equation (5.19), respectively, (5.20). Compact convergence follows by Dini's theorem due to the monotonicity. Finally, $p(l, T) \nearrow +\infty$ implies $\lim_{l\to\infty} c_{\min}(l, T) = +\infty$.

(ii) The inequalities in (3.8) follow directly from Lemma 5.6 and (i). Furthermore, the compact convergence of $C(l, t)$ (and $\tilde{C}(l, t)$) on $[0, T)$ and the fact that $C(l, \cdot)$ solves

the Matrix Differential Equation (3.7) implies

$$C(t) = \lim_{l\to\infty} C(l,t) = \lim_{l\to\infty} \int_0^t (C\Lambda^{-1}C + C\tilde{C}C - \alpha\Sigma)(l,s)ds + \lim_{l\to\infty} C(l,0)$$
$$= \int_0^t (C\Lambda^{-1}C + C\tilde{C}C - \alpha\Sigma)(s)ds + C(0)$$

and hence that C solves the Differential Equation (3.7) with boundary condition $\lim_{s\to T-} c_{\min}(s) = +\infty$ (as $\lim_{s\to T-} p(s) = +\infty$). □

5.5. Proof of Theorem 3.8

We start by proving the following bounds for $X^*(l)$ and X^*.

LEMMA 5.7. *Let $t \in [0,T)$ and $x \in \mathbb{R}^n$ be the portfolio position at time t.*

(i) Let $l > l_0$ for l_0 as in equation (3.3) and $s \in [t,T)$. Then (cf. Notation 5.3 (i))

$$X^*(l,s)^\top \Lambda X^*(l,s)$$
$$\leq x^\top \Lambda x \exp\left(-2\int_t^s p(l,u)du\right) \prod_{t\leq\tau_i\leq s} \frac{q(l,\tau_i)}{p(l,\tau_i)} \tag{5.21}$$
$$\leq x^\top \Lambda x \exp(\theta(s-t)) \frac{\left(T-s+\frac{2\lambda_{\max}}{2l+\theta\lambda_{\max}}\right)^2}{\left(T-t+\frac{2\lambda_{\max}}{2l+\theta\lambda_{\max}}\right)^2} \prod_{t\leq\tau_i\leq s} \frac{q(l,\tau_i)}{p(l,\tau_i)} \quad a.s., \tag{5.22}$$

where p and q are as in equations (5.7) and (5.8), respectively as in (5.9).

(ii) Let $s \in [t,T)$. Then

$$X^*(s)^\top \Lambda X^*(s) \leq x^\top \Lambda x \exp(\theta(s-t)) \frac{(T-s)^2}{(T-t)^2} \prod_{t\leq\tau_i\leq s} \frac{q(\tau_i)}{p(\tau_i)} \quad a.s.,$$

where p and q are as in equations (5.18) and (5.19), respectively, as in (5.20).

Proof. We prove (i); (ii) follows by exactly the same line of reasoning with the respective bounds. Let $i \in \mathbb{N}$. Observe that on $\{\tau_i < T\}$,

$$\frac{\partial}{\partial s}\left(X^*(l,s)^\top \Lambda X^*(l,s)\right) = -2X^*(l,s)^\top C(l,s) X^*(l,s)$$
$$\leq -2p(l,s) X^*(l,s)^\top \Lambda X^*(l,s)$$

for $s \in [\tau_i, \tau_{i+1}\wedge T)$ by Theorem 3.2. Gronwall's inequality implies

$$X^*(l,s)^\top \Lambda X^*(l,s) \leq X^*(l,\tau_i)^\top \Lambda X^*(l,\tau_i) \exp\left(-2\int_{\tau_i}^s p(l,r)dr\right),$$

in particular

$$X^*(l,(\tau_{i+1}\wedge T)-)^\top \Lambda X^*(l,(\tau_{i+1}\wedge T)-)$$
$$\leq X^*(l,\tau_i)^\top \Lambda X^*(l,\tau_i) \times \exp\left(-2\int_{\tau_i}^{\tau_{i+1}\wedge T} p(l,r)dr\right). \tag{5.23}$$

This implies (cf. Inequalities (5.12) and (5.13) in the proof of Lemma 5.4)

$$\begin{aligned} X^*(l,\tau_{i+1})^\top \Lambda X^*(l,\tau_{i+1}) &\le \frac{1}{p(l,\tau_{i+1})} X^*(l,\tau_{i+1}-)^\top C(l,\tau_{i+1}) X^*(l,\tau_{i+1}-) \\ &\le \frac{q(l,\tau_{i+1})}{p(l,\tau_{i+1})} X^*(l,\tau_{i+1}-)^\top \Lambda X^*(l,\tau_{i+1}-) \quad \text{a.s.} \end{aligned} \tag{5.24}$$

Using Inequalities (5.23) and (5.24), we obtain Inequality (5.21) inductively as before. Inequality (5.22) follows from

$$\frac{1}{T-r+\dfrac{2\lambda_{\max}}{2l+\theta\lambda_{\max}}} - \frac{\theta}{2} \le p(l,r)$$

which is a direct consequence of Corollary B.2. □

The main step in the proof of the admissibility of u^* is to show that the liquidation constraint holds (cf. Definition 2.3 (iv)). This is (in particular) accomplished in the following proposition.

PROPOSITION 5.8. *Let $t \in [0,T)$ and $x \in \mathbb{R}^n$ be the portfolio position at time t.*

(i) $X^*(l,\cdot) \xrightarrow{l\to\infty} X^*(\cdot)$ *a.s. compactly on* $[t,T)$.

(ii) $l \cdot \|X^*(l,T)\|_2^2 \xrightarrow{l\to\infty} 0$ *a.s. and in* L^1 *and* $X^*(l,T) \xrightarrow{l\to\infty} X^*(T) = \lim_{s\to T-} X^*(s) = 0$ *a.s.*

Proof.

(i) The spectral norm $\|\cdot\|_{2,2}$ is equivalent to the matrix maximum norm, and therefore the element-wise convergence results from Theorem 3.6 (i) transfer to the corresponding results for the spectral norm.

Let $t \le T' < T$. On $\{\tau_i < T'\}$, X^* and $X^*(l)$ solve the respective ordinary differential equations

$$X' = -\xi^*(\cdot, X) = -\Lambda^{-1} C X \quad \text{and} \quad X' = -\xi^*(l,\cdot,X) = -\Lambda^{-1} C(l) X$$

on the interval $[\tau_i, \tau_{i+1} \wedge T')$. We prove that the assertion follows from the continuous dependence of solutions of ordinary differential equations on the right-hand side and initial values. To this end, we first require some preliminary observations.

For $s \in [t,T']$ and $x, y \in \mathbb{R}^n$, we have

$$\|\Lambda^{-1}C(s)x - \Lambda^{-1}C(s)y\|_2 \le \max_{s\in[t,T']} \|Q(s)\|_{2,2} \|x-y\|_2 =: L\|x-y\|_2$$

(cf. Theorem 3.6 (ii)), that is, for all $s \in [t,T']$, $\xi^*(s,\cdot)$ is Lipschitz continuous on $\mathbb{R}^n$ with Lipschitz constant $L = L(T')$ independent of s. Furthermore, there exits a constant $K_1 \ge 1$ such that for $s \in [t,T']$, $\|\bar{C}(s)C(s)\|_{2,2} \le K_1$. We now show by induction on $i \in \mathbb{N}$ that for all $\epsilon > 0$, there exits an $l_i > l_0$ such that $l_i \ge l_{i-1}$ and for all $l \ge l_i$, $s \in [t, \tau_i \wedge T']$, $\|X^*(l,s) - X^*(s)\|_2 < \epsilon$.
The assertion is clear for $i = 0$. Let $i > 0$ and $\epsilon > 0$. By the induction hypothesis, there exists $l_{i-1} > l_0$ such that for $l > l_{i-1}$,

$$\|X^*(l,\tau_{i-1}) - X^*(\tau_{i-1})\|_2 < \epsilon \frac{e^{-L(T'-t)}}{6K_1}. \tag{5.25}$$

Note that on $\{\tau_{i-1} \geq T'\}$ the induction step is trivial. We therefore fix some $\omega \in \{\tau_{i-1} < T'\}$. Let now $l_i \geq l_{i-1}$ such that for $l > l_i$, $s \leq \tau_i \wedge T'$ (recall the uniform convergence of $(C(l,s))_l$ on $[t, T']$, Theorem 3.6 (i)),

$$\begin{aligned} \|\Lambda^{-1}C(l,s) - \Lambda^{-1}C(s)\|_{2,2} &\leq \|\Lambda^{-1}\|_{2,2}\|C(l,s) - C(s)\|_{2,2} \\ &< \epsilon \frac{e^{-L(T'-t)}}{6(T'-t)K_1} \end{aligned} \tag{5.26}$$

and

$$\|\bar{C}(l,s)C(l,s) - \bar{C}(s)C(s)\|_{2,2} < \frac{\epsilon}{3K^2} \tag{5.27}$$

for K as in Lemma 5.4 (ii). By the continuous dependence of solutions of systems of ordinary differential equations on the right-hand side and initial values, we have for $s \in [\tau_{i-1}, \tau_i \wedge T')$ (by Inequalities (5.25) and (5.26)),

$$\|X^*(l,s,\omega) - X^*(s,\omega)\|_2 \leq \left(\epsilon \frac{e^{-L(T'-t)}}{6K_1} + (T'-t)\epsilon \frac{e^{-L(T'-t)}}{6(T'-t)K_1}\right) e^{L(T'-t)} = \frac{\epsilon}{3K_1},$$

in particular

$$\|X^*(l,(\tau_i \wedge T')-) - X^*((\tau_i \wedge T')-)\|_2 \leq \frac{\epsilon}{3K_1}. \tag{5.28}$$

We can conclude by using the Inequalities (5.27) and (5.28):

$$\begin{aligned} &\|X^*(l,\tau_i(\omega)\wedge T',\omega) - X^*(\tau_i \wedge T',\omega)\|_2 \\ &\quad= \|X^*(l,(\tau_i(\omega)\wedge T')-,\omega) - \bar{C}(l,\tau_i(\omega)\wedge T')C(l,\tau_i(\omega)\wedge T')X^*(l,(\tau_i(\omega)\wedge T')-,\omega) \\ &\qquad - X^*((\tau_i(\omega)\wedge T')-,\omega) + \bar{C}(\tau_i(\omega)\wedge T')C(\tau_i(\omega)\wedge T')X^*((\tau_i(\omega)\wedge T')-,\omega)\|_2 \\ &\quad\leq \underbrace{\|X^*(l,(\tau_i(\omega)\wedge T')-,\omega) - X^*((\tau_i(\omega)\wedge T')-,\omega)\|_2}_{\leq \epsilon/(3K_1) \leq \epsilon/3 \text{ by Inequality (43)}} \\ &\qquad + \underbrace{\|X^*(l,(\tau_i(\omega)\wedge T')-,\omega)\|_2 \|\bar{C}(l,\tau_i(\omega)\wedge T')C(l,\tau_i(\omega)\wedge T') - \bar{C}(\tau_i(\omega)\wedge T')C(\tau_i(\omega)\wedge T')\|_{2,2}}_{\leq \epsilon/3 \text{ by Inequality (42)}} \\ &\qquad + \underbrace{\|\bar{C}(\tau_i(\omega)\wedge T')C(\tau_i(\omega)\wedge T')\|_2 \|X^*(l,(\tau_i(\omega)\wedge T')-,\omega) - X^*((\tau_i(\omega)\wedge T')-,\omega)\|_2}_{< \epsilon/3 \text{ by Inequality (43)}} \\ &\quad< \epsilon \end{aligned}$$

as required.

(ii) For fixed $\omega \in \Omega$, we have by Lemma 5.7 (i) that

$$\begin{aligned} &l \cdot \|X^*(l,T,\omega)\|_2^2 \\ &\leq \frac{1}{\lambda_{\min}} x^\top \Lambda x \exp(\theta(T-t)) \frac{l\left(\frac{2\lambda_{\max}}{2l+\theta\lambda_{\max}}\right)^2}{\left(T-t+\frac{2\lambda_{\max}}{2l+\theta\lambda_{\max}}\right)^2} \prod_{t \leq \tau_i \leq s} \frac{q(l,\tau_i)}{p(l,\tau_i)}. \end{aligned} \tag{5.29}$$

Furthermore,

$$\mathbb{E}\left[\prod_{t \leq \tau_i \leq T} \frac{q(l,\tau_i)}{p(l,\tau_i)}\right] < K_2 < +\infty$$

for some constant K_2 independent of l; thus for almost all $\omega \in \Omega$, there exists a constant $K_3(\omega)$ such that

$$\prod_{t \le \tau_i(\omega) \le T} \frac{q(l, \tau_i(\omega))}{p(l, \tau_i(\omega))} < K_3(\omega).$$

Therefore, Inequality (5.29) implies

$$\mathbb{E}\left[l \cdot \| X^*(l, T, \omega)\|_2^2\right] \xrightarrow{l \to \infty} 0 \quad \text{and} \quad l \cdot \| X^*(l, T, \omega)\|_2^2 \xrightarrow{l \to \infty} 0 \quad \text{a.s.}$$

Finally, Lemma 5.7 (ii) implies that $\lim_{s \to T-} \| X^*(s)\|_2 = 0$ a.s., finishing the proof. □

We are now able to prove that u^* is indeed an admissible liquidation strategy. The main step toward this goal is accomplished by Proposition 5.8 (ii). It remains thus to show that u^* fulfills the moment conditions in Definition 2.3 (ii).

Proof of Theorem 3.8. Definition 2.3 (i) and (iii) are clear and (iv) follows from Proposition 5.8 (ii).

Furthermore, we have

$$\mathbb{E}\left[\int_t^T \| \xi(s, X^*(s))\|_2^4 ds\right] \le \mathbb{E}\left[\int_t^T q(s)^4 \| X^*(s)\|_2^4 ds\right] < +\infty$$

by Lemma 5.4 (ii) and Proposition 5.8 (i). Finally, $\bar{C}(s) \le nC(s)^{-1}$ by Corollary 5.2 and thus, as $\frac{q(s)}{p(s)}$ admits a continuous extension on $[t, T]$, there exists a constant $\bar{K}$ independent of s such that

$$\| \bar{C}(s)C(s)\|_{2,2}^8 \le n^8 \| C^{-1}\|_{2,2}^8 \| C\|_{2,2}^8 \le \frac{n^8 \lambda_{\max}^8 q(s)^8}{\lambda_{\min}^8 p(s)^8} \le \bar{K}. \tag{5.30}$$

Using Lemma 5.4 (ii) and Proposition 5.8 (i) again, we can deduce from Inequality (5.30) that

$$\mathbb{E}\left[\int_t^T \| \eta(s, X^*(s-)\|_2^8 ds\right] \le \mathbb{E}\left[\int_t^T \| \bar{C}(s)C(s)\|_2^8 \| X^*(s-)\|_2^8 ds\right] < +\infty.$$

□

5.6. Proof of Theorem 3.9

We can directly deduce compact convergence of the optimal trading intensity in the primary venue from Theorem 3.6 (i) and Proposition 5.8 (i).

COROLLARY 5.9. *Let $t \in [0, T)$ and $x \in \mathbb{R}^n$ be the portfolio position at time t. Then*

$$\xi^*(l, \cdot, X^*(l, \cdot)) \longrightarrow \xi^*(\cdot, X^*(\cdot)) \quad \textit{a.s. compactly on } [t, T)$$

as $l \to \infty$.

This enables us to finally prove the main result of the article.

Proof of Theorem 3.9. We fix $t \in [0, T)$ and $x \in \mathbb{R}^n$. Note first that we have

$$v(t, x) \ge \lim_{l \to \infty} \tilde{v}(l, t, x). \tag{5.31}$$

For the converse inequality, let $A := \{\pi(T) = \pi(t)\}$ be the set of scenarios without any dark pool execution in $[0, T]$ and $K : (l_0, +\infty) \times \Omega \longrightarrow [0, +\infty]$ be the following cost function:

$$K(l, \omega) := \int_t^T (\xi^*(l, s, X^*(l, s, \omega))^\top \Lambda \xi^*(l, s, X^*(l, s, \omega)) + \alpha X^*(l, s, \omega)^\top \Sigma X^*(l, s, \omega)) ds + l \| X^*(l, T, \omega)\|_2^2.$$

Then $\mathbb{P}[A] > 0$ and for $\omega \in A$, $K(l, A) := K(l, \omega)$ is independent of the specific scenario ω almost surely. By optimality of $u^*(l)$ (Theorem 3.5), $K(l, A)$ is an upper bound for $K(l, \cdot)$ almost surely. As $\lim_{l\to\infty} \tilde{v}(l, t, x)$ is bounded and $\mathbb{P}[A] > 0$, there exists a constant K such that for all $l > l_0$, $K(l, A) \leq K$. By the dominated convergence theorem, this implies

$$\lim_{l\to\infty} \tilde{v}(l, t, x) = \lim_{l\to\infty} \mathbb{E}[K(l)] = \mathbb{E}[\lim_{l\to\infty} K(l)].$$

By Proposition 5.8 and Corollary 5.9, the limit in the last expression exists, so Fatou's lemma yields

$$\begin{aligned}
&\lim_{l\to\infty} \tilde{v}(l, t, x) \\
&\quad \geq \mathbb{E}\left[\int_t^T \lim_{l\to\infty} (\xi^*(l, s, X^*(l, s))^\top \Lambda \xi^*(l, s, X^*(l, s)) \right. \\
&\qquad \left. + \alpha X^*(l, s)^\top \Sigma X^*(l, s)) ds + \lim_{l\to\infty} l \cdot \| X^*(l, T)\|_2^2 \right] \\
&\quad = \mathbb{E}\left[\int_t^T \left(\xi^*(s, X^*(s))^\top \Lambda \xi^*(s, X^*(s)) + \alpha X^*(s)^\top \Sigma X^*(s)\right) ds\right] \\
&\quad \geq v(t, x). \qquad (5.32)
\end{aligned}$$

The Inequalities (5.31) and (5.32) establish that u^* solves the Optimization Problem (OPT) and that the value function is given by v. For uniqueness, let $u = (\xi, \eta), \tilde{u} = (\tilde{\xi}, \tilde{\eta}) \in \mathbb{A}(t, x)$, and $\mu \in (0, 1)$. We define the convex combination $\bar{u} = (\bar{\xi}, \bar{\eta})$:

$$\bar{\xi}(s) = \mu\xi(s) + (1 - \mu)\tilde{\xi}(s), \quad \bar{\eta}(s) = \mu\eta(s) + (1 - \mu)\tilde{\eta}(s)$$

for $s \in [t, T)$. Thus, $X^{\bar{u}}(s) = \mu X^u(s) + (1 - \mu) X^{\tilde{u}}(s)$ and $\bar{u} \in \mathbb{A}(t, x)$. Notice that

$$\mathbb{P} \otimes \lambda [u \neq \tilde{u}] > 0 \quad \text{implies} \quad \mathbb{P} \otimes \lambda \left[\xi \neq \tilde{\xi}\right] > 0 \qquad (5.33)$$

as else $\mathbb{P}[\lim_{s\to T-} X^u(s) \neq \lim_{s\to T-} X^{\tilde{u}}(s)] > 0$, a contradiction to Definition 2.3 (iv). Hence,

$$\begin{aligned}
J(t, x, \bar{u}) &= \mathbb{E}\left[\int_t^T f\left(\bar{\xi}(r), X^{\bar{u}}(r)\right) dr\right] \\
&\leq \mathbb{E}\left[\int_t^T \mu f(\xi(r), X^u(r)) + (1 - \mu) f\left(\tilde{\xi}(r), X^{\tilde{u}}(r)\right) dr\right] \qquad (5.34) \\
&= \mu J(t, x, u) + (1 - \mu) J(t, x, \tilde{u}),
\end{aligned}$$

where Inequality (5.34) follows from the convexity of f. We have equality in Inequality (5.34) if and only if $u = \tilde{u}\ \mathbb{P} \otimes \lambda$ - a.s. by the strict convexity of f in the first argument and (5.33). □

6. PROOFS OF THE RESULTS OF SECTION 4

6.1. Proofs of the Results of Section 4.1

We first require the following elementary result.

LEMMA 6.1. *Let* $0 < a < b$, $x > 0$. *Then*

$$0 > \frac{d}{dx}\frac{\sinh(ax)}{\sinh(bx)} > (a-b)\frac{\sinh(ax)}{\sinh(bx)}.$$

Proof. Note first that

$$\frac{d}{dx}\frac{\sinh(ax)}{\sinh(bx)} = \frac{a\cosh(ax)\sinh(bx) - b\cosh(bx)\sinh(ax)}{\sinh^2(bx)}. \tag{6.1}$$

The result follows from elementary calculus (see Kratz 2011 for details). □

Proof of Proposition 4.1. We let $t \in [0, T)$ and compute

$$\begin{aligned}
\frac{\partial}{\partial\theta}C(t;\theta) &= \frac{\Lambda\theta\coth\left(\frac{\tilde{\theta}}{2}(T-t)\right)}{2\tilde{\theta}} - \frac{\Lambda\theta(T-t)}{4\sinh^2\left(\frac{\tilde{\theta}}{2}(T-t)\right)} - \frac{\Lambda}{2} \\
&\leq \frac{\Lambda\theta\left(\cosh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\sinh\left(\frac{\tilde{\theta}}{2}(T-t)\right) - \frac{\tilde{\theta}}{2}(T-t) - \sinh^2\left(\frac{\tilde{\theta}}{2}(T-t)\right)\right)}{2\tilde{\theta}\sinh^2\left(\frac{\tilde{\theta}}{2}(T-t)\right)} \\
&< 0
\end{aligned}$$

for $\theta > 0$ since

$$\begin{aligned}
&\cosh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\sinh\left(\frac{\tilde{\theta}}{2}(T-t)\right) - \frac{\tilde{\theta}}{2}(T-t) - \sinh^2\left(\frac{\tilde{\theta}}{2}(T-t)\right) \\
&= \sinh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\left(\cosh\left(\frac{\tilde{\theta}}{2}(T-t)\right) - \sinh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\right) - \frac{\tilde{\theta}}{2}(T-t) \\
&= \left(\frac{1-\exp(-\tilde{\theta}(T-t))}{2}\right) - \frac{\tilde{\theta}}{2}(T-t) < 0
\end{aligned}$$

(note that $\frac{1}{2}(1-\exp(-2x)) - x < 0$ for $x > 0$). This establishes the first and the second assertion directly; the third assertion follows from the first equality in equation (4.1). For

the proof of (iv), we note first that

$$\mathbb{E}[X^*(t;\theta)] = \mathbb{P}[\pi(t)=0]\cdot \tilde{X}(t;\theta) = \frac{\sinh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\exp\left(-\frac{\theta}{2}t\right)}{\sinh\left(\frac{\tilde{\theta}}{2}T\right)}x. \tag{6.2}$$

We compute for $\theta > 0$,

$$\begin{aligned}\frac{\partial}{\partial\theta}\mathbb{E}[X^*(t;\theta)] &= \frac{x}{\sinh^2\left(\frac{\tilde{\theta}}{2}T\right)}\left(\frac{\theta(T-t)}{2\tilde{\theta}}\sinh\left(\frac{\tilde{\theta}}{2}T\right)\cosh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\exp\left(-\frac{\theta}{2}t\right)\right.\\ &\quad -\frac{t}{2}\sinh\left(\frac{\tilde{\theta}}{2}T\right)\sinh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\exp\left(-\frac{\theta}{2}t\right) - \frac{\theta T}{2\tilde{\theta}}\cosh\left(\frac{\tilde{\theta}}{2}T\right)\\ &\quad \left.\times\ \sinh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\exp\left(-\frac{\theta}{2}t\right)\right)\\ &< \frac{\theta\exp\left(-\frac{\theta}{2}t\right)x}{2\tilde{\theta}\sinh^2\left(\frac{\tilde{\theta}}{2}T\right)}\left((T-t)\sinh\left(\frac{\tilde{\theta}}{2}T\right)\cosh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\right.\\ &\quad \left.-\ T\cosh\left(\frac{\tilde{\theta}}{2}T\right)\sinh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\right) < 0\end{aligned}$$

by Lemma 6.1 (cf. also equation (6.1)), finishing the proof of (iv).

We have

$$\mathbb{E}[X^*(t;\theta)^2] = \mathbb{P}[\pi(t)=0]\cdot \tilde{X}(t;\theta)^2 = \frac{\sinh\left(\frac{\tilde{\theta}}{2}(T-t)\right)}{\sinh\left(\frac{\tilde{\theta}}{2}T\right)}x^2.$$

This term is differentiable and strictly decreasing in θ by Lemma 6.1 (note that $\tilde{\theta}$ is strictly increasing in θ). Thus by Fubini's theorem,

$$\frac{\partial}{\partial\theta}\mathbb{E}\left[\int_0^T X^*(t;\theta)^2 dt\right] = \frac{\partial}{\partial\theta}\int_0^T \mathbb{E}[X^*(t;\theta)^2]dt = \int_0^T \frac{\partial}{\partial\theta}\mathbb{E}[X^*(t;\theta)^2]dt < 0,$$

establishing (v).

Finally, we note that

$$\mathbb{E}[\xi^*(t,\, X^*(t;\theta);\theta)^2] = \mathbb{P}[\pi(t)=0]\cdot\frac{C(t;\theta)^2}{\Lambda^2}\tilde{X}(t;\theta)^2 = \frac{C(t;\theta)^2}{\Lambda^2}\mathbb{E}[X^*(t;\theta)^2]$$

by equation (6.2). This term is differentiable and strictly decreasing in θ as both terms are positive and strictly increasing in θ. Similarly as before, we deduce (vi). □

Proof of Proposition 4.2. (i) follows directly from the cost functional J. For (ii), we compute

$$\frac{\partial}{\partial \Lambda}\xi^*(t;\Lambda) = \frac{1}{2}\frac{\partial \tilde{\theta}}{\partial \Lambda}\left(\frac{\cosh\left(\frac{\tilde{\theta}}{2}(T-t)\right)\sinh\left(\frac{\tilde{\theta}}{2}(T-t)\right) - \frac{\tilde{\theta}}{2}(T-t)}{\sinh^2\left(\frac{\tilde{\theta}}{2}(T-t)\right)}\right) < 0$$

as $\dfrac{\partial \tilde{\theta}}{\partial \Lambda} < 0$. Monotonicity of $\tilde{X}$ follows as in the proof of Proposition 4.1 (iii). A similar calculation yields

$$\frac{\partial}{\partial(\alpha\Sigma)}\xi^*(t;\alpha\Sigma) > 0$$

as $\frac{\partial \tilde{\theta}}{\partial(\alpha\Sigma)} > 0$, finishing the proof. □

6.2. Proofs of the Results of Section 4.2

Proof of Proposition 4.3. We have $v(t, x - \eta e_i) = (x - \eta e_i)^\top C(t)(x - \eta e_i)$, which can easily be seen to be minimized by $\eta_1^*(t, x)$. □

Proof of Proposition 4.5. We prove the case $x_1, x_2 > 0, \rho < 0$, that is, x is well diversified. Let u^* be the optimal strategy for the initial portfolio position $(x_1, -x_2)^\top$. For $(x_1, x_2)^\top$, we define the strategy $u \in \mathbb{A}(t, (x_1, x_2)^\top)$ in such a way that for $i = 1, 2, s \geq t$, $X_i^u(s) \geq 0$ and $|X_i^u(s)| = |X_i^{u^*}(s)|$; this is achieved by changing the signs of the trading intensities and by adjusting the dark pool orders appropriately if necessary. In particular, we have $|\xi_i(s, X^u(s))| = |\xi_i^*(s, X^*(s))|$ and both strategies yield the same impact costs. On the other hand, the risk costs of u are strictly smaller as $\rho < 0$. Hence,

$$v(t, (x_1, -x_2)^\top) = J(t, (x_1, -x_2)^\top, u^*) > J(t, (x_1, x_2)^\top, u) \geq v(t, (x_1, x_2)^\top)$$

as desired. The remaining cases follow accordingly. □

Proof of Proposition 4.6. We prove the case $x_1, x_2 > 0, \rho < \tilde{\rho} < 0$ and proceed similarly as in the proof of Proposition 4.5. Let $\tilde{u}$ be the optimal strategy for $\tilde{\rho}$. For ρ, we define the strategy $u \in \mathbb{A}(t, x)$ in such a way that for $i = 1, 2, s \geq t$, $X_i^u(s) \geq 0$ and $|X_i^u(s)| = |X_i^{\tilde{u}}(s)|$. As $\rho < \tilde{\rho} < 0$, this yields

$$v(t, x; \rho) \leq J(t, x, u; \rho) < J(t, x, \tilde{u}; \tilde{\rho}) = v(t, x; \tilde{\rho}).$$

The remaining cases follow in the same way. □

For the proof of Proposition 4.7, we first require the following symmetry results of the value function.

LEMMA 6.2. *Let $t \in [0, T)$, $x \in \mathbb{R}^2$, and $\rho \in [-1, 1]$. Then $v(t, x; \rho) = v(t, -x, ; \rho)$ and $v(t, (x_1, x_2)^\top; \rho) = v(t, (x_1, -x_2)^\top; -\rho)$.*

Proof. We have $J(t, x, u) = J(t, -x, -u)$ and hence the first assertion follows. The second assertion follows from $J(t, (x_1, x_2)^\top, (u_1, u_2)^\top; \rho) = J(t, (x_1, -x_2)^\top, (u_1, -u_2)^\top; -\rho)$. □

Proof of Proposition 4.7.

(i) The first assertion follows directly from $C(t) > 0$. The second assertion follows from Lemma 6.2 as for $i = 1, 2$, $c_{i,i}(t, \rho) = v(t, e_i; \rho) = v(t, e_i; -\rho) = c_{i,i}(t, -\rho)$. We directly deduce the third assertion as

$$\begin{aligned} c_{1,1}(t;\rho) + c_{2,2}(t;\rho) + 2c_{1,2}(t,\rho) &= v(t,(1,1)^\top;\rho) = v(t,(1,-1)^\top;-\rho) \\ &= c_{1,1}(t;\rho) + c_{2,2}(t;\rho) - 2c_{1,2}(t;-\rho). \end{aligned}$$

Finally, it follows for $\rho < 0$ ($\rho > 0$) by Proposition 4.5 that

$$\begin{aligned} c_{1,1}(t;\rho) + c_{2,2}(t;\rho) + 2c_{1,2}(t;\rho) &= v(t,(1,1)^\top;\rho) < (>) v(t,(1,1)^\top;-\rho) \\ &= c_{1,1}(t;\rho) + c_{2,2}(t;\rho) - 2c_{1,2}(t;\rho) \end{aligned}$$

and therefore $c_{1,2}(t;\rho) < 0$ ($c_{1,2}(t;\rho) > 0$).

(ii) For the monotonicity of $c_{1,1}(t;\cdot)$, we let $\rho < \tilde{\rho} < 0$ and proceed similarly as in the proofs of Propositions 4.5 and 4.6. Let $\tilde{u}$ be the optimal strategy for $\tilde{\rho}$. For ρ, we define the strategy $u \in \mathbb{A}(t, e_1)$ in such a way that for $i = 1, 2$, $s \geq t$, $X_i^u(s) \geq 0$ and $|X_i^u(s)| = |X_i^{\tilde{u}}(s)|$. As $\rho < \tilde{\rho} < 0$, this yields $J(t, e_1, u; \rho) \leq J(t, e_1, \tilde{u}; \tilde{\rho})$ with equality if and only if $X_2^u(s) = 0$ a.s. However, we have $\xi_2(t, e_1) = 1/\lambda_2 c_{1,2}(t; \tilde{\rho}) < 0$ and hence $X_2^{\tilde{u}} > 0$ in some neighborhood of t with positive probability. Thus,

$$c_{1,1}(t;\rho) = v(t, e_1; \rho) \leq J(t, e_1, u; \rho) < J(t, e_1, \tilde{u}; \tilde{\rho}) = v(t, e_1, \tilde{u}; \tilde{\rho}) = c_{1,1}(t; \tilde{\rho}).$$

The monotonicity of $c_{2,2}$ and the case $\rho > 0$ follow by the same line of reasoning.

Before we proceed, we remark that all symmetry properties and the monotonicity of $c_{1,1}$ and $c_{2,2}$ in ρ also hold for $C(l)$ ($l \geq l_0$) with exactly the same proofs. We now prove monotonicity of $c_{1,2}(l, t; \rho)$ ($l \geq l_0$) first; monotonicity of $c_{1,2}(t; \rho)$ then follows directly from the fact that $\lim_{l\to\infty} c_{1,2}(l, t; \rho) = c_{1,2}(t; \rho)$. A straightforward computation confirms that $c_{1,2}(l)$ fulfills the following scalar initial value problem

$$\begin{aligned} &\frac{\partial}{\partial t} c_{1,2}(l, t; \rho) \\ &\qquad = c_{1,2}(l, t; \rho) \left(\frac{c_{1,1}(l, t; \rho)}{\lambda_1} + \frac{c_{2,2}(l, t; \rho)}{\lambda_2} + \theta_1 + \theta_2 \right) - \alpha\sigma_1\sigma_2\rho, \\ &c_{1,2}(l, T; \rho) = 0 \end{aligned} \tag{6.3}$$

By the continuous differentiable dependence of $c_{1,2}$ on the parameter ρ, we can exchange differentiation with respect to t and ρ and obtain the following initial value problem for $\frac{\partial c_{1,2}}{\partial \rho}$:

$$\frac{\partial}{\partial t} \frac{\partial c_{1,2}}{\partial \rho}(l, t; \rho) = \frac{\partial c_{1,2}}{\partial \rho}(l, t; \rho) f(l, t; \rho) - g(l, t; \rho), \qquad \frac{\partial c_{1,2}}{\partial \rho}(l, T; \rho) = 0$$

for

$$\begin{aligned} f(l, t; \rho) &:= \frac{c_{1,1}(l, t; \rho)}{\lambda_1} + \frac{c_{2,2}(l, t; \rho)}{\lambda_2} + \theta_1 + \theta_2, \\ g(l, t; \rho) &:= \alpha\sigma_1\sigma_2 - c_{1,2}(l, t; \rho) \left(\frac{1}{\lambda_1} \frac{\partial c_{1,1}(l, t; \rho)}{\partial \rho} + \frac{1}{\lambda_2} \frac{\partial c_{2,2}(l, t; \rho)}{\partial \rho} \right); \end{aligned}$$

this implies

$$\frac{\partial c_{1,2}}{\partial \rho}(l,t;\rho) = \exp\left(\int_T^t f(l,s;\rho)ds\right)\int_t^T g(l,r;\rho)\exp\left(\int_r^T f(l,s;\rho)ds\right)dr > 0$$

as $g(l,t;\rho) \geq 0$ by (i) and the monotonicity of $c_{1,1}(l,t;\cdot)$ and $c_{2,2}(l,t;\cdot)$. □

Proof of Proposition 4.8.

(i) We prove the assertion for $x_1, x_2 > 0$ and $\rho < 0$. Let $j \in \mathbb{N}$. We assume that $X_i^*(\tau_j) > 0$ $(i = 1, 2)$ on $\{\tau_j < T\}$ (cf. Notation 5.3). We compare the initial value problem for the controlled process on $[\tau_j, \tau_{j+1} \wedge T)$ (cf. equations (4.3)) with the case $\rho = 0$; as $c_{1,2}(s;\rho) < c_{1,2}(s;0) = 0$ and $0 < c_{i,i}(s;\rho) < c_{i,i}(s;0)$ by Proposition 4.7, $X_i^*(s;\rho) \geq X_i^*(s;0)$ for $s \in [\tau_j, \tau_{j+1} \wedge T)$. For $\rho = 0$ the two components of X^* evolve independently according to the results of Section 4.1. In particular, the optimal asset position remains positive in $[\tau_j, \tau_{j+1} \wedge T)$. It follows that $X_i^*(s;\rho) > 0$ for all $s < \tau_{j+1} \wedge T$ and that $X_i^*(\tau_{j+1}-) > 0$ on $\{\tau_{j+1} < T\}$. Applying equations (4.4) and Proposition 4.7, we obtain on $\{\tau_{j+1} < T\}$,

$$X_1^*(\tau_{j+1}-) - \eta_1(\tau_{j+1}, X^*(\tau_{j+1}-)) = -\frac{c_{1,2}(\tau_{j+1};\rho)}{c_{1,1}(\tau_{j+1};\rho)}X_2^*(\tau_{j+1}-) > 0$$

and the respective result for X_2^*. The assertion now follows by induction on j. The remaining cases follow accordingly.

(ii) We assume $x_1, x_2, \rho > 0$. It is clear that $\tau > t$ a.s. The result follows from equations (4.4) and Proposition 4.7 as for $t \leq s < \tau$,

$$\eta_1(s, X^*(s)) = X_1^*(s) + \frac{c_{1,2}(s;\rho)}{c_{1,1}(s;\rho)}X_2^*(s) > X_1^*(s).$$

□

Proof of Proposition 4.9. The assertions follow directly from equations (4.3) and (4.4) by applying Proposition 4.7. □

For the proof of Proposition 4.10, we require the following results about the monotonicity of $c_{i,i}$ and $c_{1,2}$ in t.

LEMMA 6.3. *For $i = 1, 2$ and $t \in [0, T)$, $c_{i,i}(t)$ is increasing in t. $c_{1,2}(t)$ is decreasing in t if $\rho > 0$ and increasing in t if $\rho < 0$.*

Proof. For fixed $l \geq l_0$, we consider the Initial Value Problem (3.1); its solution $C(l)$ satisfies $C(l)' \geq Q(l)'$, where $Q(l)$ solves the initial value problem $Q' = 1/\lambda_{\min} Q^2 - \alpha\sigma_{\min} I$, $Q(T) = lI$ (cf. Notation 3.1). As in the proof of Proposition 3.4, $Q(l)$ can be computed explicitly with $Q(l)'(t) > 0$ for l large enough, say $l \geq l_1$; in particular $c'_{i,i}(l,t) > 0$ and therefore $c'_{i,i}(t) \geq 0$.

For the monotonicity of $c_{1,2}$ we assume $\rho > 0$ and consider the Initial Value Problem (6.3) for $c_{1,2}(l,t)$. Note first that for all l, $c'_{1,2}(l,T) = -\alpha\rho\sigma_1\sigma_2 < 0$. Furthermore,

$$c''_{1,2}(l,t) = c'_{1,2}(l,t)\left(\frac{c_{1,1}(l,t)}{\lambda_1} + \frac{c_{2,2}(l,t)}{\lambda_2} + \theta_1 + \theta_2\right) + \underbrace{c_{1,2}(l,t)\left(\frac{c'_{1,1}(l,t)}{\lambda_1} + \frac{c'_{2,2}(l,t)}{\lambda_2}\right)}_{\geq 0 \text{ for } l \geq l_1}.$$

This implies $c'_{1,2}(l,t) < 0$ for $l \geq l_1$ and hence $c'_{1,2}(t) \leq 0$. The case $\rho < 0$ follows accordingly. □

Proof of Proposition 4.10. We assume $\rho > 0$ and that the dark pool order for the first asset is executed at time τ_1, that is, $X_1^*(\tau_1) = -\frac{c_{1,2}(\tau_1)}{c_{1,1}(\tau_1)} X_2^*(\tau_1)$; we further assume $X_2^*(\tau_1) > 0$. We have $\xi_1^*(\tau_1) = 0$ and $\xi_2^*(\tau_1) > 0$. As $c_{1,2}(s)/c_{i,i}(s)$ is decreasing in s by Lemma 6.3, we have $\xi_2^*(s) > 0$ and $\xi_1^*(s) < 0$ (in particular $\eta_2^*(s) > 0$ and $\eta_1^*(\tau_1) < 0$) until the next jump time of π. The result follows inductively as the number of jumps is almost surely finite. The proof for the remaining cases is analog. □

Proof of Proposition 4.11. The first assertion follows directly from the definition of the cost functional J. For the second assertion, let $\lambda_i < \tilde{\lambda}_i$. Then

$$\begin{aligned}\frac{v(t,x;\lambda_i)}{\lambda_i} &= \mathbb{E}\left[\int_t^T \left(\xi_i^*(s)^2 + \sum_{j\neq i}\frac{\lambda_j}{\lambda_i}\xi_j^*(s)^2 + \frac{\alpha}{\lambda_i} X^*(s)^\top \Sigma X^*(s)\right) ds\right] \\ &\geq \mathbb{E}\left[\int_t^T \left(\xi_i^*(s)^2 + \sum_{j\neq i}\frac{\lambda_j}{\tilde{\lambda}_i}\xi_j^*(s)^2 + \frac{\alpha}{\tilde{\lambda}_i} X^*(s)^\top \Sigma X^*(s)\right) ds\right] \geq \frac{v(t,x;\tilde{\lambda}_i)}{\tilde{\lambda}_i}.\end{aligned}$$

□

Proof of Proposition 4.12.

(i) We prove the case $x_1, x_2 > 0$, $\rho < 0$, that is, x is well diversified and modify the proof of Proposition 4.5. Let u^* be the optimal strategy for the initial portfolio position $(x_1, -x_2)^\top$. For $(x_1, x_2)^\top$, we can define a strategy $u \in \mathbb{A}(t, (x_1, x_2)^\top)$ with the following properties: for $s \leq \tau_1$ (where τ_1 is the first jump time of π),

$$\begin{aligned} |\xi_i(s)| = |\xi_i^{u^*}|, \qquad |X_i^u(s)| \leq |X_i^{u^*}(s)|, \\ \operatorname{sgn}(\xi_1^{u^*}(s)) = \operatorname{sgn}(\xi_2^{u^*}(s)) \;\Rightarrow \operatorname{sgn}(\xi_1(s)) = \operatorname{sgn}(\xi_2(s)), \\ \operatorname{sgn}(X_1^{u^*}(s)) = \operatorname{sgn}(X_2^{u^*}(s)) \;\Rightarrow \operatorname{sgn}(X_1^u(s)) = \operatorname{sgn}(X_2^u(s)).\end{aligned} \tag{6.4}$$

u needs to be defined carefully by considering all possible combinations of the signs of the trading intensities; note that it can be necessary to change the signs of the positions. We adjust the dark pool orders in such a way that $X_i^u(\tau_1) \in \{\gamma X_i^{u^*}(\tau_1), -\gamma X_i^{u^*}(\tau_1)\}$ for some $\gamma \leq 1$ (cf. the inequality in (6.4)). For $s \geq \tau_1$, we proceed similarly as before by defining $\xi_i(s) = \gamma \xi_i^{u*}(s)$ (or $\xi_i(s) = -\gamma \xi_i^{u*}(s)$) and the dark pool orders as before.

This ensures that the impact costs of u are less or equal than the impact costs of u^* while the risk costs are strictly smaller; hence the assertion follows inductively. The remaining cases follow by the same line of reasoning.

(ii), (iii) As in the proof of Proposition 4.7 (i), we obtain that $\operatorname{sgn}(c_{1,2}(t;\rho,\lambda_{1,2})) = \operatorname{sgn}(\lambda_{1,2}) = \operatorname{sgn}(\rho)$. The assertions follow from the fact that η^* is as in equations (4.4) (also in the presence of cross price impact).

□

Proof of Proposition 4.13. Let $\theta_i \leq \tilde{\theta}_i$. For any matrix $C > 0$, we have $\tilde{C}(\theta_i) \geq \tilde{C}(\tilde{\theta}_i)$ (cf. equation (3.2)). Similarly as in the proof of Theorem 3.2, we obtain that the respective solutions of the Initial Value Problem fulfill $C(l, t; \tilde{\theta}_i) \leq C(l, t; \theta_i)$ for l large enough. The assertion follows by taking the limit $l \to \infty$. □

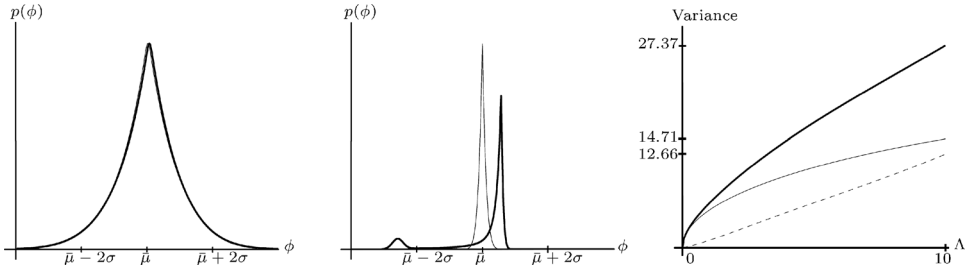

FIGURE A.1. The left and the middle picture show the distribution of the proceeds of selling ϕ (thick solid lines) and the approximation by a Gaussian mixture distribution with common mean $\bar{\phi}$ (thin solid lines). The right picture shows the overall variance $\mathrm{Var}(\phi)$ (thick solid line) and the modeled market variance $\mathrm{Var}(\bar{\phi})$ (thin solid line) depending on the temporary impact parameter Λ; the dashed line denotes the variance of liquidation proceeds due to dark pool execution risk. $x = 1$, $T = 1$, $\theta = 2.3$, $\Sigma = 100$, $\alpha = 0$ (left and middle picture), respectively, $\alpha = 1$ (right picture) and $\Lambda = 1$ (left picture), $\Lambda = 100$ (middle picture), respectively, $\Lambda \in (0, 10]$ (right picture).

APPENDIX A: MARKET AND EXECUTION RISK

Our model set-up penalizes the variance introduced by market risk but not the execution risk of the dark pool orders. Here, we investigate the impact of this model choice by analyzing the single-asset case (cf. Section 4.1). To this end, we consider the distribution of the proceeds of selling $\phi = \phi(0, x, u^*)$ (cf. Section 2.4) when the optimal liquidation strategy u^* is applied.

Liquidity in the dark pool is found at time t for the first time with probability density $w_t := \theta \exp(-\theta t)$.[25] Conditional on this event, the trades $(\xi^*(s))_{0 \le s \le t}$ and the asset positions $(X^*(s))_{0 \le s \le t}$ are deterministic (i.e., independent of the market price evolution) and are zero after time t. The proceeds are hence normally distributed with mean $\mu_t := \int_0^t \Lambda \xi^*(s)^2 ds$ and variance $\sigma_t^2 := \int_0^t \Sigma X^*(s)^2 ds$. Mixing these normal distributions with weights w_t gives the distribution of ϕ, that is, the distribution of ϕ is a Gaussian mixture distribution.

In the left and middle picture of Figure A.1, we illustrate the exact distribution of ϕ (thick solid lines) for two different parameter choices. In the left picture, the price impact Λ is small compared to the exogenous volatility Σ of the asset price; in this case, ϕ follows a unimodal distribution that is driven by market risk. In the middle picture, the price impact is large compared to the volatility and ϕ follows a bimodal distribution where the two modes correspond to the cases of presence and absence of liquidity in the dark pool.

We compare the distribution of ϕ (which mixes the distributions $N(\mu_t, \sigma_t)$) to another Gaussian mixture distribution $\bar{\phi}$ (illustrated by the thin solid lines) with the same weights w_t and the same variances σ_t^2 as in the exact distribution of ϕ; the means of the normal distributions however are changed from μ_t to $\bar{\mu} := \int_0^T w_s \mu_s ds$, that is, the same mean is applied to all underlying Gaussian distributions. This mixture of the normal distributions

[25]The case $t = T$ corresponds to not finding any liquidity in the dark pool and occurs with discrete probability $\exp(-\theta T)$.

$N(\bar{\mu}, \sigma_t)$ erased the execution risk component (i.e., the differences in μ_t), but captures the market risk component which is penalized in our model:

$$\mathrm{Var}(\bar{\phi}) = \mathbb{E}\left[\int_0^T \Sigma\, X^*(s)^2 ds\right].$$

As shown in the left picture of Figure A.1, the thick solid line (exact distribution) and the thin solid line are similar if the temporary price impact is small compared to the variance; on the other hand, they are notably different if the temporary price impact is large (middle picture). In the first case, the variance of the exact distribution is only marginally larger than the market risk component, but in the latter case, it is several orders of magnitude larger. The right picture of Figure A.1 illustrates this dependency of overall variance $\mathrm{Var}(\phi)$ (thick solid line) and modeled market variance $\mathrm{Var}(\bar{\phi})$ (thin solid line) on the size of the temporary impact parameter Λ. Depending on the relative sizes of Λ and Σ, our model might hence be missing an important component of overall variance. In reality however, price impact is almost always significantly smaller than market volatility. This is reflected, for example, in the difficulty of measuring market impact and evaluating the performance of trade execution algorithms (see, e.g., Almgren et al. 2005, Sofianos and Jeria 2008a, and Sofianos and Jeria 2008b). Under such circumstances our model appears to capture the primary risk component.

APPENDIX B: RICCATI MATRIX DIFFERENTIAL EQUATIONS

In this section, we state a well-known comparison result about matrix Riccati equations in the form in which we apply it in the proof of Theorem 3.2. A standard textbook is the one by Reid (1972). A proof for the specific form of the theorem can, for example, be found in Kratz (2011).

THEOREM B.1. *Let $A(t)$, $B_P(t)$, $C_P(t)$, $B_Q(t)$, $C_Q(t) \in \mathbb{R}^{n\times n}$ be piecewise continuous on $\mathbb{R}$. Furthermore, let $B_P(t)$, $C_P(t)$, $B_Q(t)$, $C_Q(t)$ $(t \in \mathbb{R})$, and S_P, $S_Q \in \mathbb{R}^{n\times n}$ be symmetric. Let $t_0 > t_1 \geq -\infty$ and*

$$S_Q \leq S_P, \quad 0 \leq B_Q(\cdot) \leq B_P(\cdot), \quad C_P(\cdot) \leq C_Q(\cdot)$$

on $(t_2, t_0]$. Assume that the initial value problem

$$P' = -A^\top P - PA - PB_P P + C_P, \quad P(t_0) = S_P$$

possesses a solution P on $(t_1, t_0]$. Then the initial value problem

$$Q' = -A^\top Q - QA - QB_Q Q + C_Q, \quad Q(t_0) = S_Q$$

possesses a solution Q on $(t_1, t_0]$ and $P(t) \geq Q(t)$ on $(t_1, t_0]$.

We apply the theorem to scalar Riccati equations with constant coefficients and obtain a useful lower bound for their solution.

COROLLARY B.2. *Let y be the solution of the scalar initial value problem $y' = y^2 + ay - b$, $y(T) = c$, where $a, b \geq 0$, $c > 0$, $b < c^2 + ac$, and $d := a^2/4 + b > 0$. Then for $t \in (-\infty, T]$,*

$$y(t) \geq \frac{1}{T - t + \dfrac{1}{c + a/2}} - \frac{a}{2}.$$

Proof. As $d > 0$, we have that the solution z of the initial value problem $z' = z^2 - d$, $z(T) = c + a/2$ fulfills

$$z(t) \geq \frac{1}{T - t + \dfrac{1}{c + a/2}} \quad \text{on } (-\infty, T]$$

(cf. Theorem B.1; compare z with the solution of $f' = f^2$, $f(T) = c + \frac{a}{2}$). The assertion follows directly from the fact that $y(t) = z(t) - a/2$. □

references

ALFONSI, A., A. FRUTH, and A. SCHIED (2010): Optimal Execution Strategies in Limit Order Books with General Shape Functions. *Quant. Finance* 10(2), 143–157.

ALMGREN, R. (2003): Optimal Execution with Nonlinear Impact Functions and Trading-Enhanced Risk. *Appl. Math. Finance* 10(1), 1–18.

ALMGREN, R., and N. CHRISS (2001): Optimal Execution of Portfolio Transactions. *J. Risk* 3(2), 5–39.

ALMGREN, R., and J. LORENZ (2007): Adaptive Arrival Price, in *Algorithmic Trading III: Precision, Control, Execution*, B. Bruce, ed. New York: Institutional Investor Inc., pp. 59–66.

ALMGREN, R., C. THUM, E. HAUPTMANN, and H. LI (2005): Direct Estimation of Equity Market Impact. *Risk* 18(7), 57–62.

APPLEBAUM, D. (2004): *Lévy Processes and Stochastic Calculus*. Cambridge: Cambridge University Press.

BAYRAKTAR, E., and M. LUDKOVSKI (2014): Liquidation in Limit Order Books with Controlled Intensity. *Math. Finance* 24(4), 627–650.

BERNSTEIN, D. S. (2005): *Matrix Mathematics: Theory, Facts, and Formulas with Application to Linear Systems Theory*. Princeton: Princeton University Press.

BERTSIMAS, D., and A. LO (1998): Optimal Control of Execution Costs. *J. Finan. Markets* 1(1), 1–50.

BOUCHARD, B., and N.-M. DANG (2013): Generalized Stochastic Target Problems for Pricing and Partial Hedging under Loss Constraints—Application in Optimal Book Liquidation. *Finance Stoch.* 17(1), 31–72.

BOUCHARD, B., R. ELIE, and N. TOUZI (2009): Stochastic Target Problems with Controlled Loss. *SIAM J. Control Optim.* 48(5), 3123–3150.

CARLIN, B. I., M. SOUSA LOBO, and S. VISWANATHAN (2007): Episodic Liquidity Crises: Cooperative and Predatory Trading. *J. Finance* 65(5), 2235–2274.

CARRIE, C. (2008): Illuminating the New Dark Influence on Trading and U.S. Market Structure. *J. Trading* 3(2), 40–55.

CONRAD, J., K. M. JOHNSON, and S. WAHAL (2003): Institutional Trading and Alternative Trading Systems. *J. Finan. Econ.* 70(1), 99–134.

COPPEL, W. A. (1971): *Disconjugacy*. Lecture Notes in Mathematics, Vol. 220. Berlin: Springer.

DEGRYSE, H., M. VAN ACHTER, and G. WUYTS (2009a): Dynamic Order Submission Strategies with Competition between a Dealer Market and a Crossing Network. *J. Finan. Econ.* 91(3), 319–338.

DEGRYSE, H., M. VAN ACHTER, and G. WUYTS (2009b): Shedding Light on Dark Liquidity Pools. *Institutional Investor* 2009(1), 147–155.

DÖNGES, J., and F. HEINEMANN (2013): Crossing Network versus Dealer Market: Unique Equilibria in the Allocation of Order Flow. *Eur. Econ. Rev.* 62, 41–57.

FONG, K., A. MADHAVAN, and P. L. SWAN (2004): Upstairs, Downstairs: Does the Upstairs Market Hurt the Downstairs? Preprint. Available at: available online at: http://www.finance.nsysu.edu.tw/SFM/15thSFM/14thSFM/13thSFM/12thSFM/167970922.pdf. Accessed April 26, 2013.

GANCHEV, K., M. KEARNS, Y. NEVMYVAKA, and J. WORTMAN VAUGHAN (2010): Censored Exploration and the Dark Pool Problem. *Commun. ACM* 53(5), 99–107.

GLOSTEN, L. R. and P. R. MILGROM (1985): Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders. *J. Finan. Econ.* 14(1), 71–100.

GROSSMAN, S. J., and M. H. MILLER (1988): Liquidity and Market Structure. *J. Finance* 43(3), 617–637.

GUÉANT, O., and C.-A. LEHALLE (2015): General Intensity Shapes in Optimal Portfolio Liquidation. *Math. Finance* 25(3), 457–495.

GUÉANT, O., C.-A. LEHALLE, and J. FERNANDEZ-TAPIA (2012): Optimal Portfolio Liquidation with Limit Orders. *SIAM J. Finan. Math.* 3(1), 740–764.

GUILBAUD, F., and H. PHAM (2015): Optimal High-Frequency Trading in a Pro-rata Microstructure with Predictive Information. *Math. Finance* 25(3), 545–575.

HANSON, F. B. (2007): *Applied Stochastic Processes and Control for Jump-diffusions: Modeling, Analysis, and Computation*. Philadelphia: SIAM.

HENDERSHOTT, T., and H. MENDELSON (2000): Crossing Networks and Dealer Markets: Competition and Performance. *J. Finance* 55(5), 2071–2115.

HORN, R. A., and C. R. JOHNSON (1985): *Matrix Analysis*. Cambridge: Cambridge University Press.

HÖSCHLER, M. (2011): *Limit Order Book Models and Optimal Trading Strategies*. PhD thesis, Technical University Berlin, 2011.

HUITEMA, R. (2012): Optimal Portfolio Execution Using Market and Limit Orders. Preprint. Available at: http://www.bf.uzh.ch/publikationen/pdf/3543.pdf. Accessed April 26, 2013.

KALLENBERG, O. (2002): *Foundations of Modern Probability*. New York: Springer.

KLÖCK, F., A. SCHIED, and Y. SUN (2011): Existence and Absence of Price Manipulation in a Market Impact Model with Dark Pool. Preprint, 2012. Available at arXiv: http://arxiv.org/abs/1205.4008. Accessed April 26, 2013.

KRATZ, P. (2011): *Optimal Liquidation in Dark Pools in Discrete and Continuous Time*. PhD thesis, Humboldt University Berlin.

KRATZ, P. (2013): An Explicit Solution of a Non-linear Quadratic Constrained Stochastic Control Problem with an Application to Optimal Liquidation in Dark Pools with Adverse Selection. Preprint. Available at arXiv: http://arxiv.org/abs/1204.2498. Accessed April 26, 2013.

KRATZ, P., and T. SCHÖNEBORN (2014): Optimal Liquidation in Dark Pools. *Quant. Finance* 14(9), 1519–1539.

KYLE, A. S. (1985): Continuous Auctions and Insider Trading. *Econometrica* 53(6), 1315–1336.

LARUELLE, S., C.-A. LEHALLE, and G. PAGÈS (2011): Optimal Split of Orders Across Liquidity Pools: A Stochastic Algorithm Approach. *SIAM J. Finan. Math.* 2(1), 1042–1076.

MITTAL, H. (2008): Are You Playing in a Toxic Dark Pool? A Guide to Preventing Information Leakage. *J. Trading* 3(3), 20–33.

NAUJOKAT, F., and N. WESTRAY (2011): Curve Following in Illiquid Markets. *Math. Finan. Econ.* 4(4), 299–335.

OBIZHAEVA, A., and J. WANG (2013): Optimal Trading Strategy and Supply/Demand Dynamics. *J. Finan. Markets* 16(1), 1–32.

ØKSENDAL, B. (2007): *Stochastic Differential Equations: An Introduction with Applications*. Berlin: Springer.

ØKSENDAL, B., and Y. HU (2008): Partial Information Linear Quadratic Control for Jump Diffusions. *SIAM J. Control Optim.* 47(4), 1744–1761.

ØKSENDAL, B., and A. SULEM (2007): *Applied Stochastic Control of Jump Diffusions*. Berlin: Springer.

REID, W. T. (1972): *Riccati Differential Equations*. New York: Academic Press.

ROGERS, L. C. G., and S. SINGH (2010): The Cost of Illiquidity and Its Effects on Hedging. *Math. Finan.* 20(4): 597–615.

SCHIED, A., and T. SCHÖNEBORN (2009): Risk Aversion and the Dynamics of Optimal Liquidation Strategies in Illiquid Markets. *Finance Stoch.* 13(2), 181–204.

SCHIED, A., T. SCHÖNEBORN, and M. TEHRANCHI (2010): Optimal Basket Liquidation for CARA Investors is Deterministic. *Appl. Math. Finance* 17(6): 471–489.

SCHÖNEBORN, T., and A. SCHIED (2009): Liquidation in the Face of Adversity: Stealth vs. Sunshine Trading. Preprint. Available at SSRN: http://papers.ssrn.com/sol3/papers.cfm?abstractid=1007014. Accessed April 26, 2013.

SOFIANOS, G., and D. JERIA (2008a): Quantifying the Sigma X Crossing Benefit. Equity Execution Strategies: Street Smart 31, Goldman Sachs.

SOFIANOS, G., and D. JERIA (2008b): Quantifying the Sigma X Crossing Benefit: The Sequel. *J. Trading* 3(4), 84–91.

YE, M. (2011): A Glimpse into the Dark: Price Formation, Transaction Costs, and Market Share in the Crossing Network. Preprint. Available at SSRN: http://papers.ssrn.com/sol3/papers.cfm?abstractid=1521494. Accessed April 26, 2013.

MATHEMATICAL FINANCE

An International Journal of Mathematics, Statistics and Financial Economics

Mathematical Finance, Vol. 25, No. 3 (July 2015), 545–575

OPTIMAL HIGH-FREQUENCY TRADING IN A PRO RATA MICROSTRUCTURE WITH PREDICTIVE INFORMATION

FABIEN GUILBAUD

Université Paris 7 Diderot

HUYÊN PHAM

Université Paris 7 Diderot, CREST-ENSAE, JVN Institute

We propose a framework to study optimal trading policies in a one-tick pro rata limit order book, as typically arises in short-term interest rate futures contracts. The high-frequency trader chooses to post either market orders or limit orders, which are represented, respectively, by impulse controls and regular controls. We discuss the consequences of the two main features of this microstructure: first, the limit orders are only partially executed, and therefore she has no control on the executed quantity. Second, the high-frequency trader faces the overtrading risk, which is the risk of large variations in her inventory. The consequences of this risk are investigated in the context of optimal liquidation. The optimal trading problem is studied by stochastic control and dynamic programming methods, and we provide the associated numerical resolution procedure and prove its convergence. We propose dimension reduction techniques in several cases of practical interest. We also detail a high-frequency trading strategy in the case where a (predictive) directional information on the price is available. Each of the resulting strategies is illustrated by numerical tests.

KEY WORDS: market making, limit order book, pro rata microstructure, inventory risk, marked point process, stochastic control.

1. INTRODUCTION

In most of modern public security markets, the price formation process, or price discovery, results from competition between several market agents that take part in a public auction. In particular, day trading sessions, which are also called continuous trading phases, consist of continuous double auctions. In these situations, liquidity providers continuously set bid and ask prices for the considered security, and the marketplace publicly displays a (possibly partial) information about these bid and ask prices, along with transactions prices.[1] The action of continuously providing bid and ask quotes during day trading sessions is called *market making*, and this role was traditionally performed by specialist firms. However, due to the recent increased availability of electronic trading technologies, as well as regulatory changes, a large range of investors are now able to implement such market making strategies. These strategies are part of the broader

We would like to thank both referees for all their comments, which helped to improve the first version of this paper.

Manuscript received May 2012; final revision received February 2013.

Address correspondence to Huyen Pham, University Paris 7, Probabilites et Modeles Aleatoire, 2 Place Jussieu, Paris 75251, France; e-mail: pham@math.univ-paris-diderot.fr.

[1] In this paper, we call liquidity provider any investor that currently trades with limit orders.

DOI: 10.1111/mafi.12042

category of high-frequency trading (HFT) strategies, which are characterized by the fact that they facilitate a larger number of orders being sent to the market per unit of time. HFT takes place in the continuous trading phase, and therefore in the double continuous auction context, but actual mechanisms that implement this general continuous double auction setup directly influence the price formation process and, as a consequence, HFT strategies.

In this work, we shall focus on the case where the continuous double auction is implemented by a limit order book (LOB), operated under the pro rata microstructure (see Aikin 2006; Janecek and Kabrhel 2007). This microstructure can be encountered on some derivatives markets, and especially in short-term interest rate (STIR) futures markets, also known as financial futures, traded, e.g., on LIFFE (London International Financial Futures and options Exchange) or on CME (Chicago Mercantile Exchange). This differs from the usual price/time microstructure found on most cash equity markets, and governed by the FIFO (first in first out) rule where limit orders are executed according to the first arrival at the best price. We will describe in detail the pro rata microstructure in Section 2 but the general mechanism of this microstructure is as follows: an incoming market order is dispatched on all active limit orders at the best price, with each limit order contributing to execution in proportion to its volume. In particular, we will discuss the two main consequences of this microstructure on HFT strategies which are the oversizing of the best priced slices of the LOB and the overtrading risk.

Our main goal is to construct an HFT strategy, by means of optimal stochastic control, that targets the pro rata microstructure. We allow both limit orders and market orders in this HFT strategy, modeled, respectively, as continuous and impulse controls, due to considerations about direct trading costs. From a modeling point of view, the key novelty is that we take into account partial execution for limit orders, which is crucial in the pro rata case. For this purpose, we introduce a compound Poisson model for trades processes, which can be fitted to a large class of real-world execution processes, since we make few assumptions about the distributions of execution volumes. From a practical trading point of view, we allow the HFT to input predictive information about price evolution into the strategy, so that our algorithm can be seen as an information-driven HFT strategy (this situation is sometimes called HFT with superior information; see Cartea, Jaimungal, and Ricci 2014). We derive the dynamic programming equation (DPE) corresponding to this mixed impulse/regular control problem. Moreover, we are able to reduce the number of relevant state variables to one in two situations of practical interest: first, in the simple case where the mid-price is a martingale, and second, in the case where the mid-price is a Lévy process, in particular when the HFT has predictive information on price trend, in line with recent studies (Cont and de Larrard 2013). We provide a computational algorithm for the resolution of the DPE, and prove the convergence of this scheme. We illustrate numerically the behavior of the strategy and perform a simulated data benchmarked backtest.

HFR has recently received sustained academic interest, mostly in a price/time microstructure model. The reference work for inventory-based HFR is Avellaneda and Stoikov (2008) following early work by Ho and Stoll (1981). The authors present the HFT problem as an inventory management problem and define inventory risk as the risk of holding a nonzero position in a risky asset. They also provide a closed-form approximate solution in a stylized market model where the controls are continuous. Several works are available that describe optimal strategies for HFT on cash equities or foreign exchange, e.g., Kühn and Stroh (2010), Cartea and Jaimungal (2013), Guéant, Fernandez Tapia, and Lehalle (2013), Guilbaud and Pham (2013) or Veraart (2011).

Guéant et al. (2013) provide extensive analytical treatment of the Avellaneda and Stoikov model. Veraart (2011) includes market orders (that are modeled as impulse controls) as well as limit orders in the context of FX trading. Guilbaud and Pham (2013) study market/limit order HFT strategies on stocks with a focus on the price/time priority microstructure and the bid/ask spread modeling. Cartea et al. (2014) consider an HFT strategy that takes into account influence of trades on the LOB, and give the HFT superior information about the security price evolution. A growing literature is dedicated to modeling the dynamics of the LOB itself, and its consequences for the price formation process. A popular approach is the Poisson Limit Order Book model as in Cont and de Larrard (2013). These authors are able to retrieve a predictive information on price behavior (together with other LOB features) based on the current state of the order book. Finally, in empirical literature, much work is available for cash equities, e.g., Gould et al. (2010), but very little is dedicated to markets operating under pro rata microstructure. We would like to mention the work by Field and Large (2008), which provides a detailed empirical description of such microstructure.

This paper is organized as follows: in Section 2 we detail the pro rata microstructure model and explain the HFT strategy in this context. In Section 3 we formulate the control problem, derive the corresponding DPE for the value function, and state some bounds and symmetry properties. We also simplify the DPE in two cases of practical interest, namely the case where the price is a martingale, and the case where the investor has predictive information on price trend available. In Section 4 we provide the numerical algorithm to solve the DPE, and we study the convergence of the numerical scheme, by proving the monotonicity, stability, and consistency for this scheme. We also provide numerical tests including computations of the optimal policies and performance analysis on a simulated data backtest. Finally, in Section 5 we show how to extend our model in the optimal liquidation case, i.e., when the investor's objective is to minimize the trading costs for unwinding her portfolio.

2. MARKET MODEL

Let us fix a probability space $(\Omega, \mathcal{F}, \mathbf{P})$ equipped with a filtration $\mathbb{F} = (\mathcal{F}_t)_{0\le t\le T}$ satisfying the usual conditions. It is assumed that all random variables and stochastic processes are defined on the stochastic basis $(\Omega, \mathcal{F}, \mathbb{F}, \mathbf{P})$.

Prices in a one-tick microstructure. We denote by P the mid-price, defined as a Markov process with generator $\mathcal{P}$ valued in $\mathbb{P}$. We shall assume that P is a special semimartingale with locally integrable quadratic variation process $[P]$, so that its dual predictable projection (also called sharp bracket) $\langle P\rangle$ exists (see Protter 2005). We assume that $\langle P\rangle_T$ is integrable, and that the predictable finite variation term A of the special semimartingale P satisfies the canonical structure: $dA_t \ll d\langle P\rangle_t$, with a bounded density process

$$\theta_t = \frac{dA_t}{d\langle P\rangle_t} \tag{2.1}$$

and the sharp bracket process $\langle P\rangle$ is absolutely continuous with respect to the Lebesgue measure

$$d\langle P\rangle_t = \varrho(P_t)dt \tag{2.2}$$

for some positive continuous function ϱ of P. We denote by $\delta > 0$ the tick size, and we shall assume that the spread is constantly equal to δ, i.e., the best ask (resp. bid) price is $P^a := P + \frac{\delta}{2}$ (resp. $P^b := P - \frac{\delta}{2}$). This assumption corresponds to the case of the so-called *one-tick microstructure* (Field and Large 2008), which can be encountered, e.g., on STIRs futures contracts.

Trading strategies. For most investors, brokerage costs are paid when a transaction occurs, but new limit order submission, update or cancel are free of charge. Therefore, the investor can submit or update her quotes at any time, with no costs associated to this operation: it is then natural to model the limit order strategy (*make* strategy) as a continuous-time predictable control process. On the contrary, market orders lead to immediate execution, and are costly, so that continuous submission of market orders would lead to bankruptcy. Therefore, we choose to model the market order strategy (*take* strategy) as impulse controls. More precisely, we model trading strategies by a pair $\alpha = (\alpha^{\text{make}}, \alpha^{\text{take}})$ in the form

$$\alpha^{\text{make}} = \left(L^a_t, L^b_t\right)_{t\geq 0}, \qquad \alpha^{\text{take}} = (\tau_n, \xi_n)_{n\in\mathbb{N}}.$$

The predictable processes L^a and L^b, valued in $\{0, 1\}$ represent the possible *make regimes*: when $L^a_t = 1$ (resp. $L^b_t = 1$) this means that the investor has active limit orders at the best ask price (resp. best bid price) at time t, else, if $L^a_t = 0$ (resp. $L^b_t = 0$) this means that the investor has no active order at the best ask price (resp. best bid price) at time t. A practical implementation of such rule would be, for example, to send a limit order with a fixed quantity, when the corresponding control is 1, and cancel it when it turns to 0. Another practical implementation of the rule would be to post a constant proportion of the available volume at best prices: for example, if $V^a_M(t)$ is the current offered volume at best ask, and if $L^a_t = 1$, the practical action in this situation is to post a limit order of volume $v^a(t)$ s.t. $\frac{v^a(t)}{v^a(t)+V^a_M(t)} = const$ at the best ask price. Choice of practical implementation of the limit order controls will impact the outcome of the high-frequency trader's strategy in term of executed volumes, and therefore we propose in the next paragraph an approach suitable in both cases. On the other hand, $(\tau_n)_{n\in\mathbb{N}}$ is an increasing sequence of stopping times, representing the times when the investor chooses to trade at market, and ξ_n, $n \geq 0$ are $\mathcal{F}_{\tau_n}$-measurable random variables valued in $\mathbb{R}$, representing the quantity purchased if $\xi_n \geq 0$ or sold if $\xi_n < 0$.

Execution processes in a pro rata microstructure. The pro rata microstructure (see Janecek and Kabrhel 2007, for extensive presentation and discussion) can be schematically described as follows:[2] when a market order comes in the pro rata LOB, its volume is dispatched among all active limit orders at best prices, proportionally to each limit orders volumes, and therefore create several transactions (see Figure 2.1).

This pro rata microstructure fundamentally differs from price-time microstructure (Guilbaud and Pham 2013) for two reasons: first, several limit orders at the best prices receive incoming market order flow, regardless of the time priority, and second, market makers tend to oversize their liquidity offering (that is, posting limit order with much higher volume than they actually want to trade) in order to increase their transaction

[2]For a detailed description of actual trading rules, and a general overview of STIR futures trading, we refer to Aikin (2006) and references therein.

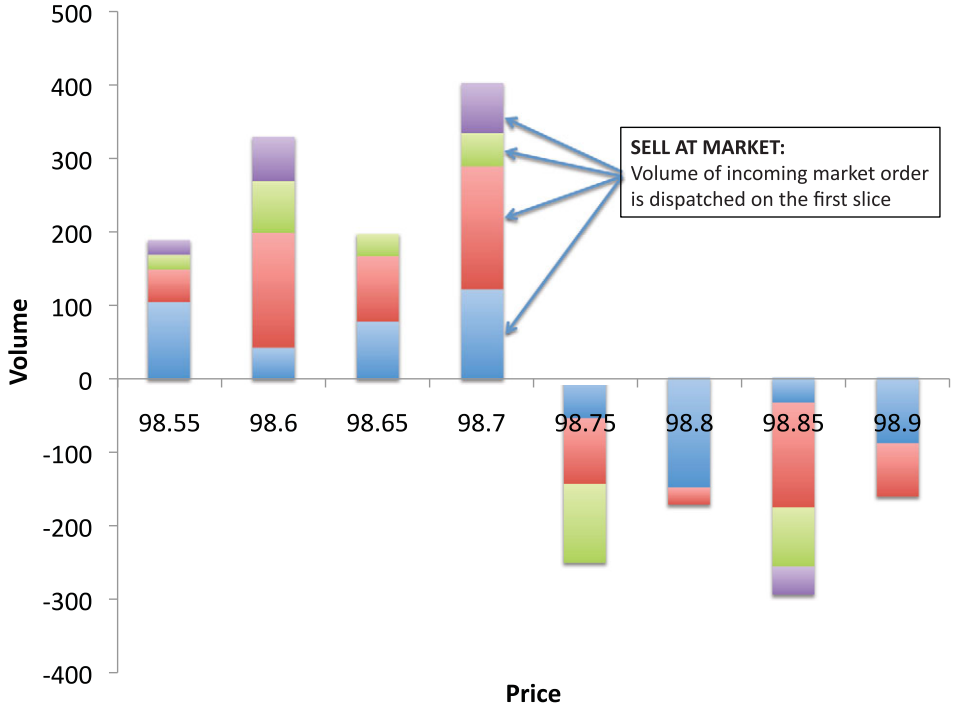

FIGURE 2.1. Schematic view of the pro rata market microstructure.

volume. For example, on the 3-month EURIBOR futures contracts, the liquidity available at the best prices is 200 times higher than the average transaction size.

Let us examine more precisely the outcomes of the two practical implementations of limit orders posting mentioned in the last paragraph. We consider the two cases where the high-frequency trader posts (1) limit orders with a fixed volume, say $V_0 = 100$ contracts, and (2) limits orders with volumes:

$$v^a(t) \text{ s.t. } \frac{v^a(t)}{v^a(t) + V_M^a(t)} = 10\%; \qquad v^b(t) \text{ s.t. } \frac{v^b(t)}{v^b(t) + V_M^b(t)} = 10\%,$$

where $V_M^a(t)$ (resp. $V_M^b(t)$) is the volume available at best ask (resp. bid) at time t. Considering an incoming market order of size V on the ask side, the high-frequency trader receives:

- in case (1), $\min\left(V, V_0 + V_M^a(t)\right) \frac{V_0}{V_0 + V_M^a(t)} \leq V_0$
- and in case (2) $10\%\ \min(V, v^a(t) + V_M^a(t)) \leq v^a(t)$.

Note that in these two cases, the volume offered by the market maker is fully executed if and only if the market order's volume V is greater or equal to the total volume offered at ask $V_0 + V_M^a(t)$, resp. $v^a(t) + V_M^a(t)$. Therefore, the probability that the high-frequency trader volume is fully executed is equal to the probability that the market order consume the first slice of the LOB in integrality. In other words, the volume $\frac{V_0}{V_0 + V_M^a(t)}$, resp. $\frac{v^a(t)}{v^a(t) + V_M^a(t)}$, that the HFT receives, never reaches the bound V_0, resp. $v^a(t)$, unless the market order consumes the first slice of the LOB in integrality.

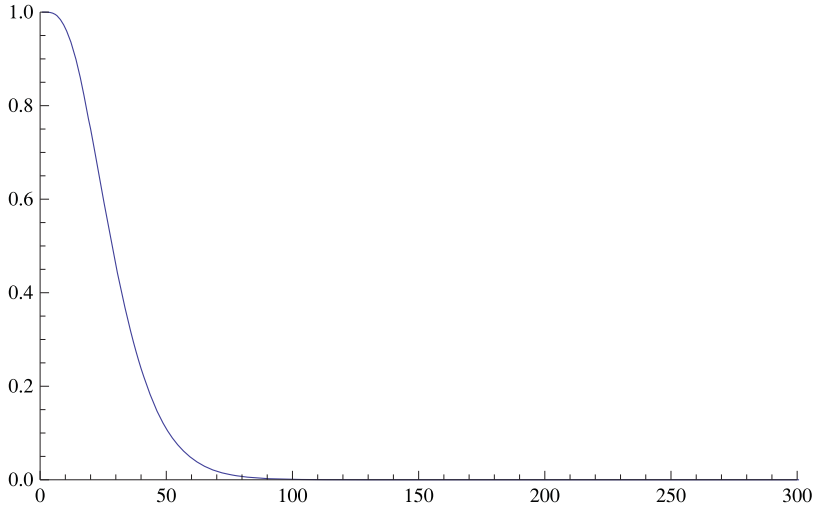

FIGURE 2.2. Probability of the HFT ask limit order V_0, resp. $v^a(t)$, to be fully executed as a function of total offered volume $V_0 + V_M^a(t)$, resp. $v^a(t) + V_M^a(t)$, when a market order of size $V \sim \text{Gamma}(4, 7.5)$ comes in the LOB at time t.

For illustration purposes, and in this discussion only, we assume that the volume of incoming market orders has a gamma distribution with shape 4 and scale 7.5 (which makes an average market order volume of 30 contracts, consistent with observations on the front 3-month EURIBOR contract, see Field and Large 2008). In Figure 2.2, we plot the probability of the HFT's limit order to be fully executed as a function of $V_0 + V_M^a(t)$, resp. $v^a(t) + V_M^a(t)$.

In this example, we see that the probability of the HFT limit order to be fully executed drops to negligible values once the total offered volume is greater than 100, which is about three times the average transaction size. Yet, in actual market, the average offered volume at the best priced slice is about 200 times larger than the average transaction size (Field and Large 2008), and therefore, the probability that the HFT limit orders are fully executed is negligible. For example, if we use the average volume offered on best prices on the front EURIBOR future, 6,000 contracts, the probability of such a market order consuming the first slice is 3×10^{-340}.

Therefore, our approach is to assume that the HFT's limit orders are never fully executed, and instead we model the executed volume as a random variable on which the market maker has no control.[3] This approach combines the advantages of abstracting from practical details of the strategy implementation while keeping precise information on executed volumes. In other words, we assume that the outcome of the practical implementation of the strategy, in terms of executed volume distribution, is known and can be measured in market data.

More precisely, let N^a (resp. N^b) be a Poisson process of intensity $\lambda^a > 0$ (resp. λ^b), whose jump times represent the times when execution by a market order flow occurs at

[3]Actually, the distribution of the executed volume depends on the posted volume, and we should then consider a family of distributions indexed by posted volume. In this paper, we consider as first approximation an a posteriori distribution for the executed volume, which can be obtained ex post by following a given volume posting strategy.

best ask (resp. best bid), and we assume that N^a and N^b are independent. Let $(\zeta_n^a)_{n\in\mathbb{N}^*}$ and $(\zeta_n^b)_{n\in\mathbb{N}^*}$ be two independent sequences of i.i.d. integrable random variables valued in $(0,\infty)$, of distribution laws μ^a and μ^b, which represent the transacted volume of the nth execution at best ask and best bid. We denote by $\nu^a(dt,dz)$ (resp. $\nu^b(dt,dz)$) the Poisson random measure associated to the marked point process $(N^a,(\zeta_n^a)_{n\in\mathbb{N}^*})$ (resp. $(N^b,(\zeta_n^b)_{n\in\mathbb{N}^*})$) of intensity measure $\lambda^a\mu^a(dz)dt$ (resp. $\lambda^b\mu^b(dz)dt$), which is often identified with the compound Poisson processes

$$\vartheta_t^a=\sum_{n=1}^{N_t^a}\zeta_n^a=\int_0^t\int_0^\infty z\,\nu^a(dt,dz),\qquad \vartheta_t^b=\sum_{n=1}^{N_t^b}\zeta_n^b=\int_0^t\int_0^\infty z\,\nu^b(dt,dz) \tag{2.3}$$

representing the cumulative volume of transaction at ask, and bid, assumed to be independent of the mid-price process P. Notice that these processes model only the trades in which the investor has participated.

Cash holdings and inventory. The cash holdings process X and the cumulated number of stocks Y (also called inventory) held by the investor evolve according to the following dynamics:

$$dX_t=L_t^a\left(P_{t^-}+\frac{\delta}{2}\right)d\vartheta_t^a-L_t^b\left(P_{t^-}-\frac{\delta}{2}\right)d\vartheta_t^b,\quad \tau_n\le t<\tau_{n+1}, \tag{2.4}$$

$$dY_t=L_t^b d\vartheta_t^b-L_t^a d\vartheta_t^a,\qquad \tau_n\le t<\tau_{n+1}, \tag{2.5}$$

$$X_{\tau_n}-X_{\tau_n-}=-\xi_n P_{\tau_n}-|\xi_n|\left(\frac{\delta}{2}+\varepsilon\right)-\varepsilon_0 1_{\xi_n\neq 0}, \tag{2.6}$$

$$Y_{\tau_n}-Y_{\tau_n-}=\xi_n. \tag{2.7}$$

The equations (2.4)–(2.5) model the evolution of the cash holdings and inventory under a limit order (make) strategy, whereas equations (2.6)–(2.7) describe the jump on the cash holdings and inventory when posting a market order (take) strategy, subject to a per share fee $\varepsilon>0$ and a fixed fee $\varepsilon_0>0$. In the sequel, we impose the admissibility condition that the inventory should remain within a bounded interval $[-M_Y, M_Y]$, $M_Y>0$, after the trade at market, i.e., $\xi_n\in[-M_Y-Y_{\tau_n-}, M_Y-Y_{\tau_n-}]$, $n\ge 0$, and we shall denote by $\mathcal{A}$ the set of all admissible make and take strategies $\alpha=(\alpha^{\text{make}},\alpha^{\text{take}})$.

REMARK 2.1. Let us define the process $V_t=X_t+Y_tP_t$, which represents at time t the marked-to-market value of the portfolio (or mid-price value of the portfolio). From (2.4)–(2.7), we see that its dynamics is governed by

$$dV_t=\frac{\delta}{2}\big(L_t^b d\vartheta_t^b+L_t^a d\vartheta_t^a\big)+Y_{t^-}dP_t, \tag{2.8}$$

$$V_{\tau_n}-V_{\tau_n-}=-|\xi_n|\left(\frac{\delta}{2}+\varepsilon\right)-\varepsilon_0 1_{\xi_n\neq 0}. \tag{2.9}$$

In equation (2.9), we notice that a trade at market will always diminish the marked-to-market value of our portfolio, due to the fact that we have to "cross the spread," hence trade at a least favorable price. On the other hand, in equation (2.8), the term $\int \frac{\delta}{2}(L_t^b d\vartheta_t^b + L_t^a d\vartheta_t^a)$ is always positive, and represents the profit obtained from a limit order execution, while the term $\int Y_{t^-} dP_t$ represents the portfolio value when holding shares in the stock, hence inducing an inventory risk, which one wants to reduce its variance.

3. MARKET MAKING OPTIMIZATION PROCEDURE

3.1. Control Problem Formulation

The market model in the previous section is fully determined by the state variables (X, Y, P) controlled by the limit/market orders strategies $\alpha = (\alpha^{\text{make}}, \alpha^{\text{take}}) \in \mathcal{A}$. The market maker wants to optimize her profit over a finite time horizon T (typically short term), while keeping control of her inventory risk, and to get rid of any risky asset by time T. We choose a mean-variance optimization criterion, and the goal is to

$$\text{(3.1)}\quad \text{maximize}\quad \mathbf{E}\left[X_T - \gamma \int_0^T Y_{t^-}^2 d\langle P\rangle_t\right] \quad \text{over all strategies } \alpha \in \mathcal{A}, \quad \text{s.t.}\quad Y_T = 0$$

with the convention that $\infty - \infty = -\infty$, as usually done in expected utility maximization. The integral $\int_0^T Y_{t^-}^2 d\langle P\rangle_t$ is a quadratic penalization term for holding a nonzero inventory in the stock, and $\gamma > 0$ is a risk-aversion parameter chosen by the investor. The penalty term $\gamma \mathbf{E}\big[\int_0^T Y_{t^-}^2 d\langle P\rangle_t\big]$ can further be motivated by noting that the variance of the total value of the investor's inventory in the case where P is a martingale by the Itô isometry:

$$\mathbf{Var}\left(\int_0^T Y_{t^-} dP_t\right) = \mathbf{E}\left[\int_0^T Y_{t^-}^2 d\langle P\rangle_t\right],$$

which is our penalty term, up to the scale factor γ. As pointed out by Cartea and Jaimungal (2015), this running penalty is much more effective than the terminal inventory constraint at controlling the trader profit and loss distribution.

Let us now rewrite problem (3.1) in a more standard formulation where the constraint $Y_T = 0$ on the inventory control is removed. For this, let us introduce the liquidation function

$$L(x, y, p) = x + yp - |y|\left(\frac{\delta}{2} + \varepsilon\right) - \varepsilon_0 1_{y\neq 0},$$

which represents the cash obtained after an immediate liquidation of the inventory via a (nonzero) market order. Then, problem (3.1) is formulated equivalently as

$$\text{(3.2)}\quad \text{maximize}\quad \mathbf{E}\left[L(X_T, Y_T, P_T) - \gamma \int_0^T Y_{t^-}^2 d\langle P\rangle_t\right] \quad \text{over all strategies } \alpha \in \mathcal{A}.$$

Indeed, the maximal value of problem (3.1) is clearly smaller than the one of problem (3.2) since for any $\alpha \in \mathcal{A}$ s.t. $Y_T = 0$, we have $L(X_T, Y_T, P_T, S_T) = X_T$. Conversely, given an arbitrary $\alpha \in \mathcal{A}$, let us consider the control $\tilde{\alpha} \in \mathcal{A}$, coinciding with α up to time T, and to which one add at the terminal date T the admissible market order consisting in

liquidating all the inventory Y_T if it is nonzero. The associated state process $(\tilde{X}, \tilde{Y}, P, S)$ satisfies: $\tilde{X}_t = X_t$, $\tilde{Y}_t = Y_t$ for $t < T$, and $\tilde{X}_T = L(X_T, Y_T, P_T, S_T)$, $\tilde{Y}_T = 0$. This shows that the maximal value of problem (3.2) is smaller and then equal to the maximal value of problem (3.1).

Recalling (2.2), let us then define the value function for the problem (3.2):

$$v(t, x, y, p) = \sup_{\alpha \in \mathcal{A}} \mathbf{E}_{t,x,y,p}\left[L(X_T, Y_T, P_T) - \gamma \int_t^T Y_s^2 \varrho(P_s) ds \right] \tag{3.3}$$

for $t \in [0, T]$, $(x, y, p) \in \mathbb{R}^2 \times \mathbb{P}$. Here, given $\alpha \in \mathcal{A}$, $\mathbf{E}_{t,x,y,p}$ denotes the expectation operator under which the process (X, Y, P) solution to (2.4)–(2.7) with initial state $(X_{t^-}, Y_{t^-}, P_{t^-}) = (x, y, p)$, is taken. Problem (3.3) is a mixed impulse/regular control problem in Markov model with jumps that we shall study by dynamic programming methods.

First, we state some bounds on the value function, which shows in particular that the value function is finite and locally bounded.

PROPOSITION 3.1. *There exists some constant K_P (depending only on the price process and γ) such that for all $(t, x, y, p) \in [0, T] \times \mathbb{R}^2 \times \mathbb{P}$,*

$$L(x, y, p) \le v(t, x, y, p) \le x + yp + \frac{\delta}{2}\big(\lambda^a \bar{\mu}^a + \lambda^b \bar{\mu}^b\big)(T - t) + K_P, \tag{3.4}$$

where $\bar{\mu}^a = \int_0^\infty z\mu^a(dz)$, $\bar{\mu}^b = \int_0^\infty z\mu^b(dz)$ are the mean of the distribution laws μ^a and μ^b.

Proof. The lower bound in (3.4) is derived easily by considering the particular strategy, which consists of liquidating immediately all the current inventory (if nonzero) via a market order, and then doing nothing else until the final horizon. Let us now focus on the upper bound. Observe that in the definition of the value function in (3.3), we can restrict obviously to controls $\alpha \in \mathcal{A}$ s.t.

$$\mathbf{E}\left[\int_0^T Y_{t^-}^2 d\langle P \rangle_t \right] < \infty. \tag{3.5}$$

For such strategies, we have

$$\begin{aligned}
&\mathbf{E}_{t,x,y,p}\left[L(X_T, Y_T, P_T) - \gamma \int_t^T Y_{s^-}^2 d\langle P \rangle_s \right] \\
&\quad \le \mathbf{E}_{t,x,y,p}\left[V_T - \gamma \int_t^T Y_s^2 d\langle P \rangle_s \right] \\
&\quad \le x + yp + \mathbf{E}_{t,x,y,p}\left[\frac{\delta}{2}\big(\vartheta^a_{T-t} + \vartheta^b_{T-t}\big) + \int_t^T Y_{s^-} dP_s - \gamma \int_t^T Y_{s^-}^2 d\langle P \rangle_s \right] \\
&\quad = x + yp + \mathbf{E}_{t,x,y,p}\left[\frac{\delta}{2}\big(\vartheta^a_{T-t} + \vartheta^b_{T-t}\big) + \int_t^T \big(Y_{s^-}\theta_s - \gamma Y_{s^-}^2\big) d\langle P \rangle_s \right].
\end{aligned}$$

Here, the second inequality follows from the relation (2.8), together with the fact that $L^a, L^b \le 1$, ϑ^a, ϑ^b are increasing processes, and also that jumps of V are negative by (2.9). The last equality holds true by (2.1) and the fact that $\int Y_- dM$ is a square-integrable martingale from (3.5), where M is the martingale part of the semimartingale

P. Since θ is bounded and $\gamma > 0$, this shows that for all strategies α satisfying (3.5), we have:

$$\mathbf{E}_{t,x,y,p}\left[L(X_T,\, Y_T,\, P_T) - \gamma \int_t^T Y_{s^-}^2 d\langle P\rangle_s\right] \leq x + yp + \frac{\delta}{2}\mathbf{E}\big[\vartheta^a_{T-t} + \vartheta^b_{T-t}\big] + K\mathbf{E}[\langle P\rangle_T]$$

for some positive constant K, which proves the required result by recalling the characteristics of the compound Poisson processes ϑ^a and ϑ^b, and since $\langle P\rangle_T$ is assumed to be integrable. □

REMARK 3.2. The terms of the upper bound in (3.4) have a financial interpretation. The term $x + yp$ represents the marked-to-market value of the portfolio evaluated at mid-price, whereas the term K_P stands for a bound on profit for any directional frictionless strategy on the fictive asset that is priced P. The term $\frac{\delta}{2}\big(\lambda^a\bar{\mu}^a + \lambda^b\bar{\mu}^b\big)(T-t)$, always positive, represents the upper bound on profit due to market making, i.e., the profit of the strategy participating in every trade, but with no costs associated to managing its inventory.

3.2. Dynamic Programming Equation

For any $(\ell^a, \ell^b) \in \{0, 1\}^2$, we introduce the nonlocal operator associated with the limit order control:

$$\mathcal{L}^{\ell^a,\ell^b} = \mathcal{P} + \ell^a\Gamma^a + \ell^b\Gamma^b, \tag{3.6}$$

where

$$\Gamma^a\phi(t, x, y, p) = \lambda^a \int_0^\infty \left[\phi\left(t, x + z\left(p + \frac{\delta}{2}\right), y - z, p\right) - \phi(t, x, y, p)\right]\mu^a(dz),$$

$$\Gamma^b\phi(t, x, y, p) = \lambda^b \int_0^\infty \left[\phi\left(t, x - z\left(p - \frac{\delta}{2}\right), y + z, p\right) - \phi(t, x, y, p)\right]\mu^b(dz)$$

for (t, x, y, p) $[0, T] \times \mathbb{R} \times \mathbb{R} \times \mathbb{P}$. Let us also consider the impulse operator associated with admissible market order controls, and defined by

$$\mathcal{M}\phi(t, x, y, p) = \sup_{e\in[-M_Y - y, M_Y - y]\setminus\{0\}} \phi\left(t, x - ep - |e|\left(\frac{\delta}{2} + \varepsilon\right) - \varepsilon_0, y + e, p\right).$$

The DPE associated to the control problem (3.3) is a quasi-variational inequality (QVI) in the form

$$\min\left[-\frac{\partial v}{\partial t} - \sup_{(\ell^a,\ell^b)\in\{0,1\}^2} \mathcal{L}^{\ell^a,\ell^b} v \;+\; \gamma g\;,\; v - \mathcal{M}v\right] = 0, \quad \text{on } [0, T) \times \mathbb{R}^2 \times \mathbb{P}, \tag{3.7}$$

together with the terminal condition

$$v(T, .) = L, \qquad \text{on } \mathbb{R}^2 \times \mathbb{P}, \tag{3.8}$$

where we denoted by g the function: $g(y, p) = y^2\varrho(p)$. By standard methods of dynamic programming, one can show that the value function in (3.3) is the unique viscosity solution under growth conditions determined by (3.4) to the DPE (3.7)–(3.8) of dimension three (in addition to the time variable) (see, e.g., chapter 9 in Øksendal and Sulem 2007).

3.3. Dimension Reduction in the Lévy Case

We now consider a special case on the mid-price process where the market making control problem can be reduced to the resolution of a one-dimensional variational inequality involving only the inventory state variable. We shall suppose actually that P is a Lévy process so that

$$\mathcal{P} I_{\mathbb{P}} = c_{_P}, \text{ and } \varrho \text{ is a constant,} \tag{3.9}$$

where $I_{\mathbb{P}}$ is the identity function on $\mathbb{P}$, i.e., $I_{\mathbb{P}}(p) = p$, and $\varrho > 0$, $c_{_P}$ are real constants depending on the characteristics triplet of P. Two practical examples are:

- **Martingale case**: The mid-price process P is a martingale, so that $\mathcal{P} I_{\mathbb{P}} = 0$. This martingale assumption in a high-frequency context reflects the idea that the market maker has no information on the future direction of the stock price.
- **Trend information**: To remove the martingale assumption, one can introduce some knowledge about the price trend. A typical simple example is when P follows an arithmetic Brownian motion (Bachelier model). A more relevant example is described by a pure jump process P valued in the discrete grid $\delta\mathbb{Z}$ with tick $\delta > 0$, and such that

$$\begin{aligned}
\mathbf{P}(P_{t+h} - P_t = \delta \,|\mathcal{F}_t) &= \pi^+ h + o(h) \\
\mathbf{P}(P_{t+h} - P_t = -\delta \,|\mathcal{F}_t) &= \pi^- h + o(h) \\
\mathbf{P}(|P_{t+h} - P_t| > \delta \,|\mathcal{F}_t) &= o(h),
\end{aligned}$$

  where $\pi^+, \pi^- > 0$, and $o(h)$ is the usual notation meaning that $\lim_{h\to 0} o(h)/h = 0$. Relation (3.9) then holds with $c_{_P} = \varpi\delta$, where $\varpi = \pi^+ - \pi^-$ represents a constant information about price direction, and $\varrho = (\pi^+ + \pi^-)\delta^2$. In a high-frequency context, this model is of practical interest as it provides a way to include a (predictive) information about price direction. For example, Cont and de Larrard (2013) infers the future prices movements (at the scale of a few seconds) from the current state of the LOB in a Poisson framework. In this work, as well as in our real data tests, the main quantities of interest are the volume offered at the best prices in the LOB, also known as the imbalance.

In this Lévy context, we can decompose the value function v into the form

$$v(t, x, y, p) = L_0(x, y, p) + w(t, y), \tag{3.10}$$

where $L_0(x, y, p) = x + yp - |y|\left(\frac{\delta}{2} + \varepsilon\right) = L(x, y, p) + \varepsilon_0 1_{y\neq 0}$ is the liquidation function up to the fixed fee, and where w is solution to the integral variational inequality:

$$\min\left[-\frac{\partial w}{\partial t} - yc_{_P} + \gamma\varrho y^2 - \mathcal{I}^a w - \mathcal{I}^b w \;,\; w - \tilde{\mathcal{M}} w\right] = 0, \quad \text{on } [0, T) \times \mathbb{R} \tag{3.11}$$

together with the terminal condition

$$w(T, y) = -\varepsilon_0 1_{y\neq 0}, \quad \forall y \in \mathbb{R} \tag{3.12}$$

with $\mathcal{I}^a$ and $\mathcal{I}^b$, the nonlocal integral operators

$$\mathcal{I}^a w(t,y) = \lambda^a \left(\int_0^\infty \left[w(t,y-z) - w(t,y) + z\frac{\delta}{2} + \left(\frac{\delta}{2}+\varepsilon\right)(|y|-|y-z|) \right] \mu^a(dz) \right)_+$$

$$\mathcal{I}^b w(t,y) = \lambda^b \left(\int_0^\infty \left[w(t,y+z) - w(t,y) + z\frac{\delta}{2} + \left(\frac{\delta}{2}+\varepsilon\right)(|y|-|y+z|) \right] \mu^b(dz) \right)_+,$$

and $\tilde{\mathcal{M}}$, the nonlocal operator

$$\tilde{\mathcal{M}} w(t,y) = \sup_{e \in [-M_Y - y, M_Y - y] \setminus \{0\}} \left[w(t,y+e) - \left(\frac{\delta}{2}+\varepsilon\right)(|y+e|+|e|-|y|) - \varepsilon_0 \right].$$

The interpretation of the decomposition (3.10) is the following. The term $L_0(x,y,p)$ represents the mid-price value that the investor would obtain by liquidating immediately with a market order (up to the fixed fee), and w is an additional correction term taking into account the illiquidity effects induced by the bid–ask spread and the fees, as well as the execution risk when submitting limit orders. Moreover, in the Lévy case, this correction function w depends only on time and inventory. From (3.4), we have the following bounds on the function w:

$$-\varepsilon_0 1_{y\neq 0} \le w(t,y) \le \left(\frac{\delta}{2}+\varepsilon\right)|y| + \frac{\delta}{2}\left(\lambda^a \bar{\mu}^a + \lambda^b \bar{\mu}^b\right)(T-t) + K_P, \quad \forall (t,y) \in [0,T] \times \mathbb{R}.$$

Actually, we have a sharper upper bound in the Lévy context.

PROPOSITION 3.3. *Under (3.9), we have:*

$$-\varepsilon_0 1_{y\neq 0} \le w(t,y) \le (T-t)\left[\frac{c_P^2}{4\gamma\rho} + \lambda^a(\delta+\epsilon)\bar{\mu}^a + \lambda^b(\delta+\epsilon)\bar{\mu}^b \right] \tag{3.13}$$

for all $(t,x,y,p) \in [0,T] \times \mathbb{R}^2 \times \mathbb{P}$.

Proof. For any $(x,y,p) \in \mathbb{R}^2 \times \mathbb{P}$, we notice that

$$\begin{aligned} L_0(x,y,p) - \sup_{e \in [-M_Y - y, M_Y - y]\setminus\{0\}} & L_0(x - ep - |e|\left(\frac{\delta}{2}+\varepsilon\right) - \varepsilon_0, y+e, p) \\ &= \inf_{e \in [-M_Y - y, M_Y - y]\setminus\{0\}} \left[\left(\frac{\delta}{2}+\varepsilon\right)(|e|+|y+e|-|y|) + \varepsilon_0 \right] \ge 0. \end{aligned} \tag{3.14}$$

We also observe that for all $z \ge 0$

$$\begin{aligned} L_0\left(x + z\left(p+\frac{\delta}{2}\right), y-z, p\right) - L_0(x,y,p) &= z\frac{\delta}{2} + \left(\frac{\delta}{2}+\varepsilon\right)(|y|-|y-z|) \\ &\le (\delta+\varepsilon)z, \end{aligned} \tag{3.15}$$

and similarly

$$L_0\left(x - z\left(p-\frac{\delta}{2}\right), y+z, p\right) - L_0(x,y,p) \le (\delta+\varepsilon)z. \tag{3.16}$$

Let us then consider the function $\phi(t, x, y, p) = L_0(x, y, p) + (T - t)u$, for some real constant u to be determined later. Then, $\phi(T, .) = L_0$, and by (3.15)–(3.16), we easily check that

$$-\frac{\partial \phi}{\partial t} - \sup_{(\ell^a, \ell^b) \in \{0,1\}^2} \mathcal{L}^{\ell^a, \ell^b} \phi + \gamma g$$
$$\geq u - \lambda^a(\delta + \varepsilon)\bar{\mu}^a - \lambda^b(\delta + \varepsilon)\bar{\mu}^b - y c_P + \gamma y^2 \rho.$$

The r.h.s. of this last inequality is a second-order polynomial in y and therefore it is always nonnegative iff

$$c_P^2 - 4\gamma\rho(u - \lambda^a(\delta + \varepsilon)\bar{\mu}^a - \lambda^b(\delta + \varepsilon)\bar{\mu}^b) \leq 0,$$

which is satisfied once the constant u is large enough, namely:

$$u \geq \hat{u} := \frac{c_P^2}{4\gamma\rho} + \lambda^a(\delta + \epsilon)\bar{\mu}^a + \lambda^b(\delta + \epsilon)\bar{\mu}^b.$$

For such choice of $u = \hat{u}$, and denoting by $\hat{\phi}$ the associated function: $\hat{\phi}(t, x, y, p) = L_0(x, y, p) + (T - t)\hat{u}$ we have

$$-\frac{\partial \hat{\phi}}{\partial t} - \sup_{(\ell^a, \ell^b) \in \{0,1\}^2} \mathcal{L}^{\ell^a, \ell^b} \hat{\phi} + \gamma g \geq 0,$$

which shows, together with (3.14), that $\hat{\phi}$ is a supersolution of (3.7)–(3.8). From comparison principle for this variational inequality, we deduce that

$$v \leq \hat{\phi} \quad \text{on} \quad [0, T] \times \mathbb{R}^2 \times \mathbb{P},$$

which shows the required upper bound for $w = v - L_0$. □

Finally, from (3.11)–(3.12), and in the case where $\lambda^a = \lambda^b$, $\mu^a = \mu^b$, and by stressing the dependence of w in c_P, we see that w satisfies the symmetry relation

$$w(t, y, c_P) = w(t, -y, -c_P), \quad \forall (t, y) \in [0, T] \times \mathbb{R}. \tag{3.17}$$

REMARK 3.4. Notice that in some situations, one may have $\mathcal{I}^a w(t, y) = \mathcal{I}^b w(t, y) = 0$ and $w(t, y) - \tilde{\mathcal{M}} w(t, y) > 0$, meaning that the optimal control at the point (t, y) is to do nothing (no market order, no limit orders on either sides). For example, if we have a zero inventory $y = 0$ at time $t \in [0, T]$, then by (3.13)

$$\int_0^\infty \left[w(t, 0 - z) - w(t, 0) + z\frac{\delta}{2} + \left(\frac{\delta}{2} + \varepsilon\right)(|0| - |0 - z|) \right] \mu^*(dz)$$
$$\leq (T - t)\left[\frac{c_P^2}{4\gamma\rho} + \lambda^a(\delta + \epsilon)\bar{\mu}^a + \lambda^b(\delta + \epsilon)\bar{\mu}^b \right] - \varepsilon\bar{\mu}^* =: C_{a,b}(T - t) - \varepsilon\bar{\mu}^*, \tag{3.18}$$

where we wrote $*$ for ask and bid a, b, and set $C_{a,b} = \left[\frac{c_P^2}{4\gamma\rho} + \lambda^a(\delta + \epsilon)\bar{\mu}^a + \lambda^b(\delta + \epsilon)\bar{\mu}^b\right]$. The r.h.s. of (3.18) is nonpositive for t close enough to T, namely for: $(T - t) \leq \varepsilon \min[\bar{\mu}^a, \bar{\mu}^b]/C_{a,b}$, and so: $\mathcal{I}^a w(t, y) = \mathcal{I}^b w(t, y) = 0$. Using again (3.13), we also see

that for t satisfying $(T-t) \leq \varepsilon_0/C_{a,b}$, then $\tilde{\mathcal{M}}w(t,0) < 0 \leq w(t,0)$. Consequently, when we have a zero inventory close to maturity, it is not optimal to send any limit order nor submitting market orders. Indeed, from a zero inventory, posting a limit order, say at ask, may lead to an execution (a sale), and this open short position, unlikely to be closed by another limit order execution, will have to be closed at T with a market buy order (hitting the ask side). Therefore, this operation is a net loss of the per share fee times the volume of the trade.

4. NUMERICAL RESOLUTION

In this section, we focus on the numerical resolution of the integral variational inequality (3.11)–(3.12), which characterizes the reduced value function of the market-making problem in the Lévy case.

4.1. Numerical Scheme

We provide a computational scheme for the integral variational inequality (3.11). We first consider a time discretization of the interval $[0, T]$ with time step $h = T/N$ and a regular time grid $\mathbb{T}_N = \{t_k = kh,\ k = 0, \ldots, N\}$. Next, we discretize and localize the inventory state space on a finite regular grid: for any $M > 0$ (in practice we choose $M = M_Y$) and $N_Y \in \mathbb{N}$, and denoting by $\Delta_Y = \frac{M}{N_Y}$, we set

$$\mathbb{Y}_M = \{y_i = i\Delta_Y,\ i = -N_Y, \ldots, N_Y\}.$$

We denote by $\mathrm{Proj}_M(y) := -M \vee (y \wedge M)$, and consider the discrete approximating distribution of μ^a and μ^b, defined by

$$\hat{\mu}^a = \sum_{i\in\mathbb{Z}^+} \mu^a([i\Delta_Y;(i+1)\Delta_Y))\delta_{i\Delta_Y}, \qquad \hat{\mu}^b = \sum_{i\in\mathbb{Z}^+} \mu^b([i\Delta_Y;(i+1)\Delta_Y))\delta_{i\Delta_Y}$$

with δ_x the Dirac measure at x. We then introduce the operator associated to the explicit time-space discretization of the integral variational inequality (3.11): for any real-valued function φ on $[0, T] \times \mathbb{R}$, $t \in [0, T]$, and $y \in \mathbb{R}$, we define

$$\mathcal{S}^{h,\Delta_Y,M}(t, y, \varphi) = \max[\mathcal{T}^{h,\Delta_Y,M}(t, y, \varphi)\ ;\ \tilde{\mathcal{M}}^{h,\Delta_Y,M}(t, y, \varphi)],$$

where

$$\begin{aligned}
\mathcal{T}^{h,\Delta_Y,M}(t, y, \varphi) = \varphi(t, y) &- h\gamma\varrho y^2 + hyc_P \\
&+ \lambda^a h\bigg(\int_0^\infty [\varphi(t, \mathrm{Proj}_M(y - z)) - \varphi(t, y)]\hat{\mu}^a(dz) \\
&+ \int_0^\infty \left[\frac{\delta}{2}z + \left(\frac{\delta}{2} + \varepsilon\right)(|y| - |y - z|)\right]\mu^a(dz)\bigg)_+ \\
&+ \lambda^b h\bigg(\int_0^\infty [\varphi(t, \mathrm{Proj}_M(y + z)) - \varphi(t, y)]\hat{\mu}^b(dz) \\
&+ \int_0^\infty \left[\frac{\delta}{2}z + \left(\frac{\delta}{2} + \varepsilon\right)(|y| - |y + z|)\right]\mu^b(dz)\bigg)_+,
\end{aligned}$$

and

$$\tilde{\mathcal{M}}^{h,\Delta_Y,M}(t,y,\varphi) = \sup_{e\in\mathbb{Y}_M\cap[-M_Y-y,M_Y-y]\setminus\{0\}} \left[\varphi(t,\mathrm{Proj}_M(y+e)) - \left(\frac{\delta}{2}+\varepsilon\right)(|y+e|+|e|-|y|)-\varepsilon_0\right]. \tag{4.1}$$

By recalling that $x_+ = \max_{\ell\in\{0,1\}} \ell x$, we see that the operator $\mathcal{T}^{h,\Delta_Y,M}$ may be written also as

$$\begin{aligned}\mathcal{T}^{h,\Delta_Y,M}(t,y,\varphi) = -h\gamma\varrho y^2 + hyc_P + \max_{\ell^a,\ell^b\in\{0,1\}}\Bigg[&\varphi(t,y)(1-\lambda^a h\ell^a-\lambda^b h\ell^b)\\ &+\lambda^a h\ell^a\Bigg(\int_0^\infty \varphi(t,\mathrm{Proj}_M(y-z))\hat{\mu}^a(dz)\\ &+\int_0^\infty\left[\frac{\delta}{2}z+\left(\frac{\delta}{2}+\varepsilon\right)(|y|-|y-z|)\right]\mu^a(dz)\Bigg)\\ &+\lambda^b h\ell^b\left(\int_0^\infty \varphi(t,\mathrm{Proj}_M(y+z))\hat{\mu}^b(dz)\right.\\ &\left.+\int_0^\infty\left[\frac{\delta}{2}z+\left(\frac{\delta}{2}+\varepsilon\right)(|y|-|y+z|)\right]\mu^b(dz)\right)\Bigg].\end{aligned} \tag{4.2}$$

Notice that on the boundary $y = M_Y$ (resp. $y = -M_Y$) the set of admissible market orders is $[-2y, 0]$ (resp. $[0, -2y]$) which implies that we only allow sell (resp. buy) market orders. Limit order controls can be of any type on the boundary, since we do not set a global constraint on the inventory.

We then approximate the solution w to (3.11)–(3.12) by the function $w^{h,\Delta_Y,M}$ on $\mathbb{T}_N\times\mathbb{Y}_M$ solution to the computational scheme

$$w^{h,\Delta_Y,M}(t_N,y) = -\varepsilon_0 1_{y\neq 0}, \qquad y\in\mathbb{Y}_M, \tag{4.3}$$

$$w^{h,\Delta_Y,M}(t_k,y) = \mathcal{S}^{h,\Delta_Y,M}(t_{k+1},y,w^{h,\Delta_Y,M}), \qquad k=0,\dots,N-1,\ y\in\mathbb{Y}_M. \tag{4.4}$$

This algorithm is described explicitly in backward induction by the following pseudo-code:

- Time step $t_N = T$: for each $y\in\mathbb{Y}_M$, set $w^{h,\Delta_Y,M}(t_n,y) := -\varepsilon_0 1_{y\neq 0}$
- For $k = N-1\dots 0$, from time step t_{k+1} to time step t_k, and for each $y\in\mathbb{Y}_M$:
  - Compute $\mathcal{T}^{h,\Delta_Y,M}(t_{k+1},y,w^{h,\Delta_Y,M})$ from (4.2), and store $\ell^{a,\star}$, $\ell^{b,\star}$ the argmax;
  - Compute $\tilde{\mathcal{M}}^{h,\Delta_Y,M}(t_{k+1},y,w^{h,\Delta_Y,M})$ from (4.1), and store $e^\star$ the argmax;
  - If $\mathcal{T}^{h,\Delta_Y,M}(t_{k+1},y,w^{h,\Delta_Y,M}) \geq \tilde{\mathcal{M}}^{h,\Delta_Y,M}(t_{k+1},y,w^{h,\Delta_Y,M})$ then set

    $$w^{h,\Delta_Y,M}(t_k,y) := \mathcal{T}^{h,\Delta_Y,M}(t_{k+1},y,w^{h,\Delta_Y,M})$$

    and the policy is make $(\ell^{a,\star},\ell^{b,\star})$. Otherwise, set

    $$w^{h,\Delta_Y,M}(t_k,y) := \tilde{\mathcal{M}}^{h,\Delta_Y,M}(t_{k+1},y,w^{h,\Delta_Y,M})$$

    and the policy is take $e^\star$.

4.2. Convergence of the Numerical Scheme

In this section, we study the convergence of the numerical scheme (4.3)–(4.4) by showing the monotonicity, stability, and consistency properties of this scheme. We denote by $C_b^1([0,T]\times\mathbb{R})$ the set of bounded continuously differentiable functions on $[0,T]\times\mathbb{R}$, with bounded derivatives.

PROPOSITION 4.1 (Monotonicity). *For any $h>0$ s.t. $h<\frac{1}{\lambda^a+\lambda^b}$ the operator $\mathcal{S}^{h,\Delta_Y,M}$ is nondecreasing in φ, i.e., for any $(t,y)\in[0,T]\times\mathbb{R}$ and any $\varphi,\psi\in C_b^1([0,T]\times\mathbb{R})$, s.t. $\varphi\le\psi$:*

$$\mathcal{S}^{h,\Delta_Y,M}(t,y,\varphi)\le\mathcal{S}^{h,\Delta_Y,M}(t,y,\psi).$$

Proof. From the expression (4.2), it is clear that $\mathcal{T}^{h,\Delta_Y,M}(t,y,\varphi)$, and then also $\mathcal{S}^{h,\Delta_Y,M}(t,y,\varphi)$ is monotone in φ once $1-\lambda^a h-\lambda^b h>0$. □

PROPOSITION 4.2 (Stability). *For any $h,\Delta_Y,M>0$ there exists a unique solution $w^{h,\Delta_Y,M}$ to (4.3)–(4.4), and the sequence $(w^{h,\Delta_Y,M})$ is uniformly bounded: for any $(t,y)\in\mathbb{T}_N\times\mathbb{Y}_M$,*

$$-\varepsilon_0 1_{y\neq 0}\ \le\ w^{h,\Delta_Y,M}(t,y)\le(T-t)\left[\frac{c_P^2}{4\gamma\rho}+\lambda^a(\delta+\epsilon)\bar{\mu}^a+\lambda^b(\delta+\epsilon)\bar{\mu}^b\right].$$

Proof. Existence and uniqueness of $w^{h,\Delta_Y,M}$ follows from the explicit backward scheme (4.3)–(4.4). Let us now prove the uniform bounds. We consider the function

$$\Psi^\star(t)=(T-t)\left[\frac{c_P^2}{4\gamma\rho}+\lambda^a(\delta+\epsilon)\bar{\mu}^a+\lambda^b(\delta+\epsilon)\bar{\mu}^b\right]$$

and notice that $\Psi^\star(t)\ge\mathcal{S}^{h,\Delta_Y,M}(t+h,y,\Psi^\star)$ by the same arguments as in Proposition 3.3. Moreover, we have, by definition, $w^{h,\Delta_Y,M}(T,y)=-\varepsilon_0 1_{y\neq 0}\le\Psi^\star(T)=0$, and therefore, a direct recurrence from (4.3)–(4.4) shows that $w^{h,\Delta_Y,M}(t,y)\le\Psi^\star(t)$ for all $(t,y)\in\mathbb{T}_n\times\mathbb{Y}_M$, which is the required upper bound for $w^{h,\Delta_Y,M}$.

On the other hand, we notice that $\mathcal{S}^{h,\Delta_Y,M}(t,0,\varphi)\ge\varphi(t,0)$ for any function φ on $[0,T]\times\mathbb{R}$, and $t\in[0,T]$, by considering the"diffusive" part of the numerical scheme with the particular controls $\ell^a=\ell^b=0$. Therefore, since $w^{h,\Delta_Y,M}(T,0)=0$, we obtain by induction on (4.3)–(4.4) that $w^{h,\Delta_Y,M}(t,0)\ge 0$ for any $t\in\mathbb{T}_N$. Finally, considering the obstacle part of the numerical scheme, with the particular control $e=-y$, shows that $w^{h,\Delta_Y,M}(t,y)\ge w^{h,\Delta_Y,M}(t,0)-\varepsilon_0 1_{y\neq 0}\ge-\varepsilon_0 1_{y\neq 0}$ for any $(t,y)\in\mathbb{T}_N\times\mathbb{Y}_M$, which proves the required lower bound for $w^{h,\Delta_Y,M}$. □

PROPOSITION 4.3 (Consistency). *For all $(t,y)\in[0,T)\times\mathbb{R}$ and $\varphi\in C_b^1([0,T]\times\mathbb{R})$, we have*

$$\lim_{\substack{(h,\Delta_Y,M)\to(0,0,\infty)\\(t',y')\to(t,y)}}\frac{1}{h}[\varphi(t',y')-\mathcal{T}^{h,\Delta_Y,M}(t'+h,y',\varphi)]$$
$$=-\frac{\partial\varphi}{\partial t}(t,y)-yc_P+\gamma\varrho y^2-\mathcal{I}^a\varphi(t,y)-\mathcal{I}^b\varphi(t,y) \tag{4.5}$$

and

$$\lim_{\substack{(h,\,\Delta_Y,\,M)\to(0,0,\infty)\\(t',\,y')\to(t,\,y)}} \tilde{\mathcal{M}}^{h,\Delta_Y,M}(t'+h,\,y',\,\varphi)=\tilde{\mathcal{M}}\varphi(t,\,y). \tag{4.6}$$

Proof. The consistency relation (4.6) follows from the continuity of the function $(t,\,y,\,e)\to\varphi(t,\,y+e)-\left(\frac{\delta}{2}+\varepsilon\right)(|y+e|+|e|-|y|)-\varepsilon_0$. On the other hand, we have for all $(t',\,y')\in[0,\,T)\times\mathbb{R}$,

$$\begin{aligned}\frac{1}{h}[\varphi(t',\,y')-\mathcal{T}^{h,\Delta_Y,M}(t'+h,\,y',\,\varphi)]&=\frac{1}{h}[\varphi(t',\,y')-\varphi(t'+h,\,y')]-y'c_P+\gamma\rho y'^2\\&\quad-\mathcal{I}_a^{h,\Delta_Y,M}(t'+h,\,y',\,\varphi)-\mathcal{I}_b^{h,\Delta_Y,M}(t'+h,\,y',\,\varphi),\end{aligned} \tag{4.7}$$

where

$$\begin{aligned}\mathcal{I}_a^{h,\Delta_Y,M}(t,\,y,\,\varphi)&=\lambda^a\Bigg(\int_0^\infty\left[\varphi(t,\,\mathrm{Proj}_M(y-z))-\varphi(t,\,y)\right]\hat{\mu}^a(dz)\\&\qquad+\int_0^\infty\left[\frac{\delta}{2}z+\left(\frac{\delta}{2}+\varepsilon\right)(|y|-|y-z|)\right]\mu^a(dz)\Bigg)_+\\\mathcal{I}_a^{h,\Delta_Y,M}(t,\,y,\,\varphi)&=\lambda^b\Bigg(\int_0^\infty\left[\varphi(t,\,\mathrm{Proj}_M(y+z))-\varphi(t,\,y)\right]\hat{\mu}^b(dz)\\&\qquad+\int_0^\infty\left[\frac{\delta}{2}z+\left(\frac{\delta}{2}+\varepsilon\right)(|y|-|y+z|)\right]\mu^b(dz)\Bigg)_+.\end{aligned}$$

The three first terms of (4.7) converge trivially to $-\frac{\partial\varphi}{\partial t}(t,\,y)-yc_P+\gamma\varrho y^2$ as h goes to zero and $(t',\,y')$ goes to $(t,\,y)$. Hence, in order to get the consistency relation, it remains to prove the convergence of $\mathcal{I}_a^{h,\Delta_Y,M}(t'+h,\,y',\,\varphi)$ to $\mathcal{I}^a\varphi(t,\,y)$ as $(h,\,\Delta_Y,\,M)$ goes to $(0,\,0,\,\infty)$, and $(t',\,y')$ goes to $(t,\,y)$ (an identical argument holds for $\mathcal{I}_b^{h,\Delta_Y,M}(t'+h,\,y',\,\varphi)$). By writing that $|x_+-x'_+|\le|x-x'|$, we have

$$\begin{aligned}&\left|\mathcal{I}_a^{h,\Delta_Y,M}(t'+h,\,y',\,\varphi)-\mathcal{I}^a\varphi(t,\,y)\right|\\&\qquad\le\lambda^a|\varphi(t'+h,\,y')-\varphi(t,\,y)|\\&\qquad\quad+\lambda^a\left|\int_0^\infty\varphi(t'+h,\,\mathrm{Proj}_M(y'-z))\hat{\mu}^a(dz)\;-\;\int_0^\infty\varphi(t,\,y-z)\mu^a(dz)\right|\\&\qquad\le\lambda^a|\varphi(t'+h,\,y')-\varphi(t,\,y)|\\&\qquad\quad+\lambda^a\left|\int_0^{M+y'}\varphi(t'+h,\,y'-z)\hat{\mu}^a(dz)\;-\;\int_0^{M+y'}\varphi(t,\,y-z)\mu^a(dz)\right|\\&\qquad\quad+\lambda^a\left|\int_{M+y'}^\infty\varphi(t'+h,\,-M)\hat{\mu}^a(dz)\;-\;\int_{M+y'}^\infty\varphi(t,\,y-z)\mu^a(dz)\right|\\&\qquad\le\lambda^a|\varphi(t'+h,\,y')-\varphi(t,\,y)|\\&\qquad\quad+\lambda^a\int_0^\infty|\varphi(t'+h,\,y'-\kappa(z))\;-\;\varphi(t,\,y-z)|\mu^a(dz)\\&\qquad\quad+2\lambda^a\|\varphi\|_\infty\mu^a([M+y',\infty)),\end{aligned}$$

where we denote by $\kappa(z) = \lfloor \frac{z}{\Delta_Y} \rfloor \Delta_Y$.

Since the smooth function φ has bounded derivatives, say bounded by $\|\varphi^{(1)}\|_\infty$, it follows that

$$\begin{aligned}\left|\mathcal{I}_a^{h,\Delta_Y,M}(t'+h, y', \varphi) - \mathcal{I}^a\varphi(t, y)\right| &\le \lambda^a \|\varphi^{(1)}\|_\infty (h + 2|y' - y| + \Delta_Y) \\ &\quad + 2\lambda^a \|\varphi\|_\infty \mu^a([M + y', \infty)),\end{aligned}$$

which proves that

$$\lim_{\substack{(h,\,\Delta_Y,\,M) \to (0,\,0,\,\infty) \\ (t',\,y') \to (t,\,y)}} \mathcal{I}_a^{h,\Delta_Y,M}(t'+h, y', \varphi) = \mathcal{I}^a\varphi(t, y),$$

hence completing the consistency relation (4.5). □

THEOREM 4.4 (Convergence). *The solution $w^{h,\Delta_Y,M}$ to the numerical scheme ((4.3)–(4.4)) converges locally uniformly to w on $[0, T) \times \mathbb{R}$, as (h, Δ_Y, M) goes to $(0, 0, \infty)$.*

Proof. Given the above monotonicity, stability, and consistency properties, the convergence of the sequence $(w^{h,\Delta_Y,M})$ toward w, which is the unique bounded viscosity solution to (3.11)–(3.12), follows from Barles and Souganidis (1991). We report the arguments for sake of completeness. From the stability property, the semirelaxed limits

$$\begin{aligned}w_*(t, y) &= \liminf_{\substack{(h,\,\Delta_Y,\,M) \to (0,\,0,\,\infty) \\ (t',\,y') \to (t,\,y)}} w^{h,\Delta_Y,M}(t', y'), \\ w^*(t, y) &= \limsup_{\substack{(h,\,\Delta_Y,\,M) \to (0,\,0,\,\infty) \\ (t',\,y') \to (t,\,y)}} w^{h,\Delta_Y,M}(t', y'),\end{aligned}$$

are finite lower semicontinuous and upper semicontinuous functions on $[0, T] \times \mathbb{R}$, and inherit the boundedness of $(w^{h,\Delta_Y,M})$. We claim that w_* are w^* are, respectively, viscosity super and subsolution of (3.11)–(3.12). Assuming for the moment that this claim is true, we obtain by the strong comparison principle for (3.11)–(3.12) that $w^* \le w_*$. Since the converse inequality is obvious by the very definition of w_* and w^*, this shows that $w_* = w^* = w$ is the unique bounded continuous viscosity solution to (3.11)–(3.12), hence completing the proof of convergence.

In the sequel, we prove the viscosity supersolution property of w_* (a symmetric argument for the viscosity subsolution property of w^* holds true). Let $(\bar{t}, \bar{y}) \in [0, T) \times \mathbb{R}$ and φ a test function in $C_b^1([0, T] \times \mathbb{R})$ s.t. $(\bar{t}, \bar{y})$ is a strict global minimum point of $w_* - \varphi$. Then, one can find a sequence (t_n', y_n') in $[0, T) \times \mathbb{R}$, and a sequence (h_n, Δ_Y^n, M_n) such that

$$\begin{aligned}(t_n', y_n') \to (\bar{t}, \bar{y}), \quad (h_n, \Delta_Y^n, M_n) \to (0, 0, \infty), \quad w^{h_n,\Delta_Y^n,M_n} \to w_*(\bar{t}, \bar{y}), \\ (t_n', y_n') \text{ is a global minimum point of } w^{h_n,\Delta_Y^n,M_n} - \varphi.\end{aligned}$$

Denoting by $\zeta_n = (w^{h_n,\Delta_Y^n,M_n} - \varphi)(t_n', y_n')$, we have $w^{h_n,\Delta_Y^n,M_n} \ge \varphi + \zeta_n$. From the definition of the numerical scheme $\mathcal{S}^{h_n,\Delta_Y^n,M_n}$, and its monotonicity, we then get

$$\begin{aligned}\zeta_n + \varphi(t_n', y_n') &= w^{h_n,\Delta_Y^n,M_n}(t_n', y_n') \\ &= \mathcal{S}^{h_n,\Delta_Y^n,M_n}(t_n' + h_n, y_n', w^{h_n,\Delta_Y^n,M_n})\end{aligned}$$

$$\geq \mathcal{S}^{h_n,\Delta_Y^n,M_n}\left(t_n'+h_n,\, y_n',\, \varphi+\zeta_n\right) = \mathcal{S}^{h_n,\Delta_Y^n,M_n}(t_n'+h_n,\, y_n',\, \varphi)+\zeta_n$$
$$= \max\left[\mathcal{T}^{h_n,\Delta_Y^n,M_n}(t_n'+h_n,\, y_n',\, \varphi)\ ,\ \tilde{\mathcal{M}}^{h_n,\Delta_Y^n,M_n}(t_n'+h_n,\, y_n',\, \varphi)\right]+\zeta_n,$$

which implies

$$\min\left[\frac{\varphi(t_n',\, y_n')-\mathcal{T}^{h_n,\Delta_Y^n,M_n}(t_n'+h_n,\, y_n',\, \varphi)}{h_n}\ ,\ \varphi(t_n',\, y_n')-\tilde{\mathcal{M}}^{h_n,\Delta_Y^n,M_n}(t_n'+h_n,\, y_n',\, \varphi)\right]\ \geq\ 0.$$

By the consistency properties (4.5)–(4.6), and by sending n to infinity in the above inequality, we obtain the required viscosity supersolution property:

$$\min\left[-\frac{\partial\varphi}{\partial t}(\bar{t},\, \bar{y})-\bar{y}c_{_P}+\gamma\varrho\bar{y}^2-\mathcal{I}^a\varphi(\bar{t},\, \bar{y})-\mathcal{I}^b\varphi(\bar{t},\, \bar{y})\ ,\ \varphi(\bar{t},\, \bar{y})-\tilde{\mathcal{M}}\varphi(\bar{t},\, \bar{y})\right]\geq 0.$$

□

4.3. Numerical Tests

In this section, we provide numerical results for the (reduced-form) value function and optimal policies in the martingale case and the trend information case, and a backtest on simulated data for the trend information case.

Within this section, we will denote by w^h the value function and by $\alpha^\star$ the make/take strategy associated with the backward numerical scheme (4.3)–(4.4). Given a generic controlled process Z and a control $\alpha\in\mathcal{A}$, we will denote Z^α the process controlled by α. Unless specified otherwise, such processes will be supposed to start at zero: typically, we assume that the investor starts from zero cash and zero inventory at date $t=0$ in the following numerical tests. Finally, we will write indifferently $w^h(t,\, y,\, c_{_P})$ or $w^h(t,\, y):=w^h(t,\, y,\, 0)$ to either stress or omit the dependence in $c_{_P}$.

• **The martingale case**: in the martingale case, we performed the algorithm (4.3)–(4.4) with parameters shown in Table 4.1. This set of parameters is chosen to be consistent with

TABLE 4.1
Parameters for Numerical Results in the Martingale Case

Parameter	Value
(a) Market and Risk Parameters	
δ	12.5 EUR/contract
ε	1.05 EUR/contract
ε_0	0
λ	0.05/s
μ	$\exp(1/\bar{\mu})$
$\bar{\mu}$	20 contracts
γ	2.5×10^{-5}
T	100 s
(b) Discretization Parameters	
N_Y	100
N_T	500

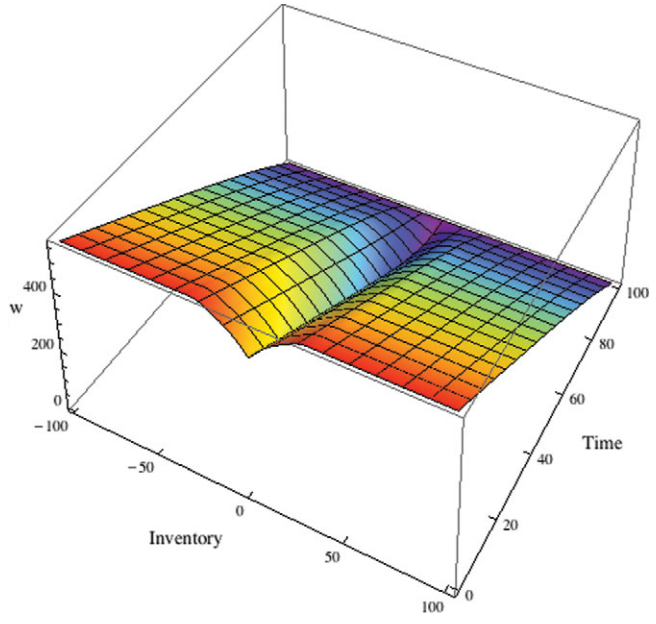

FIGURE 4.1. Reduced form value function w^h.

calibration data on the front maturity for 3-month EURIBOR future, see, for example, Field and Large (2008).

Figure 4.1 displays the reduced-form value function w^h on $[0, T] \times [-N_Y; N_Y]$. This result illustrates the linear bound (3.13) as noticed in Proposition 3.3, and also the symmetry of w^h as pointed out in (3.17). We also observe the monotonicity over $\mathbb{R}_+$ and $\mathbb{R}_-$ of the value function $w^h(t, .)$.

In Figure 4.2, we display the optimal take and make policies. The optimal take policy is represented as the volume to buy or sell with a market order, as a function of the time and inventory $(t, y) \in [0, T] \times [-N_Y; N_Y]$. We notice that a market order only occurs when the inventory becomes too large, and therefore, the take policy can be interpreted as a "stop-loss" constraint, i.e., an emergency rebalancing of the portfolio when the inventory risk is too large.

The optimal make policy is represented as the regime of limit orders posting as a function of the time and inventory $(t, y) \in [0, T] \times [-N_Y; N_Y]$. For sake of simplicity, we represented the sum of ℓ^a and ℓ^b on the map. The meaning of this surface is as follows: 0 means that there is no active limit orders on either side, 2 means that there is active limit orders on both bid and ask sides, and 1 means that there is only one active limit order either on the bid or the ask side, depending on the sign of y (if $y < 0$ only the bid side is active, and if $y > 0$ only the ask side is active). We notice that when close to maturity T, the optimal strategy tends to be more aggressive, in the sense that it will seek to get rid of any positive or negative inventory, to match the terminal liquidation constraint. Moreover, we notice that close to date 0, the dependence in t seems to be negligible, which indicates that a "stationary regime" may be attained for large T.

• **The trend information case:** in this case, we provide a backtest of the optimal strategy on simulated data in addition to the plot of the optimal policy $\alpha^\star$. We kept the same

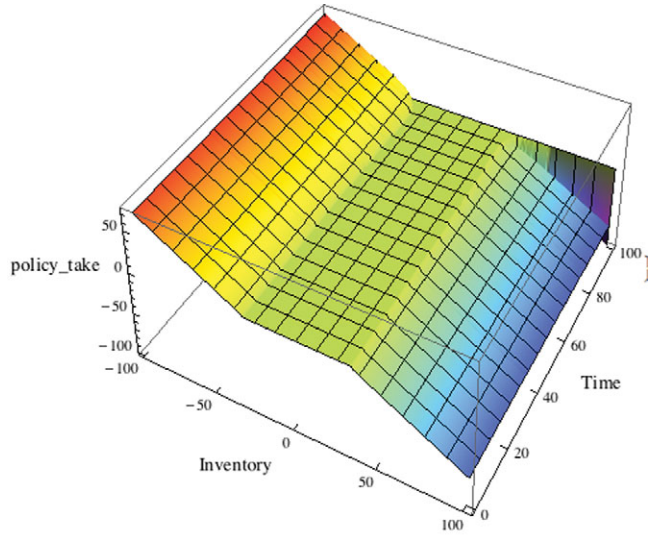

(a) Optimal take policy.

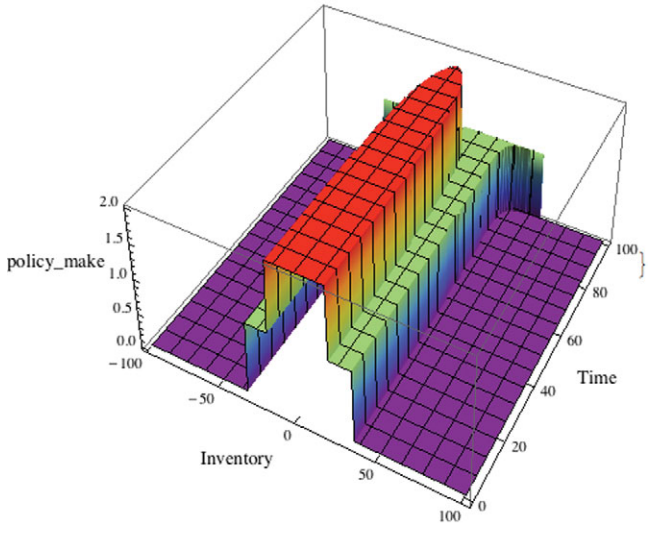

(b) Optimal make policy

FIGURE 4.2. Numerical results for the martingale case: representation of optimal make and take policies $\alpha^\star$.

parameters for execution intensity and volume, price characteristics and costs, but we chose a wider time period $T = 2{,}000$ s in order to observe multiple trade events. With this set of parameters, we expect to observe about 100 trade events of average volume 20. Note that the execution intensity $\lambda = 0.05$, a value consistent with market activity of the front quarterly EURIBOR future, is independent in our model to the trend information ϖ that we will describe below.

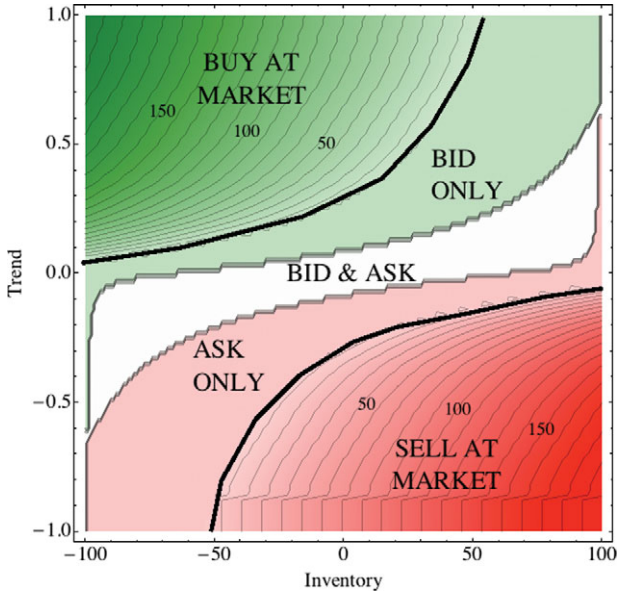

FIGURE 4.3. Optimal policy $\alpha^\star$ at date $t = 0$.

Figure 4.3 displays the optimal policy at date $t = 0$, in the plane (y, c_P). The policy has central symmetry properties as expected in (3.17), and should be read as follows: dark green zones represent a situation where a market order to buy must be sent, light green means that a limit order is active only at bid, white means that limit orders are active on both sides, light red means that a limit order is active only at ask, and dark red means that a market order to sell must be sent. Let us provide a qualitative example: assume that after the high-frequency trader acquired a positive inventory, she is informed that price is going down, i.e., negative drift $c_P < 0$. Therefore, the optimal strategy will be either to cancel the bid limit order (light red zone) and keep ask limit order active, or depending on the value of $|c_P|$, send a market order to get rid of our positive inventory (dark red zone).

We performed a benchmarked backtest on simulated data and a performance analysis in this case. The first benchmark strategy $\alpha^{\mathrm{WoMO}} = (\alpha^{\mathrm{make\,,\,WoMO}}, 0)$ corresponds to the case where we do not allow the high-frequency trader to use market orders. It is computed using the backward numerical scheme (4.3)–(4.4), but without taking into account the obstacle part, which is equivalent to setting $\varepsilon_0 = \infty$. The second benchmark strategy is made of constant controls (also known as *symmetric* or *constant* strategy):

$$\alpha^{\mathrm{cst}} := (\alpha^{\mathrm{make\,,\,cst}}, 0)$$
$$\alpha^{\mathrm{make\,,\,cst}} := (1, 1).$$

In order to make our simulated data backtest closer to reality, we chose to slightly deviate from the original price model, and use a varying price trend. We simulate a price process model given by

$$\hat{P}_t = \hat{P}_0 + \delta(N_t^+ - N_t^-),$$

where N^+ and N^- are the Euler scheme simulation of Cox processes of respective intensities π^+ and π^- defined as follows:

$$\pi^+ + \pi^- \equiv K = \varrho/\delta^2$$
$$d\pi_t^+ - d\pi_t^- := d\varpi_t = -\theta\varpi_t dt + \sigma dB_t,$$

where $K > 0, \theta > 0$, and $\sigma > 0$ are positive constants, and B is an independent Brownian motion. Note that we chose the sum $\pi^+ + \pi^-$ to be the constant K, for simplicity sake: it means that, disregarding the direction of price variation, the mean number of price change per second is assumed to be constant $\mathbf{P}(|P_{t+h} - P_t| = \delta) = Kh + o(h)$, which provides an easy way to calibrate the parameter K while reducing the dimension of the simulation. The interpretation of this simulation model is as follows: we add an exogenous risk factor B, which drives the price trend information ϖ as an Ornstein–Uhlenbeck process. Notice that this supplementary risk factor B is not taken into account in our optimization procedure and thus has a penalizing impact on the strategy's performance: therefore it does not spoil the backtest. This model choice for the process (ϖ_t) is a convenient way to simulate the real-world situation, where the high-frequency trader continuously updates her predictive information about short-term price movements, based, e.g., on the current state of the LOB.

Therefore, qualitatively speaking, our optimization procedure is consistent with this simulation model if we choose θ and σ s.t. The variation of the (reduced-form) value function w due to predictive information is very small compared to the variation of the value function due to other market events (e.g., an execution event).

This assumption is consistent with HFT practice since the HF trader is able to adapt very quickly to a modification of this predictive information. Backtest parameters involved in this simulation are shown in Table 4.2.

The interpretation of the trend information parameters is the following: independently from the trade intensity $\lambda = 0.05$, we consider the price trend, which is interpreted as the expected return of the mid-price over the next few milliseconds, and is directed by the state variable ϖ. In the stationary regime, this variable ϖ has a marginal distribution $\mathcal{L}(\varpi_t)$ which is essentially a centered normal law of standard deviation $\sigma/\sqrt{2\theta} \simeq 0.32$ with this set of parameters. Qualitatively speaking, using the 2-sigma rule, this means that the process ϖ spends most of the time in the range -0.6 to 0.6. The value $\varpi_t = 0.6$ (resp. $\varpi_t = -0.6$) represents qualitatively a 60% probability of an uptick (resp. a downtick) in the next second. Such signal can be computed, for example, using the methods developed in Cont and de Larrard (2013). Moreover, ϖ is a mean reverting process, of reversion

TABLE 4.2
Backtest Parameters

Parameter	Value
K	0.2
θ	0.2/s
σ	0.2/s
N_{MC}	50,000
N_ϖ	50

TABLE 4.3
Synthesis Table for Backtest. Categories Are, from Top to Bottom: Relative Performance Metrics, Period-Adjusted Performance Metrics, Absolute Performance Metrics and Absolute Activity Metrics

Quantity	Definition	$\alpha^\star$	α^{WoMO}	α^{cst}
Info ratio over T	$m(\hat{V}_{\dot{T}})/\sigma(\hat{V}_{\dot{T}})$	3.67	0.89	0.18
Profit per trade	$m(\hat{V}_{\dot{T}})/m(\hat{Q}^{\text{total},\cdot})$	8.06	16.31	5.57
Risk per trade	$\sigma(\hat{V}_{\dot{T}})/m(\hat{Q}^{\text{total},\cdot})$	2.19	18.31	29.56
Mean performance	$m(\hat{V}_{\dot{T}})$	31,446.4	28,246.3	21,737.2
Standard deviation of perf	$\sigma(\hat{V}_{\dot{T}})$	8,555.46	31,701.2	115,312
Skew of perf	$\text{skew}(\hat{V}_{\dot{T}})$	0.64	0.16	−0.007
Kurtosis of perf	$\text{kurt}(\hat{V}_{\dot{T}})$	3.82	3.31	7.02
Mean total executed volume	$m(\hat{Q}^{\text{total},\cdot})$	3,900.68	1,730.82	3,900.61
Mean at market volume	$m(\hat{Q}^{\text{market},\cdot})$	1,932.29	0	0
Ratio market over total exec	$m(\hat{Q}^{\text{market},\cdot})/m(\hat{Q}^{\text{total},\cdot})$	0.495	0	0

speed 0.2/s, which can be qualitatively interpreted as the timescale during which a prediction remains valid, in this case $\frac{1}{\theta_\varpi} = 5$ s. This can be viewed as the timescale on which the high-frequency trader will update her prediction about the price trend. Note that in the case of STIR futures trading, this choice of reversion speed is consistent with other market activity statistics: indeed, this reversion speed is greater than mid-price update intensity (of order 0.01/s) and smaller than order book update intensity (of order 1 to 10/s) (see Field and Large 2008, for precise statistics).

Let us denote by $\hat{\vartheta}^a$ and $\hat{\vartheta}^b$ the Euler scheme simulation of the compound Poisson processes ϑ^a and ϑ^b, with dynamics (2.3). Therefore, for $\alpha \in \{\alpha^\star, \alpha^{\text{WoMO}}, \alpha^{\text{cst}}\}$, we were able to compute the Euler scheme simulation $\hat{X}^\alpha$ (resp. $\hat{Y}^\alpha$) of X^α (resp. Y^α), starting at 0 at $t = 0$, by replacing ϑ^a (resp. ϑ^b) by $\hat{\vartheta}^a$ (resp. $\hat{\vartheta}^b$) in equation (2.4) (resp. (2.5)).

We performed N_{MC} simulation of the above processes. For each simulation $\omega \in [1 \dots N_{MC}]$ and for $\alpha \in \{\alpha^\star, \alpha^{\text{WoMO}}, \alpha^{\text{cst}}\}$, we stored the following quantities: the terminal wealth after terminal liquidation $\hat{V}_T^\alpha(\omega) := L(\hat{X}^\alpha(\omega), \hat{Y}^\alpha(\omega), \hat{P}(\omega))$, called "performance" in what follows; the total executed volume $\hat{Q}^{\text{total},\alpha}(\omega) := \sum_{[0,T]} |\hat{Y}_t^\alpha(\omega) - \hat{Y}_{t-}^\alpha(\omega)|$; and the volume executed at market $\hat{Q}^{\text{market},\alpha}(\omega) := \sum_{[0,T]} |\xi_n(\omega)^\alpha|$. Finally, we denote by $m(.)$ the empirical mean, by $\sigma(.)$ the empirical standard deviation, by skew(.) the empirical skewness, and by kurt(.) the empirical kurtosis, taken over $\omega \in [1 \dots N_{MC}]$.

Table 4.3 displays a synthesis of descriptive statistics for this backtest. We first notice that the information ratio over T of $\alpha^\star$ is more than four times that of α^{WoMO}, which itself is about four times that of α^{cst}. Second, the per trade metrics can be compared to the half-spread $\frac{\delta}{2} = 6.25$ EUR/contract, and we see that although the mean profit per trade is smaller for the optimal strategy, the risk associated to each trade is dramatically reduced compared to the benchmark.

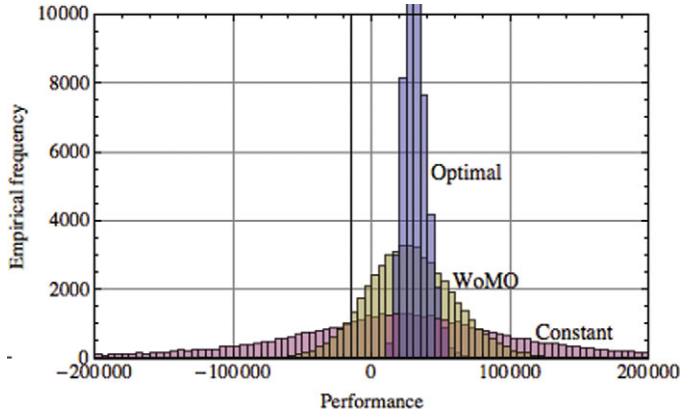

FIGURE 4.4. Empirical distribution of performance $\hat{V}_{\dot{T}}$. The graph shows the number of occurrences for each bin on $N_{MC} = 50{,}000$ simulations.

TABLE 4.4
Varying Risk-Aversion Parameter γ: Data

γ	$\alpha^\star : \sigma(\hat{V}_{\dot{T}})$	$\alpha^\star$: info ratio	$\alpha^{\text{WoMO}} : \sigma(\hat{V}_{\dot{T}})$	α^{WoMO}: info ratio
0	8,585.7	3.68	32,095.7	0.88
1×10^{-4}	8,427.4	3.74	30,613.6	0.91
2×10^{-4}	8,152.1	3.83	29,043.0	0.94
3×10^{-4}	7,880.1	3.88	27,322.3	0.98
4×10^{-4}	7,594.4	3.92	25,752.1	1.02
5×10^{-4}	7,252.7	4.00	24,300.8	1.06
1×10^{-3}	5,538.2	4.31	18,724.3	1.29
2×10^{-3}	3,730.8	4.62	13,882.4	1.51
3×10^{-3}	2,785.9	4.71	12,156.6	1.55
4×10^{-3}	2,255.1	4.65	11,428.1	1.51
5×10^{-3}	1,845.4	4.71	10,858.2	1.47
1×10^{-2}	1,014.4	4.86	9,596.6	1.28

This is confirmed by the empirical distribution of performance, also shown in Figure 4.4, where the dark blue represents the performance distribution of the optimal strategy, the light yellow represents the performance distribution of the WoMO strategy and the light purple represents the performance distribution of the benchmark strategy. We see that not only benchmark has higher standard deviation, but also higher excess kurtosis and heavy tails: this is due to the fact that inventory can be very large for the constant strategy, and therefore it bears a nonnegligible market risk (or inventory risk). Finally, we see that about 49% of the trades are done with market orders.

Our last numerical test is devoted to displaying the influence of the risk-aversion parameter γ. All other parameters remaining the same, we tested several values of γ (as indicated in Table 4.4), and characterized the performance of the corresponding strategy by the pair $(\sigma(\hat{V}_{\dot{T}}), m(\hat{V}_{\dot{T}}))$, which gives the *efficient frontier* plot displayed in Figure 4.5.

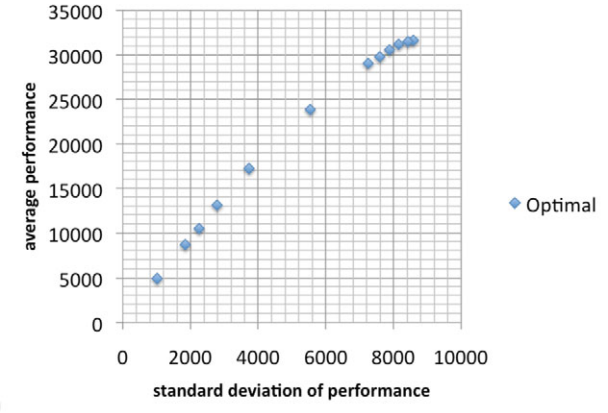

(a) Optimal

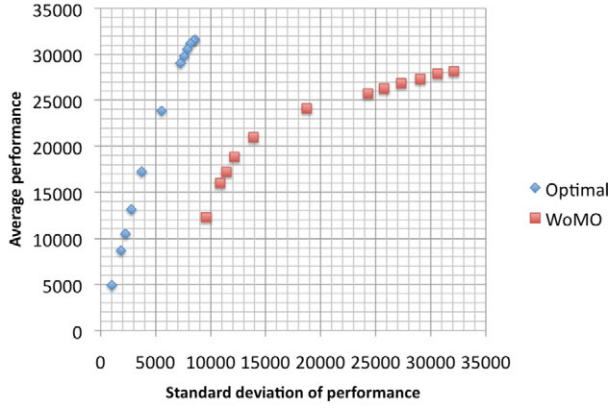

(b) Optimal and WoMO

FIGURE 4.5. Varying parameter γ for $\alpha^\star$ and α^{WoMO}. The x-axis represents the standard deviation of performance $\sigma(\hat{V}_T)$ and the y-axis the average performance $m(\hat{V}_T)$ estimated on $N_{MC} = 50{,}000$ simulations of the model.

We measure the performance of each strategy empirically, both $\alpha^\star$ and α^{WoMO}, by running N_{MC} simulations of our market model, and therefore we can observe a slight measurement error on the points $(\sigma(\hat{V}_T), m(\hat{V}_T))$. As expected, a reduction of γ increases the standard deviation of the strategy: this is due to the fact that a small γ allows for large open position, i.e., large inventory, and therefore the market risk is greater. For small γ, performance is also better since the investor can sustain large inventories, and therefore is less impatient to get rid of it: in particular, the proportion of volume executed at market is increasing in γ. In real trading condition, the value of γ should be tuned to attain the desired ratio of mean/volatility of PnL. Cartea and Jaimungal (2015) observe similar behavior and determine these frontiers analytically when market orders are absent. We also display in Figure 4.5 the plot for α^{WoMO}, that clearly exhibits a larger risk, which indicates that the market orders in our optimal strategy are not only used to gain an extra performance, but

also cut part of the risk of holding a nonzero position. Table 4.4 shows that if α^{WoMO} and $\alpha^\star$ are comparable in terms of mean performance, the standard deviation of $\alpha^\star$ is greatly reduced compared to α^{WoMO}. This is reflected in significantly better information ratios for $\alpha^\star$. The choice of our simulation parameters explains the scale of improvements in $\alpha^\star$ compared to α^{WoMO}: indeed, the choice $\lambda = 0.05/\text{s}$ implies that the typical delay between two limit orders execution is about 20 s, when our parameterization for ϖ implies that the typical delay between two price updates is about 1 s. Therefore, α^{WoMO}, using solely limit orders, is much slower (in terms of inventory adjustments) than the market's typical update speed, which is reflected in a much greater variance in α^{WoMO} than in $\alpha^\star$. For this set of simulation parameters, we observe that for $\gamma > 2 \times 10^{-2}$ (200 bps) the strategy α^{WoMO} will not trade anymore when starting in $t = 0$ with a zero inventory $y = 0$, and will have a zero mean and a zero standard deviation of performance. This is due to the fact that, in this case, ask and bid limit order posting zones in α^{WoMO} do not contain the line $(t, 0)\,,\ t \in [0, T]$. Moreover, the standard deviation of α^{WoMO} does not go under a certain bound, approximatively 9,000 reached for $\gamma = 10^{-2}$ (100 bp). Indeed, this strategy, using only limit orders, cannot immediately cut its inventory with a market order, and so any trade from the point $y = 0$ will lead to keep a nonzero position on the risky asset for a noncontrollable time period. On the contrary, $\alpha^\star$ can accept larger values of γ, up to 350 bp approximatively, before reaching the no-trade behavior, and increasing the value of γ leads to reducing the standard deviation of performance, following the trend shown in Figure 4.5.

5. BEST EXECUTION PROBLEM AND OVERTRADING RISK

In this section, we apply our market model framework to a best execution problem. The trading objective of the investor is to liquidate $Y_0 > 0$ assets over the finite time interval $[0, T]$. She is not allowed to purchase stock during the liquidation period, and may only buy back the asset in case of short position. In this context, the investor posts continuously a limit sell order (with a volume much larger that the required quantity Y_0) at the best ask price, and also runs market (sell) orders strategy until she reaches either a negative inventory or the terminal date. By doing so, she hopes to trade as much as possible at the ask price, and therefore avoiding to *cross the spread*.

Mathematically, this means that the investor uses a subset $\mathcal{A}_\ell$ of strategies $\alpha = (\alpha^{\text{make}} = (L^a, L^b), \alpha^{\text{take}})$ in $\mathcal{A}$ such that

$$(L^a_t, L^b_t) = \begin{cases} (1, 0) & \text{for } t \le \tau, \\ (0, 0) & \text{for } t > \tau, \end{cases}$$

$$\alpha^{\text{take}} = (\tau_n, \zeta_n)_n \cup (\tau, -Y_\tau), \quad \text{with } \tau_n < \tau,\ \zeta_n < 0,$$

where $\tau = \inf\{t \ge 0 :\ Y_t \le 0\} \wedge T$. The value function associated to this liquidation problem is then defined by

$$v_\ell(t, x, y, p) = \sup_{\alpha \in \mathcal{A}_\ell} \mathbf{E}_{t,x,y,p}\left[L(X_T, Y_T, P_T) - \gamma \int_t^T Y_s^2 \varrho(P_s) ds \right] \tag{5.1}$$

for $(t, x, y, p) \in [0, T] \times \mathbb{R}^2 \times \mathbb{P}$. With the notation in (3.6), the operator corresponding to the limit order in $\mathcal{A}_\ell$ is given by $\mathcal{L}^{1,0} = \mathcal{P} + \Gamma^a$, while the impulse operator associated

to the market order in $\mathcal{A}_\ell$ is defined by

$$\mathcal{M}_\ell\,\varphi(t,x,y,p) = \sup_{e\in[-M_Y-y,-(M_Y-y)_-]\setminus\{0\}} \varphi\left(t, x-ep-|e|\left(\frac{\delta}{2}+\varepsilon\right)-\varepsilon_0, y+e, p\right),$$

where $m_- = \max(-m,0)$. The DPE associated to (5.1) takes the form

$$\min\left[-\frac{\partial v_\ell}{\partial t} - \mathcal{P}v_\ell - \Gamma^a v_\ell \;+\; \gamma g,\; v_\ell - \mathcal{M}_\ell v_\ell\right] = 0, \text{ on } [0,T)\times\mathbb{R}\times(0,\infty)\times\mathbb{P}$$

together with the terminal and boundary conditions

$$v_\ell = L, \quad \text{on } (\{T\}\times\mathbb{R}\times\mathbb{R}\times\mathbb{P})\cup([0,T)\times\mathbb{R}\times\mathbb{R}_-\times\mathbb{P}).$$

The above boundary condition for nonpositive inventory is related to the *overtrading risk*, which is the risk that the investor sold too much assets via the (oversized) limit order at the best ask price. This risk occurs typically in execution problems on pro rata LOB (see Field and Large 2008).

Again, in the Lévy case (3.9), the value function v_ℓ is reduced into

$$v_\ell(t,x,y,p) = L_0(x,y,p) + w_\ell(t,y),$$

where w_ℓ is solution to the integro-variational inequality

$$\begin{aligned}\min\Bigg\{ &-\frac{\partial w_\ell}{\partial t} - yc_{_P} + \gamma\varrho y^2 \\ &-\lambda^a\int_0^\infty\left[w_\ell(t,y-z)-w_\ell(t,y)+z\frac{\delta}{2}+\left(\frac{\delta}{2}+\varepsilon\right)(|y|-|y-z|)\right]\mu^a(dz); \\ w_\ell(t,y) - &\sup_{e\in[-M_Y-y,-(M_Y-y)_-]\setminus\{0\}}\left[w_\ell(t,y+e)-\left(\frac{\delta}{2}+\varepsilon\right)(|y+e|+|e|-|y|)-\varepsilon_0\right]\Bigg\} = 0\end{aligned}$$

for $(t,y)\in[0,T)\times(0,\infty)$, together with the terminal and boundary conditions

$$w_\ell(t,y) = -\varepsilon_0 1_{y\neq 0}, \quad \forall (t,y)\in\big(\{T\}\times\mathbb{R}\big)\cup\big([0,T)\times\mathbb{R}_-\big).$$

The associated numerical scheme reads now as follows:

$$\begin{aligned}&w_\ell^h(t_N,y) = -\varepsilon_0 1_{y\neq 0}, \quad y\in\mathbb{R},\\ &w_\ell^h(t_k,y) = 0, \quad k=0,\dots,N-1,\ y\le 0,\\ &w_\ell^h(t_k,y) = \max\big[\mathcal{T}_\ell^{h,\Delta_Y,M}(t,y,\varphi)\,;\,\mathcal{M}_\ell^{h,\Delta_Y,M}(t,y,\varphi)\big],\ k=0,\dots,N-1\,,\ y\in\mathbb{Y}_M^+,\end{aligned}$$

where $\mathbb{Y}_M^+ = \mathbb{Y}_M\cap\mathbb{R}_+$,

$$\begin{aligned}\mathcal{T}_\ell^{h,\Delta_Y,M}(t,y,\varphi) = \varphi(t,y) &- h\gamma\varrho y^2 + hyc_P \\ &+\lambda^a h\Bigg(\int_0^\infty\big[\varphi(t,\mathrm{Proj}_M(y-z))-\varphi(t,y)\big]\hat{\mu}^a(dz) \\ &+\int_0^\infty\left[\frac{\delta}{2}z+\left(\frac{\delta}{2}+\varepsilon\right)(|y|-|y-z|)\right]\mu^a(dz)\Bigg)\end{aligned}$$

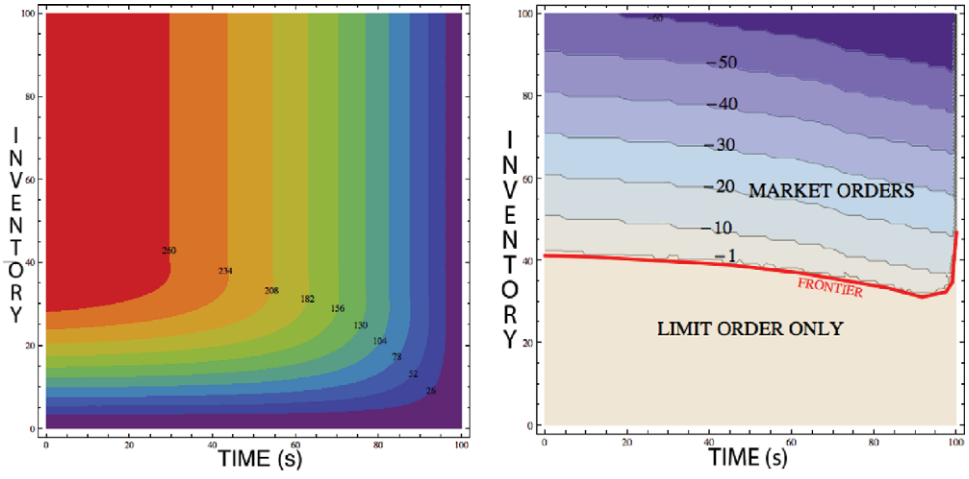

(a) Value function (b) Policy

FIGURE 5.1. Numerical results for the simple liquidation problem (for $c_P = 0$). On the left side, level lines are indicated for the value function w_ℓ. On the right side, numbers labeled represent the quantity to sell in the optimal market order control.

and

$$\mathcal{M}_\ell^{h,\Delta_Y,M}(t,y,\varphi) = \sup_{e\in\mathbb{Y}_M\cap[-M_Y-y,-(M_Y-y)_-]\setminus\{0\}} \left[\varphi(t,\mathrm{Proj}_M(y+e)) - \left(\frac{\delta}{2}+\varepsilon\right)(|y+e|+|e|-|y|) - \varepsilon_0\right].$$

In this case, the optimal policy shown in Figures 5.1 and 5.1b is simple to describe. The state space is delimited in two zones: when the inventory is small, the HFT must wait for her limit sell order to be executed; and when the inventory is large, the HFT must send a market sell order to avoid the market risk related to holding a large position.

The frontier between the two zones (indicated in bold red in Figure 5.1b) can be interpreted as an *optimal trading curve*, a concept that is extensively documented (see, e.g., Guéant et al. 2013) in the optimal execution literature. The optimal trading curve is the inventory that the investor should hold, seen as a function of time, in order to minimize overall trading costs. Therefore, in the typical setting, the execution strategy consists in trading via market orders to get as close as possible to the optimal trading curve. Similarly, in our case, we can see on Figure 5.1b that the optimal strategy will behave similarly for large inventories (i.e., when above the trading curve): indeed, we observe that the quantities to sell are such that the market orders strategy would keep the inventory close to the optimal trading curve, if no limit orders were allowed. Now, in our case, we observe two specific features of the optimal strategy: (1) the optimal trading curve does not reach 0 at maturity, and therefore the HFT has to get rid of her inventory at market at final date to match the constraint $Y_T = 0$. This is due to the fact that a supplemental gain is always achievable when the limit order is executed. Therefore, this feature leads to an execution strategy where the final trade is bigger than intermediary trades; (2) below the optimal trading curve, i.e., in the region where the HFT trades via limit orders only, the sell limit order is always active, and can lead to

an execution. Therefore, the inventory is always below the optimal trading curve, and the distance between the current inventory and the optimal trading curve equals the volume executed via limit orders. This differs from classic pattern-based best execution strategies, for example, the U-shaped execution strategy that consists in trading a large quantity at the beginning and at the end of the liquidation, and trade regularly small quantities in between. Indeed, the optimal strategy does not provide a fixed pattern for every execution, but provides the optimal action to take given the observation of the inventory that is still to be sold and the market characteristics as, e.g., the mean traded volume at ask per second $\lambda^a \bar{\mu}^a$, or trades volume distributions at ask μ^a.

Finally, let us notice that this strategy can be interpreted as a convenient way to avoid the cost of crossing the spread during the liquidation of a portfolio, but we did not take into account the impact of the market order on the transaction price. In the case of a pro rata microstructure, available volumes offered at best prices are usually about 200 times larger than the mean volume of market orders (see Field and Large 2008), and therefore it is consistent to consider that there is no impact on the price for our market orders. Yet, the model can easily be modified by adding an impact component in the obstacle operator $\mathcal{M}_\ell$ to take care of this effect. We also did not model the possibility that the intensities λ^a and λ^b of execution processes may vary, and postpone this investigation for future research.

6. CONCLUSION

In this paper, we investigate a framework to build up mixed HFR strategies in an exotic microstructure, the pro rata microstructure. This microstructure can be encountered, for example, on STIRs futures. We consider the situation of an investor willing to maximize her terminal profit over a finite time horizon, and able to trade with limit and market orders. We adopt the perspective of inventory management, which means that the investor's primary objective is to keep her position on the risky asset close to zero at all times, in order to avoid being exposed to market risk.

We provide a tractable market model that mimics the major features of our target microstructure, while being parsimonious enough to fit a large range of products. We detail the optimization procedure, by means of stochastic control, as well as the numerical scheme used to solve the resulting HJB equation. Dimension reduction techniques as well as interpretable decomposition of the profit's dynamics are described. We also discuss the practical implementation of such strategy.

In this particular microstructure, we are able to define and address two specific types of risk: the overtrading risk, which is the risk of brutal variations in the investor inventory, due to the fact that she does not control the quantity traded at limit; and the adverse selection risk, which is the risk of market reacting unfavorably to the investor quotes.

For this last purpose, we introduce a new state variable, which we interpret as a predictive price indicator, that allows us to balance our position before the price changes. This last feature also provides an extra performance on our empirical tests. We provide several examples of application of our framework, including a mixed limit/market strategy when no information is available on price, a mixed strategy with superior information on price, and a liquidation strategy without information on price. Moreover, we point out the advantages of using market orders in this setup by benchmarking the performance of our strategy against a pure limit order strategy, and we find that the mean/volatility ratio is much smaller in this last case.

REFERENCES

AIKIN, S. (2006): *Trading STIR Futures: An Introduction to Short-Term Interest Rate Futures*, Petersfield, UK: Harriman House Publishing.

AVELLANEDA, M., and S. STOIKOV (2008): High Frequency Trading in a Limit Order Book, *Quant. Finance* 8(3), 217–224.

BARLES, G., and P. SOUGANIDIS (1991): Convergence of Approximation Schemes for Fully Nonlinear Second Order Equations, *Asymptotic Anal.* 4, 271–283.

CARTEA, A., and S. JAIMUNGAL (2013): Modeling Asset Prices for Algorithmic and High Frequency Trading, *Appl. Math. Finance* 20(6), 512–547.

CARTEA, A., and S. JAIMUNGAL (2015): Risk Metrics and Fine Tuning of High Frequency Trading Strategies, *Math. Finance* 25(3), 576–611.

CARTEA, A., S. JAIMUNGAL, and J. RICCI (2014): Buy Low Sell High: A High Frequency Trading Perspective, *SIAM J. Financ. Math.* 5(1), 415–444.

CONT, R., and A. DE LARRARD (2013): Price Dynamics in a Markovian Limit Order Market, *SIAM J. Financ. Math.* 4(1), 1–25.

FIELD, J., and J. LARGE (2008): Pro-Rata Matching and One-Tick Futures Markets, Preprint.

GOULD, M. D., M. A. PORTER, S. WILLIAMS, M. MCDONALD, D. J. FENN, and S. D. HOWISON (2010): The Limit Order Book: A Survey, Preprint.

GUÉANT, O., J. FERNANDEZ TAPIA, and C.-A. LEHALLE (2011): Dealing with Inventory Risk, Preprint.

GUILBAUD, F., and H. PHAM (2013): Optimal High Frequency Trading with Limit and Market Orders, *Quant. Finance* 13(1), 79–94.

HO, T., and H. R. STOLL (1981): Optimal Dealer Pricing under Transactions and Return Uncertainty, *J. Financ. Econ.* 9, 47–73.

JANECEK, K., and M. KABRHEL (2007): *Matching Algorithms of International Exchanges*. Working Paper.

KÜHN, C., and M. STROH (2010): Optimal Portfolios of a Small Investor in a Limit Order Market: A Shadow Price Approach, *Math. Financ. Econ.* 3(2), 45–72.

ØKSENDAL, B., and A. SULEM (2007): *Applied Stochastic Control of Jump Diffusions*, Universitext, Berlin, Heidelberg, New York: Springer.

PROTTER, P. (2005): *Stochastic Integration and Differential Equations*, 2nd ed., Version 2.1, Heidelberg: Springer-Verlag.

VERAART L. A. M. (2011): Optimal Investment in the Foreign Exchange Market with Proportional Transaction Costs, *Quant. Finance* 11(4), 631–640.

MATHEMATICAL FINANCE
An International Journal of Mathematics, Statistics and Financial Economics

Mathematical Finance, Vol. 25, No. 3 (July 2015), 576–611

RISK METRICS AND FINE TUNING OF HIGH-FREQUENCY TRADING STRATEGIES

ÁLVARO CARTEA

University College London

SEBASTIAN JAIMUNGAL

University of Toronto

We propose risk metrics to assess the performance of high-frequency (HF) trading strategies that seek to maximize profits from making the realized spread where the holding period is extremely short (fractions of a second, seconds, or at most minutes). The HF trader maximizes expected terminal wealth and is constrained by both capital and the amount of inventory that she can hold at any time. The risk metrics enable the HF trader to fine tune her strategies by trading off different metrics of inventory risk, which also proxy for capital risk, against expected profits. The dynamics of the midprice of the asset are driven by information flows which are impounded in the midprice by market participants who update their quotes in the limit order book. Furthermore, the midprice also exhibits stochastic jumps as a consequence of the arrival of market orders that have an impact on prices which can give rise to market momentum (expected prices to trend up or down). The HF trader's optimal strategy incorporates a buffer to cover adverse selection costs and manages inventories to maximize the expected gains from market momentum.

KEY WORDS: algorithmic trading, high-frequency trading, momentum trading, market impact, adverse selection, risk metrics, inventory risk.

1. INTRODUCTION

Computerized trading has revolutionized the way in which financial markets work. Nowadays, all major exchanges operate on electronic platforms where market participants trade using powerful hardware and customized software. A general class of computerized trading is known as algorithmic trading (AT) which refers to algorithms that are designed to make trading decisions and to manage inventories without human intervention. It is reported that over 60% in the United States, and over 40% in Europe, of the trading volume in equities is due to AT, a figure that is expected to increase (CFTC and SEC 2010). Within the AT class there is a more narrowly defined set of strategies, known as high-frequency (HF) trading, that relies on ultra-high speed to process information in

The authors would like to thank Rob Almgren, Tomasz Bielecki, Adrien De Larrard, Jason Ricci, and participants at the SIAM Conference on Financial Mathematics and Engineering 2012, Young Researchers Workshop on Finance Tokyo 2012, and Universidad Carlos III, Madrid. As well, the authors thank two anonymous referees for their comments which ultimately improved this paper. Finally, SJ thanks NSERC and Mprime for partially funding this work.

Manuscript received March 2012; final revision received September 2012.

Address correspondence to Sebastian Jaimungal, Department of Statistics, University of Toronto, 100 St. George St., Toronto, Ontario M5S3G3, Canada; e-mail: sebastian.jaimungal@utoronto.ca.

DOI: 10.1111/mafi.12023

order to post, cancel, or amend buy and sell orders; and where inventories are managed to ensure that no positions are carried overnight after markets close (see Cvitanic and Kirilenko 2010).

Setting up a proprietary HF trading desk or becoming an HF trader requires high initial sunk costs, but more importantly, it requires skills to develop strategies and maintain them so that they stay competitive and profitable (see Cartea and Penalva 2012). Thus, understandably, there is a great deal of secrecy around what it is that HF traders do and how they do it. Although for some electronic markets it is possible to obtain detailed information of all messages sent to the limit order book (LOB), these publicly available data do not include trader identity, hence it is very difficult to reverse engineer HF strategies, measure HF traders' trading profits, or assess the impact that HF trading has on the quality of the markets.

Furthermore, an important state variable which is impossible to track with publicly available data is the inventories that HF traders hold. The exception in the literature are two studies that have access to the audit-trail tape which includes the identity of each trader submitting orders to the LOB (Kirilenko et al. 2010 and Brunetti, Kirilenko and Mankad 2011). Brunetti et al. use audit-trail data for all transactions in the September 2009 E-mini S&P 500 futures contract during the month of August 2009. Over that period they find that only 11 HF traders are active in the market (out of 20,000 traders) and find that the net position (aggregate inventory) of these HF traders starts at zero at the beginning of the trading day, ends at zero at the end of the trading day, and it is always between $-3{,}000$ and 3,000 contracts.[1]

The picture that emerges from the current literature on HF trading is that profits are made by turning over positions over extremely short periods of time, where speed, which is paramount to process information and to make trading decisions, is at the heart of their competitive advantage. Another, much less explored trait of HF trading, is how are inventories managed. Apart from the fact that HF traders aim at holding no inventories at the end of the trading day, very little is known about how HF traders manage the accumulated positions while both maximizing expected profits and obeying risk controls on their inventory exposure.

In this paper we develop an HF trading strategy to post limit buy and sell orders which are continuously updated based on the dynamics of the midquote price of the asset, the shape of the LOB, and the trader's accumulated inventories. The HF strategy results from solving an optimal control problem where the HF trader maximizes expected terminal wealth while optimally managing inventories and by taking into account trades where she might be adversely selected. To this end, the optimal limit orders include a buffer to offset losses that arise from trading with counterparties who are better informed about the direction of subsequent price moves or losses that result from trades that have a permanent (and adverse) price impact. Moreover, inventories are continuously managed to ensure that at any point in time the strategy does not build large positions and to prevent triggering unfavorable price movements when positions are unwound at the end of the trading day.

Early work on optimal postings by a securities dealer is that of Ho and Stoll (1981). More recently Avellaneda and Stoikov (2008) study the optimal HF submission strategies of bid and ask limit orders where the authors perform a small inventory asymptotics in

[1]The positions held by the HF traders are very low compared to those of the other traders. For instance, buyers and sellers start the day with zero inventories and by the end of the day they have accumulated around 30,000 and $-20{,}000$ contracts, respectively.

a symmetric model and assume that market orders arrive at Poisson times and the fundamental price is independent of order arrivals. In general, very little is known about the details of the strategies that are employed by AT desks or the more specialized proprietary HF trading desks. Algorithms are designed for different purposes and to seek profits in different ways (Bouchard, Dang, and Lehalle 2011). For example, there are algorithms that are designed to find the best execution prices for investors who wish to minimize the price impact of large buy or sell orders, e.g., Bertsimas and Lo (1998), Almgren (2003), Kharroubi and Pham (2010), and Bayraktar and Ludkovski (2012). Moreover, there are HF strategies that specialize in arbitraging across different trading venues, and that seek to profit from market making and from short-term deviations in stock prices, e.g., Cartea and Jaimungal (2013), Guilbaud and Pham (2013), Cartea, Jaimungal and Ricci (2014), and Guéant, Lehalle and Fernandez Tapia (2011).

In this paper we assume that the HF trader maximizes expected terminal wealth and is inventory and capital constrained. The inventory constraint is a consequence of the time scale over which the HF strategies are designed to profit from buying and selling assets. HF traders use their superior speed to process information and act ahead of other slower traders. For example, some of these strategies include arbitraging the same stock in different markets, thus absolute speed is all that matters; or to trade quicker than other market participants to take advantage of short-lived price deviations from the fundamental value of the asset. Admittedly, there are a great deal of HF strategies and all we know is that their success depends on being able to profit from roundtrip trades (buy followed by a sell or vice-versa). Therefore, because HF traders' competitive edge is speed, their strategies seek opportunities to enter and exit the market very quickly (seconds or minutes) and, as a result, holding periods are extremely short. In other words, speed cannot be employed only to execute one leg of a roundtrip trade and wait until there is an opportunity to close the position. On the contrary, HF traders only take positions that present a high chance of being closed very quickly and at a profit. Consequently, HF traders' inventories should never be too "far" away from the optimal inventory level. The profitability of strategies that operate over "long-time" intervals depends on more attributes than just speed and, for that reason, HF traders lose their competitive advantage to more traditional professional traders.

In addition, capital requirements also constrain the amount of assets the HF trader is willing (or allowed) to hold at any point in time. For example, if the HF trader builds a large position in a stock, she will need more capital to be able to hold the position, whether required by the exchange or by her own risk controls. Similarly, if at the end of the trading day the HF trader has a large position to be carried overnight, she not only faces the risk that the market will move against her before trading commences the following business day, but the exchange will require collateral. Finally, traders know that unwinding large positions is costly because their own trades may trigger adverse price movements and this effect might be exacerbated at the end of the trading day.

Therefore, here we assume that the control problem solved by the HF trader maximizes expected terminal wealth where inventories enter in the optimization program in three different ways. First, the amount of assets that the HF trader can hold at any point in time is between an upper and lower bound. Second, deviations of inventories from zero along the entire path of the strategy, i.e., throughout the whole trading horizon, are penalized. And third, at maturity (at most one trading day for HF traders) the inventory is liquidated at worse prices than the best available quotes because HF traders have to execute aggressive market orders to ensure that no positions are carried over after the market closes. The use of market orders is costly because there is an exchange fee for

taking liquidity as well as a high chance of causing adverse price changes when there is not enough depth at the best quotes; the HF will receive less (pay more) than the best bid (best offer) when she unwinds her positions at maturity.

Furthermore, our paper contributes to the literature in two more directions. First, we propose risk metrics which are designed for HF trading strategies. Current off-the-shelf risk metrics are intended for the more traditional investment strategies and fail to provide the right decision-making tools and risk controls for HF trading. Risk metrics for HF trading should encapsulate information on risk exposures that derive from risk profiles over long (hours) and short (seconds or minutes) time scales. We propose risk metrics which are based on the path of inventories throughout the entire period of the trading strategy. Our risk metrics embed information about: the lifetime for every inventory level; how quickly inventories return to their optimal level; how much time is the strategy exposed to undesired levels of inventories; and what is the exposure to terminal inventories. Thus, our risk metrics allow HF traders to fine tune their strategies to obey risk controls and to trade off inventory risk against expected profits.

Second, in our model the innovations in the fundamental value of the asset are driven by information flows and the arrival of market orders. Under normal market circumstances, information is gradually impounded in the fundamental price of the asset by market participants who are continuously updating their quotes in the LOB. Market orders, on the other hand, cause prices to jump up or down by a random amount depending on the arrival of news, noise traders coming to the market with large orders, informed traders taking advantage of their superior or privileged information, and/or because market orders arrive at times when there is not enough depth at the best quotes in the LOB to fully absorb the orders.

We illustrate the behavior and performance of the strategies by showing how an HF trader, who maximizes expected terminal wealth and faces capital and inventory constraints, trades off expected profits against the variance and distribution of terminal Profit and Loss (PnL), terminal inventory, and lifetime of inventories. Our analysis shows that final inventory penalties (which proxies the cost of unwinding positions at expiry of the strategy) affect the distribution of inventories at the expiry of the strategy, but have a limited effect on the distribution of the PnL of the strategy. On the other hand, our results show that penalizing inventories throughout the entire strategy (the penalization affects inventory management but does not alter the PnL directly) has an important effect on the three risk metrics we propose and provides a useful tool to trade off expected profits against risks. For example, we show that there are ranges of the penalization parameter where a small increase in the parameter causes a considerable reduction in the risks borne by the HF trader (as given by our risk metrics), while expected profits are hardly affected.

Our results also show that during times when markets become one-sided, i.e., there are more sell than buy market orders (or vice versa), the HF trader adjusts her strategy to take advantage of the market momentum. For example, if it is expected that the price of the asset will drift upwards, the optimal strategy consists in quickly building a long position and then proceed to post limit buy and sell orders to make profits on the realized spread. The optimal postings are such that inventories revert to an optimal time-dependent position which is positive most of the time, but very close to (and at) expiry the optimal inventory level becomes negative.

The rest of this paper is organized as follows: Section 2 formulates and solves the HF's optimization problem. Section 3 illustrates the behavior of the optimal limit orders posted by the HF trader as a function of the model parameters and inventories. Section 4 proposes risk metrics for HF trading strategies. Section 5 shows how the HF trader

employs the risk metrics to depict the tradeoff between expected profits and the exposure to risks and it also discusses how the HF trader fine tunes her strategies to maximize expected profits while obeying risk controls. Section 6 discusses further properties of the optimal strategy. Finally, Section 7 concludes and in the Appendix we collect proofs of our results.

2. THE HIGH-FREQUENCY TRADER'S OPTIMIZATION PROBLEM

2.1. Formulation of the High-Frequency Market Making Problem

The HF strategy we develop here is based on posting limit buy and sell orders to profit from the realized spread which is the expected gain from a roundtrip trade. The HF trader is continuously updating her positions in the LOB to reflect changes in the midprice of the traded asset as well as other state variables, such as inventories and shape of the LOB, that we discuss below. The midprice of the asset is the average of the best bid and best ask and we assume that market participants that populate the LOB are constantly updating their quotes as a result of new information arriving in the market. We assume that the midprice (or fundamental price) S_t of the traded asset satisfies

$$dS_t = \sigma\, dW_t + \int_{\mathbb{R}_+^2} y_1[\mu^+(d\mathbf{y}, dt) - \mu^-(d\mathbf{y}, dt)], \tag{2.1}$$

where $(W_t)_{0\le t\le T}$ is a standard Brownian motion, σ is a positive constant, and $\mu^\pm(d\mathbf{y}, dt)$ are Poisson random measures[2] on $\mathbb{R}_+^2 \times [0, T]$, independent of one another and the Brownian motion $(W_t)_{0\le t\le T}$, with mean measures $\nu^\pm(d\mathbf{y}, dt) = \lambda^\pm\, F^\pm(dy_1)\, G^\pm(dy_2)\, dt$ where the distribution $F^\pm$ satisfies $\int_0^\infty F(dy_1) = 1$, $\int_0^\infty y_1^2\, F(dy_1) < +\infty$ (i.e., with finite second moments) and $G^\pm(dy_2) = \frac{1}{\kappa^\pm} e^{-\kappa^\pm y_2} dy_2$ with $\kappa^\pm > 0$. The distributions $F^\pm$ determine the size of the impact of trades on the midprice, while the distributions $G^\pm$ are related to the probability of limit orders at price y being filled. Below we discuss the role of $G^\pm$ further. The counting processes corresponding to the arrival of sell and buy market orders are $M_t^- = \int_0^t \int_{\mathbb{R}_+^2} \mu^-(d\mathbf{y}, ds)$ and $M_t^+ = \int_0^t \int_{\mathbb{R}_+^2} \mu^+(d\mathbf{y}, ds)$, respectively, and are therefore Poisson processes with intensities λ^- and λ^+. We often find it convenient to work with the random variables $\epsilon^\pm$ whose distribution functions are $F^\pm$. As usual, we work on a completed filtered probability space $(\Omega, \mathbb{F}, \mathcal{F} = \{\mathcal{F}_t\}_{0\le t\le T}, \mathbb{P})$ where $\mathcal{F}$ is the natural filtration generated by the midprice $(S_t)_{0\le t\le T}$ and the counting processes $(N_t^\pm)_{0\le t\le T}$ which count the number of filled sell/buy limit orders that the trader posts. We define $N^\pm$ more precisely below.[3]

Intuitively, we can view the dynamics of the midprice as the sum of two components. The first component, the Brownian motion on the right-hand side of (2.1), captures the changes in the midprice that are due to information flows that reach all or some market participants who subsequently update their quotes. The other component, the jump process with increments $\epsilon^+ dM_t^+ - \epsilon^- dM_t^-$, represents the changes in the midprice caused by the arrival of market orders that have a permanent price impact. Market orders may come at times when there is enough liquidity in the market—hence prices remain unchanged or change by a negligible amount; they may arrive at times when

[2]See Kühn and Stroh (2012) for a thorough development of Poisson random measures for limit order book modeling.

[3]For stock price models with HF data where the interarrival times of trades are not exponential, see Cartea and Meyer-Brandis (2010) and Cartea (2013).

liquidity is thin or the orders sent by traders with superior information and these trades have a permanent impact on prices. The impact of trading on the midprice may also be viewed as the action of informed traders. If an informed trader purchases (sells) shares, they will only do so if the asset price is known to be going up (down). The resulting increase (decrease) of the midprice following informed trading can be approximated by an immediate, and permanent, price impact as we model here. Note we also assume that the size of the jump in the midprice at the time of a market-order is stochastically independent of the depth at which the market order eats into the book.

In practice market orders also have a temporary price impact which depends on the type of information that triggered the trade and the resiliency of the LOB. For example, if a large market order (or a series of one-side market orders) arrives it will eat into one side of the LOB and it will take some time before other market participants update their quotes. If the market order is a pure liquidity trade which is not motivated by private information about the fundamental value of the asset, the price impact is temporary and, everything else equal, the midquote will revert to the value prior to the arrival of the market order. The time it takes for the midquote to mean-revert is determined by how long it takes market participants to process market information (i.e., find out that the trade is purely a liquidity trade) and repost resting orders to replenish the book. In our model, market orders have either a permanent price impact or no price impact, but it is possible to extend it so that market orders have a transitory price impact.

In our setting HF traders maximize expected terminal wealth and are constrained by the amount of inventories they are willing to hold. As mentioned above, HF traders' main competitive advantage is speed. HF strategies that profit from the realized spread are designed to enter and exit the market over very short-term horizons (seconds, or at most minutes) because the shorter the time scales the more likely are HF traders to take advantage of speed to process information and execute roundtrip trades at a profit. In our framework speed comes into play in two ways. One, we assume that the HF trader is able to process information and send instructions to the exchange in microseconds. Speed enables her to place optimal buy and sell limit orders in the LOB, i.e., how far away from the midquote S_t is it optimal to post a buy limit order $S_{t-} - \delta_t^-$ and a sell limit order $S_{t-} + \delta_t^+$ (with $\delta_t^{\pm} > 0$), before other slower traders can make decisions or reach the exchange. And two, because HF traders exploit their speed advantage to open and close positions, optimal strategies generate a high turnover of inventories. Everything else equal, HF traders prefer strategies with very short expected holding periods because they are constrained by the amount of inventories they can hold long or short.

Thus, in the optimization problem that the HF trader solves, inventories q_t are penalized and managed as follows: First, inventories receive a penalty throughout the entire path of the strategy, and, at maturity final inventory q_T is liquidated using market orders which are executed at worse prices than the best available quotes. The use of market orders is costly because there is an exchange fee for taking liquidity as well as an adverse impact on prices which is increasing, and generally convex, in the size of the inventories; a market order will eat into the LOB away from the best quote. Finally, to avoid accumulating large positions, the HF trader imposes a cap on the number of shares that she can hold long or short at any point in time.

To formalize the investor's problem we need to introduce some more notation. Let N_t^- and N_t^+ denote the counting processes for the market orders that filled the HF trader's buy and sell limit orders. Thus, the total inventory of the HF trader is given by $q_t = N_t^- - N_t^+$ and so $q_t \in \mathbb{Z}$ for each t. Upon a buy or sell order being filled, the HF trader pays $(S_{t-} - \delta_t^-)$ or receives $(S_{t-} + \delta_t^+)$. We assume that a limit order at price level

$S_{t-} \pm \delta_t^{\pm}$ is filled (given that a market order arrives) with probability $e^{-\kappa^{\pm}\delta_t^{\pm}}$. Hence, $\kappa^{\pm}$ can be interpreted as the exponential decay factor for the fill rate of orders placed away from the midprice, and the fill rate function can be written as $\Lambda_t^{\pm} = \lambda^{\pm} e^{-\kappa_{\pm}\delta_t^{\pm}}$ which is the rate of execution of a limit order. For queuing models of the LOB we refer the reader to Cont, Talreja, and Stoikov (2010). For simplicity we assume that the LOB is instantly replenished after a market order arrives and that the spread is zero.[4]

In terms of the Poisson random measures $\mu^{\pm}(d\mathbf{y}, dt)$ introduced earlier, the counting processes corresponding to the trader's filled limit orders are written as $N_t^{\pm} = \int_0^t \int_{\mathbb{R}_+ \times [\delta_s^{\pm}, +\infty)} \mu^{\pm}(d\mathbf{y}, ds)$. Note that $N_t^{\pm}$ counts those events which lie above $y_2 > \delta_t^{\pm}$ (representing market orders large enough to eat into that level of the limit order book) regardless of the impact size y_1 on the midprice. Recall that the filtration $\mathcal{F}$ is the natural one generated by both $(S_t)_{0\le t\le T}$ and $(N_t^{\pm})_{0\le t\le T}$ and the jumps of S_t coincide with those of $M_t^{\pm}$, therefore the collection of compensated counting processes $(\widehat{N}_t^{\pm}, \widehat{M}_t^{\pm})_{0\le t\le T}$ where $\widehat{N}_t^{\pm} = N_t^{\pm} - \int_0^t \lambda^{\pm} e^{-\kappa_{\pm}\delta_s^{\pm}}\, ds$ and $\widehat{M}_t^{\pm} = M_t^{\pm} - \lambda^{\pm}\, t$ are $\mathcal{F}$-martingales.

Finally, the HF trader's wealth X_t satisfies the stochastic differential equation

$$dX_t = \left(S_{t-} + \delta_t^{+}\right) dN_t^{+} - \left(S_{t-} - \delta_t^{-}\right) dN_t^{-}, \tag{2.2}$$

and the investor seeks the strategy $(\delta_s^{\pm})_{t\le s\le T}$ which maximizes expected terminal wealth while penalizing and constraining inventories. Therefore, the HF trader solves the control problem

$$H(t, x, q, S) = \sup_{(\delta_s^{\pm})_{t\le s\le T}\in\mathcal{A}} \mathbb{E}\Bigg[X_T + q_T(S_T - \alpha\, q_T) - \phi\,\sigma^2 \int_t^T q_s^2\, ds \,\Bigg| \, X_{t-} = x, q_{t-} = q, S_{t-} = S \Bigg], \tag{2.3}$$

where $q_t \in \mathbb{Z}$ is the inventory position at time t, the terminal date of the strategy is $T > t$, and $\mathcal{A}$ denotes the set of admissible strategies—strategies that are $\mathcal{F}_t$-predictable processes such that inventories are bounded above by $\overline{q} > 0$ and below by $\underline{q} < 0$ both finite. Here, $\phi \ge 0$ penalizes deviations of q_t from zero along the entire path of the strategy, while the term $-\alpha\, q_T^2$ only penalizes deviations away from zero at maturity itself.[5] Moreover, admissible strategies must obey upper and lower bounds on inventories: $\underline{q} \le q_s \le \overline{q}$, $t \le s \le T$, where $\underline{q}, \overline{q} \in \mathbb{Z}$ with $-\infty < \underline{q} < 0 < \overline{q} < +\infty$. In particular, if $q_t = \overline{q}$, then the investor places only limit sell orders (i.e., she sets $\delta^- = +\infty$) and if $q_t = \underline{q}$, then the investor places only limit buy orders (i.e., she sets $\delta^+ = +\infty$).

The penalty term $\phi\sigma^2\mathbb{E}[\int_t^T q_s^2\, ds|\mathcal{F}_t]$ can further be motivated by noting that the total value of the investor's inventory (assuming a mark-to-market value at midprice and ignoring jumps) is $V = \int_t^T q_s\, dS_t^c$, where S^c denotes the continuous component of the midprice S. Hence its variance is

$$\mathbb{E}\left[\left(\int_t^T q_s\, dS_s^c \right)^2 \middle| \mathcal{F}_t \right] = \mathbb{E}\left[\int_t^T q_s^2\, d[S^c, S^c]_s \middle| \mathcal{F}_t \right] = \sigma^2\, \mathbb{E}\left[\int_t^T q_s^2\, ds \middle| \mathcal{F}_t \right]$$

[4]Our model could be extended so that the shape and depth of the book undergoes temporary changes in response to the arrival or market orders that eat into the book and to reflect the update of quotes due to the arrival of news. Moreover, a constant nonzero bid-ask spread can be incorporated, however, the qualitative results do not change and essentially result in shifting the optimal drift by the half-spread.

[5]In the work of Guéant et al. (2011) final inventory is penalized with $-\alpha|q_T|$. In principle we can include any liquidation penalty as a function of q or any inventory integrated penalty that is an even function of q.

which, up to the scale factor ϕ, is the penalty we employ. In particular $\phi \geq 0$ reflects how severe the inventory controls are.

2.2. The Optimal Strategy

To solve the optimal control problem described above, a dynamic programming principle holds and the value function (2.3) is the unique viscosity solution of the Hamilton-Jacobi-Bellman (HJB) equation (see, e.g., Fleming and Soner 2006, chapter V):

$$
\begin{aligned}
0 = {} & \partial_t H + \frac{1}{2}\sigma^2 \partial_{SS} H - \phi\, \sigma^2 q^2 \\
& + \lambda^+ \sup_{\delta^+} \Big\{ e^{-\kappa^+ \delta^+} \mathbb{E}(H(t, x + (S + \delta^+), q - 1, S + \epsilon^+) - H) \\
& \qquad + (1 - e^{-\kappa^+ \delta^+}) \mathbb{E}(H(t, x, q, S + \epsilon^+) - H) \Big\} \\
& + \lambda^- \sup_{\delta^-} \Big\{ e^{-\kappa^- \delta^-} \mathbb{E}(H(t, x - (S - \delta^-), q + 1, S - \epsilon^-) - H) \\
& \qquad + (1 - e^{-\kappa^- \delta^-}) \mathbb{E}(H(t, x, q, S - \epsilon^-) - H) \Big\} ,
\end{aligned} \tag{2.4}
$$

subject to the terminal condition

$$
H(T, x, q, S) = x + q(S - \alpha q) , \tag{2.5}
$$

and where x, S, and q are the quantities at $t-$ (and not t) and the expectation is over the random variables $\epsilon^{\pm}$ and not over x, S, or q. Moreover, recall that the set of admissible strategies imposes bounds on q_t, this means that when $q_t = \overline{q}$ ($\underline{q}$) the optimal strategy is to post one-sided limit orders which are obtained by solving (2.4) with the term proportional to λ^- (λ^+) absent. Alternatively, one can view these boundary cases as imposing $\delta^- = +\infty$ ($\delta^+ = +\infty$) when $q = \overline{q}$ ($\underline{q}$).

Intuitively, the various terms in the HJB equation represent the arrival of market orders that may be filled by limit orders together with the diffusion of the asset price through the term $\frac{1}{2}\sigma^2 \partial_{SS} H$ and the effect of penalizing deviations of inventories from zero along the entire path of the strategy which is captured by the term $\phi\sigma^2 q^2$. The sup over δ^+ contain the terms due to the arrival of a market buy order (which is filled by a limit sell order). The first term represents the expected change in the value function H due to the arrival of a market order which fills the limit order and the midquote S_t jumps up by the random amount ϵ^+; and the second term represents the arrival of a market order which does not reach the limit order's price level (but still causes a random jump in the midquote). Similarly, the sup over δ^- contain the analogous terms for the market sell orders which are filled by limit buy orders.

To solve the HJB equation we use the terminal condition (2.5) to make an ansatz for H. In particular, write

$$
H(t, x, q, S) = x + q(S - \alpha\, q) + h_q(t) , \tag{2.6}
$$

which leads to a coupled system of ordinary differential equations (ODEs) for the functions $h_q(t)$ and allows us to solve for the optimal feedback controls. Before proceeding, note that the function $h_q(t)$ is related to the so-called *reservation price* of the asset as in Ho and Stoll (1981). In particular, the reservation bid price r_t^- satisfies $H(t, x - r_t^-, q + 1, s) = H(t, x, q, s)$ and the reservation ask price r_t^+ satisfies

$H(t, x + r_t^+, q - 1, s) = H(t, x, q, s)$. In other words, the reservation bid (ask) price is the price at which the HF trader is indifferent between buying (selling) and the inventory increasing (decreasing) by one unit or not trading at all.

Let, $\varepsilon^\pm = \mathbb{E}[\epsilon^\pm]$ be the mean jump size in asset price when a market-order arrives. Then, substituting (2.6) in the HJB equation results in $h_q(t)$ satisfying

(2.7)
$$\phi\sigma^2 q^2 = \begin{cases} \partial_t h_q + \lambda^+ \sup_{\delta^+}\{\varepsilon^+ q + e^{-\kappa^+\delta^+}(\delta^+ - \varepsilon^+ - \alpha(1-2q) + h_{q-1} - h_q)\} \\ \quad + \lambda^- \sup_{\delta^-}\{-\varepsilon^- q + e^{-\kappa^-\delta^-}(\delta^- - \varepsilon^- - \alpha(1+2q) + h_{q+1} - h_q)\}, & q \neq \underline{q},\ \overline{q}, \\ \partial_t h_q + \lambda^+ \sup_{\delta^+}\{\varepsilon^+ q + e^{-\kappa^+\delta^+}(\delta^+ - \varepsilon^+ - \alpha(1-2q) + h_{q-1} - h_q)\} - \varepsilon^-\lambda^- q, & q = \overline{q}, \\ \partial_t h_q + \lambda^- \sup_{\delta^-}\{-\varepsilon^- q + e^{-\kappa^-\delta^-}(\delta^- - \varepsilon^- - \alpha(1+2q) + h_{q+1} - h_q)\} + \varepsilon^+\lambda^+ q, & q = \underline{q}, \end{cases}$$

subject to $h_q(T) = 0$ which allows us to solve for the optimal feedback controls, in terms of $h_q(t)$, as shown in the proposition below.

PROPOSITION 2.1 (Optimal Feedback Controls). *The optimal feedback controls of the HJB equation are given by*

$$\delta^{+*}(t, q) = \frac{1}{\kappa^+} + \varepsilon^+ + \alpha(1 - 2q) - h_{q-1}(t) + h_q(t), \quad q \neq \underline{q}, \tag{2.8a}$$

and

$$\delta^{-*}(t, q) = \frac{1}{\kappa^-} + \varepsilon^- + \alpha(1 + 2q) - h_{q+1}(t) + h_q(t), \quad q \neq \overline{q}\,. \tag{2.8b}$$

Proof. Applying the first-order conditions directly to each sup term in the HJB equation (2.7) leads immediately to the result. □

To understand the intuition behind the feedback controls we first note that the optimal $\delta^\pm$ can be decomposed into four terms. The first component, $1/\kappa^\pm$, is the optimal strategy that a risk-neutral agent would employ in the absence of both terminal date restrictions (i.e., $T = \infty$, below in Proposition 2.6 we study the limiting case when $T \to \infty$) and inventory constraints. To see this, note that the expected gains from buying (selling) the asset at the midprice S_t, followed by selling (buying) it using a limit order at $S_{t-} + \delta^+$ ($S_{t-} - \delta^-$), is given by $\delta^\pm e^{-\kappa^\pm\delta^\pm}$, which is maximized if the limit order is posted $1/\kappa^\pm$ away from the midprice.

The second term shows that the HF trader incorporates the expectation of the jump in prices, conditioned on a market order arriving, by posting limit orders which are $\varepsilon^\pm = \mathbb{E}[\epsilon^\pm]$ further away from the midprice. In this way the HF trader recovers, on average, the losses she incurs when her limit orders are filled by trades that have a permanent price impact—for example, losses arising from trading with more informed traders (adverse selection).

Furthermore, the term $\alpha(1 \pm 2q)$ ensures that positions neither become too large nor too small, relative to the optimal inventory level, because at time T there are liquidation costs. For example, the time T optimal spreads are given by $\delta^\pm = 1/\kappa^\pm + \varepsilon^\pm + \alpha(1 \mp 2q)$ and one can see that when $q > 0$ ($q < 0$), the limit sell (buy) order is posted nearer

to the midprice to increase the likelihood of selling (buying) the stock, which induces inventories to revert to zero.

Finally, the term $-h_{q-1}(t) + h_q(t)$, controls for inventories through time. As expected, if inventories are long, then the strategy consists in posting limit orders that increase the probability of sell orders being hit. Moreover, the function $h_q(t)$ also induces mean reversion to an optimal inventory level as a result of penalizing accumulated inventories throughout the entire trading horizon and the strategy approaching T as well as the other parameters of the model, including $\phi\sigma^2$.

PROPOSITION 2.2 (Solving Reduced Value Function). *If $\kappa^+ = \kappa^- = \kappa$, then write $h_q(t) = \frac{1}{\kappa}\ln\omega_q(t)$ and stack $\omega_q(t)$ into a vector $\boldsymbol{\omega}(t) = [\omega_{\overline{q}}(t), \omega_{\overline{q}-1}(t), \ldots, \omega_{\underline{q}}(t)]'$. Further, let $\mathbf{A}$ denote the $(\overline{q} - \underline{q} + 1)$-square matrix whose rows are labeled from $\overline{q}$ to $\underline{q}$ and whose entries are given by*

$$\mathbf{A}_{i,q} = \begin{cases} q\,\kappa\,(\varepsilon^+\,\lambda^+ - \varepsilon^-\,\lambda^-) - \phi\sigma^2\kappa\,q^2, & i = q, \\ \lambda^+\,e^{-\kappa\left(\frac{1}{\kappa}+\varepsilon^+ + \alpha(1-2q)\right)}, & i = q-1, \\ \lambda^-\,e^{-\kappa\left(\frac{1}{\kappa}+\varepsilon^- + \alpha(1+2q)\right)}, & i = q+1, \\ 0, & \textit{otherwise.} \end{cases} \tag{2.9}$$

Then,

$$\boldsymbol{\omega}(t) = e^{\mathbf{A}(T-t)}\mathbf{1}\,, \tag{2.10}$$

where $\mathbf{1} = [1, \ldots, 1]'$ is a $(\overline{q} - \underline{q} + 1)$-dim vector of ones.

Proof. See Appendix A1. □

As a direct consequence of Proposition 2.2, when the market is symmetric the spread on the buy side with q shares equals the spread on the sell side with $-q$ shares.

COROLLARY 2.3 (Symmetry in Spreads). *When $\kappa^\pm = \kappa$, $\lambda^\pm = \lambda$, $\varepsilon^\pm = \varepsilon$, the optimal spreads are symmetric in the sense that $\delta^{\pm *}(t, q) = \delta^{\mp *}(t, -q)$.*

Proof. See Appendix A2.

Since we now have a classical solution of the HJB equation (2.7) which is also bounded above and below, it is easy to see that the feedback controls in (2.8) are admissible and the verification theorem below follows.

THEOREM 2.4 (Verification). *Let $h_q(t)$ be a classical solution to (2.7). Then $H(t, w, S, q) = x + q(S - \alpha q) + h_q(t)$ is the value function of the trader's control problem (2.3) and the optimal controls are given by (2.8) in feedback form.*

Proof. See Appendix B.

In the following proposition we show that the optimal HF strategy $\delta^{*\pm}(t, q)$ induces mean reversion in inventories—below in the simulations section we show that the speed of mean reversion is increasing in the penalty parameter ϕ.

COROLLARY 2.5 (Mean Reversion in Inventories). *Given the pair of optimal strategies* $\delta^+(t,q)$, $\delta^-(t,q)$, *the expected drift in inventories* q_t *is given by*

$$\mu_t(q) \triangleq \lim_{s\downarrow t} \frac{1}{s-t}\mathbb{E}[q_s - q_t | q_t = q] = \lambda^- e^{-\kappa^-\delta^-(t,q)} - \lambda^+ e^{-\kappa^+\delta^+(t,q)}. \tag{2.11}$$

Proof. It is straightforward to see that the expected drift in inventories is given by the difference in the arrival rates of filled orders. □

There are two remaining interesting limiting cases to study. The first is in the limit of being "far" from the investor's maturity and the second is in the limit of being "near" to the investor's maturity. Since the matrix **A** has elements which are linear in the activity rates, the time scale on which the "far" and "near" limits are valid is of the order $\min(1/\lambda^+, 1/\lambda^-)$. The key mathematical results are recorded below.

PROPOSITION 2.6 (Limiting Spreads). *Assume that* $\kappa^\pm = \kappa$ *and let* $\delta_0^\pm(q) = \frac{1}{\kappa} + \varepsilon^\pm + \alpha(1 \mp 2q)$. *Then,*

1. *As* $T \to \infty$, *the limiting spreads are given by the (time independent) values*

$$\delta^+(t,q) = \delta_0^+(q) - c^+(q), \quad q \neq \underline{q}, \quad \text{and} \quad \delta^-(t,q) = \delta_0^-(q) - c^-(q), \quad q \neq \overline{q}\,, \tag{2.12}$$

   where $c^\pm(q) = \mathbf{v}_{q\mp1} - \mathbf{v}_q$ *and the vector* $\mathbf{v} = \frac{1}{\kappa}\ln(\mathbf{U}[\mathbf{U}^{-1}\mathbf{1}]_\ell)$ *where the logarithm is taken componentwise,* ℓ *denotes the index of the largest eigenvalue of the matrix* **A**, *the notation* $[\cdot]_\ell$ *denotes the vector obtained by setting all entries to zero except for the* ℓ^{th} *entry, and* **U** *denotes the matrix whose columns are the eigenvectors of* **A**.
2. *As* $t \uparrow T$, *the limiting spreads for* $q \neq \overline{q}, \underline{q}$ *are given by*

$$\delta^+(t,q) = \delta_0^+(q) + \Bigg[\frac{2}{\kappa}\left(\lambda^+ e^{-\kappa(\frac{1}{\kappa}+\varepsilon^+-2\alpha q)} - \lambda^- e^{-\kappa(\frac{1}{\kappa}+\varepsilon^- +2\alpha(q-1))}\right)\sinh(\kappa\alpha) + (\lambda^+\varepsilon^+ - \lambda^-\varepsilon^-) - \phi\sigma^2\kappa(2q-1)\Bigg](T-t) + o(T-t), \tag{2.13a}$$

$$\delta^-(t,q) = \delta_0^-(q) + \Bigg[\frac{2}{\kappa}\left(-\lambda^+ e^{-\kappa(\frac{1}{\kappa}+\varepsilon^+-2\alpha(q+1))} + \lambda^- e^{-\kappa(\frac{1}{\kappa}+\varepsilon^- +2\alpha q)}\right)\sinh(\kappa\alpha) - (\lambda^+\varepsilon^+ - \lambda^-\varepsilon^-) + \phi\sigma^2\kappa(2q+1)\Bigg](T-t) + o(T-t). \tag{2.13b}$$

Proof. See Appendix A3. □

As direct consequence of the previous result we have the following limiting behavior of the strategies:

COROLLARY 2.7 (Terminal Strategy Approach). *When* $\kappa^\pm = \kappa$, $\lambda^\pm = \lambda$, $\varepsilon^\pm = \varepsilon$ *and* $\phi = 0$, *we have for* $q \neq \overline{q}, \underline{q}$

$$(2.14)\quad 0 > \begin{cases} \lim_{t\uparrow T} \partial_t \delta^+(t,q), & \forall q > 0, \\ -\lim_{t\uparrow T} \partial_t \delta^+(t,q), & \forall q \le 0, \end{cases} \quad \textit{and} \quad 0 < \begin{cases} \lim_{t\uparrow T} \partial_t \delta^-(t,q), & \forall q \le 0, \\ -\lim_{t\uparrow T} \partial_t \delta^-(t,q), & \forall q < 0. \end{cases}$$

This result tells us whether the strategy approaches its terminal strategy of $\delta_0^{\pm}(q)$ from above or below. For example, if $q \le 0$ then $\delta^+(t,q)$ approaches $\delta_0^+(q)$ from below, while if $q > 0$, then $\delta^+(t,q)$ approaches $\delta_0^+(q)$ from above.

3. THE BEHAVIOR OF THE STRATEGY

In this section we illustrate several aspects of the behavior of the optimal strategy as a function of $q, \phi, \alpha, \lambda^{\pm}, \overline{q}, -\underline{q}$. In all the examples of this section the terminal date of the strategy is $T = 10$ seconds, $\kappa^{\pm} = 25$, $\sigma = 0.1$, and upon the arrival of market orders, the midprice jumps up or down by the amount $1/25$, with probability $1/10$, or 0, with probability $9/10$, so that $\mathbb{E}[\epsilon^{\pm}] = \varepsilon^{\pm} = 1/250$.

3.1. Liquidation Penalty and the Optimal Spreads as a Function of Time for Various Inventory Levels

Figure 3.1 shows the behavior of the optimal spreads as a function of the liquidation penalty $\alpha = \{0, 0.01\}$ for different inventory levels. In the examples the arrival rate of market orders is $\lambda^{\pm} = 2$ (there are on average 2 buy and 2 sell market orders per second), $\overline{q} = -\underline{q} = 3$, and $\phi = 0$. In this subsection we discuss the results for the sell spread, the buy spread has a similar interpretation. In panel (a) the HF strategy is not penalized for unwinding inventories at expiry, $\alpha = 0$, and we show the optimal sell postings δ^+, i.e., upon the arrival of a market buy order the HF trader is willing to sell one unit of the asset at the price $S_t + \delta^+$. Thus, when $\alpha = 0$ and inventories are close to the allowed minimum, the optimal posting is furthest away from the midquote because only at a very "high" price is the HF trader willing to decrease her inventories further. However, as the strategy approaches T and $q_t < 0$, the optimal spread δ^+ decreases. To understand the intuition behind the optimal strategy note that if the terminal inventory $q_T < 0$ is

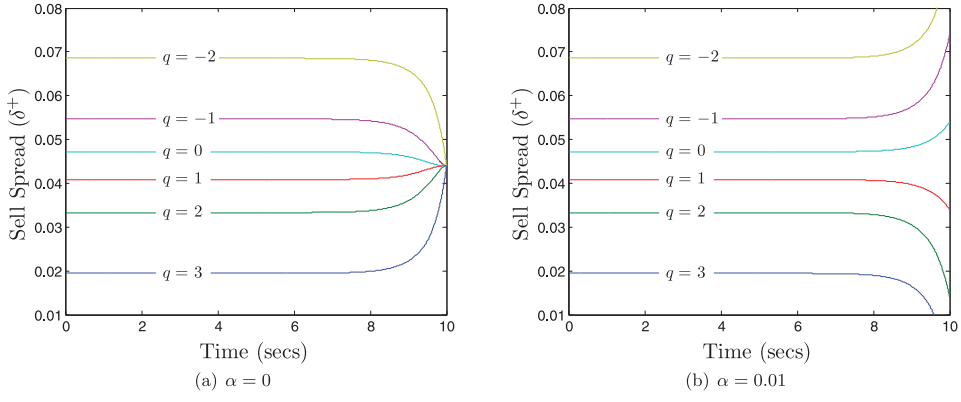

FIGURE 3.1. The optimal spreads as a function of time for various inventory levels and $T = 10$. The remaining model parameters are: $\lambda^{\pm} = 2$, $\kappa^{\pm} = 25$, $\varepsilon^{\pm} = \frac{1}{250}$, $\overline{q} = -\underline{q} = 3$, $\phi = 0$, $\sigma = 0.1$. Only the sell-side is shown as the buy-side is similar.

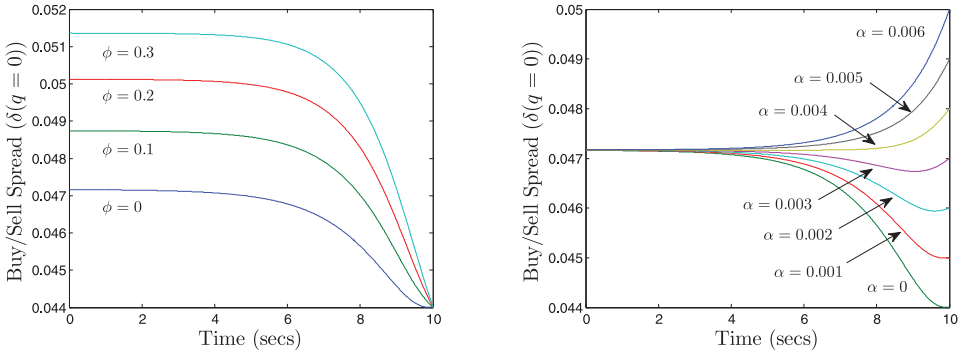

FIGURE 3.2. The optimal spread at $q = 0$ as a function of time for various levels of ϕ and α. The remaining model parameters are: $T = 10$, $\lambda^{\pm} = 2$, $\kappa^{\pm} = 25$, $\varepsilon^{\pm} = \frac{1}{250}$, $-\underline{q} = \overline{q} = 3$, and $\sigma = 0.1$. In the left panel $\alpha = 0$ and in the right panel $\phi = 0$.

liquidated at the price $S_T - \alpha q_T$, then when α is sufficiently low and $q_t < 0$, as well as being fractions of a second away from expiry, it is optimal to post nearer the midprice to increase the chances of being filled (i.e., selling one more unit of the asset) because the price is not expected to move too much before expiry and the entire position will be unwound at the midprice—making a profit on the last unit of the asset that was sold.

When $q_t > 0$ and the strategy is far away from T it is optimal to post low spreads because there is enough time before expiry to turn over positions. On the other hand, when $q_t \geq 0$ and the strategy approaches T, it is optimal to post further away from the midquote. Initially this result appears to be counterintuitive because one would expect that it would be best to post nearer the midprice, rather than further, to increase the probability of selling the asset at a price slightly above S_t to reduce the long position instead of closing the position at S_T. Note however, that if the strategy is one second away from expiry, the HF trader knows that on average there will be 2 buy/sell market incoming orders and that there is a tradeoff between posting nearer or further away from the midprice. The tradeoff consists in comparing what is the change in expected wealth from posting a limit sell order far from the midprice compared to the change in the expected wealth from posting close to the midprice. The former has a lower probability of being filled than the latter, but overall the HF trader is better off, in expectation, if she posts the sell limit order deeper in the LOB as shown by panel (a) in the figure.

When the penalty for liquidating terminal inventory is increased, as in panel (b), we observe a change in the qualitative behavior of the strategy as T is approached. For instance, panel (b) shows that when $q_t < 0$ and the strategy approaches T, the spreads increase—if the HF trader sells another asset close to expiry, it is at a price that ensures that it would cover the liquidation penalty of αq_T per share. Similarly, when $q_t > 0$ and the strategy approaches T, it is optimal to post very close to the midprice to increase the probability of unwinding the long position with a limit order before T rather than picking up the liquidation penalty.

3.2. The Optimal Spread at Zero Inventory as a Function of Time for Various Levels of ϕ and α

Figure 3.2 shows how the optimal buy/sell spread changes as a function of the liquidation penalty parameter α and the inventory penalty ϕ when $q_t = 0$. We recall that

due to the symmetry of both the arrival rate of market orders and the size of jumps in the midprice, the buy and sell spreads are the same when $q_t = 0$ (as shown in Corollary 2.3). In our discussion we focus on the sell spread, the interpretation of the buy spread is similar. In the left panel, $\alpha = 0$ and we see that when $t < T$ the higher is the inventory penalty ϕ, the further away from the midprice is the optimal sell spread, but as the strategy approaches T, all spreads, as a function of ϕ, converge to the same terminal value of $\frac{1}{\kappa^{\pm}} + \varepsilon^{\pm}$. The monotonic decreasing behavior is natural because, here, there is no terminal penalty and the running penalty (driven by ϕ) does not have time to constrain the strategy.

The picture on the right-hand shows how the liquidation penalty affects the optimal sell strategy. We observe that for low values of the liquidation parameter α the optimal spread starts to decrease a few seconds before T and then increases. The intuition behind this behavior is the following. When the strategy approaches expiry, and $q_t = 0$, it is optimal to lower the sell spread in the hope of acquiring a unit of the asset and have enough time left for other market orders to arrive on the other side of the LOB to be able to unwind the position. However, if the strategy is approaching T then there is not enough time left to execute a roundtrip trade and this explains why the spread turns and increases very close to T. This can be seen as a direct consequence of (i) Corollary 2.7 which shows that the strategy (for $q = 0$) must approach its terminal limit δ_0^+ from below, and (ii) the fact that the terminal limit lies below the far from maturity limit $\delta^+(t, q)$. On the other hand, for high values of α we see that as the strategy approaches expiry, the spread increases because if the offer is lifted by an incoming buy market order, the HF trader has very little time left to unwind her position before expiry and the spread has to be high enough to cover the high-liquidation penalty.

3.3. The Optimal Spreads as a Function of Time for Various Inventory Levels and Maximal Positions

Figure 3.3 shows how the optimal sell spread depends on the maximum amount of shares q (long or short) that the HF trader is allowed to hold. We assume that $T = 10$ seconds, $\lambda^{\pm} = 2$, $\kappa^{\pm} = 25$, $\varepsilon^{\pm} = \frac{1}{250}$, $\phi = 0$, $\sigma = 0.1$, and $\alpha = 0.001$. As expected, as

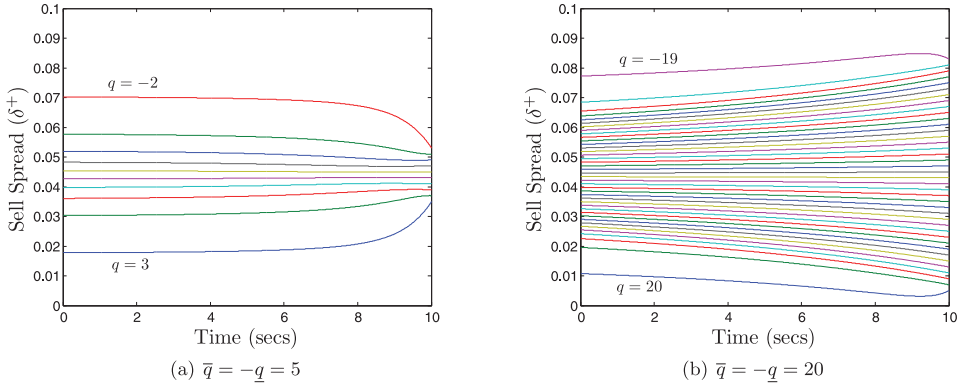

(a) $\overline{q} = -\underline{q} = 5$ (b) $\overline{q} = -\underline{q} = 20$

FIGURE 3.3. The optimal spreads as a function of time for two inventory levels and maximal positions with $T = 10$. The remaining model parameters are: $\lambda^{\pm} = 2$, $\kappa^{\pm} = 25$, $\varepsilon^{\pm} = \frac{1}{250}$, $\phi = 0$, $\sigma = 0.1$, and $\alpha = 0.001$.

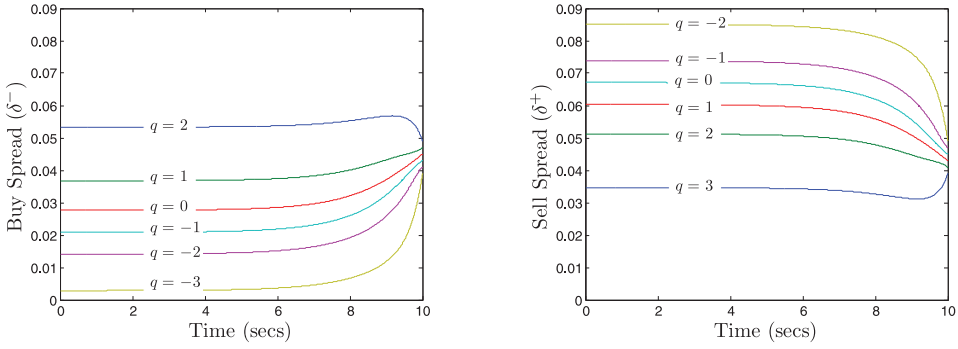

FIGURE 3.4. The optimal spreads as a function of time for various inventory levels when buy market orders are more likely $\lambda^+ = 6$, $\lambda^- = 4$. The remaining model parameters are: $T = 10$, $\kappa^\pm = 25$, $\varepsilon^\pm = \frac{1}{250}$, $\overline{q} = -\underline{q} = 3$, $\phi = 0$, $\sigma = 0.1$, and $\alpha = 0.001$.

inventories get closer to $\underline{q} + 1$, which are given by the lines with highest spread in each panel, the sell spread increases because only if the expected realized spread is higher will the HF trader be willing to increase her short position in the asset. The figure also shows that when the strategy is approaching T the optimal spread decreases or increases when inventory constraints are $\underline{q} + 1$ or $\overline{q}$, respectively. Finally, it is interesting to see that as $\overline{q}$ ($\underline{q}$) increases (decreases), the maximum and minimum spread, which correspond to inventories reaching the largest position (long or short) that can be held, do not vary much when inventory bounds increase from 5 to 20. In fact, we observe that the effect of a larger q is to reduce the difference between optimal spreads that only differ by one unit of inventory.

3.4. The Optimal Spreads as a Function of Time for Various Inventory Levels When Buy Market Orders Are More Likely

Figure 3.4 shows the behavior of the optimal buy and sell spread when the arrival of market orders is not symmetric. We assume that the intensity of arrival of buy market orders is higher than that of sell orders, $\lambda^+ = 6$, $\lambda^- = 4$, and that $T = 10$, $\kappa^\pm = 25$, $\varepsilon^\pm = \frac{1}{250}$, $\overline{q} = -\underline{q} = 3$, $\phi = 0$, $\sigma = 0.1$, and $\alpha = 0.001$. The figure shows that for all levels of inventory, the postings around the midprice will show a larger spread on the sell than on the buy side because the HF trading strategy incorporates the fact that the value of the asset is drifting upwards as a consequence of the arrival of more buy than sell market orders. Thus, when the strategy is not close to expiry, posting lower spreads on the buy side means that the HF trader expects to build a long inventory position because inventories will on average appreciate in value due to the upward pressure in prices. Finally, as the strategy approaches expiry, spreads on both sides of the book become more symmetric to ensure that inventories are unwound.

3.5. Mean Reversion of Inventories: Speed and Level

In Figure 3.5 we assume that the strategy is far away from expiry and we illustrate how the inventory drift (see Proposition 2.5) varies as a function of inventory and how

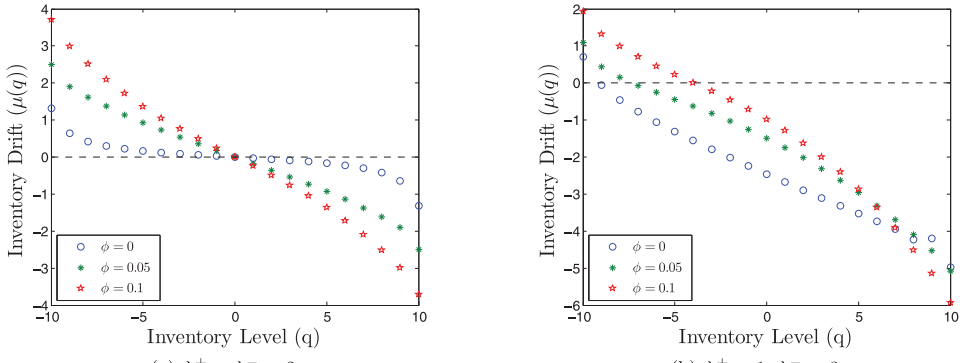

(a) $\lambda^+ = \lambda^- = 2$ (b) $\lambda^+ = 1$, $\lambda^- = 3$

FIGURE 3.5. Inventory drift $\mu(q)$ for each level of inventory with equal and unequal arrival rate of market orders, $\overline{q} = -\underline{q} = 10$, $\alpha = 0$, $\varepsilon^{\pm} = \frac{1}{250}$, and $\kappa^{\pm} = 25$, and the strategy is far away from expiry.

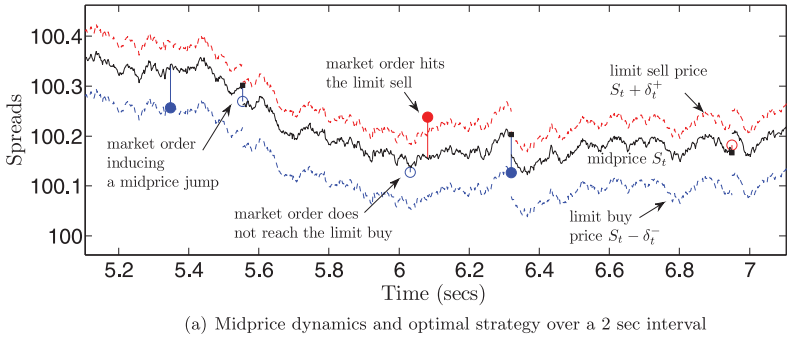

(a) Midprice dynamics and optimal strategy over a 2 sec interval

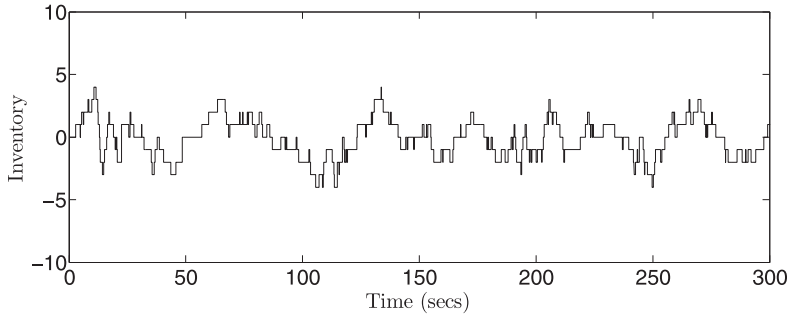

(b) Inventory path over 300 sec interval

FIGURE 3.6. Strategy and inventory sample paths over a 2- and 300-second interval, respectively. The model parameters are $\kappa^{\pm} = 25$, $\lambda^{\pm} = 2$, $\epsilon^{\pm} = \frac{1}{25}\mathfrak{a}^{\pm}$ where $\mathfrak{a}^{\pm}$ are independent Bernoulli r.v. with success prob of $\frac{1}{10}$, $\sigma = 0.1$, $\overline{q} = -\underline{q} = 10$, $S_0 = 100$, $\phi = 0.1$, and $\alpha = 0.001$. Filled circles indicate market orders that hit/lift the limit order, unfilled circles indicate orders that did not reach the post limit order.

this rate depends on the mean arrival rate of market orders and the parameter ϕ that penalizes deviations of inventories from zero. In the left panel, the arrival rates of market buy and sell orders are equal and we observe that the drift changes sign at zero inventory level (for positive inventories the drift is negative, while for negative inventories the drift is positive), thus the HF trader's inventory mean-reverts to zero. The figure also shows that the higher the penalty ϕ, the quicker will inventories mean revert to zero.

In the right panel of Figure 3.5 the arrival rate of market sell orders is much larger than that of buy orders. We observe that the drift changes sign at negative inventory levels, and consequently, the HF trader's inventory mean reverts to holding a short position in the asset in anticipation of the downward pressure on prices due to the excess arrival of sell market orders. Recall that the midquote price process is given by (2.1) so the HF trader knows that the midprice is drifting downward, $\lim_{\Delta t\downarrow 0} \frac{1}{\Delta t}\mathbb{E}[S_{t+\Delta t} - S_t|\mathcal{F}_t] = -\frac{2}{250}$, in which case it is optimal to always hold short positions, when far away from expiry, because the value of the asset holdings is expected to decrease. Moreover, if the penalty ϕ is increased, inventories mean revert to values which are closer to zero.

In panel (a) of Figure 3.6 we show a 2-second snippet of one simulation of the midprice where market buy and sell orders arrive with intensity $\lambda^{\pm} = 2$, and we also show how these market orders cross with the HF trader's limit orders. In panel (b) of the same figure we show the dynamics of the HF trader's inventory position over a 300-second window—the picture clearly shows that inventories are quickly mean reverting to zero and that they are always far away from the allowed upper and lower bound $\overline{q} = -\underline{q} = 10$.

4. RISK METRICS FOR HIGH-FREQUENCY TRADING STRATEGIES

HF trading is relatively new and relies on strategies which, as any other type of investment in the stock market, seek to profit from buying and selling assets. However, one of the differences between HF trading and any other type of activities that we observe in the exchanges, is the speed at which HF traders are able to process information and make trading decisions. Moreover, because speed is at the core of their competitive advantage, it is logical that HF traders invest in skills to develop trading strategies based on completing roundtrip trades as quickly as possible which is one of the reasons that explain why inventories revert to an optimal level of inventory so rapidly.[6]

Therefore, everything else equal, HF traders prefer strategies which are less exposed to inventory risk. Holding larger positions than the optimal one is risky for HF traders, and holding smaller positions than the optimal one for a long period of time is also risky. Ideally, HF traders prefer to hold small positions (relative to what other types of market participants are holding) for a short period of time because the longer the holding period, the riskier it is for them because they lose their competitive advantage to the more traditional market participants—being able to take decisions in milliseconds becomes less effective and less valuable.

Most risk metrics have been designed with the more traditional market participants in mind. For instance, the variance of the PnL is a widely used risk metric to assess an investment strategy that maximizes expected profits. However, this risk metric fails to address, in a direct rather than indirect way, the tradeoff between inventory risk and expected profits which is fundamental for HF trading. In the rest of this section we

[6]This is another striking difference between the behavior of HF traders and other investors, such as traditional market makers, who are willing (and required by the exchange) to hold inventories for much longer periods of time.

propose risk metrics for HF trading strategies and illustrate them using the model we developed above. We point out that the risk metrics proposed here have been designed to be applied to all HF trading strategies, not only the ones discussed in this paper. The risk metrics provide a tool to understand and quantify exposure to inventory risk (which also proxies for capital constraints), as well as to enable HF traders to fine tune their strategies by trading off inventory risk against expected profits.

4.1. Risk Metrics for High-Frequency Trading

In this section we discuss three risk metrics designed to assess the risk of HF strategies which profit from making the realized spread in the intraday market. Thus, the risk metrics are for HF market makers that meet all or some of the following: cannot or do not wish to hold inventories that deviate from the optimal level for a long period of time; employ strategies which turn over a vast amount of positions in the intraday market but do not want to hold large inventories; are capital constrained; operate in markets where liquidity varies throughout the trading day.

The first risk metric quantifies the variance of the PnL of the strategy, and the other two metrics quantify the exposure that HF trading strategies have to inventory risk. Below q_t^* and X_t^* denote the path of inventory and cash holdings respectively, followed by a general HF strategy—i.e., not limited to the one developed above. We let $\mathbb{E}_t[\cdot]$ and $\mathbb{V}_t[\cdot]$ denote the expectation and variance operators conditional on $X_{t-} = x$, $q_{t-} = q$ and $S_{t-} = S$.

1. Variance of terminal PnL:

$$\mathcal{R}^1(t, x, q, S) \triangleq \mathbb{V}_t\big[X_T^* + q_T^*(S_T - \alpha q_T^*)\big]. \tag{4.1}$$

2. Variance of terminal inventory:

$$\mathcal{R}^2(t, q) \triangleq \mathbb{V}_t\big[q_T^*\big]. \tag{4.2}$$

3. Variance of lifetime inventory:

$$\mathcal{R}^3(t, q) \triangleq \sum_{n \in \{\underline{q}, \ldots, \overline{q}\}} (n - \langle q^* \rangle_{t,q})^2 \, \mathbb{Q}_n(t, q), \tag{4.3}$$

where $\langle q^* \rangle_{t,q} = \sum_{n \in \{\underline{q}, \ldots, \overline{q}\}} n \, \mathbb{Q}_n(t, q)$ and

$$\mathbb{Q}_n(t, q) \triangleq \mathbb{E}_t \left[\frac{1}{T - t} \int_t^T \mathbb{I}\{q_s^* = n\} \, ds \right], \tag{4.4}$$

denotes the percentage of time (from t to T) that the strategy holds an inventory of n given that the strategy begins with q-shares. Moreover, note that $\mathbb{Q}_n(t, q)$ defines a distribution over the discrete set $\{\underline{q}, \ldots, \overline{q}\}$ of possible inventory levels and that risk metric $\mathcal{R}^3(t, q)$ assesses the strategy over the entire time horizon and not just at maturity.

For the particular optimal HF strategy derived in this paper we can compute risk metrics $\mathcal{R}_1, \mathcal{R}_2$, and $\mathcal{R}_3$ in closed-form in terms of time ordered exponentials. It is possible to further obtain an approximation of the three risk metrics by setting the strategy to the limiting strategy of Proposition 2.6 where the spreads are not time dependent. As

we saw in the previous section, for the particular parameters we have chosen to study, the effect of the terminal date T starts to show only a couple of seconds before expiry. In general the scale on which the strategy approaches its asymptotic value is very short, consequently, using the limiting strategy provides us with an excellent approximation to the strategy of Proposition 2.1. In the proposition below, we use the limiting strategy to calculate in closed-form risk metrics $\mathcal{R}_1$, $\mathcal{R}_2$, and $\mathcal{R}_3$. In A4 we show the fully time dependent risk metrics.

PROPOSITION 4.1 (Risk Metrics). *Suppose the agent follows the strategy specified by the limiting optimal spreads*[7] $\delta_q^{\pm *}$ *given by Proposition 2.6. Let* $\mathbf{B}$ *denote the matrix whose elements (labeled from* $\overline{q}$ to $\underline{q}$*) are*

$$\mathbf{B}_{i,q} = \begin{cases} -\left(\lambda^- e^{-\kappa\,\delta_q^{-*}} + \lambda^+ e^{-\kappa\,\delta_q^{+*}}\right), & i = q, \\ \lambda^+ e^{-\kappa\,\delta_q^{+*}}, & i = q-1, \\ \lambda^- e^{-\kappa\,\delta_q^{-*}}, & i = q+1, \\ 0, & \textit{otherwise,} \end{cases} \tag{4.5}$$

and let $\mathbf{V}$ *denote the matrix of eigenvectors stacked columnwise in decreasing order of its eigenvalues. Then we have the following:*

(i) $\mathbf{B}$ *contains exactly one zero eigenvalue, and all other eigenvalues are negative.*

(ii) *The terminal variance risk metric is given by*

$$\begin{aligned} \mathcal{R}^1(t,x,q) = {} & \mathbf{a}(t,q) - \mathbf{m}_1(t,q)^2 + 2\alpha q^2 \mathbf{m}_1(t,q) + (\mathbf{b}(q) - 2\mathbf{m}_1(t,q))x \\ & + (\mathbf{c}(t,q) - 2q\mathbf{m}_1(t,q))\,S \end{aligned} \tag{4.6}$$

where $\mathbf{m}_1(t,q)$ *is provided in equation (A.5),* $\mathbf{a} = \mathbf{V}\widetilde{\mathbf{a}}$, $\mathbf{b} = \mathbf{V}\widetilde{\mathbf{b}}$, *and* $\mathbf{c} = \mathbf{V}\widetilde{\mathbf{c}}$ *with* $\widetilde{\mathbf{a}}$, $\widetilde{\mathbf{b}}$, *and* $\widetilde{\mathbf{c}}$ *provided in equations (A.10a), (A.10b), and (A.10c). Moreover, we have (with* $\tau = T - t$*)*

$$\begin{aligned} \widetilde{\mathbf{a}}_j(t,q) &= r_1\,\tau^3\,(1 + o(\tau)), \quad \widetilde{\mathbf{b}}_j(t,q) = \widetilde{\boldsymbol{\chi}}_1\,\tau\,(1 + o(\tau)), \\ \widetilde{\mathbf{c}}_j(t,q) &= \frac{1}{2}\widetilde{\boldsymbol{\Psi}}_{11}\widetilde{\boldsymbol{\chi}}_1\,\tau^2\,(1 + o(\tau)), \end{aligned} \tag{4.7}$$

for $j = 1$ *and equal zero otherwise. Here, the constants* $\widetilde{\boldsymbol{\chi}}_1$ *and* $\widetilde{\boldsymbol{\Psi}}_{11}$ *are given in (A.11a) and* r_1 *is given in (A.11h).*

(iii) *The limiting terminal inventory risk metric is given by*

$$\lim_{T\to+\infty} \mathcal{R}^2(t,q) = (\mathbf{V}[\mathbf{V}^{-1}\,\boldsymbol{\beta}]_1)_q - \left((\mathbf{V}[\mathbf{V}^{-1}\,\boldsymbol{\gamma}]_1)_q\right)^2 \tag{4.8}$$

where $\boldsymbol{\beta}$ *and* $\boldsymbol{\gamma}$ *are* $(\overline{q} - \underline{q} + 1)$*-dim vectors with elements* $\boldsymbol{\beta}_q = q^2$ *and* $\boldsymbol{\gamma}_q = q$*, respectively.*

(iv) *For each inventory level, the limiting lifetime is given by*

$$\lim_{T\to+\infty} \mathbb{Q}_n(t,q) = (\mathbf{V}\,[\mathbf{V}^{-1}\,\mathbf{1}_n]_1)_q\,, \tag{4.9}$$

where $\mathbf{1}_n$ *is a* $(\overline{q} - \underline{q} + 1)$ *vector of zeros except for the* n^{th} *entry which equals* 1.

[7]Recall these are independent of time and depend only on q.

5. EXPECTED PROFIT FRONTIERS: FINE TUNING HIGH-FREQUENCY STRATEGIES

Here we show the tradeoff between expected profits and exposure to volatility of terminal PnL ($\mathcal{R}^1$), volatility of terminal inventory ($\mathcal{R}^2$), and volatility of lifetime inventory ($\mathcal{R}^3$). Unless otherwise stated, in this section, the model parameters are $T = 300$ seconds, $\kappa^{\pm} = 25$, $\lambda^{\pm} = 2$, $\sigma = 0.1$, $\overline{q} = -\underline{q} = 10$, $S_0 = 100$, and upon the arrival of a market order prices jump by $\epsilon^{\pm} = \frac{1}{25}$ with a probability of $\frac{1}{10}$ and do not jump with probability $\frac{9}{10}$.

5.1. Expected Profit Frontiers and Risk Metrics for High-Frequency Traders

HF traders seek to maximize expected profits but also take into account the tradeoff between profits and their exposure to inventory risk. In Figure 5.1 we depict three investment frontiers that show the risk-return profile for the three risk metrics developed above as the penalty parameter ϕ and the liquidation cost parameter α vary. To depict

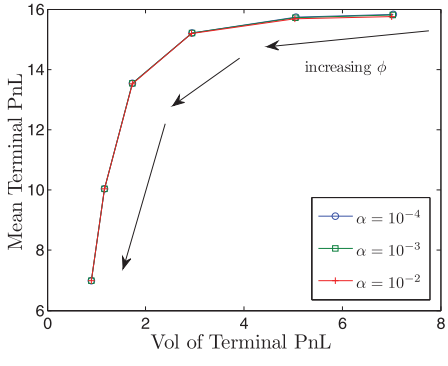

(a) Risk Measure 1

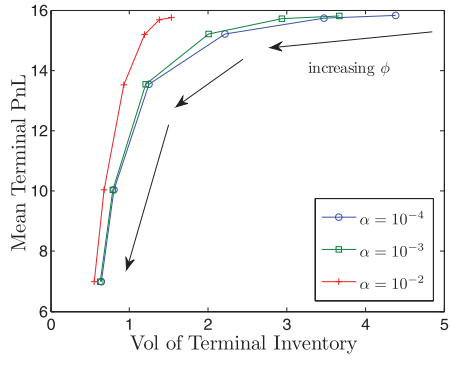

(b) Risk Measure 2

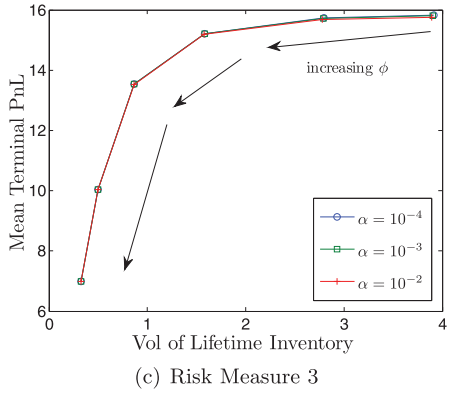

(c) Risk Measure 3

FIGURE 5.1. The effects of the penalty parameter ϕ and the terminal liquidation cost α on the risk-return profiles. Each point corresponds to $\phi = \{0, 0.01, 0.1, 1, 5, 10\}$. The investment horizon $T = 300$ seconds and the remaining model parameters are $\kappa^{\pm} = 25$, $\lambda^{\pm} = 2$, $\epsilon^{\pm} = \frac{1}{25}\mathfrak{a}^{\pm}$ where $\mathfrak{a}^{\pm}$ are independent Bernoulli r.v. with success prob of $\frac{1}{10}$, $\sigma = 0.1$, $\overline{q} = -\underline{q} = 10$, and $S_0 = 100$.

the frontiers we use the optimal strategy in Proposition 2.1 and calculate the risk metrics by numerical integrating the system of ODEs (A.3) and (A.6).[8]

Panels (a), (b), and (c) show the tradeoff between the mean of the PnL and the volatility of the PnL, volatility of terminal inventory, and volatility of lifetime inventory, respectively. The three panels show that increasing the inventory penalty parameter ϕ reduces both the mean PnL and the exposure to risk as measured by $\mathcal{R}_1$, $\mathcal{R}_2$, and $\mathcal{R}_3$. Moreover, it is interesting to see that there are ranges where fine tuning ϕ can significantly reduce the risk of the strategy without a significant reduction in expected profits.

The panels also show how the expected profit frontiers depend on the liquidation penalty α. In panels (a) and (c) we observe that the liquidation penalty has virtually no effect on the tradeoff between the mean of the PnL and both the volatility of the PnL and the volatility of lifetime inventory. On the other hand, panel (b) shows that the frontier that depicts the tradeoff between expected profits and the variance of final inventory is mildly affected by low values of the liquidation penalty, but when the penalty is increased to $\alpha = 0.01$ the frontier shifts to the left by an economically significant amount.[9] A shift to the left implies that by increasing α (and not varying ϕ) the HF trader can achieve the same level of expected profits and considerably reduce her exposure to final inventory risk.

Overall, Figure 5.1 shows that HF traders who maximize terminal expected wealth and are constrained by inventories and capital requirements can fine tune their strategies to find an optimal point in the expected profit frontiers. The constraints are either self-imposed or required by a regulator or the exchange and the severity of these constraints will determine what is the HF trader's optimal choice.

5.2. Distribution of PnL, Terminal Inventory, and Lifetime Inventory

5.2.1. Symmetric arrival of market orders. Moreover, to further understand how HF traders can fine tune their strategies to maximize expected profits, while obeying risk controls, we employ 10,000 simulations[10] and illustrate the effects that the penalties $\phi = \{0, 0.1, 2\}$ and $\alpha = \{0, 0.001, 0.01\}$ have on the distribution of the PnL, terminal inventory q_T, and inventory lifetime when the optimal strategy derived in Proposition 2.1 is employed. We recall that the model parameters are $\lambda^{\pm} = 2$ (i.e., symmetric arrival of market orders), $\kappa^{\pm} = 25$, $\overline{q} = -\underline{q} = 10$, $\sigma = 0.1$, $T = 300$ seconds and upon the arrival of a market order prices jump by $\epsilon^{\pm} = \frac{1}{25}$ with a probability of $\frac{1}{10}$ and do not jump with probability $\frac{9}{10}$. The results are reported in Panel A of Table 5.1.

[8]Note that the various matrices and vectors appearing in these ODEs will be time-dependent when the optimal strategy from Proposition 2.1 is inserted—which is why numerical integration is necessary.

[9]We note that for panels (a) and (c) the expected profit frontiers that result from employing the optimal limiting strategy in Proposition 2.6, and calculating the risk metric using the results in Proposition 4.1, coincide with the numerical solution of the optimal strategy in Proposition 2.1 presented here—the effect of the terminal date T on the expected profit frontiers is negligible.

[10]One path of the simulation is obtained by generating the time of the next market order (occurring at a rate of $\lambda^+ + \lambda^-$); generating a uniform random variable to decide if it is a buy/sell (it is a buy with probability $\frac{\lambda^+}{\lambda^+ + \lambda^-}$); generating a uniform random variable to decide if the limit order on the appropriate side of the book is filled (occurs with probability $e^{-\kappa^{\pm}\delta_t^{*\pm}}$). Moreover, the cash and inventories are tracked and the process repeated until maturity at which point any remaining inventory is liquidated (accounting for the liquidation cost). It is important to note that this process effectively prevents the optimal spreads from being negative, because a negative spread implies that market orders arrive faster. However, here market orders always arrive at the specified rate. See Section 6 for further discussion of this and related points.

TABLE 5.1

Mean and Standard Deviations of the PnL, Lifetime Inventory, and Terminal Inventory from Following the Optimal Policy with $\kappa^{\pm} = 25$, $\epsilon^{\pm} = \frac{1}{25}\mathfrak{a}^{\pm}$ where $\mathfrak{a}^{\pm}$ Are Independent Bernoulli r.v. with Success Prob of $\frac{1}{10}$, $\overline{q} = -\underline{q} = 10$, $\sigma = 0.1$, and $T = 300$ Seconds as the Penalty Parameter ϕ and Liquidation Cost Parameter α Vary

Panel A

Symmetric market: $\lambda^+ = 2, \lambda^- = 2$

	PnL			Lifetime inventory			Terminal Inventory		
	$\phi = 0$	$\phi = 0.1$	$\phi = 2$	$\phi = 0$	$\phi = 0.1$	$\phi = 2$	$\phi = 0$	$\phi = 0.1$	$\phi = 2$
$\alpha = 0$	15.776	15.210	12.392	0.008	−0.000	−0.000	0.002	−0.017	−0.006
	(7.10)	(2.97)	(1.52)	(3.92)	(1.59)	(0.74)	(4.75)	(2.20)	(0.95)
$\alpha = 0.001$	15.802	15.170	12.375	0.007	−0.000	−0.000	0.005	−0.017	−0.006
	(7.13)	(2.94)	(1.54)	(3.89)	(1.59)	(0.74)	(3.54)	(1.98)	(0.93)
$\alpha = 0.01$	15.716	15.201	12.368	0.007	−0.000	−0.000	−0.008	0.001	−0.003
	(7.02)	(2.96)	(1.52)	(3.85)	(1.58)	(0.74)	(1.51)	(1.26)	(0.80)

Panel B

Asymmetric market: $\lambda^+ = 1, \lambda^- = 3$

	PnL			Lifetime inventory			Terminal inventory		
	$\phi = 0$	$\phi = 0.1$	$\phi = 2$	$\phi = 0$	$\phi = 0.1$	$\phi = 2$	$\phi = 0$	$\phi = 0.1$	$\phi = 2$
$\alpha = 0$	31.596	22.338	10.524	−8.640	−3.886	−0.127	−7.143	−1.523	0.259
	(15.58)	(7.52)	(1.54)	(1.45)	(1.61)	(0.74)	(1.89)	(2.16)	(0.94)
$\alpha = 0.001$	31.547	22.415	10.566	−8.631	−3.883	−0.127	−6.023	−1.288	0.244
	(15.69)	(7.54)	(1.54)	(1.46)	(1.61)	(0.74)	(2.05)	(1.96)	(0.93)
$\alpha = 0.01$	31.317	22.361	10.534	−8.592	−3.874	−0.128	−2.272	−0.620	0.163
	(15.70)	(7.49)	(1.52)	(1.52)	(1.61)	(0.74)	(1.63)	(1.28)	(0.81)

The table shows the mean and standard deviations of the PnL, lifetime inventory, and terminal inventory from following the optimal policy. We observe that for any level of liquidation cost α, the effect of an increase in the penalty ϕ reduces both the expected profit and its standard deviation.

It is also clear, that everything else equal, the HF trader will prefer strategies where the lifetime inventory is concentrated around the optimal level of inventories which in this case is zero because $\lambda^+ = \lambda^-$. In the table we see that penalizing asset holdings along the entire path of the strategy with $\phi = 2$ ensures that inventories revert to zero very quickly.

Panel A clearly shows that the higher is the liquidation penalty and the higher is the penalization of inventories throughout the entire strategy, the terminal inventory is more concentrated around zero. This information will be used by the HF trader to choose among strategies with similar PnL distributions by picking the ones that are least exposed to terminal inventory.

Thus, given the liquidation costs (which proxy for the price impact of large market orders as well as capturing the fees charged by the exchange for taking liquidity) the HF trader will fine tune her strategy by looking at Table 5.1 and picking a preferred ϕ and/or by looking at the expected profit. For example, an HF trader that is highly constrained by the amount of inventory that she can hold at any point in time, will select an HF strategy with a high ϕ at the expense of low expected profits—the HF trader must ensure that the expected profits are sufficiently high to allow her to recover the fixed and variable costs from being in the business.

Finally, we remark that the liquidation penalty α seems to be useful only to reduce the volatility of the final inventory. On the other hand, ϕ plays a prominent role in the distribution of the PnL, terminal inventory, and lifetime inventory. The inventory penalty parameter ϕ is self-imposed by the HF trader and captures how binding are the constraints on intraday inventory exposures for the HF trader which could be different depending on the market—and particular asset—she is trading.

5.2.2. Asymmetric arrival of market orders. Above we assumed that the arrival rate of market buy and sell orders was the same, on average there were four market orders per second and half of them were buy and the other half were sell. Here we proceed as above but assume that the arrival of market orders is asymmetric, $\lambda^+ = 1$ and $\lambda^- = 3$, so that on average there are four market orders per second but there are much more sell than buy orders (market momentum). We perform 10,000 simulations and the remaining parameters are as above: $\phi = \{0, 0.1, 2\}$, $\alpha = \{0, 0.001, 0.01\}$, $\kappa^{\pm} = 25$, $\overline{q} = -\underline{q} = 10$, $\sigma = 0.1$, $T = 300$ seconds and upon the arrival of a market order prices jump by $\epsilon^{\pm} = \frac{1}{25}$ with a probability of $\frac{1}{10}$ and do not jump with probability $\frac{9}{10}$. We report the results in Panel B of Table 5.1.

Panel B shows that when inventories are not penalized throughout the entire strategy, the mean PnL in the asymmetric market is higher than the mean PnL resulting from a symmetric market. This occurs because in the former the strategy benefits from roundtrip trades as well as inventory appreciation. When the strategy is far away from expiry it is optimal to quickly build a short inventory position, and then start executing roundtrips around that short optimal level. The excess arrival of market sell orders exerts a downward pressure on the midquote which means that on average the HF trader earns a profit on short positions.

Panel B also shows that even though the HF trader can profit from the expected appreciation of her short position, if $\phi = 0.1$ her inventory levels are "forced" to be

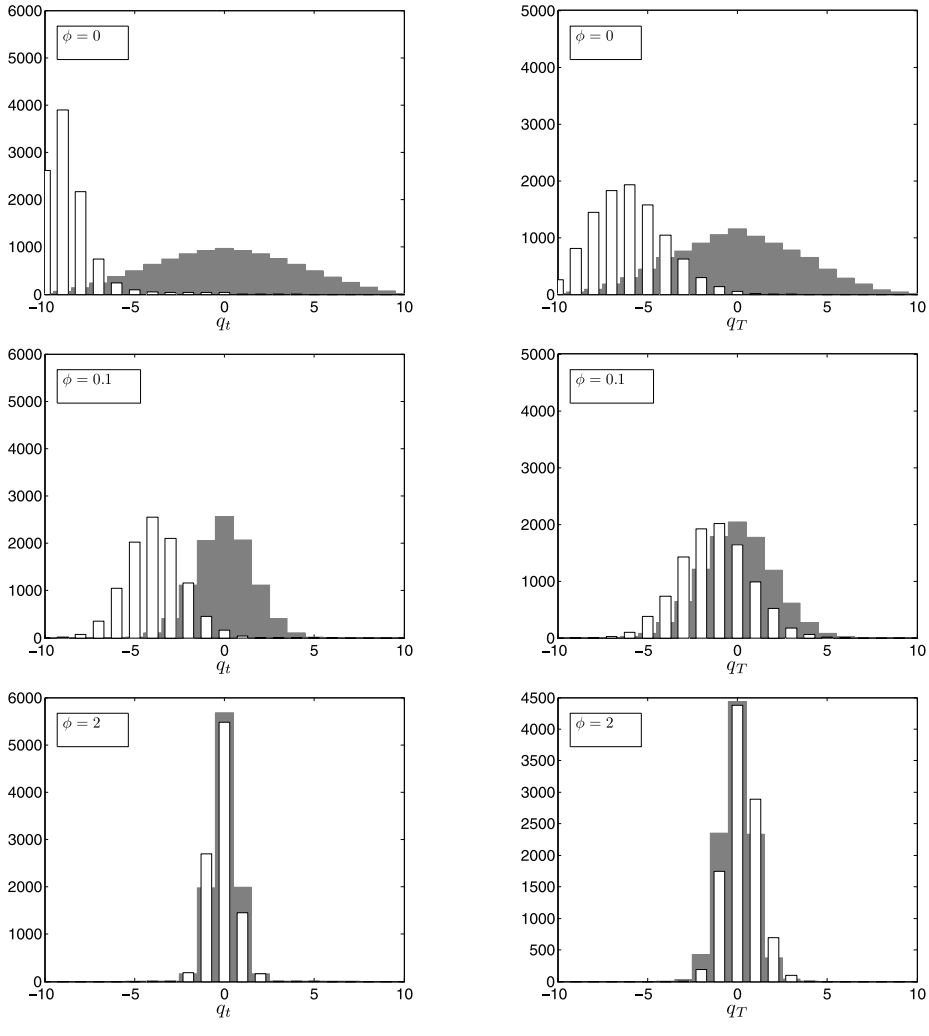

(a) Lifetime Inventory Distribution (b) Terminal Inventory Distribution

FIGURE 5.2. Distribution of inventory using lifetime inventory and terminal inventory levels from following the optimal policy with $\lambda^+ = 1$, $\lambda^- = 3$ (white bars), and $\lambda^\pm = 2$ (gray bars), $\kappa^\pm = 25$, $\epsilon^\pm = \frac{1}{25}\mathfrak{a}^\pm$ where $\mathfrak{a}^\pm$ are independent Bernoulli r.v. with success prob of $\frac{1}{10}$, $\overline{q} = -\underline{q} = 10$, $\sigma = 0.1$, $T = 300$ seconds, $\alpha = 0.001$, and $\phi = 0$, 0.1, and 2 (from top to bottom).

closer to zero and these appreciation gains are forgone. Furthermore, if the penalization of inventory is further increased to $\phi = 2$, then the distribution of the PnL in the asymmetric market shifts to the left of that of a symmetric market. The intuition is that in an asymmetric market the expected appreciation gains are an important component of the PnL, but these are considerably curbed by ϕ.

Figure 5.2 shows histograms of lifetime inventory, column (a), and terminal inventories, column (b), for both the symmetric (gray bars) and asymmetric (white bars) markets. Column (a) shows that when the penalty parameter ϕ is increased, the distribution of

lifetime inventory becomes more symmetric and centered at zero. When $\phi = 0$ the optimal strategy profits from the expected decline in prices by building a short position in the asset and then performing as many roundtrips as possible, but always reverting to an optimal short inventory position—the strategy reaches the minimum allowed $q = -10$ and spends most of the time between $q = -9$ and $q = -10$. As the inventory penalties are increased, the optimal strategy relies less on gains via inventory appreciation and more on roundtrips with inventories reverting to zero.

Column (b) shows that when $\phi = 0$ the distribution of terminal inventory is negatively skewed and to the left of that of the symmetric market. As the penalty parameter ϕ is increased the distribution of terminal inventory becomes skewed to the right. Initially this result may seem counterintuitive because one would expect that in a market with negative momentum it is always optimal to keep a short position which is either unwound fractions of a second before T or liquidated at time T. But on the contrary, the figure shows that the histogram is positively skewed and that the mean final inventory level is also positive. To understand the intuition behind this result assume that we are at expiry or milliseconds before expiry. In this case the chances of completing a roundtrip trade are essentially zero. Thus, the HF trader knows that if anything she can only execute the first leg of a trade to then liquidate it at the midprice and pick up the liquidation penalty. If we assume, for simplicity, that the liquidation penalty is $\alpha = 0$ and we have reached expiry, the best strategy is to maximize the probability of being filled on either side of the book—the same argument that follows still applies if the strategy is microseconds before expiry. In this case the optimal deltas are given by $\delta^{\pm} = \delta = 1/\kappa + \varepsilon$ (see Proposition 2.1) and the fill rates are $\Lambda^{\pm} = \lambda^{\pm} e^{-\kappa\delta}$. Obviously $\Lambda^{-} > \Lambda^{+}$, because $\lambda^{-} > \lambda^{+}$, and one can verify that although at the beginning of the strategy it is optimal to build short positions to gain from the expected fall in prices (see panel (b) in Figure 3.5), as we approach expiry the optimal postings will be such that the HF trader is more likely to be buying rather than selling the asset. This explains why we see the positive skew in the distribution shown in the last picture of column (b) in Figure 5.2 for $\phi = 2$. We remark that although this effect is always present at times close to expiry, in our examples it only becomes visible in the last picture of column (b) occurs because a high ϕ "forces" the strategy not to deviate from zero inventories before T, but at the last moment the optimal limit orders are more likely to be filled by sell, rather than buy, market orders. Therefore, at the end of the strategy there is an upward pressure on inventories and the distribution of terminal inventory is shifted and skewed to the right.

Finally, if we look at the last pictures of columns (a) and (b) we corroborate the intuition given above to explain the positive skew in terminal inventories. To this end, note that although the distribution of lifetime inventory is left skewed, the distribution of the final inventory is skewed to the right, which confirms that very close to and at expiry the optimal strategy is to increase inventories as much as possible—this last upward pressure on inventories is the same for all cases, but only becomes visible for high ϕ.

6. CONSTRAINED STRATEGIES

As a final note, we discuss the issue of negative spreads ($\delta_t^{\pm} < 0$). Recall that conditional on the arrival of a market order, the probability of a limit order at price $S_t \pm \delta_t^{\pm}$ being filled is assumed to be $e^{-\kappa^{\pm}\delta_t^{\pm}}$ resulting in a rate of arrival of filled limit orders of $\Lambda^{\pm} = \lambda^{\pm} e^{-\kappa^{\pm}\delta_t^{\pm}}$. Consequently, negative spreads ($\delta_t^{\pm} < 0$) can be interpreted as increasing the rate of arrival of market orders. However, this may not be a desirable feature. To correct

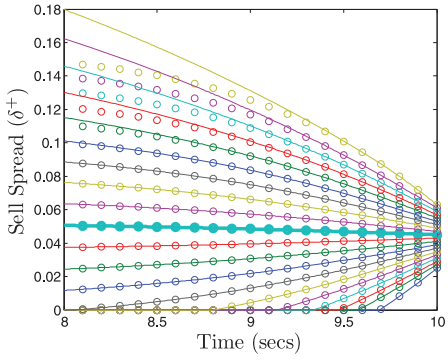

(a) $\lambda^+ = \lambda^- = 2$

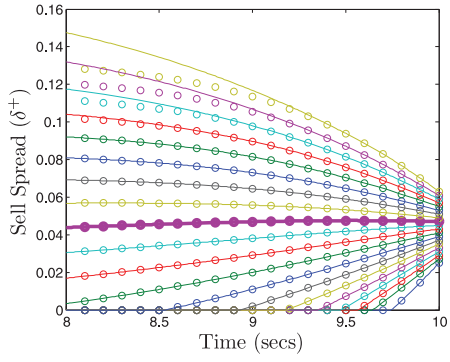

(b) $\lambda^+ = 1,\ \lambda^- = 3$

FIGURE 6.1. Comparison between the exact constrained spreads (solid lines) and the spreads computed by constraining afterwards (circles). The remaining model parameters are $T = 10, \alpha = 0.001, \phi = 0.1, \kappa^+ = \kappa^- = 25, \varepsilon^+ = \varepsilon^- = \frac{1}{250}, \overline{q} = -\underline{q} = 10$, and $\sigma = 0.1$. The inventories decrease from 10 to -9 from the smallest spreads to the largest. The thick lines show the spreads at the mean-reverting inventory level: top panels $q = 0$, bottom panels $q = -1$.

this effect we can restrict the spreads to be nonnegative. In this case, the optimal controls are given by

$$\text{(6.1a)}\quad \delta^+(t,q) = \max\left(\frac{1}{\kappa^+} + \varepsilon^+ + \alpha(1-2q) - h_{q-1}(t) + h_q(t),\ 0\right), \quad q \neq -\underline{q},$$

and

$$\text{(6.1b)}\quad \delta^-(t,q) = \max\left(\frac{1}{\kappa^-} + \varepsilon^- + \alpha(1+2q) - h_{q+1}(t) + h_q(t),\ 0\right) \quad q \neq \overline{q}\,,$$

rather than (2.8). Upon substituting the feedback controls into the HJB equation the resulting coupled system of ODEs no longer admits a closed-form solution. However, it is straightforward to numerically integrate this system.

Figure 6.1 compares the spreads computed by numerically solving the ODEs to the spreads computed by ignoring the constraint but then replacing $\delta_t^\pm \to \max(\delta_t^\pm, 0)$ after the fact. This after-the-fact truncation of spreads to nonnegative values is how our simulations are carried out. Notice that for most of the inventory levels corresponding to small/moderate spreads the exact constrained spreads and the spreads constrained afterwards are essentially identical. Only at large spreads, corresponding to short inventory levels on the sell side of book, are the solutions different. However, these inventories are only rarely attained and the impact on the overall strategy appears to be minimal. The highlighted spreads correspond to the mean-reverting inventory levels—measured by the highest lifetimes of inventory. Only the sell side is shown here, the buy side is similar except that $\overline{q}$ and $-\underline{q}$ are interchanged in the sell side analog to Figure 6.1(a)—as required by Corollary 2.3—while in the analog to Figure 6.1(b), the symmetry is only approximate.

7. CONCLUSIONS

We propose risk metrics for HF trading strategies that profit from market making at very high speeds. The HF trader maximizes expected terminal wealth, and is capital and inventory constrained. These constraints enter into the optimization problem by penalizing accumulated inventories throughout the entire strategy, imposing a financial penalty for liquidating terminal inventories, and by imposing a cap on the amount of accumulated inventories that the HF trader can hold, long or short, at any point in time.

The risk metrics calculate the variance of terminal PnL, variance of terminal inventory, and the variance of lifetime inventory. We show that these metrics provide information that enables HF traders to fine tune their strategies to trade off expected profits against exposure to different metrics of inventory risk (which also proxy for capital constraints). We remark that the risk metrics we propose are for HF trading strategies in general and not limited to the particular ones discussed here.

For the optimal strategies we develop here, our analysis shows that final inventory penalties (which proxies the cost of unwinding positions at expiry of the strategy) affect the distribution of inventories at the expiry of the strategy, but have a very limited effect on the distribution of the PnL of the strategy. On the other hand, penalizing inventories throughout the entire strategy (the penalization affects inventory management but does not alter the PnL directly) has an important effect on the three risk metrics we propose. For example, we show that there are ranges of the penalization parameter where a very small increase in the parameter causes a considerable reduction in the risks borne by the HF trader (as given by our risk metrics), while expected profits are hardly affected.

Furthermore, the strategies include a component to recover adverse selection costs and induce mean reversion of inventories to an optimal level. This optimal level can change sign throughout the life of the strategy. When markets are asymmetric (arrival of sell market orders is different from that of buy orders) the optimal strategy benefits from both inventory appreciation and roundtrip trades. In a market with positive momentum—the arrival rate of buy market orders is higher than the arrival rate of market sell orders—the optimal strategy starts by quickly building positive inventories and is immediately followed by executing roundtrip trades while reverting to a time-dependent positive inventory level. However, as expiry is approached, the optimal postings are such that the optimal inventory level decreases and changes sign from positive to negative. This happens because very close to or at expiry the chances of executing roundtrip trades are virtually zero, and the expected gains from inventory appreciation (due to market momentum) are also very low, so the optimal strategy is to maximize the probability of being filled, and then liquidate the position, which in a market with positive momentum means that the optimal inventory level very close to or at expiry is negative.

Finally, we also discuss other possible risk metrics that are closely connected to the risk metrics we propose. In particular we show that the distribution of the lifetime inventory provides important information for HF traders who are capital and inventory constrained. For example, everything else equal, an HF trader prefers a strategy where inventories are turned over very quickly but the lifetime inventory is highest for the optimal level of inventory (which is not necessarily zero as discussed above). Thus, to provide further insights into the fine tuning of the optimal strategies derived here, we show the trade off between expected profits and the distribution of lifetime inventory as a function of the parameter ϕ that penalizes inventories throughout the entire strategy. In particular, we show that small increases in ϕ have an important effect on the distribution

of lifetime inventory while the distribution of the PnL (expected profits) is much less affected by this parameter—this provides a useful tool to fine tune HF strategies.

APPENDIX A: PROOFS

A.1. Proof of Proposition 2.2

Inserting the controls (2.8) into the HJB equation (2.7) and writing $h_q(t) = \frac{1}{\kappa}\ln\omega_q(t)$, after some straightforward computations, one finds that $\omega_q(t)$ satisfy the coupled system of equations

$$\partial_t\boldsymbol{\omega}(t) + \mathbf{A}\boldsymbol{\omega}(t) = \mathbf{0}. \tag{A.1}$$

The boundary condition $h_q(T) = 0$ translates into $\omega_q(T) = 1$. This is a first-order coupled system of ODEs. The solution to (A.1) is trivially given by (2.10).

A.2. Proof of Corollary 2.3

Let $\widetilde{\boldsymbol{\omega}}_q(t) = e^{\kappa\alpha q^2}\boldsymbol{\omega}_q(t)$, then through straightforward computations we see that $\partial_t\widetilde{\boldsymbol{\omega}} + \widetilde{\mathbf{A}}\boldsymbol{\omega} = \mathbf{0}$ subject to $\widetilde{\boldsymbol{\omega}}_q(T) = e^{\kappa\alpha q^2}$ where

$$\widetilde{\mathbf{A}}_{i,q} = \begin{cases} -\phi\sigma^2\kappa\, q^2, & i = q, \\ \lambda\, e^{-\kappa(\frac{1}{\kappa}+\varepsilon)}, & i = q-1, q+1 \\ 0, & \text{otherwise.} \end{cases}$$

Clearly, this system is symmetric under $q \to -q$, therefore $\widetilde{\boldsymbol{\omega}}_q(t) = \widetilde{\boldsymbol{\omega}}_{-q}(t)$. Furthermore, $\boldsymbol{\omega}_q(t) = \boldsymbol{\omega}_{-q}(t)$ and hence $h_q(t) = h_{-q}(t)$. Finally, the optimal feedback control $\delta^{+*}(t,q) = \frac{1}{\kappa} + \varepsilon + \alpha(1-2q) - h_{q-1}(t) + h_q(t) = \frac{1}{\kappa} + \varepsilon + \alpha(1+2(-q)) - h_{-(q-1)}(t) + h_{-q}(t) = \delta^{-*}(t,-q)$. □

A.3. Proof of Proposition 2.6

Part 1: Clearly, $e^{\mathbf{A}(T-t)} = \mathbf{U}e^{\mathbf{D}(T-t)}\mathbf{U}^{-1}$ where $\mathbf{D}$ is the diagonal matrix of eigenvalues of $\mathbf{V}$. We then have $\exp\{\mathbf{A}(T-t)\}\mathbf{1} = \mathbf{U}\exp\{[\mathbf{D}]_\ell(T-t)\}\mathbf{U}^{-1}\mathbf{1} + o(\exp\{\mathbf{D}_{\ell\ell}(T-t)\})$, where $[\mathbf{D}]_\ell$ denotes the diagonal matrix whose entries are all zero except for the ℓ^{th} entry. Next, notice that $\exp\{[\mathbf{D}]_\ell(T-t)\}\mathbf{U}^{-1}\mathbf{1} = [\mathbf{U}^{-1}\mathbf{1}]_\ell\, e^{\mathbf{D}_{\ell\ell}(T-t)}$. Consequently, $\mathbf{h}(t) = \frac{\mathbf{D}_{\ell\ell}}{\kappa}(T-t) + \frac{1}{\kappa}\ln(\mathbf{U}[\mathbf{U}^{-1}\mathbf{1}]_\ell) + o(T-t)$ and the result immediately follows since the linear term is independent of q and applying Proposition 2.1.

Part 2: We have $e^{\mathbf{A}(T-t)} = \mathbb{I} + \mathbf{A}(T-t) + o(T-t)$, therefore $\boldsymbol{\omega}(t) = \mathbf{1} + \mathbf{A}\mathbf{1}(T-t) + o(T-t)$ and $\mathbf{h}(t) = \frac{1}{\kappa}\mathbf{A}\mathbf{1}\,(T-t) + o(T-t)$. Moreover, for $q \neq \overline{q}, \underline{q}$ we have

$$(\mathbf{A}\mathbf{1})_q = \lambda^+ e^{-\kappa(\frac{1}{\kappa}+\varepsilon^+-\alpha(1+2q))} + \lambda^- e^{-\kappa(\frac{1}{\kappa}+\varepsilon^-+\alpha(2q-1))} + q\kappa(\lambda^+\epsilon^+ - \lambda^-\epsilon^-) - \phi\sigma^2\kappa q^2\,.$$

Using Proposition 2.1, after some tedious computations we arrive at the result.

A.4. Proof of Proposition 4.1

Part 1: Since the rows of $\mathbf{B}$ sum to zero, and nondiagonal entries are positive, it represents a generator matrix of a Markov chain. Consequently, $\mathbf{1}$ is an eigenvector with eigenvalue zero and its remaining eigenvalues are negative. This proves part 1.

Part 2: We compute $M_1(t, x, q, S_t) \triangleq \mathbb{E}_t[X_T + q_T(S_T - \alpha\, q_T)]$ and $M_2(t, x, q, S_t) \triangleq \mathbb{E}_t[(X_T + q_T(S_T - \alpha\, q_T))^2]$ under the strategy δ_t^* separately. Clearly, $M_j(t, X_t, q_t, S_t)$ are martingales under any strategy, therefore they must each satisfy the following PIDE:

(A.2)
$$\begin{aligned} 0 = {} & \partial_t M_j + \frac{1}{2}\sigma^2 \partial_{SS} M_j \\ & + \lambda^+ \Big\{ e^{-\kappa\delta^{*+}(q)} \mathbb{E}(M_j(t, x + S + \delta^{*+}(q), q - 1, S + \epsilon^+) - M_j(t, x, q, S)) \\ & \qquad + \big(1 - e^{-\kappa\delta^{*+}(q)}\big)\mathbb{E}(M_j(t, x, q, S + \epsilon^+) - M_j(t, x, q, S))\Big\} \\ & + \lambda^- \Big\{ e^{-\kappa\delta^{*-}(q)} \mathbb{E}(M_j(t, x - S + \delta^{*-}(q), q + 1, S - \epsilon^-) - M_j(t, x, q, S)) \\ & \qquad + \big(1 - e^{-\kappa\delta^{*-}(q)}\big)\mathbb{E}(M_j(t, x, q, S - \epsilon^-) - M_j(t, x, q, S))\Big\}, \end{aligned}$$

for $q = \underline{q}, \underline{q} + 1 \ldots, \overline{q}$. Here, and in the remainder of this proof, we have written $\delta^{+*}(\underline{q}) = \delta^{-*}(\overline{q}) = +\infty$ to avoid having to treat these boundary cases separately.

Solving for M_1: Using the ansatz $M_1 = x + q(S - \alpha q) + m_1(t, q)$ then after some tedious computations, $m_1(q)$ satisfies

$$\partial_t \mathbf{m}_1 + \mathbf{B}\,\mathbf{m}_1 + \boldsymbol{\theta} = \mathbf{0}, \tag{A.3}$$

where the vector $\boldsymbol{\theta}$ has elements

$$\begin{aligned} \boldsymbol{\theta}_q = {} & (\varepsilon^+\widetilde{\lambda}^+ - \varepsilon^-\widetilde{\lambda}^-)\, q + e^{-\kappa\delta^{*+}(q)}\,\widetilde{\lambda}^+ \left(\delta^{*+}(q) - \varepsilon^+ - \alpha(1 - 2q)\right) \\ & + e^{-\kappa\delta^{*-}(q)}\,\widetilde{\lambda}^- \left(\delta^{*-}(q) - \varepsilon^- - \alpha(1 + 2q)\right) \end{aligned} \tag{A.4}$$

subject to $\mathbf{m}_1(T) = \mathbf{0}$. Delicate treatment of the zero eigenvalue of the matrix $\mathbf{B}$ is required, but it is not difficult to show that the solution to (A.3) is given by $\mathbf{m}_1 = \mathbf{V}\widetilde{\mathbf{m}}_1$ where the components of $\widetilde{\mathbf{m}}_1$ are

$$\widetilde{\mathbf{m}}_{1,j} = \begin{cases} (\mathbf{V}^{-1}\boldsymbol{\theta})_1\, \tau, & j = 1, \\ (\mathbf{V}^{-1}\boldsymbol{\theta})_j\, \dfrac{e^{d_j \tau} - 1}{d_j}, & j > 1. \end{cases} \tag{A.5}$$

Solving for M_2: Using the ansatz $M_2 = (x + q(S - \alpha q))^2 + m_2(t, x, q, S)$ with $m_2(t, x, q, S) = a(t, q) + b(t, q)\, x + c(t, q)\, S$, we find that $\mathbf{a}$, $\mathbf{b}$, and $\mathbf{c}$ satisfy the coupled system of equations

$$\partial_t \mathbf{a} + \mathbf{B}\,\mathbf{a} + \boldsymbol{\Upsilon}\,\mathbf{b} + \boldsymbol{\Omega}\,\mathbf{c} + \boldsymbol{\upsilon} = \mathbf{0}, \tag{A.6a}$$

$$\partial_t \mathbf{b} + \mathbf{B}\,\mathbf{b} + \boldsymbol{\chi} = \mathbf{0}, \tag{A.6b}$$

$$\partial_t \mathbf{c} + \mathbf{B}\,\mathbf{c} + \boldsymbol{\Psi}\,\mathbf{b} + \boldsymbol{\psi} = \mathbf{0}, \tag{A.6c}$$

where the elements of the matrices $\boldsymbol{\Upsilon}$, $\boldsymbol{\Omega}$, and $\boldsymbol{\Psi}$ are

$$\boldsymbol{\Upsilon}_{i,q} = \begin{cases} e^{-\kappa\,\delta^{*+}(q)}\,\widetilde{\lambda}^{+}\,\delta^{*+}(q), & i = q-1, \\ e^{-\kappa\,\delta^{*-}(q)}\,\widetilde{\lambda}^{-}\,\delta^{*-}(q), & i = q+1, \\ 0 & \text{otherwise}, \end{cases} \tag{A.7a}$$

$$\boldsymbol{\Omega}_{i,q} = \begin{cases} \left\{\left(1 - e^{-\kappa\,\delta^{*+}(q)}\right)\widetilde{\lambda}^{+} - \left(1 - e^{-\kappa\,\delta^{*-}(q)}\right)\widetilde{\lambda}^{-}\right\}\varepsilon, & i = q, \\ e^{-\kappa\,\delta^{*+}(q)}\,\widetilde{\lambda}^{+}\,\varepsilon, & i = q-1, \\ -e^{-\kappa\,\delta^{*-}(q)}\,\widetilde{\lambda}^{-}\,\varepsilon, & i = q+1, \\ 0 & \text{otherwise}, \end{cases} \tag{A.7b}$$

$$\boldsymbol{\Psi}_{i,q} = \begin{cases} e^{-\kappa\,\delta^{*+}(q)}\,\widetilde{\lambda}^{+}, & i = q-1, \\ -e^{-\kappa\,\delta^{*-}(q)}\,\widetilde{\lambda}^{-}, & i = q+1, \\ 0 & \text{otherwise}, \end{cases} \tag{A.7c}$$

and the vectors $\boldsymbol{\upsilon}$, $\boldsymbol{\chi}$, and $\boldsymbol{\Psi}$ have elements

$$\begin{aligned} \boldsymbol{\upsilon}_q = {} & \widetilde{\lambda}^{+}\left\{q^2\left(\varepsilon_2^{+} - 2\alpha q\varepsilon^{+}\right) + e^{-\kappa\delta^{*+}}(\widehat{\delta}^{+}(q) + 2\alpha q)(\widehat{\delta}^{+}(q) - 2\alpha(q^2 - q) + 2q\varepsilon^{+})\right\} \\ & + \widetilde{\lambda}^{-}\left\{q^2\left(\varepsilon_2^{-} + 2\alpha q\varepsilon^{-}\right) + e^{-\kappa\delta^{*-}}(\widehat{\delta}^{-}(q) - 2\alpha q)(\widehat{\delta}^{-}(q) - 2\alpha(q^2 + q) - 2q\varepsilon^{-})\right\}, \end{aligned} \tag{A.8a}$$

$$\boldsymbol{\chi}_q = 2\widetilde{\lambda}^{+}\left\{(\widehat{\delta}^{+}(q) + 2\alpha q)e^{-\kappa\delta^{*+}(q)} + \varepsilon^{+}\,q\right\} + 2\widetilde{\lambda}^{-}\left\{(\widehat{\delta}^{-}(q) - 2\alpha q)e^{-\kappa\delta^{*-}(q)} - \varepsilon^{-}\,q\right\}, \tag{A.8b}$$

$$\boldsymbol{\psi}_q = 2\widetilde{\lambda}^{+}\left\{q(\widehat{\delta}^{+}(q) + 2\alpha q)e^{-\kappa\delta^{*+}(q)} + q^2\varepsilon^{+}\right\} + 2\widetilde{\lambda}^{-}\left\{q(\widehat{\delta}^{-}(q) - 2\alpha q)e^{-\kappa\delta^{*-}(q)} - q^2\varepsilon^{-}\right\}, \tag{A.8c}$$

where

$$\widehat{\delta}^{\pm}(q) = \delta^{*\pm}(q) - \alpha - \varepsilon^{\pm}, \quad \text{and} \quad \varepsilon_2^{\pm} = \mathbb{E}[(\epsilon^{\pm})^2]. \tag{A.8d}$$

The existence of the zero eigenvalue in the matrix **B** introduces some complexity in the solution to ODEs. This system of ODEs is solved by (i) solving (A.6b) by diagonalization (using the eigenvectors of **B**). The solution to the equations corresponding to nonzero(negative) eignevalues decays exponentially. The solution to the equations corresponding to the zero eigenvalue grows linearly. (ii) Substituting the solution for **b** into (A.6c), diagonalizing once again and carefully separating out the zero-eigenvalues which now interact with the remaining eigenvalues. The zero eignevalue solution now contains linear combinations of τ, τ^2, and $e^{-d_k\tau}$, while the remaining eigenvalue solutions are linear combinations of τ, $\tau e^{-d_k\tau}$, and $e^{-d_k\tau}$. (iii) Substituting the solution for both **b** and **c** into (A.6a), diagonalizing again and treating the interaction of the zero and nonzero eigenvalues. The zero eigenvalue solution is now a linear combination of τ, τ^2, τ^3, $e^{-d_k\tau}$, and $\tau e^{-d_k\tau}$, while the nonzero eigenvalue solutions are linear combinations of τ, τ^2, $e^{-d_k\tau}$, $\tau e^{-d_k\tau}$, and $\tau^2 e^{-d_k\tau}$. After some tedious computations we find that

(A.9) $$\mathbf{a} = \mathbf{V}\widetilde{\mathbf{a}}, \quad \mathbf{b} = \mathbf{V}\widetilde{\mathbf{b}}, \quad \mathbf{c} = \mathbf{V}\widetilde{\mathbf{c}},$$

where,

(A.10a) $$\widetilde{\mathbf{a}}_j = \begin{cases} i_1\,\tau + m_1\,\tau^2 + r_1\,\tau^3 + \sum_{k\neq 1} n_{1k}\,(e^{d_k\,\tau} - 1) + \sum_{k\neq 1} o_{1k}\,\tau\,e^{d_k\,\tau}, & j = 1\,, \\ i_j\,\tau + m_j\,\tau^2 + h_j(1 - e^{d_j\,\tau}) + \sum_{k\neq j,1} n_{jk}\,(e^{d_k\,\tau} - e^{d_j\,\tau}) & \\ \quad + \sum_{k\neq j,1} o_{jk}\,\tau\,e^{d_k\,\tau} + q_j\,\tau^2\,e^{d_j\tau}, & j > 1\,, \end{cases}$$

(A.10b) $$\widetilde{\mathbf{b}}_j = \begin{cases} \widetilde{\boldsymbol{\chi}}_1\,\tau, & j = 1\,, \\ \widetilde{\boldsymbol{\chi}}_j\,\dfrac{e^{d_j\,\tau} - 1}{d_j}, & j > 1\,, \end{cases}$$

(A.10c) $$\widetilde{\mathbf{c}}_j = \begin{cases} -\Theta_1\,\tau + \dfrac{1}{2}\widetilde{\boldsymbol{\Psi}}_{11}\widetilde{\boldsymbol{\chi}}_1\,\tau^2 + \sum_{k\neq 1}\dfrac{\widehat{\boldsymbol{\Psi}}_{1k}}{d_k}(e^{d_k\tau} - 1), & j = 1\,, \\ -\Theta_j\,\tau + \eta_j\,\dfrac{e^{d_j\,\tau} - 1}{d_j} + \sum_{k\neq j,1}\dfrac{\widehat{\boldsymbol{\Psi}}_{jk}}{d_k - d_j}(e^{d_k\,\tau} - e^{d_j\,\tau}) + \widehat{\boldsymbol{\Psi}}_{jj}\,\tau\,e^{d_j\,\tau}, & j > 1, \end{cases}$$

and the various coefficients are given by

(A.11a)
$$\begin{aligned} &\widetilde{\boldsymbol{\chi}}_j = (\mathbf{V}^{-1}\boldsymbol{\chi})_j, && \widetilde{\boldsymbol{\psi}}_j = (\mathbf{V}^{-1}\boldsymbol{\psi})_j, && \widetilde{\boldsymbol{\upsilon}}_j = (\mathbf{V}^{-1}\boldsymbol{\upsilon})_j, \\ &\widetilde{\boldsymbol{\Psi}}_{jk} = (\mathbf{V}^{-1}\boldsymbol{\Psi}\mathbf{V})_{jk}, && \widetilde{\boldsymbol{\Omega}}_{jk} = (\mathbf{V}^{-1}\boldsymbol{\Psi}\boldsymbol{\Omega})_{jk}, && \widetilde{\boldsymbol{\Upsilon}}_{jk} = (\mathbf{V}^{-1}\boldsymbol{\Psi}\boldsymbol{\Upsilon})_{jk}, \\ &\widehat{\boldsymbol{\Psi}}_{jk} = \frac{\widetilde{\boldsymbol{\Psi}}_{jk}\widetilde{\boldsymbol{\chi}}_k}{d_k}, && \widehat{\boldsymbol{\Omega}}_{jk} = \frac{\widetilde{\boldsymbol{\Omega}}_{jk}\,\eta_k^c}{d_k}, && \widehat{\boldsymbol{\Upsilon}}_{jk} = \frac{\widetilde{\boldsymbol{\Upsilon}}_{jk}\widetilde{\boldsymbol{\chi}}_k}{d_k}, \\ &\overline{\boldsymbol{\Omega}}_{jk} = \widetilde{\boldsymbol{\Omega}}_{jk}\,\widehat{\boldsymbol{\Psi}}_{kk}, \end{aligned}$$

(A.11b) $$\Theta_j = \begin{cases} \sum_{k\neq 1}\widehat{\boldsymbol{\Psi}}_{1k} - \widetilde{\boldsymbol{\psi}}_1, & j = 1, \\ \dfrac{\widetilde{\boldsymbol{\Psi}}_{j1}\widetilde{\boldsymbol{\chi}}_1}{d_j}, & j > 1, \end{cases}$$

(A.11c) $$\eta_j = \widetilde{\boldsymbol{\psi}}_j + \Theta_j - \sum_{k\neq 1}\widehat{\boldsymbol{\Psi}}_{jk},$$

(A.11d) $$i_j = (2m_j - \beta_j)/d_j, \quad m_j = \frac{1}{2}\widetilde{\boldsymbol{\Omega}}_{j1}\widetilde{\boldsymbol{\Psi}}_{11}\widetilde{\boldsymbol{\chi}}_1/d_j, \quad h_j = (i_j - \delta_j)/d_j,$$

(A.11e) $$\delta_j = \widetilde{\boldsymbol{\upsilon}}_j - \sum_{k\neq 1}\zeta_{jk}, \quad \zeta_{jk} = \frac{1}{d_k}\left(\widetilde{\boldsymbol{\Upsilon}}_{jk}\widetilde{\boldsymbol{\chi}}_1 + \widetilde{\boldsymbol{\Omega}}_{jk}\eta_k + \widetilde{\boldsymbol{\Omega}}_{j1}\widehat{\boldsymbol{\Psi}}_{1k}\right),$$

(A.11f)
$$\beta_j = \widetilde{\boldsymbol{\Upsilon}}_{j1}\widetilde{\boldsymbol{\chi}}_1 - \sum_{k\neq 1}\widetilde{\boldsymbol{\Omega}}_{jk}\Theta_k, \quad \xi_{jk} = \zeta_{jk} - \sum_{l\neq k,1}(g_{kl}-g_{lk}), \quad g_{kl} = \sum_{m\neq l,1}\frac{\widetilde{\boldsymbol{\Omega}}_{kl}\widetilde{\boldsymbol{\Psi}}_{lm}}{d_m - d_l},$$

(A.11g)
$$n_{jk} = \begin{cases} \dfrac{\xi_{1k}-o_{1k}}{d_k}, & j=1\,, \\ (\xi_{jk}-o_{jk})/(d_k-d_j), & j\neq 1\,, \end{cases} \qquad o_{jk} = \begin{cases} \widetilde{\boldsymbol{\Omega}}_{1k}\widehat{\boldsymbol{\Psi}}_{kk}/d_k, & j=1\,, \\ \xi_{jj}+d_j n_{jj}, & k=j\neq 1, \\ \widetilde{\boldsymbol{\Omega}}_{jk}\widehat{\boldsymbol{\Psi}}_{kk}/(d_k-d_j), & k\neq j,1\,, \end{cases}$$

(A.11h)
$$q_j = \frac{1}{2}\widetilde{\boldsymbol{\Omega}}_{jj}\widehat{\boldsymbol{\Psi}}_{jj}\,, \quad r_1 = \frac{1}{6}\widetilde{\boldsymbol{\Omega}}_{11}\widetilde{\boldsymbol{\Psi}}_{11}\widetilde{\boldsymbol{\chi}}_1\,.$$

Large T asymptotics: Since the matrix $\mathbf{B}$ is a generator matrix of a Markov chain, its nonzero eigenvalues are negative. Consequently, for large τ, only the polynomial terms[11] of the functions $\widetilde{\mathbf{a}}$, $\widetilde{\mathbf{b}}$, and $\widetilde{\mathbf{c}}$ contribute. This leads to the stated result.

Part 3: For risk metric $\mathcal{R}^2$, note that $\mathbb{E}_t[q_T^2]$ and $\mathbb{E}_t[q_T]$ are given by Proposition A.1 with $\boldsymbol{\ell}=0$ and $\varphi(q)=q^2$ and q, respectively. Consequently, we have

(A.12)
$$\mathcal{R}^2(t,\cdot) = \mathbf{V}\begin{pmatrix} (\mathbf{V}^{-1}\boldsymbol{\beta})_1 \\ e^{\mathbf{d}_2(T-t)}(\mathbf{V}^{-1}\boldsymbol{\beta})_2 \\ \vdots \\ e^{\mathbf{d}_{\overline{q}-\underline{q}+1}(T-t)}(\mathbf{V}^{-1}\boldsymbol{\beta})_{\overline{q}-\underline{q}+1} \end{pmatrix} - \left(\mathbf{V}\begin{pmatrix} (\mathbf{V}^{-1}\boldsymbol{\gamma})_1 \\ e^{\mathbf{d}_2(T-t)}(\mathbf{V}^{-1}\boldsymbol{\gamma})_2 \\ \vdots \\ e^{\mathbf{d}_{\overline{q}-\underline{q}+1}(T-t)}(\mathbf{V}^{-1}\boldsymbol{\gamma})_{\overline{q}-\underline{q}+1} \end{pmatrix}\right)^{\odot 2}.$$

Here, the notation $A^{\odot 2}$ represents componentwise squaring. The limiting risk metric (as $T\to+\infty$) in (4.8) immediately follows from the fact that the nonzero eignevalues are negative.

Part 4: Finally, for $\mathcal{R}^3(t,q)$ note that $\mathbb{E}_t[\int_t^T \mathbb{I}\{q_s=n\}\,ds]$ is given by Proposition A.1 with $\boldsymbol{\ell}(q)=\mathbb{I}\{q=n\}$ and $\varphi(q)=0$. Consequently, we have,

(A.13)
$$\mathbb{Q}_n(t,q) = \frac{1}{T-t}\,\mathbf{V}\begin{pmatrix} (T-t)(\mathbf{V}^{-1}\mathbf{1}_n)_1 \\ \dfrac{e^{\mathbf{d}_2(T-t)}-1}{\mathbf{d}_2}(\mathbf{V}^{-1}\mathbf{1}_n)_2 \\ \vdots \\ \dfrac{e^{\mathbf{d}_{\overline{q}-\underline{q}+1}(T-t)}-1}{\mathbf{d}_{\overline{q}-\underline{q}+1}}(\mathbf{V}^{-1}\mathbf{1}_n)_{\overline{q}-\underline{q}+1} \end{pmatrix}.$$

Once again, the limiting case (as $T\to+\infty$) in (4.9) follows from the fact that the nonzero eignevalues are negative. □

PROPOSITION A.1 *Let $(\delta_t^{\pm *})_{0\le t\le T}$ denote the limiting optimal spreads*[12] *given by Proposition 2.6 and $(q_t^*)_{0\le t\le T}$ the corresponding inventory. Define, for arbitrary bounded functions*

[11] Although only these terms contribute for large τ, it is necessary to construct the full solution because the boundary conditions depend on the exponential terms and feed into the constant term of the polynomial.

[12] Recall these are independent of time and depend only on q.

$\varphi(q)$ and $\ell(q)$ for $q \in \{\underline{q}, \dots, \overline{q}\}$,

$$g_q(t) \triangleq \mathbb{E}\left[\varphi\left(q_T^*\right) + \int_t^T \ell\left(q_s^*\right)\, ds\right], \tag{A.14}$$

and let $\mathbf{g}$ *denote the collection of these functions, i.e.,* $\mathbf{g} = [g_{\overline{q}}(t),\ g_{\overline{q}-1}(t), \dots, g_{\underline{q}}(t)]'$. *Then,* $\mathbf{g} = \mathbf{V}\widetilde{\mathbf{g}}$ *where*

$$\widetilde{\mathbf{g}}_k(t) = \begin{cases} (\mathbf{V}^{-1}\boldsymbol{\varphi})_1 + (T-t)(\mathbf{V}^{-1}\boldsymbol{\ell})_1, & k = 1\,, \\ (\mathbf{V}^{-1}\boldsymbol{\varphi})_k\, e^{\mathbf{d}_k(T-t)} + \dfrac{e^{\mathbf{d}_k(T-t)} - 1}{\mathbf{d}_k}(\mathbf{V}^{-1}\boldsymbol{\ell})_k, & \text{otherwise.} \end{cases} \tag{A.15}$$

Proof. Clearly, $g_q(t) + \int_0^t \ell(q_s)\, ds$ is a martingale, and therefore g satisfies the coupled system of ODEs

$$\begin{cases} \partial_t g_q(t) + \lambda^+ e^{-\kappa^+\delta^{*+}(q)}(g_{q-1}(t) - g_q(t)) + \lambda^- e^{-\kappa^-\delta^{*-}(q)}(g_{q+1}(t) - g_q(t)) + \ell_q = 0, \\ g_q(T) = \varphi(q)\,, \end{cases} \tag{A.16}$$

which can compactly be written as $\partial_t\widetilde{\mathbf{g}}(t) + \mathbf{B}\widetilde{\mathbf{g}}(t) + \boldsymbol{\ell} = 0$ subject to $\widetilde{\mathbf{g}}_q(T) = \varphi(q)$. Diagonalizing, and integrating, treating the zero eigenvalue separately, leads to the stated result. □

APPENDIX B: PROOF OF THEOREM 2.4

Here we prove the verification theorem for the value function in (2.7). Let $(\delta_s^{\pm})_{t\le s\le T} \in \mathcal{A}$ denote an arbitrary admissible control, $\delta_s = (\delta_s^+, \delta_s^-)$ denote the pair of controls, $(X_s^{x,\delta})_{t\le s\le T}$ (solving (2.2)) and $(N_s^{\pm,\delta})_{t\le s\le T}$ denote the trader's wealth and counting process for filled limit sell/buy orders, respectively, obtained by following the control with $X_{t-} = x$. Proceeding in the usual manner, write the tentative value function as $\widehat{H}(t, x, S, q) = x + q(S - \alpha q) + h(t, q)$ (here, we write $h(t, q) = h_q(t)$ to ease readability). Next, observe that

$$\begin{aligned} \int_t^T dX_s + \int_t^T S_{s-}\, dq_s &= \int_t^T \int_{\mathbb{R}_+\times[\delta_s^+,+\infty)} \delta_s^+\, \mu^+(d\mathbf{y}, ds) \\ &\quad + \int_t^T \int_{\mathbb{R}_+\times[\delta_s^-,+\infty)} \delta_s^-\, \mu^-(d\mathbf{y}, ds), \\ [S, q]_{T-} - [S, q]_t &= \int_t^T \int_{\mathbb{R}_+\times[\delta_s^+,+\infty)} y_1\, \mu^+(d\mathbf{y}, ds) \\ &\quad - \int_t^T \int_{\mathbb{R}_+\times[\delta_s^-,+\infty)} y_1\, \mu^-(d\mathbf{y}, ds), \quad \text{and} \\ q_{T-}^2 - q_t^2 &= \int_t^T \int_{\mathbb{R}_+\times[\delta_s^+,+\infty)} (1 - 2q_{s-})\, \mu^+(d\mathbf{y}, ds) \\ &\quad + \int_t^T \int_{\mathbb{R}_+\times[\delta_s^-,+\infty)} (1 + 2q_{s-})\, \mu^-(d\mathbf{y}, ds)\,. \end{aligned} \tag{B.1}$$

Note that only $\delta_s^{\pm}$ appear in the above and not $\delta_{s-}^{\pm}$ since the class of admissible strategies are taken to be $\mathcal{F}$-predictable. Also, recall that $\delta_s^{\pm}$ are bounded from below and $\delta_s^{+} = +\infty$ ($\delta_s^{-} = +\infty$) whenever $q_s = \underline{q}$ ($q_s = \overline{q}$). Using the above together with Ito's lemma for Poisson random measures and the fact that $h(t, q)$ given by Proposition 2.2 is differentiable in t, we then have

(B.2)
$$\begin{aligned}
\widehat{H}(T, X_{T-}^{x,\delta}, S_{T-}, q_{T-}^{\delta}) &= \widehat{H}(t, x, S, q) + \int_t^T \partial_t h(s, q_s^{q,\delta})\, ds \\
&+ \sigma \int_t^T q_s\, dW_s + \int_t^T \int_{\mathbb{R}^2} q_{s-}\, y_1 \left(\mu^+(d\mathbf{y}, ds) - \mu^-(d\mathbf{y}, ds)\right) \\
&+ \int_t^T \int_{\mathbb{R}\times[\delta_s^+, +\infty)} \left[\delta_s^+ + y_1 - \alpha(1 - 2q_{s-}) + h\big(s, q_{s-}^{q,\delta} - 1\big) - h\big(s, q_{s-}^{q,\delta}\big)\right] \mu^+(d\mathbf{y}, ds) \\
&+ \int_t^T \int_{\mathbb{R}\times[\delta_s^-, +\infty)} \left[\delta_s^- - y_1 - \alpha(1 + 2q_{s-}) + h\big(s, q_{s-}^{q,\delta} + 1\big) - h\big(s, q_{s-}^{q,\delta}\big)\right] \mu^-(d\mathbf{y}, ds)\,.
\end{aligned}$$

Next, note that, since the processes $\delta^{\pm}$ are bounded below (say $-\Delta^{\pm} \le \delta_s^{\pm}\ \forall s \in [0, T]$), then

(B.3)
$$\begin{aligned}
\mathbb{E}\left[\int_0^T \int_{\mathbb{R}\times[\delta_s^{\pm}, +\infty)} \left|\delta_s^{\pm}\right|^2 \nu^{\pm}(d\mathbf{y}, ds)\right] &\le (\Delta^{\pm})^2 \mathbb{E}\left[\int_0^T \int_{\mathbb{R}\times[-\Delta^{\pm}, +\infty)} \mathbb{I}_{\{\delta_s^{\pm} < 0\}} \nu^{\pm}(d\mathbf{y}, ds)\right] \\
&\quad + \mathbb{E}\left[\int_0^T \int_{\mathbb{R}\times[\delta_s^{\pm}, +\infty)} (\delta_s^{\pm})^2 \mathbb{I}_{\{\delta_s^{\pm} \ge 0\}}\, \nu^{\pm}(d\mathbf{y}, ds)\right] \\
&\le \left(\lambda(\Delta^{\pm})^2 e^{\kappa^{\pm}\Delta^{\pm}} + \frac{4e^{-2}\lambda}{(\kappa^{\pm})^2}\right) T < +\infty\,,
\end{aligned}$$

Furthermore, the process q is bounded in the interval $[\underline{q}, \overline{q}]$ and finally, from Proposition 2.2, $|h(s, q_s^{q,\delta})| < C$ for some $C < +\infty$, therefore from (B.3) we have

(B.4)
$$\begin{aligned}
&\mathbb{E}_{t,x,q,S}\left[\widehat{H}\big(T, X_{T-}^{x,\delta}, S_{T-}, q_{T-}^{q,\delta}\big)\right] \\
&\quad = \widehat{H}(t, x, S, q) + \mathbb{E}_{t,x,q,S}\left[\int_t^T \left\{\partial_t h\big(s, q_s^{q,\delta}\big) + q_{s-}(\varepsilon^+ - \varepsilon^-)\right.\right. \\
&\qquad + \lambda^+(\delta_s^+ + \varepsilon^+ - \alpha(1 - 2q_{s-}) + h\big(s, q_{s-}^{q,\delta} - 1\big) - h\big(s, q_{s-}^{q,\delta}\big) \\
&\qquad \left.\left. + \lambda^-(\delta_s^- + \varepsilon^+ - \alpha(1 + 2q_{s-}) + h\big(s, q_{s-}^{q,\delta} + 1\big) - h\big(s, q_{s-}^{q,\delta}\big)\right\} ds\right] \\
&\quad \le \widehat{H}(t, x, S, q) + \phi\, \sigma^2\, \mathbb{E}_{t,x,q,S}\left[\int_t^T \big(q_s^{q,\delta}\big)^2\, ds\right], \qquad \text{by (2.7).}
\end{aligned}$$

The above inequality holds for any admissible strategy $\delta \in \mathcal{A}$ and is binding for the optimal strategy δ^*. Furthermore,

$$\begin{aligned}
\mathbb{E}_{t,x,q,S}\big[\widehat{H}\big(T, X_{T-}^{x,\delta}, S_{T-}, q_{T-}^{q,\delta}\big)\big] &= \mathbb{E}_{t,x,q,S}\big[\widehat{H}\big(T, X_T^{x,\delta}, S_T, q_T^{q,\delta}\big)\big] \\
&= \mathbb{E}_{t,x,q,S}\big[X_T^{x,\delta} + q_T^{q,\delta}\big(S_T - \alpha\, q_T^{q,\delta}\big)\big].
\end{aligned}$$

Hence we have,

(B.5)
$$\sup_{\delta\in\mathcal{A}} \mathbb{E}_{t,x,q,S}\left[X_T^{x,\delta} + q_T^{q,\delta}\left(S_T - \alpha\, q_T^{q,\delta}\right) - \phi\,\sigma^2 \int_t^T \left(q_s^{q,\delta}\right)^2\, ds \right]$$
$$\leq \widehat{H}(t,x,S,q) = \mathbb{E}_{t,x,q,S}\left[X_T^{x,\delta^*} + q_T^{q,\delta^*}\left(S_T - \alpha\, q_T^{q,\delta^*}\right) - \phi\,\sigma^2 \int_t^T \left(q_s^{q,\delta^*}\right)^2\, ds \right]$$

which proves that $\widehat{H}(t, x, S, q)$ is indeed the value function of the original control problem (2.8). Moreover, the feedback controls given by (2.8) are clearly admissible. □

REFERENCES

Almgren, R. (2003): Optimal Execution with Nonlinear Impact Functions and Trading-Enhanced Risk, *Appl. Math. Finance* 10(1), 1–18.

Avellaneda, M., and S. Stoikov (2008): High-Frequency Trading in a Limit Order Book, *Quant. Finance* 8, 217–224.

Bayraktar, E., and M. Ludkovski (2012): Liquidation in Limit Order Books with Controlled Intensity, *Math. Finance* 21(4), 681–701.

Bertsimas, D., and A. W. Lo (1998): Optimal Control of Execution Costs, *J. Financ. Markets* 1(1), 1–50.

Bouchard, B., N.-M. Dang, and C.-A. Lehalle (2011): Optimal Control of Trading Algorithms: A General Impulse Control Approach, *SIAM J. Financ. Math.* 2, 404–438.

Brunetti, C., A. Kirilenko, and S. Mankad (2011): Identifying High-Frequency Traders in Electronic Markets: Properties and Forecasting, unpublished manuscript.

Cartea, Á. (2013): Derivatives Pricing with Marked Point Processes Using Tick-by-Tick Data, *Quant. Finance* 13(1), 111–123.

Cartea, Á., and S. Jaimungal (2013): Modeling Asset Prices for Algorithmic and High Frequency Trading, *Appl. Math. Finance* 20(6), 512–547.

Cartea, Á., S. Jaimungal, and J. Ricci (2014): Buy Low Sell High: A High Frequency Trading Perspective, *SIAM J. Financ. Math.* 5(1), 415–444.

Cartea, Á., and T. Meyer-Brandis (2010): How Duration between Trades of Underlying Securities Affects Option Prices, *Rev. Finance* 14(4), 749–785.

Cartea, Á., and J. Penalva (2012): Where Is the Value in High Frequency Trading? *Q. J. Finance* 2(3), 1250014.

CFTC and SEC (2010): Findings Regarding the Market Events of May 6, 2010. Report, SEC.

Cont, R., R. Talreja, and S. Stoikov (2010): A Stochastic Model for Order Book Dynamics, *Oper. Res.* 58, 217–224.

Cvitanic, J., and A. A. Kirilenko (2010): High Frequency Traders and Asset Prices, SSRN eLibrary, http://ssrn.com/paper=1569067.

Fleming, W. H., and H. M. Soner (2006): *Controlled Markov Processes and Viscosity Solutions*, 2nd edn. New York: Springer.

Guéant, O., C.-A. Lehalle, and J. Fernandez Tapia (2011): Dealing with the Inventory Risk, Quant. Finance Papers 1105.3115, arXiv.org.

Guilbaud, F., and H. Pham (2013): Optimal High Frequency Trading with Limit and Market Orders, *Quant. Finance* 13(1), 79–94.

HO, T., and H. R. STOLL (1981): Optimal Dealer Pricing under Transactions and Return Uncertainty, *J. Financ. Econ.* 9, 47–73.

KHARROUBI, I., and H. PHAM (2010): Optimal Portfolio Liquidation with Execution Cost and Risk, *SIAM J. Financ. Math.* 1, 897–931.

KIRILENKO, A. A., A. P. S. KYLE, M. SAMADI, and T. TUZUN (2010): The Flash Crash: The Impact of High Frequency Trading on an Electronic Market, SSRN eLibrary, http://ssrn.com/paper=1686004.

KÜHN, C., and M. STROH (2012): Continuous Time Trading of a Small Investor in a Limit Order Market, Preprint.

MATHEMATICAL FINANCE
An International Journal of Mathematics, Statistics and Financial Economics

Mathematical Finance, Vol. 25, No. 3 (July 2015), 612–639

OPTIMAL EXECUTION OF A VWAP ORDER: A STOCHASTIC CONTROL APPROACH

CHRISTOPH FREI

Department of Mathematical and Statistical Sciences, University of Alberta

NICHOLAS WESTRAY

Deutsche Bank AG, London

We consider the optimal liquidation of a position of stock (long or short) where trading has a temporary market impact on the price. The aim is to minimize both the mean and variance of the order slippage with respect to a benchmark given by the market volume-weighted average price (VWAP). In this setting, we introduce a new model for the relative volume curve which allows simultaneously for accurate data fit, economic justification, and mathematical tractability. Tackling the resulting optimization problem using a stochastic control approach, we derive and solve the corresponding Hamilton–Jacobi–Bellman equation to give an explicit characterization of the optimal trading rate and liquidation trajectory.

KEY WORDS: optimal trade execution, VWAP, HJB equation, gamma bridge.

1. INTRODUCTION

In investment banks today algorithmic trading is rapidly becoming the preferred method for clients to acquire and liquidate positions of stock. Typically, a computer-based algorithm is used to buy (or sell) a position while attempting to stick to a client-selected benchmark. One of the oldest and most popular of these algorithms is volume-weighted average price (VWAP). The popularity of the VWAP benchmark for both brokers and clients stems from several reasons. First, it is very simple to calculate, facilitating easy posttrade reporting. Second, it encourages the splitting of larger orders into smaller orders, reducing demand for large liquidity and hence market impact/volatility. Finally, given a time interval, it is considered a "fair" benchmark price; in the language of Berkowitz, Logue, and Noser (1988), VWAP is a price which

> " . . . is an unbiased estimate of prices that could be achieved by any randomly selected nonstrategic trader."

Beating market VWAP would thus be considered as a "good" execution, see Madhavan (2002) for further detailed discussion.

We thank three anonymous referees for very useful comments, which enabled us to significantly improve the paper. We thank Mark DiBattista and Christian Hesse for helpful comments and suggestions. Christoph Frei gratefully acknowledges financial support by the Natural Sciences and Engineering Research Council of Canada. The opinions and ideas expressed in this article are those of the authors alone, and do not necessarily reflect the views of Deutsche Bank AG, its subsidiaries, or affiliates.

Manuscript received August 2012; final revision received April 2013.

Address correspondence to Christoph Frei, Department of Mathematical and Statistical Sciences, University of Alberta, Edmonton, AB T6G 2G1, Canada; e-mail: cfrei@ualberta.ca.

DOI: 10.1111/mafi.12048

This paper proposes a stochastic control approach to tackle the question of how a broker should optimally schedule a VWAP-benchmarked trade. In reality, the client specifies either a buy (or sell) quantity as well as a start and end time and the broker must then acquire (liquidate) the position attempting to minimize the mean and variance of the difference of the volume-weighted price achieved with the market VWAP over the order lifetime (slippage). The main motivation for the present research is attempting to improve execution efficiency. Due to the huge notional volumes being traded algorithmically, small gains arising from the application of results obtained can lead to substantial increase in profits for both brokers and their clients.

The question of optimal execution with an arrival price benchmark is well studied in the literature going back to Bertsimas and Lo (1998), see also the seminal papers of Almgren and co-authors (Almgren and Chriss 2000; Almgren et al. 2005; Almgren 2012). In contrast, perhaps due to the stochastic nature of the benchmark, there is significantly less literature related to the present problem. The first work in this area is the article by Konishi (2002) who derives the optimal execution trajectory for single and basket VWAP executions when the price is given by a Brownian motion. The strategy is then assessed against actual trade data from the Tokyo Stock Exchange. Following this article, VWAP tracking has been attacked using a variety of different methods. McCulloch and Kazakov (2012) view it as a quadratic hedging problem under partial information, whereas Kakade et al. (2004) and Białkowski et al. (2008) use online learning and dynamic volume approaches and Humphery-Jenner (2011) gives a VWAP trading rule, which takes intraday noise into consideration. Finally, Bouchard and Dang (2013) formulate it as a stochastic target problem and derive a viscosity solution characterization of the value function. Note that the above articles (excluding Bouchard and Dang 2013) do not take into account the market impact of a trade and none of them impose any parametric structure on the intraday volume curve.

This paper has several contributions. First, we extend previous literature in this area by allowing for a linear temporary market impact model. More general models of price impact have been studied both theoretically, e.g., by Gatheral (2010), and empirically, e.g., by Bouchaud et al. (2004) (see Gatheral and Schied 2013 for a good overview); however, the linear model leads to a tractable problem. Second, we provide a parametric model for relative volume, which fits real data well, reflecting meaningful underlying economic assumptions and simultaneously being tractable enough to perform optimization. Finally, although the optimization problem is involved due to the use of VWAP as benchmark, we are able to explicitly characterize the optimal control thus providing a closed-form solution for the optimal trading rate. This final result opens up, for the first time, a rigorous mathematical approach to the determination of commission for guaranteed VWAP trades, similar to that done for Implementation Shortfall in Almgren and Chriss (2000).

The paper is organized as follows. In the next section, we introduce and justify our model, which uses VWAP as a benchmark in optimal trade execution. We present the main result in Section 3, Theorem 3.1, which provides the explicit solution for the optimal trading rate. Deferring the proof to Section 5, we explain in Section 3 two crucial properties of the optimal trading rate. Its sign can switch only once from negative to positive and never the other way, and the optimal trading rate can be decomposed into two parts, with one being a deterministic TWAP (time-weighted average price) strategy and the other reflecting the adjustment necessary due to jumps in the relative volume

curve. In Section 4, we show and discuss how well the parametric model fits to trading volume.

2. A FRAMEWORK FOR USING A VWAP BENCHMARK

Here, we describe the model formulation and the key assumptions. We begin with our trading strategies, without loss of generality we consider a buy program for Y shares. The situation of a sell can be considered by reversing the time.

2.1. Trading Costs

We are given a start and end time by the client, which we assume (without loss of generality) to be given by $t_0 = 0$ and T, respectively. We must complete the purchase of stock by T and will be benchmarked to the market VWAP over the period $[0, T]$, for simplicity, the reader may think of $T = 1$, corresponding to a day VWAP order. As is standard in the literature, we work on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0\le t\le T}, P)$ satisfying the usual hypotheses of right continuity and completeness, and we denote by $X^u(t)$ our share holdings at t where the trading strategy u is an adapted and integrable process typically referred to as trading rate. In particular, our holdings evolve according to

$$\mathrm{d}X^u(t) = u(t)\,\mathrm{d}t, \quad X^u(0) = 0, \quad X^u(T) = Y.$$

The asset price process $(P(t))_{0\le t\le T}$ is assumed to be an arithmetic Brownian motion

$$P(t) = P(0) + \sigma W(t),$$

where $\sigma > 0$ represents the daily volatility in dollars and $(W(t))_{0\le t\le T}$ is a standard Brownian motion on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0\le t\le T}, P)$. Using an arithmetic as opposed to a geometric Brownian motion in the present setup is possible since we are considering an intraday trading horizon, so that $P(t)$ is negative with an extremely small probability and there are negligible differences between the two models, see Gatheral and Schied (2011).

As is well known, when we trade, we do not realize the above price $P(t)$ (which could be thought of as the mid quote) but instead we pay

$$P^u(t) = P(t) + \kappa u(t),$$

where κ is the coefficient of the linear (temporary) market impact model. It represents an instantaneous linear premium on the price due to how fast we trade. This model (among many more complicated ones) is studied in detail in the literature, see Gatheral (2010) or Bouchaud et al. (2004). The linear form is necessary as it leads to quadratic trading costs, which then produces a tractable problem. However, it is also important to note that given the extremely low predictive accuracy of market impact models (typically $<$ 5% R^2), the cost of increased complexity arising from moving away from a linear model would outweigh any gains from better describing market impact. Moreover, a linear price impact is supported in the recent empirical study by Cont, Kukanov, and Stoikov (2014). The above will introduce some dependence on impact into the solution, which is its purpose.

Similarly to section 1.1 of Almgren (2012), our total expenditure TE^u to buy the shares Y using a control u is thus given by

$$\mathrm{TE}^u = \int_0^T P^u(t)\,\mathrm{d}X^u(t) = YP(0) + \sigma\int_0^T W(t)u(t)\,\mathrm{d}t + \kappa\int_0^T u^2(t)\,\mathrm{d}t,$$

where we applied the relation $\int_0^T u(t)\,\mathrm{d}t = X^u(T) - X^u(0) = Y$. Using $X^u(T) = Y$ and that X^u is of finite variation, the product rule yields

$$\int_0^T W(t)u(t)\,\mathrm{d}t = \int_0^T W(t)\,\mathrm{d}X^u(t) = -\int_0^T X^u(t)\,\mathrm{d}W(t) + YW(T),$$

so that

$$\mathrm{TE}^u = YP(0) - \sigma\int_0^T X^u(t)\,\mathrm{d}W(t) + \sigma YW(T) + \kappa\int_0^T u^2(t)\,\mathrm{d}t. \tag{2.1}$$

2.2. The VWAP Benchmark

Given a series of prices $(P_i)_{i=1,\dots,N}$ together with volumes $(V_i)_{i=1,\dots,N}$ executed at those prices, the VWAP is defined to be

$$\mathrm{VWAP} = \frac{\sum_{i=1}^N V_i P_i}{\sum_{i=1}^N V_i}.$$

If we define $\tilde{V}_i = \sum_{j=1}^i V_j$ as the cumulative volume (with $\tilde{V}_0 = 0$), we have

$$\mathrm{VWAP} = \frac{\sum_{i=1}^N (\tilde{V}_i - \tilde{V}_{i-1})P_i}{\tilde{V}_N} \approx \int_0^T P(t)\frac{\mathrm{d}\tilde{V}(t)}{\tilde{V}(T)}.$$

In particular, to model VWAP, we need a continuous-time process for $\frac{\tilde{V}(t)}{\tilde{V}(T)}$. This will be nondecreasing and satisfy $\frac{\tilde{V}(T)}{\tilde{V}(T)} = 1$ as well as $\frac{\tilde{V}(0)}{\tilde{V}(T)} = 0$. A natural process to use is the gamma bridge.

DEFINITION 2.1. (1) A gamma process $(L(t))_{0\le t\le T}$ is a process with independent and identically distributed increments such that $L(0) = 0$ and $L(t)$ is gamma distributed with mean $mt\theta$ and variance $mt\theta^2$ for some $m > 0, \theta > 0$. (2) For a gamma process $(L(t))_{0\le t\le T}$, the gamma bridge $(\gamma(t))_{0\le t\le T}$ is defined by $\gamma(t) = L(t)\big/L(T)$.

There are several reasons for our choice of modeling the intraday relative volume curve by a gamma bridge. First, we will see in Section 4 that our model fits well to real stock data provided the stock is sufficiently liquidly traded as well as being finite variation, like real data. Second, we can think of the cumulative trading volume as analogous to the accumulation of dam rain, similarly to Gani (1957) considering the arrival of insurance claims as analogous to the accumulation of dam rain. The latter can be modeled by a gamma process as pointed out by Moran (1956) so that similarly the relative amount will be a gamma bridge. Finally, we can prove that the intraday volume curve must be a gamma bridge if we assume that trading volume is independent and stationarily distributed through the day and the relative intraday volume is independent of the total volume. This link is based on the following theoretical result on gamma processes.

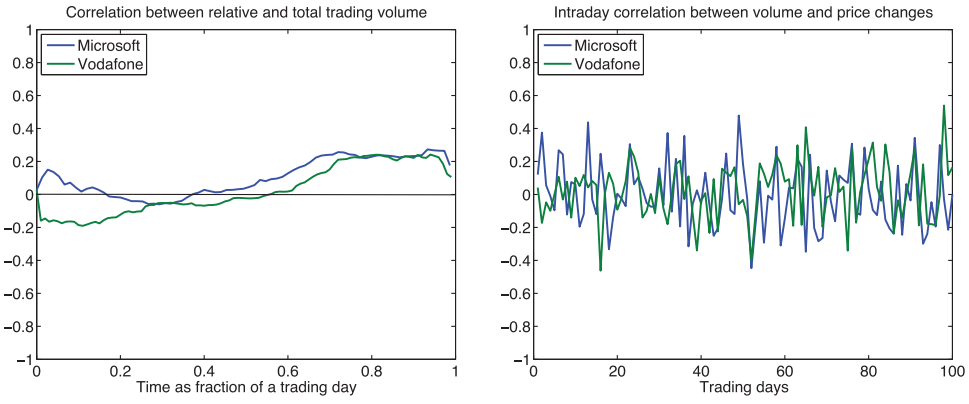

FIGURE 2.1. Correlation between relative and total trading volume (left panel) and intraday correlation between volume and price changes (right panel) for the stocks of Vodafone and Microsoft during the first 100 trading days of 2012, using 5-minute intraday data.

PROPOSITION 2.2. *Let* $(L(t))_{0\le t\le T}$ *be a Lévy process with* $L(T) > 0$ *a.s. and nondeterministic (i.e.,* $P[L(T) = c] < 1$ *for all* c*). Then the following are equivalent:*

(i) L *is a gamma process;*
(ii) there exists $t \in (0, T)$ *such that* $L(T)$ *and* $L(t)/L(T)$ *are independent;*
(iii) for all $t \in [0, T]$, $L(T)$ *and* $L(t)/L(T)$ *are independent.*

In particular, the proposition shows that the gamma process is the only positive Lévy process whose intermediate relative values are independent of the terminal value. We immediately get from Proposition 2.2 the following application to the relative volume curve.

COROLLARY 2.3. *Assume that the cumulative trading volume has independent and stationary increments, its terminal value is nondeterministic and strictly positive (i.e., trading happens a.s.), and the relative volume curve is independent of the total trading volume. Then the cumulative volume is a gamma process and the relative volume curve is a gamma bridge.*

The assumption that the relative volume curve is independent of the total volume is not too unrealistic; it is hard to imagine that the relative volume curve on a given day depends significantly on the total traded volume on that day. This is supported by the left panel of Figure 2.1 for a major US and European stock. A consequence of modeling volume with a gamma process and prices with a Brownian motion is their mutual independence; see lemma 15.6 of Kallenberg (2002). The right panel of Figure 2.1 shows that the intraday correlation between volume and price changes varies a lot but is small on average. Hence, it is difficult to incorporate in a model, leading to only a small potential increase in performance and thus justifying the independence assumption.

Looking at the historical volume data and noting its episodic nature, the hypothesis that volume is independent and stationarily distributed throughout the day is idealistic. However, we can think of the gamma bridge as a zero-order approximation to the true relative volume curve, which would be exact should such an assumption hold. This then provides a tractable basis upon which one can incorporate more realistic features. As a

final point, let us note that the fit to actual data is sufficiently good to conclude that the effect of deviation from such a hypothesis is (perhaps surprisingly) not too severe.

Proof of Proposition 2.2. The implication "$(i) \Longrightarrow (iii)$" is shown in proposition 3.2 of Brody, Hughston, and Macrina (2008), and (iii) clearly implies (ii) so that it remains to show "$(ii) \Longrightarrow (i)$." Let $t \in (0, T)$ be such that $L(T)$ and $L(t)/L(T)$ are independent. Since L is a Lévy process, it is enough to show that $L(t)$ is gamma distributed. If we set $A = L(t)$ and $B = L(T) - L(t)$, we have that A, B are strictly positive, nondeterministic, and independent. By assumption, $A + B$, $A/(A + B)$ are independent, hence so are $A + B$, A/B by using the measurable mapping $x \mapsto x/(1 - x)$. It follows from Lukacs' proportion-sum independence theorem (Lukacs 1955) that both A and B have gamma distributions. □

Since we will only consider a gamma bridge but not its underlying gamma process, by scaling we can (and henceforth do) set $\theta = 1$ without loss of generality. Moreover, in the later simulations, we fix $T = 1$ and use m as the model parameter. Under the assumption of gamma bridge for the relative volume curve, the expenditure (in dollars) to buy Y shares of market VWAP is given by

$$\begin{aligned}\text{VWAP} &= \int_0^T P(t)\,\mathrm{d}(Y\gamma(t)) = YP(0) + \sigma Y \int_0^T W(t)\,\mathrm{d}\gamma(t) \\ &= YP(0) - \sigma Y \int_0^T \gamma(t-)\,\mathrm{d}W(t) + \sigma YW(T),\end{aligned}$$

where we argued similarly to the derivation of (2.1). The VWAP benchmark should be thought of as an average market price over the lifetime of the order and it ignores the market impact because the mid quote represents a reasonable estimate of the current fair market price. Observing that we can scale by $1/Y$, we thus can (and do) assume that the client wants to buy 1 share.

One may object to the model formulation in that $L(T)$ is unknown. While this is a valid objection, it misunderstands the objective of this paper. All major brokers who engage in algorithmic trading require models for the intraday relative volume curve. The primary application for such models is in the execution of VWAP and PVol (percentage of volume) orders. When evaluating the viability of such a model in a trading application, one can compare the performance when using the new model against that in the "perfect information" case, modeled here by a gamma bridge. The framework presented here allows closed-form solutions for this case and thus allows new volume curve models to be assessed in a way relevant to their use rather than just a pure goodness-of-fit test. A second benefit is in posttrade analysis, the solution presented here provides a broker-independent benchmark for the execution of VWAP trades allowing the scope for relative comparison as well as the common absolute performance. We will address the issues of comparing and evaluating different VWAP strategies in future work. In summary, the focus here is not a real-time (i.e., based on an adapted estimator of relative volume) strategy for trading VWAP but rather an effort to set up a framework in which one can get closed-form and implementable solutions to the VWAP trading problem, which are of use for setting upper bounds on performance and providing broker-independent benchmarks. As a consequence, the filtration $(\mathcal{F}_t)_{0\le t\le T}$ we are working with is such that the gamma bridge $(\gamma(t))_{0\le t\le T}$ and the Brownian motion $(W(t))_{0\le t\le T}$ are adapted to it.

The reader will note that our trading is not taken into account when calculating the relative volume curve, so that it becomes exogenous. This is clearly a simplification and

is effectively equivalent to assuming that the (unnormalized) size Y is small relative to the total volume traded, i.e., that we are dealing with small orders. Since typical algorithmically traded VWAP orders are traded with an effective (when scaled to a day) size of 10–15% of the daily volume, this assumption is not too restrictive and hugely simplifies the problem under consideration. As explained in Hu (2007), who analyzes the difference in VWAP with and without own trading, it might even be desirable to exclude our own trading from the market VWAP calculation in order to compare our performance with a benchmark, which is not affected by us. The interesting case of large VWAP trades (>50% daily volume) will be considered elsewhere.

REMARK 2.4. It is important to be precise about what is meant here by "volume." In recent years the number of trading facilities with visible order books (lit venues) has increased substantially. For example, in 2008 one could only trade Vodafone Group PLC (VOD) on the London Stock Exchange (its primary exchange) whereas today one can trade this on BATS, CHI-X, and Turquoise (among others). Since most modern smart order routers are able to access these MTFs (multilateral trading facilities), it makes sense to consider everything in a consolidated view, so that (for example) day volume in this paper should be understood as the total volume traded in a day on all (lit) venues.

2.3. Algorithm Performance

Now that the quantity of shares is normalized to $Y = 1$, our costs can be thought of as per share in dollars. Our first aim is to minimize the expected slippage, defined as

$$\text{slip}^u = \text{TE}^u - \text{VWAP} = \sigma \int_0^T (\gamma(t-) - X^u(t))\,\mathrm{d}W(t) + \kappa \int_0^T u^2(t)\,\mathrm{d}t.$$

From a broker's perspective, simply having small expected slippage, while good, is not the sole goal. From a posttrade perspective, it will be necessary to explain to clients why slippage was significantly far from that expected, should that occur. Our second aim is, therefore, to attempt to minimize the variance of the slippage. We use the approximation

$$\text{Var}[\text{slip}^u] \approx E\left[\left(\sigma \int_0^T (\gamma(t-) - X^u(t))\,\mathrm{d}W(t)\right)^2\right] = \sigma^2 E\left[\int_0^T (\gamma(t) - X^u(t))^2\,\mathrm{d}t\right],$$

where we assumed for the last equation that $\int X^u\,\mathrm{d}W$ is a square integrable martingale and used that γ is bounded and has only countably many jumps.

The main reason for this approximation is tractability, without this the difficulty would be increased significantly. The mathematical reason for the difficulty of the original problem is its time inconsistency. As explained in section 1.2 of Björk and Murgoci (2010), time inconsistence of mean–variance problems of the form $E[X] + \lambda\text{Var}[X]$ is caused by the term $(E[X])^2$ in $\text{Var}[X] = E[X^2] - (E[X])^2$. While standard time-consistent problems are allowed to have expected values of nonlinear functions, such as $E[X^2]$, the term $(E[X])^2$ is a nonlinear function of the expected value and not an expected value of a nonlinear function. However, in our particular problem, the time inconsistency is mild in the sense that the value of $\text{Var}[\text{slip}^u]$ is close to the variance of $\int_0^T (\gamma(t) - X^u(t))\,\mathrm{d}W(t)$, which has zero mean, hence leading to a time-consistent formulation. Therefore, such an approximation is appropriate in our situation.

Hence, we study the minimization of

$$\kappa\, E\bigg[\int_0^T u^2(t)\,\mathrm{d}t\bigg] + \lambda\,\sigma^2 E\bigg[\int_0^T (\gamma(t) - X^u(t))^2\,\mathrm{d}t\bigg],$$

as an approximation of the mean–variance problem

$$\inf_u (E[\mathrm{slip}^u] + \lambda \mathrm{Var}[\mathrm{slip}^u]),$$

for a given mean–variance trade-off parameter $\lambda > 0$. In further justification of this approximation, we note that similar to Almgren (2012) and Konishi (2002), it is the case that the main driver of $\mathrm{Var}[\mathrm{slip}^u]$ is typically the volume curve γ and not the trading rate u, so that we capture the dominant term.

To get an idea of the approximation error, we used a Monte Carlo simulation to compare $\mathrm{Var}[\mathrm{slip}^u]$ and $\sigma^2 E[\int_0^T (\gamma(t) - X^u(t))^2\,\mathrm{d}t]$. Our calculation showed that, when using the optimal strategy $\hat{u}$ for the approximated problem from Theorem 3.1, the relative error

$$\frac{\mathrm{Var}[\mathrm{slip}^{\hat{u}}] - \sigma^2 E[\int_0^T (\gamma(t) - X^{\hat{u}}(t))^2\,\mathrm{d}t]}{\mathrm{Var}[\mathrm{slip}^{\hat{u}}]},$$

was always very small. For example, choosing $\sigma = 0.01$, $\kappa = 10^{-8}$ (see Remark 2.5), $\lambda = 1$, $m = 25$, and $T = 1$, the relative error was less than $10^{-3} = 0.1\,\%$.

REMARK 2.5. It is a simple exercise to calibrate a market impact model and determine κ, however, as a first-order approximation we simply use the trading rule of thumb, see Gatheral (2010) (among others), that trading one day's volume costs approximately one day's volatility in basis points. This is quite accurate in most cases and suffices for our purposes. In particular, this implies the following relation for our linear impact model:

$$\kappa \int_0^T u^2(t)\,\mathrm{d}t \approx \sigma\, V,$$

where V is the daily volume (shares) and σ is again the volatility in (\$). Assuming a linear execution for V (and noting $T = 1$), this reduces to $\kappa \approx \frac{\sigma}{V} \approx 10^{-8}$ (for a stock like Vodafone).

Dynamically formulated, the *value function* for the optimization problem is given by

$$v(t, x, \gamma) = \inf_u E\bigg[\kappa \int_t^T u^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_t^T (\gamma(s) - X^u(s))^2\,\mathrm{d}s\bigg], \tag{2.2}$$

where $(\gamma(s))_{t\le s\le T}$ is a gamma bridge with $\gamma(t) = \gamma$ and the infimum is over all adapted and integrable u with

$$\mathrm{d}X^u(s) = u(s)\,\mathrm{d}s, \quad X^u(t) = x, \quad X^u(T) = 1.$$

In conclusion, we are considering a mean–variance formulation of minimizing slippage from VWAP.

REMARK 2.6. The framework here is similar to the seminal paper of Almgren and Chriss (2000). The key extra technical difficulty in the present case comes from the switch

from an arrival price benchmark ($P(0)$) to a VWAP benchmark (VWAP). This introduces extra stochastic complexity coming from the gamma bridge.

Under a gamma bridge $(\gamma(s))_{t\le s\le T}$ with $\gamma(t)=\gamma$, we understand a process of the form $\gamma(s)=\gamma+\frac{L(s)-L(t)}{L(T)-L(t)}(1-\gamma)$ for a gamma process L so that starting with γ at t, we take the remaining part $1-\gamma$ proportional to the remaining relative portion of L. Note that $\frac{L(s)-L(t)}{L(T)-L(t)}$, $t\le s\le T$, is again a gamma bridge and independent of $\frac{L(t)}{L(T)}$ (Émery and Yor 2004, p. 673).

The reader will note that we have not considered auctions in the current model. Indeed, we have formulated the VWAP tracking problem for orders executed in the so-called "continuous" (nonauction) trading phase. For US stocks, it would be a reasonable approximation to simply ignore auctions, since the average volumes traded there are a small percentage of total volume. In Europe where auction volumes are typically much higher, one could imagine a small modification to the above framework where a dynamic model would be used to estimate the fractions executed in the open and close auction and the above model could then be used to execute the remainder. Since the focus of this paper is not the prediction of auction volumes, we assume that the historical mean values have been used and the total amount to be traded has been reduced accordingly in a preliminary step so that the results presented here apply equally to US and European stocks.

3. MAIN RESULT

Our main result is an explicit characterization of the value function v and the optimal control in (2.2).

THEOREM 3.1. *The value function v is given by*

$$v(t,x,\gamma)=a(t)x^2+b(t)\gamma x+c(t)x+d(t)\gamma^2+f(t)\gamma+g(t)$$

for $t\in[0,T)$, $x\ge 0$ and $\gamma\in[0,1]$ with the functions $a,b,c,d,f,g:[0,T)\to\mathbb{R}$ defined by

$$a(t)=\sqrt{\kappa\lambda\sigma^2}\,\frac{e^{2T\sqrt{\lambda\sigma^2/\kappa}}+e^{2t\sqrt{\lambda\sigma^2/\kappa}}}{e^{2T\sqrt{\lambda\sigma^2/\kappa}}-e^{2t\sqrt{\lambda\sigma^2/\kappa}}},\quad b(t)=-2a(t)+\frac{2\kappa}{T-t},\quad c(t)=-\frac{2\kappa}{T-t},$$

$$d(t)=\int_t^T\left(\lambda\sigma^2-\frac{1}{4\kappa}b^2(s)\right)\frac{T-s}{T-t}\,\frac{T-s+1/m}{T-t+1/m}\,\mathrm{d}s,$$

$$f(t)=\frac{1}{T-t}\int_t^T\left(b(s)+2d(s)\frac{T-s}{T-s+1/m}\right)\mathrm{d}s,$$

$$g(t)=a(t)-(T-t)\lambda\sigma^2-\int_t^T\left(\frac{b(s)}{2}\left(\frac{a(s)}{\kappa}(T-s)+1\right)-f(s)-\frac{d(s)/m}{T-s+1/m}\right)\frac{1}{T-s}\,\mathrm{d}s.$$

The optimal control $\hat{u}$ and the corresponding share holdings $\hat{X}$ are given by

$$\hat{u}(s)=-\frac{v_x}{2\kappa}(s,\hat{X}(s),\gamma(s))=-\frac{2a(s)\hat{X}(s)+b(s)\gamma(s)+c(s)}{2\kappa},\tag{3.1}$$

$$\hat{X}(s) = x\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} - \frac{1}{2\kappa}\int_t^s (b(z)\gamma(z) + c(z))\exp\left(-\frac{1}{\kappa}\int_z^s a(r)\,\mathrm{d}r\right)\mathrm{d}z.$$

We postpone the proof of Theorem 3.1 to Section 5 and instead give a nonrigorous motivation.

Suppose we replace (2.2) by the approximate problem,

$$v^n(t,x,\gamma) = \inf_u E\left[\kappa\int_t^T u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^T (\gamma(s) - X^u(s))^2\,\mathrm{d}s + n(X^u(T)-1)^2\right], \tag{3.2}$$

where the infimum is over all adapted and integrable u with

$$\mathrm{d}X^u(s) = u(s)\,\mathrm{d}s, \quad X^u(t) = x.$$

Observe now that X^u has no fixed terminal condition. Using that $\gamma(./m)$ is a gamma bridge on $[tm, Tm]$ with underlying mean growth rate equal to 1, it follows from corollary 1 of Émery and Yor (2004) that the infinitesimal generator of the gamma bridge γ is given by

$$m\int_0^1 (f(\gamma(t) + (1-\gamma(t))z) - f(\gamma(t)))(1-z)^{Tm-tm-1}\frac{1}{z}\,\mathrm{d}z,$$

where f is a function on $\mathbb{R}_+$ with bounded variation on compacts. If we assume that v^n is sufficiently regular, it should satisfy

$$\begin{aligned} &v_t^n + \lambda\sigma^2(\gamma - x)^2 + \inf_{u\in\mathbb{R}}(v_x^n u + \kappa u^2) \\ &\qquad + m\int_0^1 (v^n(t,x,\gamma + (1-\gamma)z) - v^n(t,x,\gamma))(1-z)^{Tm-tm-1}\frac{1}{z}\,\mathrm{d}z = 0, \\ &v^n(T,x,\gamma) = n(x-1)^2. \end{aligned}$$

Due to the quadratic structure, we propose the ansatz

$$v^n(t,x,\gamma) = a^n(t)x^2 + b^n(t)\gamma x + c^n(t)x + d^n(t)\gamma^2 + f^n(t)\gamma + g^n(t).$$

After some algebra, we derive that $a^n, \ldots, g^n$ should solve

$$\begin{aligned} a_t^n - \frac{1}{\kappa}(a^n)^2 + \lambda\sigma^2 = 0, &\qquad a^n(T) = n, \\ b_t^n - \frac{1}{\kappa}a^n b^n - 2\lambda\sigma^2 - b^n\varphi_0 = 0, &\qquad b^n(T) = 0, \\ c_t^n + b^n\varphi_0 - \frac{1}{\kappa}a^n c^n = 0, &\qquad c^n(T) = -2n, \\ d_t^n - 2d^n\varphi_0 + d^n\varphi_1 - \frac{1}{4\kappa}(b^n)^2 + \lambda\sigma^2 = 0, &\qquad d^n(T) = 0, \\ f_t^n - \frac{1}{2\kappa}b^n c^n + 2d^n\varphi_0 - f^n\varphi_0 - 2d^n\varphi_1 = 0, &\qquad f^n(T) = 0, \\ g_t^n - \frac{1}{4\kappa}(c^n)^2 + f^n\varphi_0 + d^n\varphi_1 = 0, &\qquad g^n(T) = n, \end{aligned}$$

where $\varphi_0(t) = \frac{1}{T-t}$ and $\varphi_1(t) = \frac{1}{T-t} - \frac{1}{T-t+1/m}$. It is intuitively sensible to expect that $v^n \uparrow v$ as well as that $a^n \to a$, $b^n \to b$, etc. However, if we let $n \uparrow \infty$ we would end up with difficult singular conditions in c and g. To avoid this, observe that the ratio $\frac{a^n(T)}{c^n(T)}$ and the difference $a^n(T) - g^n(T)$ are all constant and independent of n. Due to the convergence we would expect this to hold in the limit, this is precisely what the functions a, c, and g satisfy; see Lemma 5.1 below. A second justification is to see that as $n \uparrow \infty$, the terminal condition behaves like a quadratic function in x with one root at 1, which is infinite elsewhere. That is to say we have

$$\lim_{t \nearrow T}(a(t)x^2 + c(t)x + g(t)) = \lim_{t \nearrow T} v(t, x, 0) = \infty$$

for all $x \neq 1$. We thus expect that g behaves like a, and c behaves like $-2a$ for $t \nearrow T$, again consistent with Lemma 5.1.

REMARK 3.2. Note that our main result gives explicit formulae for both the optimal trading rate $\hat{u}$ and the optimal holdings $\hat{X}$. This desirable formula for $\hat{X}$ is due to our requirement that the holdings be absolutely continuous with respect to t. In contrast, this is not enforced in McCulloch and Kazakov (2012) so that one may not directly compare their results with ours.

Observe that the optimal trading rate depends on the volume curve but not on the price process. We give some intuition as to why this is a natural consequence of using a VWAP optimization criterion together with a Brownian price process. When comparing the VWAPs of two strategies, only price movements but not the absolute level of the price process are relevant. Since in our model the price movements are independent from past prices (they are given by Brownian increments), information about past prices is not included in the optimal strategy.

Let us now describe in further detail the structure of the optimal control. It is intuitively clear that there should be a buy and sell region; more precisely from (3.1) we can see that the sign of $\hat{u}(s)$ depends on $\hat{X}(s)$. In particular, $\hat{u}(s)$ is positive if and only if $\hat{X}(s) < \frac{-b(s)\gamma(s)-c(s)}{2a(s)}$. Indeed, if we have low holdings $\hat{X}(s)$, we will make purchases to come closer to our target while for high $\hat{X}(s)$, it can be beneficial to temporarily reduce the holdings to come closer to $\gamma(s)$. This leads us to define the *frontier* ζ by

$$\zeta(t, \gamma) = \frac{-b(t)\gamma - c(t)}{2a(t)}, \quad t \in [0, T], \quad \gamma \in [0, 1]$$

so that we have

$$\begin{aligned}
\hat{u}(t) &< 0 \ \text{ on } \ \hat{X}(t) > \zeta(t, \gamma(t)),\\
\hat{u}(t) &= 0 \ \text{ on } \ \hat{X}(t) = \zeta(t, \gamma(t)),\\
\hat{u}(t) &> 0 \ \text{ on } \ \hat{X}(t) < \zeta(t, \gamma(t)).
\end{aligned}$$

A key practical requirement is that for the original buy program (i.e., when we start with 0 shares) we should not sell. Indeed, most clients would typically be unhappy if during their execution the stock holdings were not monotone increasing. In addition, in US markets such behavior is actually prohibited by the regulators.[1] The following proposition shows

[1]See FINRA directive 5310 on Best Execution and Interpositioning, http://finra.complinet.com/en/display/display_main.html?rbid=2403&element_id=10455 .

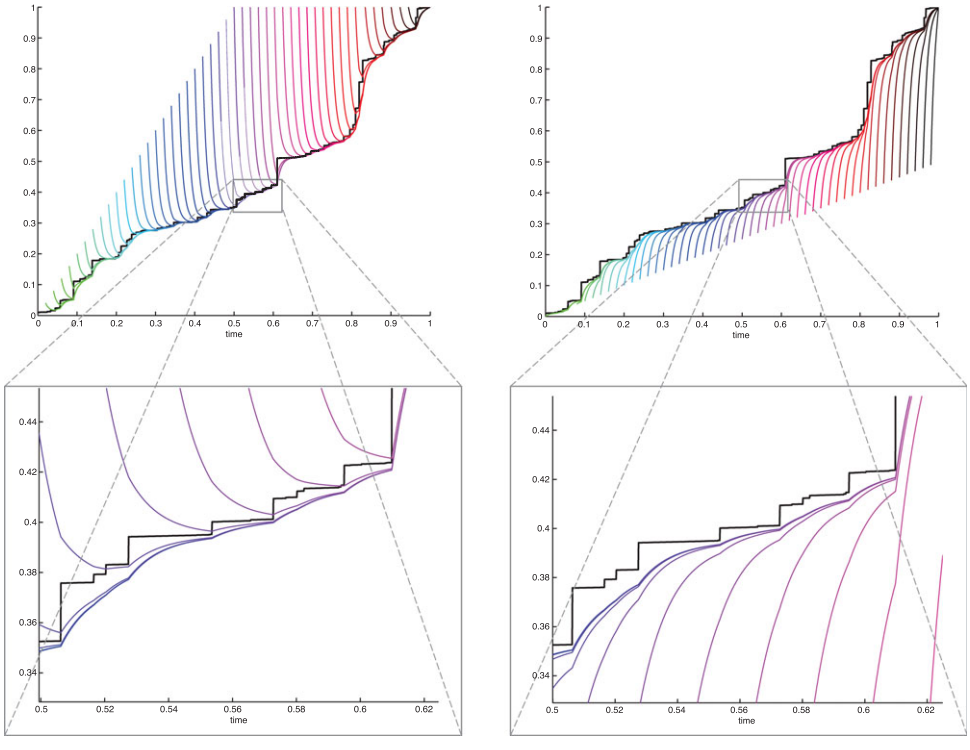

FIGURE 3.1. Simulations of $\hat{X}$: each colored curve corresponds to and starts at one initial condition $(t, \hat{X}(t))$ while the path of the gamma bridge is always the same. $\hat{X}$ can cross the frontier ζ (black curve) only from above (left panels) but not from below (right panels). The parameters are $\kappa = 10^{-8}$ (compare Remark 2.5), $\sigma = 0.01$, $\lambda = 1$, $m = 25$, and $T = 1$.

that in our formulation, for any parameter values, this is indeed the case, underlining the model's applicability and compliance with this important regulatory aspect.

COROLLARY 3.3. *Both partial derivatives of ζ are positive on $[0, T) \times [0, 1)$ and it holds that $\zeta(0, 0) > 0$. Starting with $\hat{X}(0) = 0$ and $\gamma(0) = 0$, the process $\hat{u}(t)$ is nonnegative for all t.*

We provide the proof in Section 5.3. The idea of the second part is that when the process $\hat{X}(t)$ is below $\zeta(t, \gamma(t))$, it will not cross $\zeta(s, \gamma(s))$ at a later time point due to the properties of the frontier ζ. Figure 3.1 illustrates this behavior of $\hat{X}$, that it can cross ζ only from above but not from below. We can also see that for different starting points, the paths of $\hat{X}$ are on very similar trajectories after only a short time. This is due to the multiplication by $e^{-\frac{1}{\kappa}\int_t^s a(r)\,dr}$ of x in the definition (3.1) of $\hat{X}$, which makes the impact of the starting value x vanish soon since a diverges to ∞.

Recalling that the gamma bridge satisfies $E[\gamma(t)] = \frac{t}{T}$, a natural question to ask is how our current optimization is related to the deterministic problem obtained by replacing $\gamma(t)$ by its mean, namely

$$\inf_u \left(\kappa \int_t^T u^2(s)\,ds + \lambda\sigma^2 \int_t^T \left(\frac{s}{T} - X^u(s) \right)^2 ds \right), \tag{3.3}$$

where the infimum is over all integrable u with

$$\mathrm{d}X^u(s) = u(s)\,\mathrm{d}s, \quad X^u(t) = \frac{t}{T}, \quad X^u(T) = 1.$$

The optimal control for the problem (3.3) is constant and since time is measured in calendar units in our model, this corresponds to a TWAP execution. A nice consequence of attempting to relate (3.3) to our original problem is that we are able to show that the solution of the original problem can be decomposed into two components.

COROLLARY 3.4. *The optimal trading rate $\hat{u}$ and the holdings $\hat{X}$ of the optimization problem (2.2) can be decomposed as*

$$\hat{u} = \hat{u}_1 + \hat{u}_2 \quad \text{and} \quad \hat{X} = \hat{X}_1 + \hat{X}_2, \tag{3.4}$$

where $\hat{u}_1(s) = 1/T$, $\hat{X}_1(s) = s/T$ for $s \geq t$ is the solution to (3.3). The processes $\hat{X}_2$ and $\hat{u}_2$ are given by

$$\hat{X}_2(s) = \left(x - \frac{t}{T}\right) \mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} + \frac{1}{2\kappa}\int_t^s b(z)\left(\frac{z}{T} - \gamma(z)\right) \exp\left(-\frac{1}{\kappa}\int_z^s a(r)\,\mathrm{d}r\right) \mathrm{d}z$$

and $\hat{u}_2(s) = \hat{X}_2'(s) = -\frac{1}{2\kappa}(2a(s)\hat{X}_2(s) + b(s)(\gamma(s) - \frac{s}{T}))$.

The proof of Corollary 3.4 is provided in Section 5.3. This decomposition of the optimal control can be interpreted as follows: the first part corresponds to the TWAP execution. The second part shows how we need to deviate from this deterministic strategy due to the randomness and jumps of $\gamma(s)$. Such a result corresponds nicely to heuristic algorithm design where a "historical" or average amount is always executed with an adaptive correction due to volume spikes.

Using $a > 0$ and the form of $\hat{u}_2$, we see that $\hat{X}_2$ has a zero-reversion property: if $\hat{X}_2$ equals to a big positive (negative) value, $\hat{u}_2$ will become negative (positive) so that $\hat{X}_2$ will be decreasing (increasing).

We also note that for $\kappa T \gg 1$, the optimal strategy $\hat{u}$ may be approximated by $\hat{u}_1$ because $\hat{u}_2$ depends on $\frac{1}{\kappa}(\gamma(s) - \frac{s}{T})$, whose variance vanishes as $\kappa T \to \infty$. This behavior is inline with section 1.2 of Almgren (2012) and to be expected from the original optimization problem (2.2): if κ is huge, the first term dominates the second, which leads to an approximately uniform distribution of the stock purchases due to the assumption of a linear market impact. Similarly, a huge T means that $\gamma(s)$ will be close to s/T (low variance of $\gamma(s)$) and hence the second term will again have a minor impact for a linear $\hat{X}(t)$.

4. FITTING THE MODEL TO DATA

To fit our gamma bridge-based model to data, one could attempt to fit m, using, for example, least squares. We note that since the mean increase of the gamma bridge over the day is independent of m due to the fact that $E[\gamma(t)] = \frac{t}{T}$, this fit must be done on the variance or standard deviation. In reality, it is known that trading volume is U-shaped (higher volume at the beginning and end of the day), which implies an approximate cubic shape for cumulative relative volume, see Figure 4.3. To incorporate this feature in our model, we make a deterministic time change given by a polynomial

$$t \mapsto G(t) = at^3 + bt^2 + ct + d,$$

for some constants a, b, c, d. We require that G be an increasing bijection of $[0, T]$, so we have $d = 0$ and $c = 1 - aT^2 - bT$. In the next subsection, we discuss how the parameters a, b, and m can be chosen to fit data. In Section 4.2, we explain how this time change affects our model and results.

4.1. Estimating the Model Parameters

We exemplify the estimation of the parameters on the stocks of Vodafone Group PLC (VOD) and Microsoft Corp (MSFT).[2] We also analyzed other liquid European and US stocks and the results were similar. We used intraday volume observed at 5-minute frequencies, a reasonable duration for a volume prediction, although the below parameters were not sensitive to this choice. We applied a method of moments estimator computed using nonlinear regression and based on the first and second moments of the intraday volume curves for the first 60 trading days of 2012. Since for our problem, the relative and not the absolute volume is relevant, we study the fit to the relative volume. Figure 4.1 displays the resulting curves, which led to the estimations:

	$\hat{a}$	$\hat{b}$	$\hat{c}$
VOD	1.3538	−1.6467	45.2344
MSFT	1.0739	−1.8151	84.9270

.

We see in Figure 4.1 that we get overall a good fit for the mean and standard deviation of the intraday volume curve, given that we only have three model parameters. To verify that the model does not exhibit any extreme seasonal dependence, we calibrated the parameters a, b, and m with a 60-day rolling window for the next 30 days. The time series are shown in Figure 4.2. Note that the variance of the gamma bridge is proportional to $\frac{1}{mT+1}$ so that the large range of m is not as severe as might be first assumed.

We next discuss not only the fit in the moments but the general fit. In a first qualitative analysis, we can compare the intraday volume curves of the 60 trading days with a sample of 60 trajectories based on the gamma bridge with a time change using the estimated parameters. We see in Figure 4.3 that the patterns of the true and sample trajectories look similar although the simulated paths have a more erratic behavior.

To make a more quantitative statement, we consider two goodness-of-fit tests. We again use the estimated parameters $\hat{a}$, $\hat{b}$, and $\hat{m}$, given in the above table and based on the first 60 trading days of 2012. We then take the following 30 days as an out-of-sample test data set. Under our assumption of a gamma bridge for the relative volume, we have, at each time t, 30 i.i.d. observations from a beta distribution with parameters

$$\alpha(t) = \hat{m}(\hat{a}t^3 + \hat{b}t^2 + (1 - \hat{a} - \hat{b})t) \text{ and } \beta(t) = \hat{m}T - \alpha(t).$$

The first test uses that the sample mean is approximately normally distributed with mean $\frac{\alpha(t)}{\alpha(t)+\beta(t)}$ and variance $\frac{\alpha(t)\beta(t)}{30(\alpha(t)+\beta(t))^2(\alpha+\beta+1)}$ by the central limit theorem. Performing a z-test leads to the P-values shown in Figure 4.4.

The second test, displayed also in Figure 4.4, is the well-known Kolmogorov–Smirnov test, which assesses the null hypothesis that the sample is drawn from a Beta($\alpha(t)$, $\beta(t)$)-distribution. Observe that we now consider the whole empirical cumulative distribution and not just the sample mean. Since there is dependence across different time points,

[2]All data are used with the permission of Bloomberg L.P.

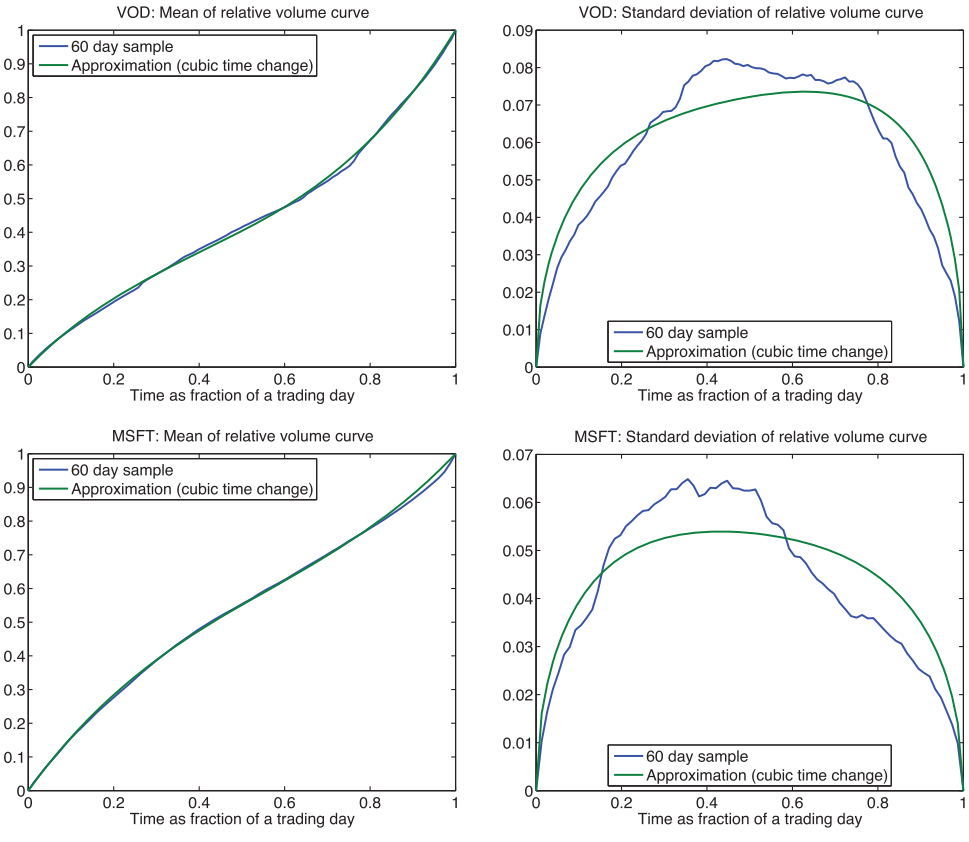

FIGURE 4.1. Example fits for the mean and standard deviation for one US and one EU stock.

one should interpret Figure 4.4 in the following way: if one selects a time interval during the day at random, the corresponding y-values are the P-values of the two tests at that given time. Generally, in the morning we do not reject the null hypothesis whereas in parts of the afternoon we do. This conclusion of an imperfect fit in parts of the day is not very surprising. Indeed, relative volume is not distributed as a gamma bridge, being subject to idiosyncratic factors that the model does not (and cannot) fully capture. The model proposed here aims to capture the main stylized features while being tractable enough for optimization, compare with the use of Brownian motion as a model for the price despite the non-Gaussianity of returns as well as the Black–Scholes framework for option pricing. Taking these considerations into account, we conclude that the model is suitable for our purposes.

One final point related to fitting the model to data concerns the choices of the parameters κ and λ. A calculation shows that one consequence of our model formulation is that as $\kappa \downarrow 0$, since the costs vanish, we end up being able to track VWAP perfectly. In a real trading environment, this is of course not possible since each execution of a market order incurs a per share cost of one-half the spread (vs. mid quote). One way to resolve this issue would be a reformulation of the impact costs to include a fixed cost of one-half the spread. This naturally leads to an impulse control formulation which is beyond the

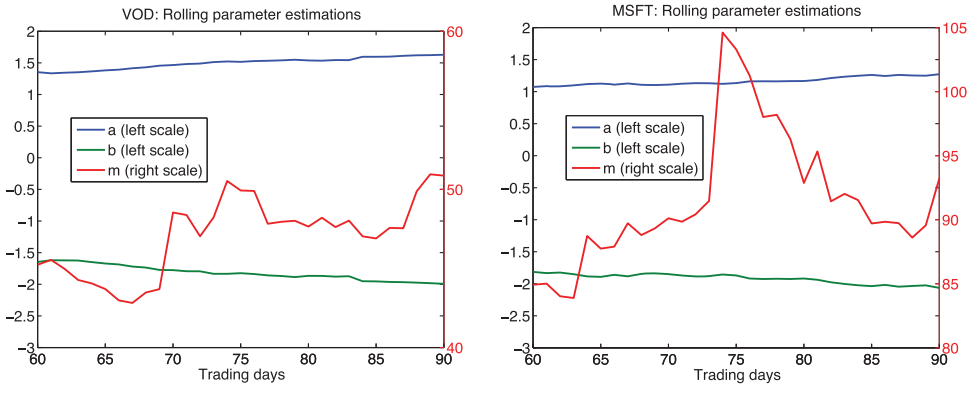

FIGURE 4.2. Rolling parameter estimations of a, b, and m.

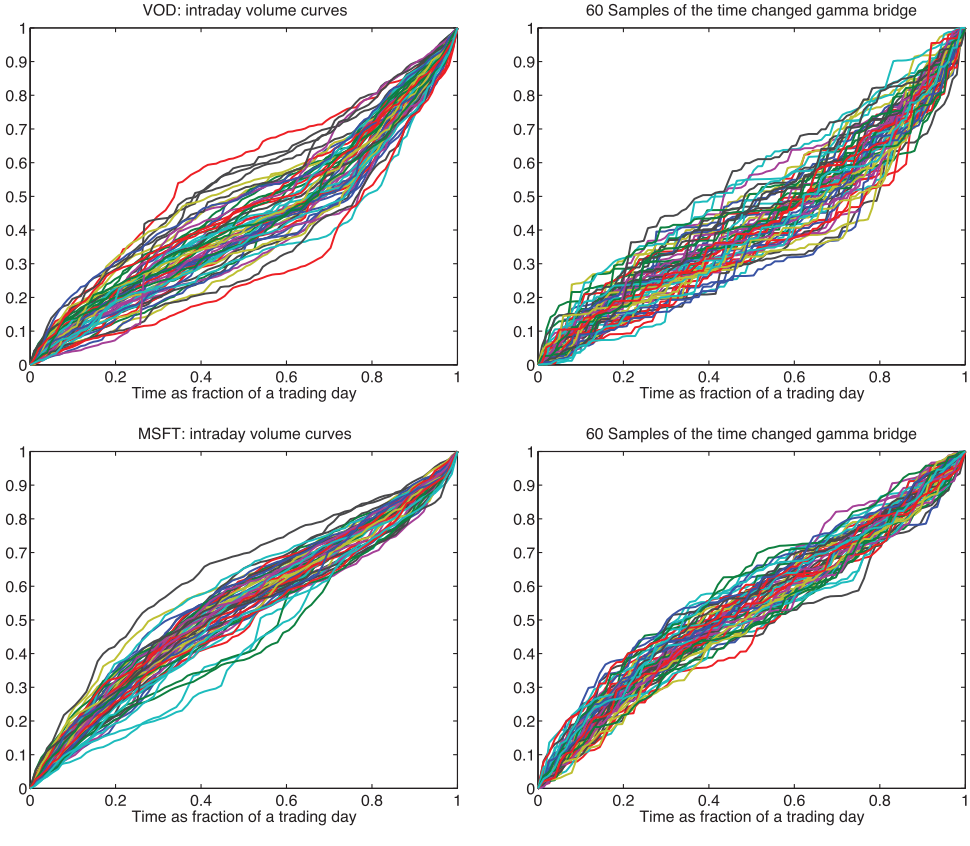

FIGURE 4.3. Comparison of the intraday volume curves with sample paths.

scope of this paper. A more pragmatic solution observes that the key parameter driving the optimal control/trajectory is $\frac{\kappa}{\lambda\sigma^2}$. In practice, one can therefore simply choose a value of λ so that this ratio is in an appropriate range for the model to be applicable. Broadly speaking, one is choosing an "effective λ" to take into account spread costs, similarly as

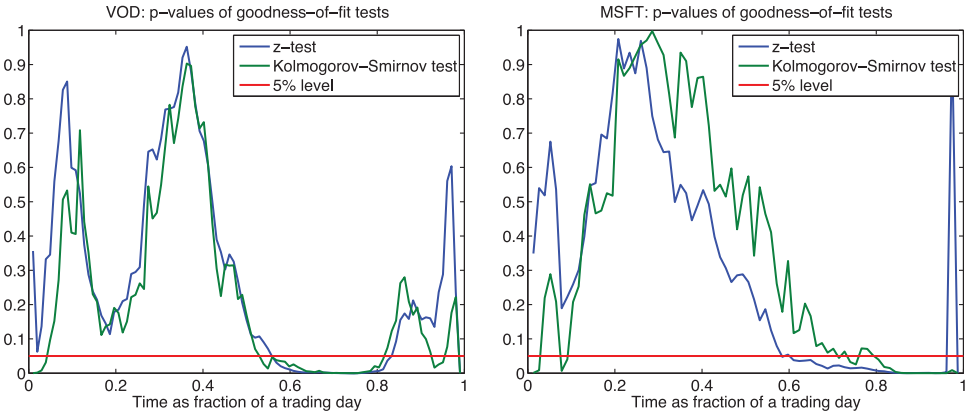

FIGURE 4.4. *P*-values based on a 30-day out-of-sample data set.

is done in trading systems employing the Implementation Shortfall framework described in Almgren and Chriss (2000).

4.2. Integrating the Time Change into the Model

As we have seen in the above discussion, the gamma bridge fits well to the relative volume curve, but only after a deterministic time transformation. We next see that our main result still holds when the model undergoes such a time change. We first explain the reason in a simple example and give afterward the mathematical argument. Assume that one has calculated the optimal control and the value function in the model of Section 2. Now it is given that the model fits well except that the expected trading frequency in the second-half of the day is double that of the first-half. How should the optimal trading strategy be modified? A natural answer is to scale the original strategy accordingly by simply changing the speed of trading execution. In the morning, it is reduced by a third so that in the middle of the day one has the same position as one would have had after one-third of the day with the original strategy. In the afternoon, the trading frequency is then correspondingly increased. In summary, the new strategy is just the time-changed original strategy scaled by the derivative of the time change.

It is also sensible to expect that the value of the optimization problem will not change because the algorithm performance (minimal slippage from VWAP) should not change under a deterministic time change. Indeed, we may take the time change into consideration when choosing our strategy so that there is no additional information present. Of course, these arguments use that the time change is deterministic and hence known in advance, however, this is acceptable for most practical purposes.

Let us now make the above arguments rigorous by adapting the model of Section 2. The time change is assumed to be given by a differentiable deterministic function $G : [0, T] \to [0, T]$ with $G(0) = 0$, $G(T) = T$, and $G'(t) \geq C$ for some constant $C > 0$ and all $t \in [0, T]$. The relative volume curve is modeled by $\eta(t) = \gamma(G(t))$, where $(\gamma(t))_{0 \leq t \leq T}$ is a gamma bridge whose underlying gamma process has a general parameter m. The time change means that trading frequency depends on the time of the day, e.g., higher at the beginning and closing of the day.

There is a well-known link between traded volume and prices, we thus expect that the varying trading frequency will also affect the price process so that the asset price at time t equals

$$\tilde{P}(t) = P(0) + \sigma W(G(t)).$$

Observe that now we have that for two time points t_0 and t_1, the volatility of returns is proportional to

$$\sqrt{E\left[(\tilde{P}(t_1) - \tilde{P}(t_0))^2\right]} = \sigma \sqrt{\int_{t_0}^{t_1} G'(s)\mathrm{d}s}.$$

Assuming that the difference $t_1 - t_0$ is small and fixed, this can be approximated by $\sigma\sqrt{t_1 - t_0}\sqrt{G'(t_0)}$. Since G' reflects the expected intraday trading frequency, it is typically U-shaped and hence $\sqrt{G'}$ is also U-shaped (or V-shaped). Indeed, a U-shaped G' means that its derivative G'' is negative at the beginning and increases to become positive at the end, and this property translates to $\frac{\mathrm{d}}{\mathrm{d}t}\sqrt{G'} = \frac{G''}{2\sqrt{G'}}$ as well. Therefore, we expect the instantaneous volatility to be U-shaped. Apart from being a mathematical consequence of making a time change to our model, this phenomenon is very well documented in the empirical finance literature across many different stock markets, dating back to Wood, McInish, and Ord (1985) and Harris (1986). Also the link between the U-shaped forms of intraday volume and volatility is well known and goes back, at least, to Harris (1987). The common theoretical explanation is that the patterns of both volume and volatility are related to the flow of information, which is not constant over time. For example, in a market with asymmetrically informed participants, trading volume itself conveys information so that a U-shaped volume will lead to a U-shaped flow of information. This nonconstant flow of information is captured in our model by the time change.

Similarly, the coefficient of market impact at time t is now $\tilde{\kappa}(t) = \kappa/G'(t)$ because a decrease (or increase) in G' means a slowdown (acceleration) in overall market trading frequency so that the market impact of our trades is increased (decreased).

Proceeding along the same lines as in Section 2 and additionally using $\mathrm{d}\langle W \circ G\rangle_s = \mathrm{d}G(s) = G'(s)\,\mathrm{d}s$, we see that the new value function to the optimization problem is given by

$$w(\tau, x, \gamma) = \inf_y E\left[\int_\tau^T \tilde{\kappa}(s) y^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_\tau^T (\eta(s) - X^y(s))^2 G'(s)\,\mathrm{d}s\right],$$

where $\eta(s) = \gamma(G(s))$ for a gamma bridge $(\gamma(s))_{G(\tau)\le s\le T}$ with $\gamma(G(\tau)) = \gamma$ and the integrable y is such that

$$\mathrm{d}X^y(s) = y(s)\,\mathrm{d}s, \quad X^y(\tau) = x, \quad X^y(T) = 1$$

and y is adapted to the time-changed filtration $(\mathcal{F}_{G(s)})_{0\le s\le T}$. We associate such a y with a process u by $u(s) = \frac{y(G^{-1}(s))}{G'(G^{-1}(s))}$, which is adapted to the original filtration $(\mathcal{F}_s)_{0\le s\le T}$. We then have

$$\begin{aligned} &E\left[\int_\tau^T \tilde{\kappa}(s) y^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_\tau^T (\eta(s) - X^y(s))^2 G'(s)\,\mathrm{d}s\right] \\ &\quad = E\left[\int_t^T \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_t^T (\gamma(s) - X^u(s))^2\,\mathrm{d}s\right] \end{aligned}$$

for $t = G(\tau)$ and

$$X^y(G^{-1}(s)) = x + \int_\tau^{G^{-1}(s)} y(r)\,\mathrm{d}r = x + \int_t^s u(r)\,\mathrm{d}r = X^u(s).$$

Since this holds for any such y, we obtain

$$w(\tau, x, \gamma) = v(G(\tau), x, \gamma),$$

where v is the original value function from (2.2), which is characterized in Theorem 3.1. We can also see that the optimal control of the time-changed problem is given by

$$\begin{aligned}\hat{y}(\tau, x, \gamma) &= \hat{u}(G(\tau), x, \gamma)G'(\tau)\\ &= -\frac{2a(G(\tau))x + b(G(\tau))\gamma + c(G(\tau))}{2\kappa}G'(\tau)\\ &= -\frac{2a(G(\tau))x + b(G(\tau))\gamma + c(G(\tau))}{2\tilde{\kappa}(t)},\end{aligned}$$

again by using Theorem 3.1.

5. PROOF OF THE MAIN RESULT

We split the proof of Theorem 3.1 into two parts: we first show some properties of the functions a, b, c, d, f, g, and the candidate for the value function, and then we verify that this candidate indeed gives rise to the value function v. Finally, we will provide the proofs of Corollaries 3.3 and 3.4.

5.1. Properties of the Auxiliary Functions

LEMMA 5.1. *The functions $a, b, c, d, f, g : [0, T) \to \mathbb{R}$ are well defined and there exist constants k_1, k_2, k_3 such that*

$$\lim_{s \nearrow T} a(s)(T - s) = \kappa, \tag{5.1}$$

$$|b(t)| + |d(t)| + |f(t)| \leq k_1(T - t), \tag{5.2}$$

$$\left|\frac{c(t)}{a(t)} + 2\right| \leq k_2(T - t)^2, \tag{5.3}$$

$$|g(t) - a(t)| \leq k_3(T - t), \tag{5.4}$$

for all $t \in [0, T)$.

Proof. We consider $a, \dots, g$ and their limiting behavior sequentially.

(1) *The functions a, b, c.* The functions a, b, c are well defined for $t \in [0, T)$, and an application of L'Hôpital's rule yields (5.1). To study the behavior of b, we use a series expansion

$$\begin{aligned} b(t) &= -2a(t) + \frac{2\kappa}{T-t} \\ &= -2\sqrt{\kappa\lambda\sigma^2}\,\frac{\mathrm{e}^{2(T-t)\sqrt{\lambda\sigma^2/\kappa}}+1}{\sum_{n=1}^{\infty} 2^n (T-t)^n (\lambda\sigma^2/\kappa)^{n/2}/n!} + \frac{2\kappa}{T-t} \\ &= \frac{2\kappa}{T-t}\,\frac{\sum_{n=0}^{\infty} 2^{n+1}(T-t)^n(\lambda\sigma^2/\kappa)^{n/2}/(n+1)! - \mathrm{e}^{2(T-t)\sqrt{\lambda\sigma^2/\kappa}} - 1}{\sum_{n=0}^{\infty} 2^{n+1}(T-t)^n(\lambda\sigma^2/\kappa)^{n/2}/(n+1)!} \\ &= 2\kappa(T-t)\frac{\sum_{n=2}^{\infty} 2^n (T-t)^{n-2}(\lambda\sigma^2/\kappa)^{n/2}(2/(n+1)! - 1/n!)}{\sum_{n=0}^{\infty} 2^{n+1}(T-t)^n(\lambda\sigma^2/\kappa)^{n/2}/(n+1)!}. \end{aligned}$$

Using $2/(n+1)! - 1/n! \le 0$ for all n, we see that

$$0 > b(t) \ge -C(T-t) \tag{5.5}$$

for all $t < T$ and some constant C. For the function c, we obtain

$$\begin{aligned} \frac{c(t)}{a(t)} + 2 &= -2\frac{\sum_{n=0}^{\infty} 2^{n+1}(T-t)^n(\lambda\sigma^2/\kappa)^{n/2}/(n+1)!}{\mathrm{e}^{2(T-t)\sqrt{\lambda\sigma^2/\kappa}}+1} + 2 \\ &= 2(T-t)^2\frac{\sum_{n=2}^{\infty} 2^n (T-t)^{n-2}(\lambda\sigma^2/\kappa)^{n/2}(1/n! - 2/(n+1)!)}{\mathrm{e}^{2(T-t)\sqrt{\lambda\sigma^2/\kappa}}+1}, \end{aligned}$$

which shows (5.3).

(2) *The function d.* Since b is bounded by (5.5) and we have $\frac{T+1/m-s}{T+1/m-t} \le 1$ and $\frac{T-s}{T-t} \le 1$ for $s \in [t, T]$, the integrand in

$$d(t) = \int_t^T \left(\lambda\sigma^2 - \frac{1}{4\kappa}b^2(s)\right)\frac{T-s}{T-t}\,\frac{T+1/m-s}{T+1/m-t}\,\mathrm{d}s$$

is bounded so that d is well defined and we can deduce $|d(t)| \le C(T-t)$ for some constant C.

(3) *The function f.* From (5.5) and the boundedness of d, we derive that

$$\left|b(s) + 2d(s)\frac{T-s}{T-s+1/m}\right| \le C(T-s)$$

for some constant C. Hence, f is well defined and

$$|f(t)| \le \frac{1}{T-t}\int_t^T \left|b(s) + 2d(s)\frac{T-s}{T-s+1/m}\right|\mathrm{d}s \le \frac{C}{2}(T-t).$$

(4) *The function g.* Thanks to (5.1), the function $s \mapsto a(s)(T-s)$ is continuous and bounded on $[0, T)$. Together with (5.2), this yields

$$\left|\frac{b(s)}{2}\left(\frac{a(s)}{\kappa}(T-s)+1\right) - f(s) - \frac{d(s)/m}{T-s+1/m}\right| \le C(T-s)$$

for some constant C. Therefore, the function g is well defined and

$$\begin{aligned}&|g(t)-a(t)|\\&\quad=\left|T-t+\int_t^T\left(\frac{b(s)}{2}\left(\frac{a(s)}{\kappa}(T-s)+1\right)-f(s)-\frac{d(s)/m}{T-s+1/m}\right)\frac{1}{T-s}\,\mathrm{d}s\right|\\&\quad\le k_3(T-t)\end{aligned}$$

for some constant k_3, which shows (5.4). □

LEMMA 5.2. *The function* $\varphi:[0,T)\times\mathbb{R}_+\times[0,1]\to\mathbb{R}$ *defined by*

$$\varphi(t,x,\gamma)=a(t)x^2+b(t)\gamma x+c(t)x+d(t)\gamma^2+f(t)\gamma+g(t) \tag{5.6}$$

satisfies

$$\begin{aligned}&\varphi_t+\lambda\sigma^2(\gamma-x)^2-\frac{\varphi_x^2}{4\kappa}\\&\qquad+m\int_0^1(\varphi(t,x,\gamma+(1-\gamma)z)-\varphi(t,x,\gamma))(1-z)^{Tm-tm-1}\frac{1}{z}\,\mathrm{d}z=0.\end{aligned} \tag{5.7}$$

Proof. We can write

$$\begin{aligned}&m\int_0^1(\varphi(t,x,\gamma+(1-\gamma)z)-\varphi(t,x,\gamma))(1-z)^{Tm-tm-1}\frac{1}{z}\,\mathrm{d}z\\&\quad=(b(t)+f(t)+2\gamma d(t))(1-\gamma)\varphi_0(t)+d(t)(1-\gamma)^2\varphi_1(t),\end{aligned} \tag{5.8}$$

where

$$\begin{aligned}\varphi_0(t)&=m\int_0^1(1-z)^{Tm-tm-1}\,\mathrm{d}z=\frac{1}{T-t},\\\varphi_1(t)&=m\int_0^1 z(1-z)^{Tm-tm-1}\,\mathrm{d}z=\frac{1}{T-t}-\frac{1}{T-t+1/m}.\end{aligned}$$

We also have

$$-\frac{\varphi_x^2}{4\kappa}=\frac{1}{\kappa}(-a^2x^2-abx\gamma-acx)-\frac{1}{4\kappa}(b^2\gamma^2+2bc\gamma+c^2), \tag{5.9}$$

where we write a for $a(t)$, etc. By means of a straightforward calculation, we can check that

$$\begin{aligned}a_t&=\frac{1}{\kappa}a^2-\lambda\sigma^2, & d_t&=2d\varphi_0-d\varphi_1+\frac{1}{4\kappa}b^2-\lambda\sigma^2,\\b_t&=\frac{1}{\kappa}ab+2\lambda\sigma^2+b\varphi_0, & f_t&=\frac{1}{2\kappa}bc-2d\varphi_0+f\varphi_0+2d\varphi_1,\\c_t&=\frac{1}{\kappa}ac-b\varphi_0, & g_t&=\frac{1}{4\kappa}c^2-f\varphi_0-d\varphi_1.\end{aligned}$$

Using this in calculating φ_t together with (5.8) and (5.9), we derive (5.7). □

5.2. Verification

We next relate the candidate φ to the value function v by using the properties (5.1)–(5.4). We start with two auxiliary results.

LEMMA 5.3. *The candidate optimal process $\hat{X}$ given by*

$$\mathrm{d}\hat{X}(t) = -\frac{1}{2\kappa}(2a(t)\hat{X}(t) + b(t)\gamma(t) + c(t))\,\mathrm{d}t, \quad \hat{X}(0) = x \tag{5.10}$$

is bounded.

Proof. Solving (5.10) for $\hat{X}$ yields

$$\hat{X}(t) = x\mathrm{e}^{-\frac{1}{\kappa}\int_0^t a(r)\,\mathrm{d}r} + \frac{1}{2\kappa}\int_0^t (-b(s)\gamma(s) - c(s))\exp\left(-\frac{1}{\kappa}\int_s^t a(r)\,\mathrm{d}r\right)\mathrm{d}s.$$

Recalling that b is bounded by (5.2), $a > 0$ and γ is a gamma bridge, we see that it is enough to show that

$$\int_0^t |c(s)|\exp\left(-\frac{1}{\kappa}\int_s^t a(r)\,\mathrm{d}r\right)\mathrm{d}s \text{ is bounded uniformly in } t.$$

Because of $a > 0$ and $|c(s)| \le a(s)(k_2(T-s)^2 + 2)$ by (5.3), this follows from

$$\begin{aligned}\int_0^t a(s)\exp\left(-\frac{1}{\kappa}\int_s^t a(r)\,\mathrm{d}r\right)\mathrm{d}s &= \kappa\exp\left(-\frac{1}{\kappa}\int_s^t a(r)\,\mathrm{d}r\right)\Big|_{s=0}^{s=t}\\ &= \kappa - \kappa\exp\left(-\frac{1}{\kappa}\int_0^t a(r)\,\mathrm{d}r\right)\\ &\le \kappa.\end{aligned}$$

□

We next establish some a priori estimates for φ.

LEMMA 5.4. *There exists a constant K such that*

$$-K(T-t)(x+1) \le \varphi(t,x,\gamma) - a(t)(x-1)^2 \le K(T-t)(x+1)$$

for all $t \in [0, T)$, $x \ge 0$, and $\gamma \in [0, 1]$. In particular, for every compact set $M \subset \mathbb{R}_+$, there exists a constant $\overline{K}$ such that

$$-\overline{K}(T-t) \le \varphi(t,x,\gamma) - a(t)(x-1)^2 \le \overline{K}(T-t)$$

for all $t \in [0, T)$, $x \in M$, and $\gamma \in [0, 1]$.

Lemma 5.4 shows that $\varphi(t,x,\gamma) - a(t)(x-1)^2$ loses the x-dependence in the limit behavior $t \nearrow T$. This can be interpreted as canceling the term $n(X^u(T)-1)^2$ in the auxiliary optimization problem (3.2) in the limit $n \to \infty$.

Proof. We deduce from (5.2)–(5.4) that

$$\begin{aligned}&|\varphi(t,x,\gamma) - a(t)(x-1)^2|\\ &\quad = |b(t)\gamma x + d(t)\gamma^2 + f(t)\gamma + x(c(t) + 2a(t)) + g(t) - a(t)|\\ &\quad \le (T-t)(k_1(x+1) + k_3 + k_2 a(t)(T-t)).\end{aligned}$$

From (5.1), it follows that $a(t)(T-t)$ is bounded, which implies the first claim. For the second part, it is enough to set $\overline{K} := K \sup_{x\in M}(1+x)$. □

We are now in position to prove Theorem 3.1.

Proof of Theorem 3.1. We first show that $\varphi \le v$, where φ is defined in (5.6) and v is the value function. This claim is equivalent to

$$\varphi(t,x,\gamma) \le E\bigg[\int_t^T \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_t^T (X^u(s)-\gamma(s))^2\,\mathrm{d}s\bigg], \tag{5.11}$$

for all controls u such that $X^u(t)=x$ and $X^u(T)=1$ a.s., $\gamma(t)=\gamma$, where $t\in[0,T]$, $x\ge 0$, $\gamma\in[0,1]$. Without loss of generality, we can assume that $E[\int_t^T \kappa u^2(s)\,\mathrm{d}s] < \infty$ since otherwise, the result holds trivially. Using Itô's formula, we have, for any stopping time τ_ℓ valued in $[t,T]$,

$$\begin{aligned}
&\varphi(\tau_\ell\wedge T, X^u(\tau_\ell\wedge T), \gamma(\tau_\ell\wedge T)) - \varphi(t,x,\gamma)\\
&\quad= \int_t^{\tau_\ell\wedge T} (\varphi_t + u\varphi_x)(s, X^u(s), \gamma(s))\,\mathrm{d}s\\
&\qquad+ \sum_{s\in(t,\tau_\ell\wedge T]} (\varphi(s, X^u(s), \gamma(s)) - \varphi(s, X^u(s), \gamma(s-)))\,.
\end{aligned}$$

Applying proposition 4 of Émery and Yor (2004) separately to ψ^+ and ψ^-, where

$$\psi(s\,m,\omega,y) := \varphi(s, X^u(s)(\omega), y+\gamma(s-)(\omega)) - \varphi(s, X^u(s)(\omega), \gamma(s-)(\omega)),$$

we deduce that the process

$$\begin{aligned}
&\sum_{s\in(tm,m\tau_\ell\wedge\cdot]} \psi\left(s,\omega,\Delta\gamma\left(\frac{s}{m}\right)(\omega)\right)\\
&\quad-\int_{tm}^{\cdot}\int_0^1 \psi\left(s,\omega,\left(1-\gamma\left(\frac{s-}{m}\right)(\omega)\right)z\right)\frac{(1-z)^{Tm-s-1}}{z}\,\mathrm{d}z\mathrm{d}s
\end{aligned}$$

is a local martingale on the interval $[tm, Tm]$ in the time-changed filtration $(\mathcal{F}_{\frac{s}{m}})_{tm\le s\le Tm}$ to which $(\gamma(\frac{s}{m}))_{tm\le s\le Tm}$ is adapted. Equivalently, the process

$$\begin{aligned}
&\sum_{s\in(t,\tau_\ell\wedge\cdot]} \psi(s\,m,\omega,\Delta\gamma(s)(\omega))\\
&\quad-m\int_t^{\cdot}\int_0^1 \psi(s\,m,\omega,(1-\gamma(s-)(\omega))z)\frac{(1-z)^{Tm-sm-1}}{z}\,\mathrm{d}z\mathrm{d}s
\end{aligned} \tag{5.12}$$

is a local martingale on $[t,T]$ in the standard filtration $(\mathcal{F}_s)_{t\le s\le T}$ (to which $(\gamma(s))_{t\le s\le T}$ is adapted). Using Lemma 5.2 together with $-\frac{\varphi_x^2}{4\kappa} \le \varphi_x u + \kappa u^2$ for all $u\in\mathbb{R}$, we derive

$$\begin{aligned}
&\varphi_t + \lambda\sigma^2(\gamma-x)^2 + \varphi_x u + \kappa u^2\\
&\qquad+ m\int_0^1 (\varphi(t,x,\gamma+(1-\gamma)z) - \varphi(t,x,\gamma))(1-z)^{Tm-tm-1}\frac{1}{z}\,\mathrm{d}z \ge 0.
\end{aligned}$$

In particular, choosing a sequence $(\tau_\ell)_{\ell\in\mathbb{N}}$ of stopping times such that $\tau_\ell \nearrow T$ and the stopped process (5.12) is a true martingale, we obtain

$$\varphi(t,x,\gamma) \le E\bigg[\int_t^{\tau_\ell\wedge T} \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^{\tau_\ell\wedge T}(X^u(s)-\gamma(s))^2\,\mathrm{d}s\bigg] + E[\varphi(\tau_\ell\wedge T,\, X^u(\tau_\ell\wedge T),\, \gamma(\tau_\ell\wedge T))]. \tag{5.13}$$

Monotone convergence implies

$$\lim_{\ell\to\infty} E\bigg[\int_t^{\tau_\ell\wedge T} \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^{\tau_\ell\wedge T}(X^u(s)-\gamma(s))^2\,\mathrm{d}s\bigg] = E\bigg[\int_t^{T} \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^{T}(X^u(s)-\gamma(s))^2\,\mathrm{d}s\bigg].$$

Applying Lemma 5.4, we have

$$E[\varphi(\tau_\ell\wedge T,\, X^u(\tau_\ell\wedge T),\, \gamma(\tau_\ell\wedge T))] \le KE\left[(T-\tau_\ell\wedge T)\left(\sup_{s\in[t,T]} X^u(s)+1\right)\right] + E\big[a(\tau_\ell\wedge T)(X^u(\tau_\ell\wedge T)-1)^2\big].$$

The first integrand converges to zero by dominated convergence, using

$$\sup_{s\in[t,T]}|X^u(s)| \le x + \int_t^T |u(s)|\,\mathrm{d}s \le x + T - t + \int_t^T u^2(s)\,\mathrm{d}s \in L^1.$$

For the second term, we observe that

$$(X^u(\tau_\ell\wedge T)-1)^2 = (X^u(\tau_\ell\wedge T)-X^u(T))^2 \le \left(\int_{\tau_\ell\wedge T}^T |u(s)|\,\mathrm{d}s\right)^2 \le (T-\tau_\ell\wedge T)\int_{\tau_\ell\wedge T}^T u^2(s)\,\mathrm{d}s,$$

by Hölder's inequality. Using that $a(\tau_\ell\wedge T)(T-\tau_\ell\wedge T)$ is uniformly bounded by (5.1), we obtain, for some constant C,

$$E[a(\tau_\ell\wedge T)(X^u(\tau_\ell\wedge T)-1)^2] \le CE\bigg[\int_{\tau_\ell\wedge T}^T u^2(s)\,\mathrm{d}s\bigg] \to 0,$$

as $m\to\infty$ by dominated convergence. Together this shows that

$$\varphi(t,x,\gamma) \le E\bigg[\int_t^{T} \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^{T}(X^u(s)-\gamma(s))^2\,\mathrm{d}s\bigg],$$

and concludes the proof of (5.11). To prove the reverse inequality $\varphi\ge v$, we deduce similarly to (5.13) that

$$\varphi(t,x,\gamma) = E\bigg[\int_t^{\tau_\ell\wedge T} \kappa \hat{u}^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^{\tau_\ell\wedge T}(\hat{X}(s)-\gamma(s))^2\,\mathrm{d}s\bigg] + E[\varphi(\tau_\ell\wedge T,\, \hat{X}(\tau_\ell\wedge T),\, \gamma(\tau_\ell\wedge T))], \tag{5.14}$$

for the candidate optimal control $\hat{u}$ corresponding to $\hat{X}$ given in (5.10)—this time we have an equality from Lemma 5.2 using $\hat{u}(s) = -\frac{\varphi_x}{2\kappa}(s, \hat{X}(s), \gamma(s))$. By Lemma 5.3, $X^{\hat{u}}$ is bounded and hence we obtain from Lemma 5.4 that

$$\begin{aligned} &E[\varphi(\tau_\ell \wedge T, \hat{X}(\tau_\ell \wedge T), \gamma(\tau_\ell \wedge T))] \\ &\geq E[a(\tau_\ell \wedge T)(\hat{X}(\tau_\ell \wedge T) - 1)^2] - \overline{K} E[T - \tau_\ell \wedge T]. \end{aligned} \tag{5.15}$$

Since the first term in (5.15) is nonnegative and the second converges to zero, monotone convergence yields

$$\varphi(t, x, \gamma) \geq E\bigg[\int_t^T \kappa \hat{u}^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_t^T (\hat{X}(s) - \gamma(s))^2\,\mathrm{d}s\bigg].$$

Using the admissible control $\bar{u}(s) = \frac{1-x}{T-t}$, one can see that

$$\varphi(t, x, \gamma) \leq v(t, x, \gamma) \leq \kappa \frac{(1-x)^2}{T-t} + M' < \infty, \tag{5.16}$$

for some constant M'. We conclude the proof by showing $\hat{X}_T = 1$ a.s. Using (5.14), (5.15), and (5.16), we deduce

$$\kappa \frac{(1-x)^2}{T-t} + M' \geq E[a(\tau_\ell \wedge T)(\hat{X}(\tau_\ell \wedge T) - 1)^2] - \overline{K} E[T - \tau_\ell \wedge T].$$

Since a is increasing, it follows, for every $t_0 \in [t, T)$, that

$$\kappa \frac{(1-x)^2}{T-t} + M' \geq a(t_0) E[(\hat{X}(\tau_\ell \wedge T) - 1)^2 1_{\{\tau_\ell \geq t_0\}}] - \overline{K} E[T - \tau_\ell \wedge T].$$

Using that $\hat{X}$ is bounded, dominated convergence yields

$$\kappa \frac{(1-x)^2}{T-t} + M' \geq a(t_0) E[(\hat{X}(T) - 1)^2],$$

for all $t_0 \in [t, T)$, and hence

$$\kappa \frac{(1-x)^2}{T-t} + M' \geq \lim_{t_0 \nearrow T} a(t_0) E[(\hat{X}(T) - 1)^2].$$

Since $\lim_{t_0 \nearrow T} a(t_0) = \infty$, this can only hold if $E[(\hat{X}(T) - 1)^2] = 0$, which means $\hat{X}(T) = 1$ a.s. The uniqueness of the optimal control is a consequence of the linearity of X^u in u and the strict convexity of the optimization problem in the control. □

5.3. Proofs of Corollaries 3.3 and 3.4

We conclude by proving Corollaries 3.3 and 3.4.

Proof of Corollary 3.3. From $b(t) < 0$ for $t < T$ by (5.5) and $a(t) > 0$, it follows $\zeta_\gamma(t,\gamma) = \frac{-b(t)}{2a(t)} > 0$ for all $t < T$. We next write

$$\zeta(t,\gamma) = \frac{-b(t)\gamma - c(t)}{2a(t)} = \gamma + \kappa(1-\gamma)\frac{1}{(T-t)a(t)}.$$

To prove $\zeta_t(t,\gamma) > 0$, it is enough to show that $\frac{1}{(T-t)a(t)}$ is increasing or, equivalently, that $(T-t)a(t)$ is decreasing. Hence, we consider

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}t}(T-t)a(t) &= -a(t) + (T-t)a_t(t) \\
&= -\sqrt{\kappa\lambda\sigma^2}\,\frac{\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} + \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}}}{\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} - \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}}} + (T-t)\frac{4\lambda\sigma^2\mathrm{e}^{2(T+t)\sqrt{\lambda\sigma^2/\kappa}}}{(\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} - \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}})^2} \\
&= \frac{\mathrm{e}^{2(T+t)\sqrt{\lambda\sigma^2/\kappa}}(-\sqrt{\kappa\lambda\sigma^2}(\mathrm{e}^{2(T-t)\sqrt{\lambda\sigma^2/\kappa}} - \mathrm{e}^{-2(T-t)\sqrt{\lambda\sigma^2/\kappa}}) + 4\lambda\sigma^2(T-t))}{(\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} - \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}})^2} \\
&< \frac{\mathrm{e}^{2(T+t)\sqrt{\lambda\sigma^2/\kappa}}(-\sqrt{\kappa\lambda\sigma^2}\,4(T-t)\sqrt{\lambda\sigma^2/\kappa} + 4\lambda\sigma^2(T-t))}{(\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} - \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}})^2} = 0,
\end{aligned}$$

using $\mathrm{e}^x - \mathrm{e}^{-x} = 2\sinh(x) = 2\sum_{n=0}^{\infty}\frac{x^{2n+1}}{(2n+1)!} > 2x$ for all $x > 0$. This proves $\zeta_t(t,\gamma) > 0$. Moreover, we have

$$\zeta(0,0) = \kappa\frac{\varphi_0(0)}{2a(0)} = \sqrt{\kappa}\,\frac{\mathrm{e}^{2T/\sqrt{\kappa}} - 1}{2T(\mathrm{e}^{2T/\sqrt{\kappa}} + 1)} > 0.$$

For the second part of the lemma, we define a process $\xi(t) = \zeta(t,\gamma(t)) - \hat{X}(t)$, which satisfies $\xi(0) = \zeta(0,\gamma(0)) = \zeta(0,0) > 0$. We have $\xi(t) \geq 0$ a.s. for all t by the following reason. Fix t and consider the event $\xi(t) < 0$. Since $\hat{X}$ is continuous and $\zeta_\gamma > 0$ and γ has only positive jumps, there needs to exist a random $\tau < t$ such that $\xi(\tau) = 0$ and $\xi(s) \leq 0$ for all $s \in [\tau, t]$. However, we have

$$\xi(t) = \underbrace{\xi(\tau)}_{=0} + \underbrace{\zeta(t,\gamma(t)) - \zeta(\tau,\gamma(\tau))}_{\geq 0} - \int_\tau^t \underbrace{\hat{u}(s)}_{\leq 0}\,\mathrm{d}s \geq 0,$$

since $\zeta(s,\gamma(s))$ is increasing and $\hat{u}(s) \leq 0$ if $\xi(s) \leq 0$. Therefore, $\xi(t)$ cannot become negative and $\hat{u}(t)$ will be nonnegative for all t. □

Proof of Corollary 3.4. Let $\hat{X}_2, \hat{u}_2$ be defined as in the corollary and set $\hat{u}_1(s) = \hat{X}_1'(s)$ and

$$\hat{X}_1(s) = \frac{t}{T}\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} + \frac{1}{2\kappa}\int_t^s \left(-b(z)\frac{z}{T} - c(z)\right)\exp\left(-\frac{1}{\kappa}\int_z^s a(r)\,\mathrm{d}r\right)\mathrm{d}z,$$

so that (3.4) is satisfied by construction; compare with (5.10). By definition, we have

$$-b(z)\frac{z}{T} - c(z) = a(z)\frac{2z}{T} - \frac{2z\kappa}{(T-z)T} + \frac{2\kappa}{T-z} = a(z)\frac{2z}{T} + \frac{2\kappa}{T},$$

so that

$$\begin{aligned}\hat{X}_1(s) &= \frac{t}{T}\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} + \frac{1}{2\kappa}\int_t^s \left(a(z)\frac{2z}{T} + \frac{2\kappa}{T}\right)\exp\left(-\frac{1}{\kappa}\int_z^s a(r)\,\mathrm{d}r\right)\mathrm{d}z \\ &= \frac{t}{T}\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} + \frac{s}{T} - \frac{t}{T}\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} = \frac{s}{T},\end{aligned}$$

by integration by parts. Finally, we check that $\hat{u}_1(t) = \hat{X}_1'(t) = \frac{1}{T}$ is the optimizer to (3.3). To this end, we calculate

$$\begin{aligned}&\kappa\int_t^T u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^T \left(\frac{s}{T} - X^u(s)\right)^2\mathrm{d}t \geq \kappa\int_t^T u^2(s)\,\mathrm{d}s \\ &\quad\geq \frac{\kappa}{T-t}\left(\int_t^T u(s)\,\mathrm{d}s\right)^2 = \frac{\kappa}{T-t}\left(1-\frac{t}{T}\right)^2 = \kappa\frac{T-t}{T^2},\end{aligned}$$

by Jensen's inequality using the probability measure $\frac{1}{T-t}\,\mathrm{d}s$. Equality holds for the choice $u = \hat{u}_1 = \frac{1}{T}$ with corresponding $X^u(s) = \hat{X}_1(s) = \frac{s}{T}$. This shows that $\hat{u}_1$ is indeed the minimizer of (3.3). □

REFERENCES

ALMGREN, R. (2012): Optimal Trading with Stochastic Liquidity and Volatility, *SIAM J. Finan. Math.* 3, 163–181.

ALMGREN, R., and N. CHRISS (2000): Optimal Execution of Portfolio Transactions, *J. Risk* 3, 5–39.

ALMGREN, R., C. THUM, E. HAUPTMANN, and H. LI (2005): Equity Market Impact, *Risk* 18, 57–62.

BERKOWITZ, S., D. LOGUE, and E. NOSER (1988): The Total Cost of Transactions on the NYSE, *J. Finance* 43, 97–112.

BERTSIMAS, D., and A. LO (1998): Optimal Control of Execution Costs, *J. Finan. Markets* 1, 1–50.

BIAŁKOWSKI, J., S. DAROLLES, and G. LE FOL (2008): Improving VWAP Strategies: A Dynamic Volume Approach, *J. Banking Finance* 32, 1709–1722.

BJÖRK, T., and A. MURGOCI (2010): A General Theory of Markovian Time Inconsistent Stochastic Control Problems, Preprint. Available at: SSRN http://ssrn.com/abstract=1694759. Accessed date March 2013.

BOUCHARD, B., and N.-M. DANG (2013): Generalized Stochastic Target Problems for the Pricing and Partial Hedging Under Loss Constraints—Application in Optimal Book Liquidation. *Finance Stoch.* 17(1), 31–72.

BOUCHAUD, J.-P., Y. GEFEN, M. POTTERS, and M. WYART (2004): Fluctuations and Response in Financial Markets: The Subtle Nature of "Random" Price Changes, *Quant. Finance* 4, 176–190.

BRODY, D., L. HUGHSTON, and A. MACRINA (2008): Dam Rain and Cumulative Gain, *Proc. R. Soc. A* 464, 1801–1822.

CONT, R., A. KUKANOV, and S. STOIKOV (2014): The Price Impact of Order Book Events. *J. Finan. Econ*. 12(1), 47–88.

ÉMERY, M., and M. YOR (2004): A Parallel between Brownian Bridges and Gamma Bridges, *Publ. RIMS, Kyoto Univ.* 40, 669–688.

GANI, J. (1957): Problems in the Probability Theory of Storage Systems, *J. R. Statist. Soc. B* 19, 181–206.

GATHERAL, J. (2010): No-Dynamic-Arbitrage and Market Impact, *Quant. Finance* 10, 749–759.

GATHERAL, J., and A. SCHIED (2011): Optimal Trade Execution under Geometric Brownian Motion in the Almgren and Chriss Framework, *Int. J. Theoretical Appl. Finance* 14, 353–368.

GATHERAL, J., and A. SCHIED (2013): Dynamical Models of Market Impact and Algorithms for Order Execution, in *Handbook on Systemic Risk*, J.-P. Fouque and J. Langsam, eds. Cambridge. UK: Cambridge University Press, pp. 579–602.

HARRIS, L. (1986): A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns, *J. Finan. Econ.* 16, 99–117.

HARRIS, L. (1987): Transaction Data Test of the Mixture of Distribution Hypothesis, *J. Finan. Quant. Anal.* 22, 127–141.

HU, G. (2007): VWAP Cost Excluding Own Trades, *J. Trading* 2, 30–34.

HUMPHERY-JENNER, M. (2011): Optimal VWAP Trading under Noisy Conditions, *J. Bank. Finance* 35, 2319–2329.

KAKADE, S., M. KEARNS, Y. MANSOUR, and L. ORTIZ (2004): Competitive Algorithms for VWAP and Limit Order Trading, in *Proceedings of the 5th ACM Conference on Electronic Commerce* May 17–20, 2004, New York, NY, pp. 189–198.

KALLENBERG, O. (2002): Foundations of Modern Probability, in *Probability and Its Applications*, 2nd ed., J. Gani, C.C. Heyde and T.G. Kurtz, eds. New York: Springer.

KONISHI, H. (2002): Optimal Slice of a VWAP Trade, *J. Finan. Markets* 5, 197–221.

LUKACS, E. (1955): A Characterization of the Gamma Distribution, *Ann. Math. Stat.* 26, 319–324.

MADHAVAN, A. (2002): VWAP Strategies, *Trading* 1, 32–39.

MCCULLOCH, J., and V. KAZAKOV (2012): Mean Variance Optimal VWAP Trading, Preprint. Available at SSRN: http://ssrn.com/abstract=1803858http://ssrn.com/abstract=1803858. Accessed date March 2013.

MORAN, P. (1956): A Probability Theory of a Dam with a Continuous Release, *Quart. J. Math. Oxf.* 7, 130–137.

WOOD, R., T. MCINISH, and J. ORD (1985): An Investigation of Transaction Data for NYSE Stocks, *J. Finance* 25, 723–739.

MATHEMATICAL FINANCE

An International Journal of Mathematics, Statistics and Financial Economics

Mathematical Finance, Vol. 25, No. 3 (July 2015), 640–672

OPTIMAL EXECUTION HORIZON

DAVID EASLEY

Cornell University

MARCOS LOPEZ DE PRADO

Hess Energy Trading Company and RCC at Harvard University

MAUREEN O'HARA

Johnson Graduate School of Management, Cornell University

Execution traders know that market impact greatly depends on whether their orders lean with or against the market. We introduce the OEH model, which incorporates this fact when determining the optimal trading horizon for an order, an input required by many sophisticated execution strategies. This model exploits the trader's private information about her trade's side and size, and how it will shift the prevailing order flow. From a theoretical perspective, OEH explains why market participants may rationally "dump" their orders in an increasingly illiquid market. We argue that trade side and order imbalance are key variables needed for modeling market impact functions, and their dismissal may be the reason behind the apparent disagreement in the literature regarding the functional form of the market impact function. We show that in terms of its information ratio OEH performs better than participation rate schemes and VWAP strategies. Our backtests suggest that OEH contributes substantial "execution alpha" for a wide variety of futures contracts. An implementation of OEH is provided in Python language.

KEY WORDS: liquidity, flow toxicity, broker, VWAP, market microstructure, adverse selection, probability of informed trading, VPIN, OEH.

1. INTRODUCTION

Optimal execution strategies compute a trajectory that minimizes the shortfall cost of acquiring or disposing of a position in an asset. Well-known contributions to this subject are Perold (1998), Bertsimas and Lo (1998), Almgren and Chriss (2000), and Kissell and Glantz (2003), to cite only a few. These strategies provide an abstract and general framework to model the costs that trading imposes on a particular investor. One drawback of this abstraction is that it does not explicitly define how market impact arises from a trade's perturbation of the liquidity provision process. This omission is important as the market impact of orders generally depends on whether the orders lean with or against the market. We use asymmetric information market microstructure theory to take this feature of the provision of liquidity into account in our determination of optimal execution. A trader has private information about the size and direction of his upcoming trade.

Manuscript received February 2013; final revision received August 2012.

Address correspondence to Maureen O'Hara, Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853; e-mail: mo19@cornell.edu.

DOI: 10.1111/mafi.12045

If, for example, he plans to buy in a selling market, his trade actually provides liquidity to the market and can be executed at a lower trading cost than a sell in a similar market. His execution strategy should take this private information into account.

An alternative approach to modeling the liquidity provision process is taken by Obizhaeva and Wang (2013), which models the impact of trades on the limit order book and evolution of the book. Our approaches differ in how the impact of orders on liquidity is modeled. Our approach is based on an asymmetric information model of the behavior of market maker; their approach is based on a model of the dynamics of orders arriving to the book. The resulting optimal execution strategies differ as well.

The PIN theory (Easley et al. 1996) argues that market makers adjust the range at which they are willing to provide liquidity based on their estimates of the probability of being adversely selected by informed traders. Easley, López de Prado, and O'Hara (2012) show that, in high frequency markets, this probability can be accurately approximated as a function of the absolute order imbalance (absolute value of buy volume minus sell volume). Suppose that we are interested in selling a large amount of E-mini S&P500 futures in a market that is imbalanced toward sells. Because our order is leaning with previous orders, it reinforces market makers' fears that they are being adversely selected, and that their current (long) inventory will be harder to liquidate without incurring a loss. As a result, market makers will further widen the range at which they are willing to provide liquidity, increasing our order's market impact. Alternatively, if our order were on the buy side, market makers would narrow their trading range (the spread between the bid and ask prices) as they believe that their chances of turning over their inventory improve, in which case we would experience a lower market impact than with a sell order. Thus, order imbalance and market maker behavior set the stage for understanding how orders fare in terms of execution costs. Our goal in this paper is to apply this theory in the context of market impact and execution strategies.

Besides reducing transaction costs, there are a number of reasons why traders care about not increasing the order imbalance. First, Federal Regulation and Exchange Circulars limit a trader's ability to disrupt or manipulate market activity.[1] A trading strategy that disrupts the market can bring fines, restrictions, and sanctions.[2] Second, traders often will have to come back to the market shortly after completing the initial trade. If a trader got great fills in the previous trade at the expense of disrupting the liquidity process, the previous trade's gains may transform into losses on the successive trades.[3] Third, the position acquired will be marked-to-market, so posttrade liquidity conditions will be reflected in the unrealized P&L of the position.[4] Thus, it would be useful to determine the amount of volume needed to "conceal" a trade so that it leaves a minimum footprint on the trading range.

Many optimal execution strategies make the assumption that liquidity cost per share traded is a linear function of trading rate or of the size of the block to be traded. This

[1] We could envision a future in which regulators and exchanges limit the assets under management of an investment firm based on that market participant's technology as it relates to the disruption of the liquidity provision process.

[2] See, for example, the "stupid algo" rules recently introduced by the Deutsche Börse (discussed in *"Superfast traders feel heat as Bourses act,"* Financial Times, March 5, 2012). The CFTC and SEC have also recently introduced explicit rules relating to algorithmic impact on markets.

[3] Even if there are no successive trades, leaving a footprint leaks information that can be recovered by competitors and used against that trader in future occasions. See Easley et al. (2012c) for examples.

[4] If a buyer pushes the mid-price up at the expense of draining liquidity, she may find that the liquidation value of that position implies a loss.

seems an unrealistic assumption and it is the motivation for Almgren's (2003) study of optimal execution. The model we introduce here also incorporates a nonlinear response in trading cost to size of the trade. We are not aware of other methodologies explicitly taking into account the presence of asymmetric information in determining the optimal execution horizon, although Obizhaeva and Wang (2013) do analyze optimal execution using a model of dynamic limit order book. In this context, we find that the side of the trade is as important as its size.

Our **O**ptimal **E**xecution **H**orizon model, which for brevity we denote OEH henceforth, is related to a growing number of recent studies concerned with execution in the context of high frequency markets and tactical liquidity provision. Schied and Schöneborn (2009) point out that the speed by which the remaining asset position is sold can be decreasing in the size of the position, but increasing in the liquidity price impact. Gatheral and Schied (2011) extend the Almgren–Chriss framework to an alternative choice of the risk criterion. Forsyth (2011) formulates the trading problem as an optimal stochastic control problem, where the objective is to maximize the mean-variance tradeoff as measured at the initial time. Hendershott and Mendelson (2000), Donges and Heinemann (2006), and Ye (2011) consider the choice of how to split an order between a dark pool and a primary venue. Bayraktar and Ludkovski (2011) and Kratz and Schöneborn (2014) propose optimal trade execution schemes for dark pools.

OEH does not replace or supersede minimum market impact strategies. On the contrary, it is complementary to them. OEH does not address the question of how to slice the orders (create a "trading schedule"), which has been studied by Hasbrouck and Schwartz (1988), Berkowitz et al. (1988), Grinold and Kahn (1999), Konishi and Makimoto (2001) among others previously cited. Our main concern is with understanding how transaction costs are derived from the impact that a trade has on liquidity providers. There are some connections, however, between OEH and the previous studies. Like earlier models, OEH minimizes the impact on liquidity subject to a timing risk. From the standard VWAP (Madhavan 2002) to sophisticated nonlinear approaches (Almgren 2003; Dai and Zhong 2012), many execution strategies require as an input the *trading horizon*, which is typically exogenous to those models. One of our contributions is to provide a framework for determining this critical input variable from a new perspective: *Informational leakage*. In doing so, we stress the importance of modeling the order side and its asymmetric impact on the order imbalance.

The current paper is organized as follows. Section 2 reviews how order imbalance and trading range are related. Section 3 explains why our trading actions leave a footprint on the market makers' trading range. Section 4 incorporates timing risk and risk aversion to our analysis. Section 5 develops the OEH algorithm that determines the optimal horizon over which to execute a trade. Section 6 presents three numerical examples, illustrating several stylized cases. Section 7 argues why the apparently irrational behavior of some market participants during the "*flash crash*" may be explained by OEH. Section 8 compares OEH's performance with that of trading schemes that target a volume participation rate. Section 9 discusses how our model relates to alternative functional forms of the market impact function. Section 10 provides backtests of the performance of OEH, and compares it with a VWAP execution strategy. Section 11 summarizes our conclusions. Appendix 1 provides a proof. Appendix 2 offers a generalization of our approach. Appendix 3 contains an implementation of Appendices 1 and 2 in Python language. Appendix 4 provides a procedure for estimating market imbalance.

2. TRADING RANGE AND FLOW TOXICITY

We begin by summarizing a standard sequential-trade market microstructure approach to determining the trading range for an asset. In a series of papers, Easley et al. (1992a, 1992b, 1996) demonstrate how a microstructure model can be estimated for individual assets using trade data to determine the probability of information-based trading, PIN. This microstructure model views trading as a game between liquidity providers and traders (position takers) that is repeated over trading periods $i = 1, \ldots, I$. At the beginning of each period, nature chooses whether an information event occurs. These events occur independently with probability α. If the information is good news, then informed traders know that by the end of the trading period the asset will be worth $\bar{S}$; and, if the information is bad news, that it will be worth $\underline{S}$, with $\bar{S} > \underline{S}$. Good news occurs with probability $(1-\delta)$ and bad news occurs with probability δ. After an information event occurs or does not occur, trading for the period begins with traders arriving according to Poisson processes. During periods with an information event, orders from informed traders arrive at rate μ. These informed traders buy if they have seen good news, and sell if they have seen bad news. Every period, orders from uninformed buyers and uninformed sellers each arrive at rate ε.

Easley, Kiefer, O'Hara, and Paperman (1996) argue that, for the natural case that $\delta = \frac{1}{2}$, the trading range at which market makers are willing to provide liquidity is

$$\Sigma = \frac{\alpha\mu}{\alpha\mu + 2\varepsilon}[\bar{S} - \underline{S}]. \tag{2.1}$$

This trading range gives the market maker's targeted profit per portfolio turnover. The first term in equation (2.1) is the probability that a trade is information based. It is known as PIN. Easley, Engle, O'Hara, and Wu (2008) show that expected absolute trade imbalance can be used to approximate the numerator of PIN. They demonstrate that, letting V^S and V^B represent sell and buy volume, respectively, $E[|V^S - V^B|] \approx \alpha\mu$ for a sufficiently large μ. Easley, López de Prado, and O'Hara [2011a, 2012] argue that in volume-time space, PIN can be approximated as[5]

$$PIN \equiv \frac{\alpha\mu}{\alpha\mu + 2\varepsilon} \approx VPIN \equiv E[|OI|] \tag{2.2}$$

where we have grouped trades in equal volume buckets of size $V^B + V^S = V = \alpha\mu + 2\varepsilon$, and $OI = \frac{V^B - V^S}{V}$ represents the order imbalance within V.[6] This volume-time approximation of PIN, known as VPIN, has been found to be useful in a number of settings (see Easley, López de Prado, and O'Hara 2012; Bethel et al. 2011; Abad and Yagüe 2012; or Menkveld and Yueshen 2013, for example). In the next section we will show how VPIN's expectations play a role in modeling the liquidity component of an execution strategy.

3. THE LIQUIDITY COMPONENT

For the reasons argued in the introduction, traders are mindful of the footprint that their actions leave on the trading range, Σ. According to equation (2.2) the trading range

[5]The expectation of the absolute value of order imbalance, $E[|V^B - V^S|]$ is approximated by the expected arrival rate of informed trades, $\alpha\mu$. This approximation is discussed in Easley, Engle, O'Hara, and Wu (2008).

[6]For a procedure that can be used to estimate OI, see Easley, López de Prado, and O'Hara (2012b).

will change due to adjustments on VPIN expectations. In the next two subsections we discuss the impact that our trade has on the order imbalance, how that may affect VPIN expectations and consequently the transaction cost's liquidity component.

3.1. Impacting the Order Imbalance

Suppose that we are an aggressive trader for m contracts to be traded in the next volume bucket, where $m > 0$ denotes a purchase, and $m < 0$ denotes a sale.[7] For now, we set the bucket size at an exogenously given V (we explain in Section 5 how to optimize this bucket size). Our trade of m contracts will cause the bucket to fill faster than it would otherwise fill. Let us suppose that our trade pushes buys and sells from other traders out of the next bucket at equal rates. Given our knowledge of OI, we let $\widetilde{V^B}$ and $\widetilde{V^S}$ be our forecasts of buy and sell volume that would occur in the next bucket without our trade.[8] The expected order imbalance over V using our private information about our trade of m contracts can be represented by

$$\widehat{OI} \equiv \frac{\widehat{V^B} - \widehat{V^S}}{V} = \frac{\frac{\widetilde{V^B}}{V}(V - |m|) - \frac{\widetilde{V^S}}{V}(V - |m|) + m}{V} = \underbrace{(2v^B - 1)}_{OI}\left(1 - \frac{|m|}{V}\right) + \frac{m}{V} \tag{3.1}$$

where $V \geq m$, $v^B \equiv \frac{V^B}{V}$ represents the forecasted fraction of buy volume in absence of our trade and $OI \equiv 2v^B - 1$ is the order imbalance in absence of our trade.[9] The second equality follows from observing that the predicted fraction of buys from others, v^B, times the volume from others, $V - |m|$, is their predicted buy volume; a similar term occurs for their predicted sell volume; and finally adding in our trade gives the predicted order imbalance. In the nomenclature above, $\widehat{OI}$ incorporates private information and OI only public information. If $\frac{|m|}{V} \approx 0$, $\widehat{OI} \approx OI$. Alternatively, in the extreme case of $\frac{|m|}{V} \approx 1$, we have $|\widehat{OI}| \approx 1$. Generally, the impact of our trade on $\widehat{OI}$ will depend on m, but also the OI that would otherwise occur in the next volume bucket. If, for instance, $V^B < V^S$, a trade of size $m = \frac{(\widetilde{V^S} - \widetilde{V^B})V}{\widetilde{V^S} - \widetilde{V^B} + V}$ will make $\widehat{OI} = 0$.

3.2. Informational Leakage on the Liquidity Component

The previous section explained how trading m contracts interspersed with $V - m$ external volume displaces the expected order imbalance from OI to $\widehat{OI}$. Next, we discuss how that displacement triggers an updating of the market makers' expectations on the order imbalance, and thus VPIN.

Market makers adjust their estimates of VPIN as a result of the information leaked during the trading process. *Ceteris paribus*, if $|m|$ is relatively small, we expect it to convey little information to market makers. For example, if a small trade is executed

[7] Of course, m must be smaller than or equal to V in order for us to do it within one volume bucket. The meaning of "aggressive trader" is that we determine the timing of the trade (as opposed to being a "passive trader").

[8] Easley et al. (2012a) present evidence that order imbalance shows persistence over (volume) time.

[9] For a procedure that can be used to estimate v^B, see Easley, López de Prado, and O'Hara (2012b).

in block but it is however large enough to fill a small bucket of size V, market makers may expect VPIN to remain around forecasted levels rather than jump to 1. We model the expected order imbalance (leaked to market makers) during execution, $\widetilde{OI}$, as a convex combination of two extreme outcomes—no leakage ($\varphi[|m|] \to 0$) and complete ($\varphi[|m|] \to 1$) informational leakage:

$$\widetilde{OI} \equiv \varphi[|m|] \underbrace{\left[(2v^B - 1)\left(1 - \frac{|m|}{V}\right) + \frac{m}{V} \right]}_{\widehat{OI}} + (1 - \varphi[|m|]) \underbrace{(2v^B - 1)}_{OI} \tag{3.2}$$

where $\varphi[.]$ is a monotonic increasing function in $|m|$ with real values in the range (0, 1). $\widetilde{OI}$ contains private information to the extent that it has been leaked by m. That occurs as a result of the trade's size ($|m|$). The role of $\varphi[.]$ is to determine the degree by which the effective order imbalance during our execution ($\widehat{OI}$) impacts the market markers' expectations on VPIN, and consequently leads to an adjustment of the range at which they provide liquidity, Σ.

It is critical to understand that the privately known order imbalance ($\widehat{OI}$) may differ from the one inferred by the market makers ($\widetilde{OI}$) due to the trade's "footprint." Some private information may have been leaked by m, but not all. If $\varphi[|m|] \to 0$, there is no informational leakage, and $\widetilde{OI} \approx OI$, regardless of $\widehat{OI}$. In this case, it is as if m had not been traded and market makers will not adjust their behavior. However, when the order is so large that $\varphi[|m|] \to 1$, the leakage is complete and the market maker knows as much as the trader, $\widetilde{OI} \approx \widehat{OI}$. For instance, an order of 75,000 contracts on the E-mini S&P500 futures may dramatically displace the expectation of VPIN, even if blended among a volume 10 times greater (see SEC-CFTC [2010] in connection with the Waddell & Reed order). This displacement of the expectation from OI to $\widetilde{OI}$ is the *footprint* left by m in the liquidity provision process.

The extent to which private information about our trade leaks to the market during the volume bucket in which we are trading is critical. In our examples and empirical implementation, we consider both complete leakage, $\varphi[|m|] \to 1$, and leakage that is a simple function of the size of our trade. One could also consider more complex forms for the leakage function; for example, it could depend on size of the volume bucket or which way the market is leaning without our trade and the extent to which it is imbalanced. We do not consider those extensions here. They would not change our conceptual framework, but the extension to including the bucket size in the leakage function would certainly change the details of the solution to the OEH.

For simplicity, we have assumed that $E[|OI|] = |OI|$, because this expectation is based on past, public information. Similarly, $E[|\varphi[|m|]\widehat{OI}|] = \varphi[|m|]|\widehat{OI}|$ because $\varphi[|m|]\widehat{OI}$ is precisely the portion of $\widehat{OI}$ that has become public due to m's leakage. In this particular model we have not considered other sources of uncertainty in forming VPIN's expectation; thus $\widetilde{VPIN} \equiv E[|\widetilde{OI}|] = |\widetilde{OI}|$.

4. THE TIMING RISK COMPONENT

Minimizing the footprint that trading has on Σ is an important factor when deciding the execution horizon. But it is not the only factor that matters to a trader. Minimal impact may require slicing a large order into multiple small trades, resulting in a long

delay in executing the intended trade. This delay involves a risk, called *timing risk*. To model this timing risk, we introduce a simple, standard model of how the midpoint of the spread evolves. It is useful to view the trading range, Σ, as being centered on the mid-price, S, which moves stochastically as (volume) time passes. We model this process as an exogenous arithmetic random walk

$$\Delta S \equiv \hat{\sigma}\sqrt{\frac{V}{V_\sigma}}\xi \tag{4.1}$$

where $\xi \sim N(0, 1)$ i.i.d. and V_σ is the volume used to compute each mid-price change. So V/V_σ is the number of price changes sampled in a bucket of size V. The standard deviation, σ, of price changes is unknown, but it can be estimated from a sample of n equal-volume buckets as $\hat{\sigma}$, in which case $\frac{(n-1)\hat{\sigma}^2}{\sigma^2} \sim \chi^2_{n-1}$.

According to this specification, the market loss from a trade of size m can be probabilistically bounded as

$$P\left[Sgn\,(m)\,\Delta S > Z_\lambda \hat{\sigma}\sqrt{\frac{V}{V_\sigma}}\right] = 1 - \lambda \tag{4.2}$$

where $Sgn(m)$ is the sign of the trade m, λ is the probability with which we are willing to accept a loss greater than $Z_\lambda \hat{\sigma}\sqrt{\frac{V}{V_\sigma}}$, and Z_λ is the critical value from a standard normal distribution associated with $\lambda \in (0, \frac{1}{2}]$.[10] λ can also be interpreted as a *risk aversion* parameter, as it modulates the relative importance that we give to timing risk.

The specification above abstracts from any direct impact that m may have on the midpoint price process. As is well known from microstructure research, the private information in our trade can introduce permanent effects on prices. In Appendix 2, we provide a generalization of this component taking into account the possibility that m leaks part of our private information into the mid-price.

5. OPTIMAL EXECUTION HORIZON

We have argued that the impact of our trade of m on Σ will depend on the size of m relative to V. In this section we take the intended trade, m, as fixed and determine the optimal trading volume V in which to hide our trading action m without incurring excessive timing risk. In order to compute this quantity, we define a *probabilistic loss* Π which incorporates the liquidity and timing risk components discussed earlier. The probabilistic loss models how different execution horizons will impact our per unit portfolio valuation after the trade's completion. This is not an implementation cost, as that is contingent upon the execution model adopted to slice the parent order into child orders. OEH's goal is to determine the V^* that optimally conceals (in the sense of minimizing Π for a given risk aversion λ) our order m, given the prevalent market state ($\varphi[|m|]$, v^B, $[\bar{S} - \underline{S}]$, V_σ, $\hat{\sigma}$),

[10]Note that at $\lambda = 0.5$ the critical value changes sign. We restrict our attention to $\lambda \leq 0.5$ as otherwise we would be mixing losses and gains.

where

$$\Pi \equiv \underbrace{\left|\varphi[|m|]\left[(2v^B-1)\left(1-\frac{|m|}{V}\right)+\frac{m}{V}\right]+(1-\varphi[|m|])(2v^B-1)\right|\left[\bar{S}-\underline{S}\right]}_{\textit{liquidity component}} \underbrace{-Z\lambda\sqrt{\frac{V}{V_\sigma}}\hat{\sigma}}_{\textit{timing risk component}} \tag{5.1}$$

with $Z_\lambda \leq 0$, thus the subtraction. The greater V, the smaller the impact on the order imbalance, but also the larger the possible change in the center of the trading range. Appendix 1 demonstrates the solution to minimizing Π, which can be implemented through the following algorithm (see Appendix 3 for an implementation in Python language):

(1) If $(2v^B-1)|m| < m$, try $V_1 = (2\varphi[|m|][(2v^B-1)|m|-m][\bar{S}-\underline{S}]\frac{\sqrt{V_\sigma}}{Z_\lambda\hat{\sigma}})^{2/3}$ and compute the value of $\widetilde{OI}$ associated with V_1, $\widetilde{OI}[V_1]$.
   (a) If $\widetilde{OI}[V_1] > 0$ and $v^B \leq v^{B+}$, then $V^* = V_1$ is the solution.
   (b) If $\widetilde{OI}[V_1] > 0$ and $v^B > v^{B+}$, then $V^* = |m|$ is the solution.

(2) If $(2v^B-1)|m| > m$, try $V_2 = (2\varphi[|m|][m-(2v^B-1)|m|][\bar{S}-\underline{S}]\frac{\sqrt{V_\sigma}}{Z_\lambda\hat{\sigma}})^{2/3}$ and compute the value of $\widetilde{OI}$ associated with V_2, $\widetilde{OI}[V_2]$.
   (a) If $\widetilde{OI}[V_2] < 0$ and $v^B \geq v^{B-}$, then $V^* = V_2$ is the solution.
   (b) If $\widetilde{OI}[V_2] < 0$ and $v^B < v^{B-}$, then $V^* = |m|$ is the solution.

(3) If $(2v^B-1)|m| = m$, then $V_3 = |m|$ is the solution.
(4) Else, try $V_4 = \varphi[|m|](|m| - \frac{m}{2v^B-1})$.
   (a) If $v^B \geq v^{B=}$, then $V^* = V_4$ is the solution.
   (b) If $v^B < v^{B=}$, then $V^* = |m|$ is the solution.

where

$$\widetilde{OI}[V^*] = \varphi[|m|]\left[\frac{m-(2v^B-1)|m|}{V^*}+2v^B-1\right]+(1-\varphi[|m|])(2v^B-1) \tag{5.8}$$

$$v^{B+} = \frac{\frac{|m|}{m}+1}{2}+\frac{Z_\lambda\hat{\sigma}\sqrt{|m|}}{4\varphi[|m|][\bar{S}-\underline{S}]\sqrt{V_\sigma}} \tag{5.9}$$

$$v^{B-} = \frac{\frac{|m|}{m}+1}{2}-\frac{Z_\lambda\hat{\sigma}\sqrt{|m|}}{4\varphi[|m|][\bar{S}-\underline{S}]\sqrt{V_\sigma}} \tag{5.10}$$

$$v^{B=} = \frac{1}{2}\left(\frac{|m|\varphi[|m|]}{m(\varphi[|m|]-1)}+1\right) \tag{5.11}$$

$\widetilde{OI} \in [-1, 1]$ is the signed order imbalanced as a result of the informational leak that comes with trading m. v^{B-}, v^{B+}, and $v^{B=}$ set the boundaries for v^B in order to meet the constraint that $V^* \geq |m|$. The liquidity component is nonlinear in m and takes into account the side of the trade (leaning against or with the market), which we illustrate with numerical examples in Section 6.

Note that V^* has been optimized to minimize Π, and it will generally differ from the V used to compute VPIN. We reiterate the point made earlier that various uses of VPIN require different calibration procedures. In this paper, we propose a method for determining the V^* that minimizes the probabilistic loss Π, not the V that maximizes our forecasting power of v^B or short-term toxicity-induced volatility.

In practice, V^* could be estimated once before submitting the first child order and reestimated again before submitting every new slice. By doing so, we incorporate the market's feedback into the model: If our initial slices have a greater (or smaller) impact than expected on the liquidity provision, we could adjust in real-time. That is, we assume that the values observed will remain constant through the end of the liquidation period, and determine the statically optimal strategy using those values. As the input parameters change, we could recompute the stationary solution. Almgren's (2009) shows that this *"rolling horizon"* approach, although not dynamically optimal, provides reasonably good solutions to a related problem.

6. NUMERICAL EXAMPLES

In this section, we develop numerical examples using the following state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] \to 1$. We want to find the optimal horizon to buy our desired amount of 1,000, recognizing that the optimal strategy must take account of both liquidity costs and timing risk costs. This strategy will depend, in part, on the imbalance in the market, which we capture as the fraction of buy volume, v^B. Three scenarios appear relevant: $v^B < \frac{1}{2}$, $v^B = \frac{1}{2}$ and $v^B > \frac{1}{2}$. We limit our attention to $m > 0$, with the understanding that a symmetric outcome would be attained with an alternative sell order $m' < 0$ and $v^{B\prime} = 1 - v^B$ (see Figure 7.1).

Figure 6.1(a) displays the optimal volume horizons (V^*) for various values of the fraction of buy volume, v^B. As is apparent, the optimal trading horizon differs dramatically with imbalance in the market. The reason for this is illustrated in Figure 6.1(b), which demonstrates how the probabilistic loss of trading 1,000 shares is also a function of the buy imbalance v^B. The probabilistic loss is the sum of the loss from the liquidity component and the loss from the timing component. Notice that with $m > 0$, the liquidity component does not contribute to Π until shortly before the market is balanced ($v^B = \frac{1}{2}$).[11]

6.1. Scenario I: $v^B = 0.4$

In this scenario we want to buy 1,000 contracts from a market that we believe to be slightly tilted toward sales (projected buys are only 40% of the total volume). If we plot the values of $\Pi[V, .]$ for different levels of V, we obtain Figure 6.2.

[11]The liquidity component is positive at $v^B < \frac{1}{2}$ as the order $m = 1{,}000$ makes the market unbalanced toward buys for $v^B < \frac{1}{2}$.

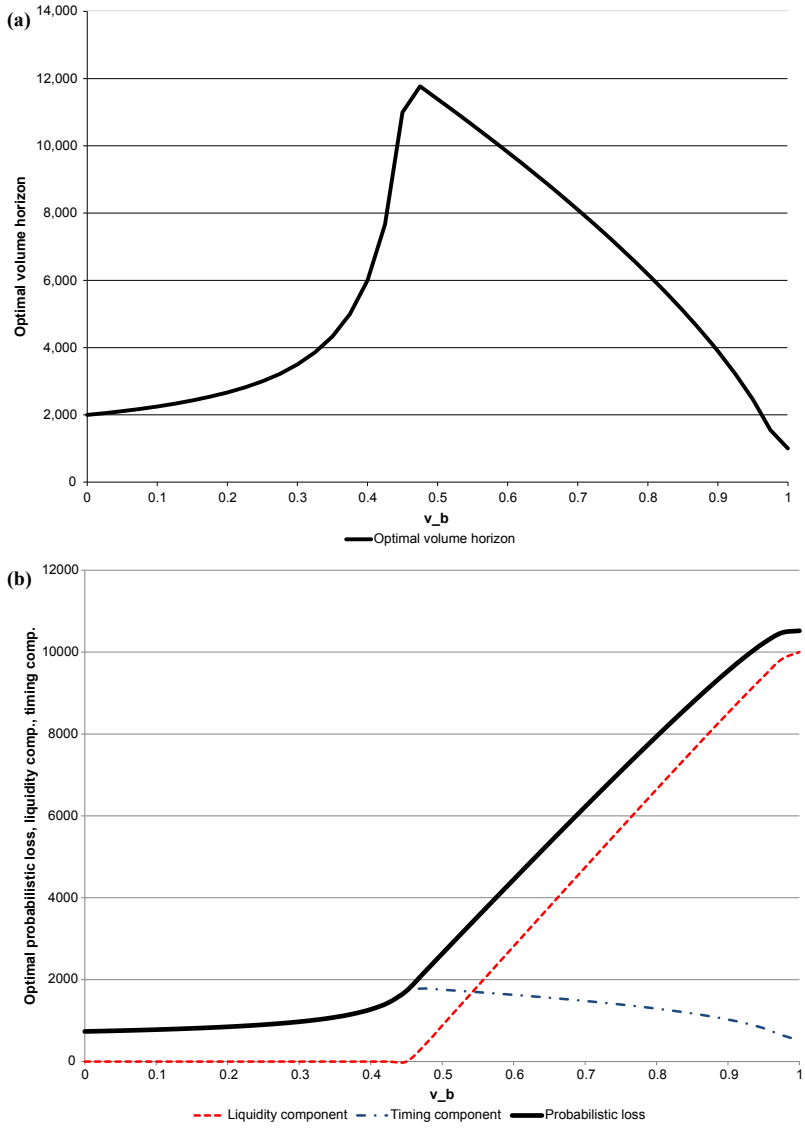

FIGURE 6.1. (a) Optimal V for various v^B on a buy order. This figure demonstrates how the optimal trading horizon for a buy order depends upon the expected fraction of buy orders in the market. When all orders are buys, v^B is 1, while if all orders are sells then v^B is 0. The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\widehat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$. (b) Optimal probabilistic loss and its components for various v^B on a buy order. This figure shows the expected total trading loss for a buy order and its relation to market imbalance. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The expected market imbalance is given by the expected fraction of buy orders in the market. When all orders are buys, v^B is 1, while if all orders are sells v^B is 0. The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$.

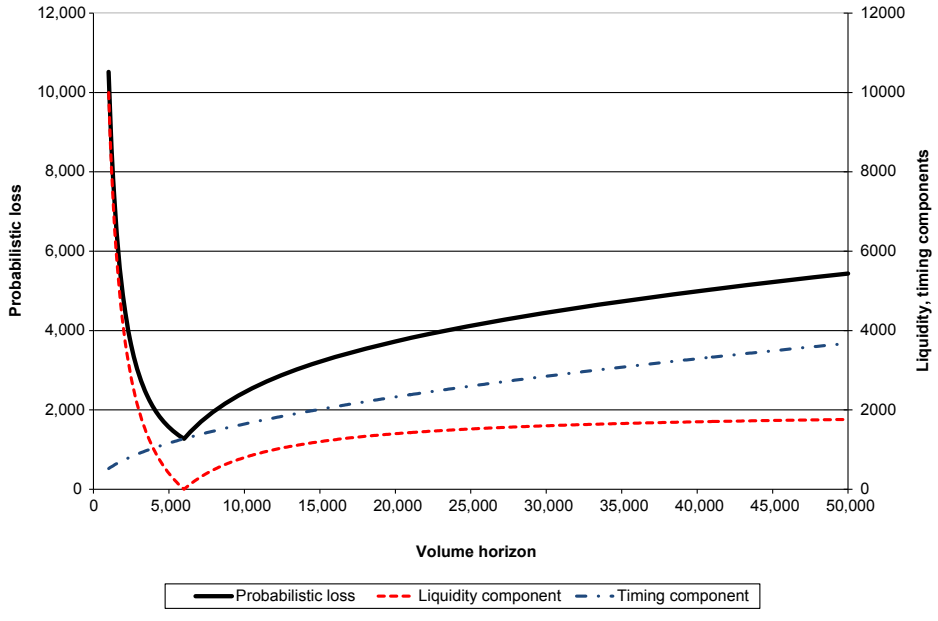

FIGURE 6.2. $\Pi(V, .)$ for different volume horizons (V), with $v^B = 0.4$. This figure demonstrates the expected total trading loss for a buy order and its relation to the volume horizon when the market is expected to have more selling activity. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$.

In this case $\widetilde{OI}[V^*] = 0$ at the optimum trading horizon of $V^* = 6,000$ contracts (of which 5,000 come from other market participants). The optimum occurs at the inflexion point where the liquidity component function changes from convex decreasing to concave increasing.[12] This may seem counter-intuitive as it is natural to expect the liquidity component to be decreasing in V. But this reasoning misses the important role played by market imbalance. Because we are buying in a selling market, there is a $v^B = 0.6$ at which we are narrowing Σ to the minimum possible and our liquidity cost is zero. Once we pass that liquidity-component optimal V^*, the trading range Σ necessarily widens again. The only way this can happen is with an increasing concave section in the liquidity component function, which explains the appearance of the inflexion point.

If we trade the desired 1,000 contracts within less than 6,000 total contracts traded our loss Π will increase because Σ increases more rapidly than the timing component of loss declines. If, on the other hand, we trade those 1,000 while more than 6,000 contracts are traded, our loss Π will increase because we will be taking on both excessive liquidity and timing costs.

One question to consider is, why does OEH allow the purchase to occur in the presence of selling flow which may result in lower future prices? Or put differently, why is it not

[12]Note that in this case the timing risk plays no role as the kink in the liquidity component overwhelms the curvature of the timing risk component. So, in this example, the value of λ is not important.

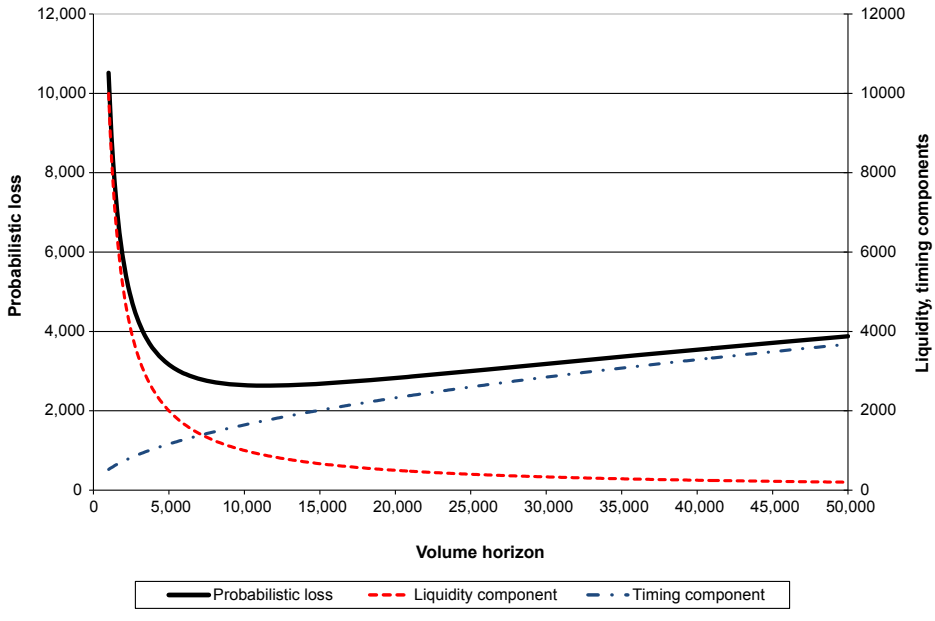

FIGURE 6.3. $\Pi(V, .)$ for different volume horizons (V), with $v^B = 0.5$. This figure shows the expected total trading loss for a buy order and its relation to the volume horizon when the market is expected to be in a balanced state. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$.

optimal to have a larger, possibly infinite execution horizon for a buy order, as long as $v^B < 0.5$? After all, prices may go lower as a result of the selling pressure, and that would give the trader a chance to buy at a better level. The reason is, this is an execution model, not an investment strategy. *The portfolio manager has decided that the trade must occur as soon as liquidity conditions allow it.* It is not the prerogative of OEH to speculate on the appropriateness of the trader's decision, which may be motivated for a variety of reasons (she holds private information, it is part of her asset management mandate, she must obey a stop loss or abide by a risk limit, a release is about to occur, etc.). The role of the OEH model is merely to determine the execution horizon that minimizes the probabilistic trading loss.

6.2. Scenario II: $v^B = 0.5$

In this scenario, the market is expected to be balanced (buys are 50% of the total volume). If we plot the values of $\Pi[V, .]$ for different levels of V, we obtain Figure 6.3.

The optimum now occurs at $V^* = 11{,}392$ contracts, with a value for $\widetilde{OI}[V^*] = 0.088$. The model recognizes that a larger volume horizon is needed to place a buy order in an already balanced market than in a market leaning against our order (note the change in shape of the liquidity component). In this scenario, our order does not narrow Σ

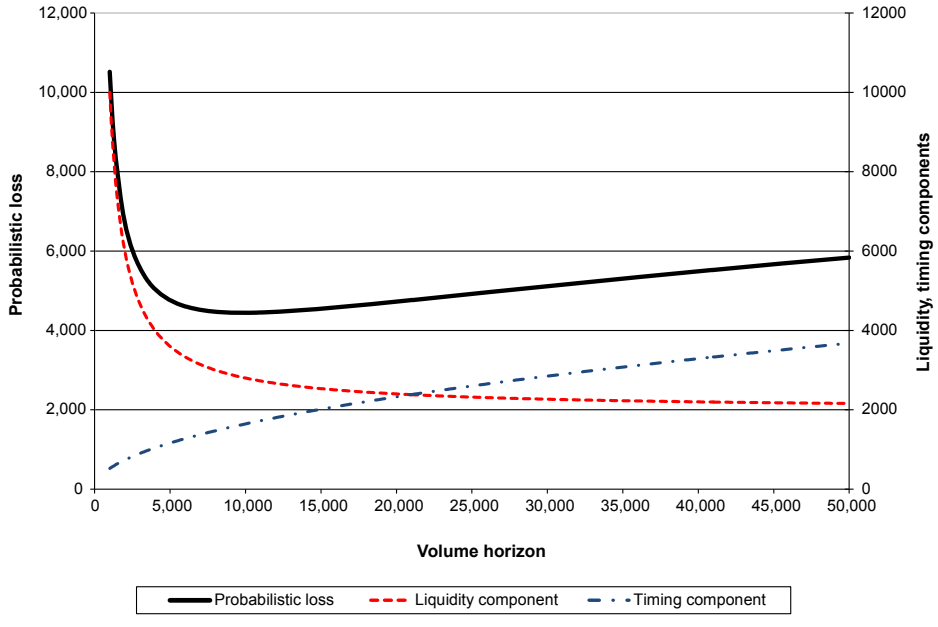

FIGURE 6.4. $\Pi(V, .)$ for different volume horizons (V), with $v^B = 0.6$. This figure demonstrates the expected total trading loss for a buy order and its relation to the volume horizon when the market is expected to have a greater buy imbalance. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$.

abruptly (as in Scenario I), and the only way to reduce Π is by substantially increasing V (i.e., concealing our order among greater market volume). The liquidity component function is now convex decreasing, without an inflexion point, because the market is not leaning against us. But the optimal V^* is still limited because of greater timing risk with increasing V.

6.3. Scenario III: $v^B = 0.6$

The market is now expected to be tilted toward buys, which represent 60% of the total volume. If we plot the values of $\Pi[V, .]$ for different levels of V, we obtain Figure 6.4.

The optimum now occurs at $V^* = 9,817$ contracts, with a value for $\widetilde{OI}[V^*] = 0.2816$. Two forces contribute to this outcome: On one hand, we are leaning with the market, which means that we are competing for liquidity ($v^B > \frac{1}{2}$), and we need a larger volume horizon than in Scenario I. On the other hand, the gains from narrowing Σ are offset by the additional timing risk, and Π eventually cannot be improved further. The relative strength of these two forces determines the optimal volume horizon. Note that this need not be a longer horizon than in a balanced market. In this example, the equilibrium between these two forces is reached at 9,817 contracts, a volume horizon between those obtained in Scenarios I and II.

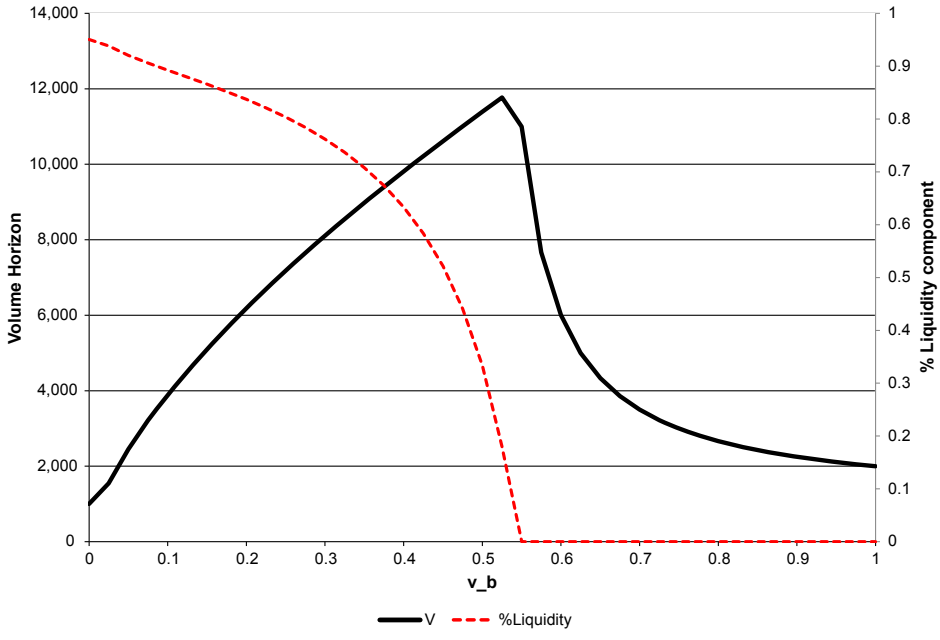

FIGURE 7.1. Optimal V for various v^B on a sell order. This figure shows how the optimal trading horizon for a sell order depends upon the expected order imbalance in the market. When all orders are buys (sells), v^B is 1, while if all orders are sells v^B is 0. The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$.

7. EXECUTION HORIZONS AND THE OSCILLATORY NATURE OF PRICES

Figure 6.1 illustrates how traders behave as predicted toxicity increases. For buyers in a market tilted toward buys, the execution horizon is a decreasing function of v^B. For sellers in a market tilted toward sells, the execution horizon is also a decreasing function. Figure 7.1 depicts the optimal horizon for a sell order. Note that, in this case, v^S decreases as we move from left to right in the graph. For both buyers and sellers, therefore, their optimal trading horizon will be influenced by the toxicity expected in the market.

We next consider the combined impact of alternative trade sizes, sides, and v^B. For simplicity, we assume that $\varphi[|m|]$ is linear in $|m|$, with $\varphi[|m|] \equiv \frac{m}{10^3}$ for $m \in (0, 10^3)$. The largest execution horizons occur in a perfectly balanced market, due to the inherent difficulty of concealing trading intentions. Again we observe in Figure 7.2 that *leaning against the market* (selling in a buying market or buying in a selling market) allows for shorter execution horizons, thanks to the possibility of achieving zero liquidity cost. *Leaning with the market* (selling in a selling market or buying in a buying market) leads to horizons that are longer than when the trade leans against the market, due to the fact that our trade is taking liquidity. The key point is that we have private information about

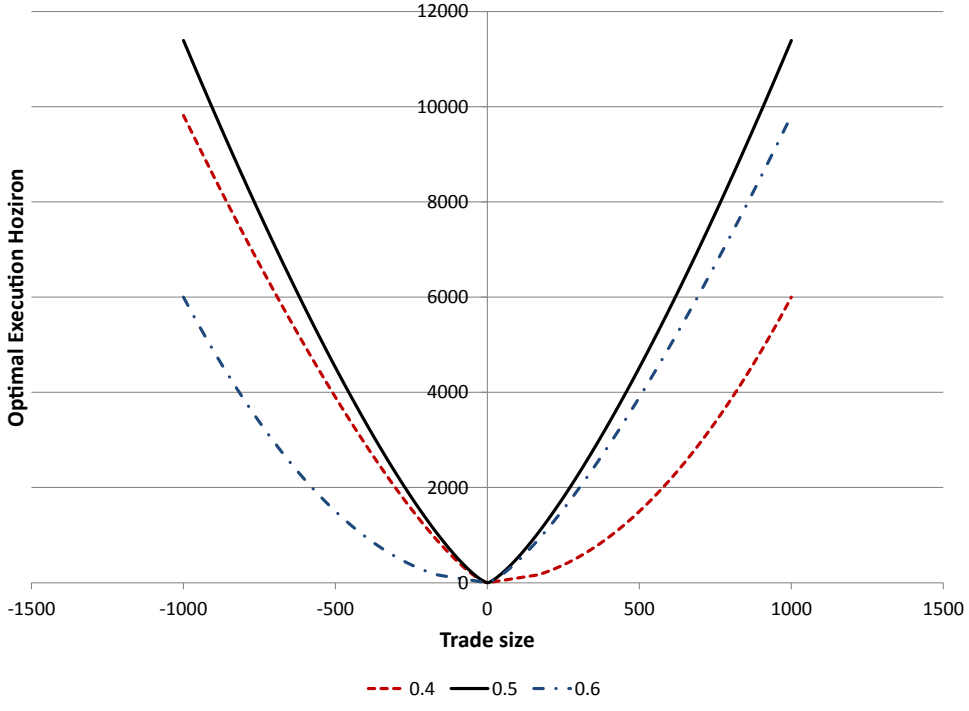

FIGURE 7.2. Optimal execution horizons for various order imbalances and trade sizes/sides. Combining alternative trade sizes and sides with our three scenarios ($v^B = 0.4$, $v^B = \frac{1}{2}$, $v^B = 0.6$) results in the optimal execution horizons displayed in the figure above. $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|]$ linear.

how our trade will disturb the liquidity provision process and we use this information to derive an optimal trading horizon.

Now consider a recent event in which toxicity played a significant role. The CFTC-SEC examination of the "flash crash" (see CFTC-SEC [2010]) indicates that, during the crash, market participants dumped orders on the market. Why would anyone reduce their execution horizon in the midst of a liquidity crisis? Perhaps the most cited example is the Waddell & Reed order to sell 75,000 E-mini S&P500 contracts. It seems at first unreasonable to execute large orders in an increasingly illiquid market. The model introduced in this paper provides a possible explanation for this behavior, and illustrates how it contributes to the oscillatory nature of prices.[13]

Suppose that v^S increases to a level sufficient to have a measurable impact on PIN. This prompts sellers to shorten their execution horizons in volume-time (Scenario III, left side of Figure 7.1) because small gains from the liquidity component come at the expense of substantial timing risks. As several sellers compete for the same gradually scarcer liquidity, an increase in the volume traded per unit of time will likely occur,

[13]We are not claiming that the Waddell & Reed order was executed optimally. Our argument instead is that trading faster in toxic markets is not necessarily irrational.

which is consistent with the observation that market imbalance accelerates the rate of trading (i.e., greater volume occurs during liquidity crises, like the "flash crash"). Prices are then pushed to lower levels, at which point buyers will have even shorter execution horizons than the sellers (Scenario I, left side of Figure 6.2). This is caused by the convex increasing section of the optimum for buyers (Figure 6.1(a)), compared to the concave decreasing section of the optimum for sellers (Figure 7.1). The increasing activity of buyers causes prices to recover some of the lost ground, v^S returns to normal levels, and execution horizons expand (Scenario II). The outcome is an oscillatory price behavior induced solely by timing reactions of traders to an initial market imbalance.[14]

8. VOLUME PARTICIPATION STRATEGIES

Many money managers conceal large orders among market volume by targeting a certain participation rate. Part of OEH's contribution is to show that, to determine an optimal participation rate, order side and order imbalance should also be taken into account. In this section, we will illustrate how a typical volume participation scheme performs in relative terms to an OEH strategy.

Suppose we apply the same parameter values used in Section 6: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$. For simplicity, consider a $\varphi[|m|]$ linear in $|m|$, with $\varphi[|m|] \equiv \frac{m}{10^4}$ for $m \in (0, 10^4)$. Figure 8.1 compares the probabilistic loss from OEH and a scheme that participates in 5% of the volume, for various buy order sizes when $v^B = \frac{1}{2}$ (i.e., when the markets is balanced between buying and selling). What is apparent is that volume participation results in a far mostly costly trading outcome. For a desired trade size of 2,000 contracts, the Volume Participation probabilistic loss is almost 50% greater than that of the OEH algorithm.

The divergent behavior of these algorithms is also influenced by market imbalance. Figures 8.2 and 8.3 present the equivalent results for $v^B = 0.4$ and $v^B = 0.6$, respectively. OEH's outperformance is particularly noticeable in those cases when the order leans against the market (i.e., buying in a selling market).

9. THE SQUARE ROOT RULE

Loeb (1983) was the first to present empirical evidence that market impact is a square root function of the order size. Grinold and Kahn (1999) justified this observation through an inventory risk model: Given a proposed trade size $|m|$, the estimated time before a sufficient number of trades appears in the market to clear out the liquidity supplier's net inventory is a linear function of $\frac{|m|}{ADV}$, where ADV is the average daily volume. Because prices are assumed to follow an arithmetic random walk, the transaction cost incorporates an inventory risk term of the form $\hat{\sigma}\sqrt{\frac{|m|}{ADV}}$.

Over the last three decades, studies have variably argued in favor of linear (Breen, Hodrick, and Korajczyk 2002; Kissell and Glantz 2003), square-root (Barra 1997) or power-law (Lillo, Farmer, and Mantegna 2003) market impact functions. A possible explanation for these discrepancies is that these studies did not control for two critical

[14]Note that in this discussion we are taking the existence of buyers and sellers as given. The effects discussed in the text arise only from optimal trading horizons. If the impacts on prices also induce changes in desired trades, the effects on market conditions may be exacerbated or moderated.

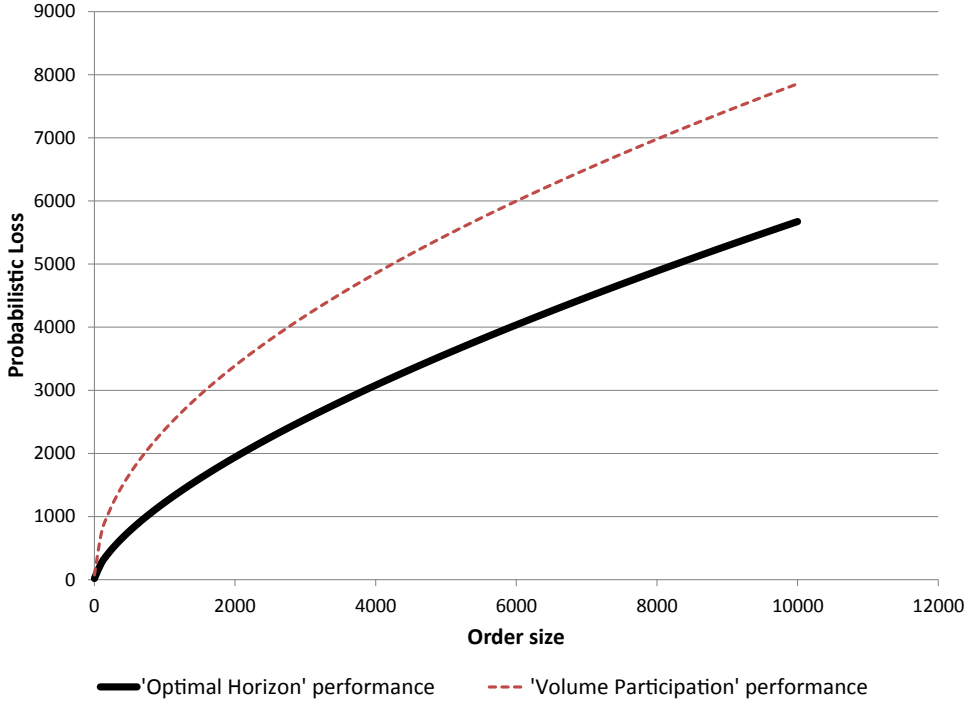

FIGURE 8.1. OEH's performance versus a volume participation strategy when $v^B = \frac{1}{2}$. This figure shows the expected total trading loss for a buy order arising from either a volume participation strategy or the OEH strategy when the market is expected to be in a balanced state. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The volume participation strategy is assumed to participate in 5% of the volume. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|]$ linear.

variables affecting transaction costs: order side and its relation to order imbalance. Consider once again the same parameter values used in Section 6: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$. Suppose that $\varphi[|m|]$ is linear in $|m|$, with $\varphi[|m|] \equiv \frac{m}{10^4}$ for $m \in (0, 10^4)$.

Figure 9.1 plots the probabilistic loss Π that results from executing a buy order of size m at optimal horizons $V^*(m)$ given $v^B = \frac{1}{2}$. A power function fits Π almost perfectly, with a power coefficient very close to the $\frac{5}{5}$ reported by Almgren, Thum, Hauptmann, and Li (2005). However, Figure 9.2 shows that, if $v^B = 0.4$, Π has a linear end and is below the levels predicted by the square root. Finally, when $v^B = 0.6$, Figure 9.3 displays Π values greater than in the other cases, as the order now competes for liquidity. These conclusions depend on our assumption of a linear $\varphi[|m|]$, but they can be generalized to other functional forms. For example, Figure 9.4 shows that, for $\varphi[|m|] \propto \sqrt{m}$ and $v^B = \frac{1}{2}$, Π fits the square root, as originally reported by Loeb (1983).

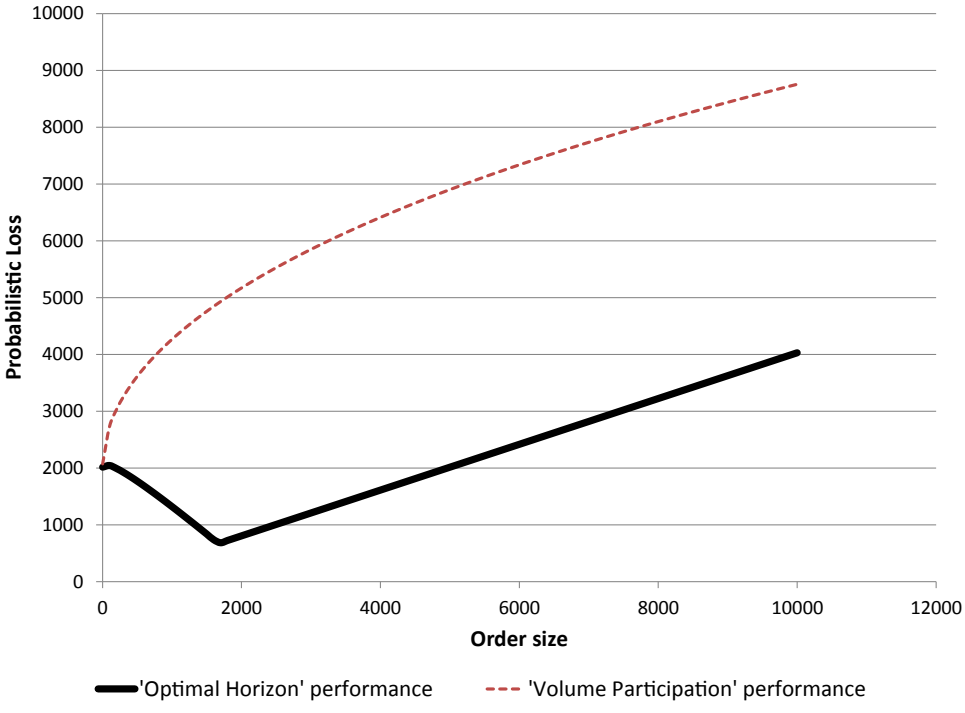

FIGURE 8.2. OEH's performance versus a volume participation strategy when $v^B = 0.4$. This figure shows the expected total trading loss for a buy order arising from either a volume participation strategy or the OEH strategy when the market is expected to have a greater sell imbalance. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The volume participation strategy is assumed to participate in 5% of the volume. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|]$ linear.

10. EMPIRICAL ANALYSIS

In the previous sections we presented a theory showing that OEH optimally uses order imbalance to determine the amount of volume needed to disguise a trade. We have also shown that OEH provides an explanation for the apparently contradictory views on the functional form of the market impact function. In this section we study the empirical performance of OEH relative to a VWAP strategy.

We consider an informed trader who wishes to trade m units of a particular futures contract. To keep the discussion as general as possible, we will consider two scenarios. In the first scenario, the trader has information about the sign of the price change over the next volume bucket, but not its magnitude. In this case, we model the desired trade by $m = Sgn(\Delta S_\tau)q$, where q is a standard trade size determined by the trader. In the second scenario, the trader has information about the sign as well as the magnitude of the price move, and in this case we model the desired trade by $m = f(\Delta S_\tau)q$. As in Section 5, we can analyze the performance of an execution strategy in terms of two components: liquidity and timing costs. The *timing* cost can be directly observed, as measured by the

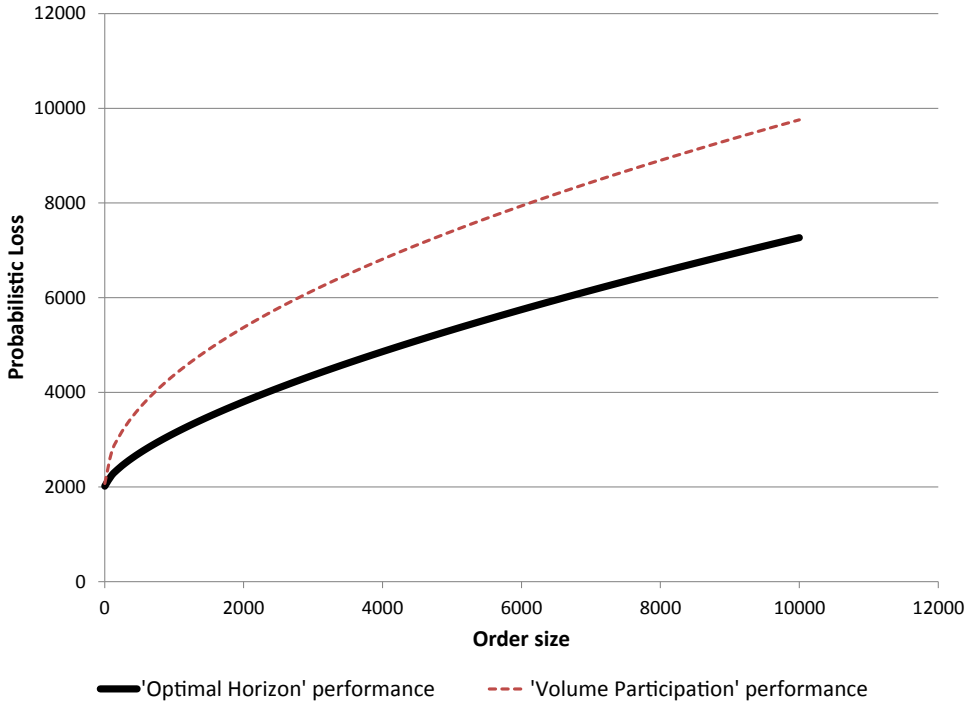

FIGURE 8.3. OEH's performance versus a volume participation strategy when $v^B = 0.6$. This figure shows the expected total trading loss for a buy order arising from either a volume participation strategy or the OEH strategy when the market is expected to have a greater buy imbalance. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The volume participation strategy is assumed to participate in 5% of the volume. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $[\bar{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|]$ linear.

price change with respect to the fill price, excluding the impact that the order has on *liquidity*. For example, in the first scenario the trader may know that prices will go up, but she may not know the magnitude of the price increase, so for a sufficiently large m the expected timing gain may turn into a loss due to the liquidity cost.

When the trader executes using OEH, she combines her private information with market estimates of σ and v^B over the next bucket. In this particular exercise, the latter is based on the procedure described in Appendix 4, and the former is the standard deviation of price changes over a given sample length. Then, for some $\varphi[|m|]$, λ, and a long-run volatility $\bar{\sigma}$, assuming $[\bar{S} - \underline{S}] = -Z_\lambda \bar{\sigma}$,[15] she can compute the OEH V^* over which to trade m contracts. Let's denote the average fill price over the horizon V^*, in absence of her trade, by $\bar{F}_{OEH}$. The realized timing component is $(S_\tau - \bar{F}_{OEH,\tau})m$. The liquidity

[15]Other specifications could be considered to model the maximum trading range at which market makers are willing to provide liquidity. Here, we express that number as a function of the long-term volatility, multiplied by a market makers' risk aversion factor. That factor may not coincide with Z_λ; however we see an advantage in keeping the model as parsimonious as possible, rather than introducing a new variable.

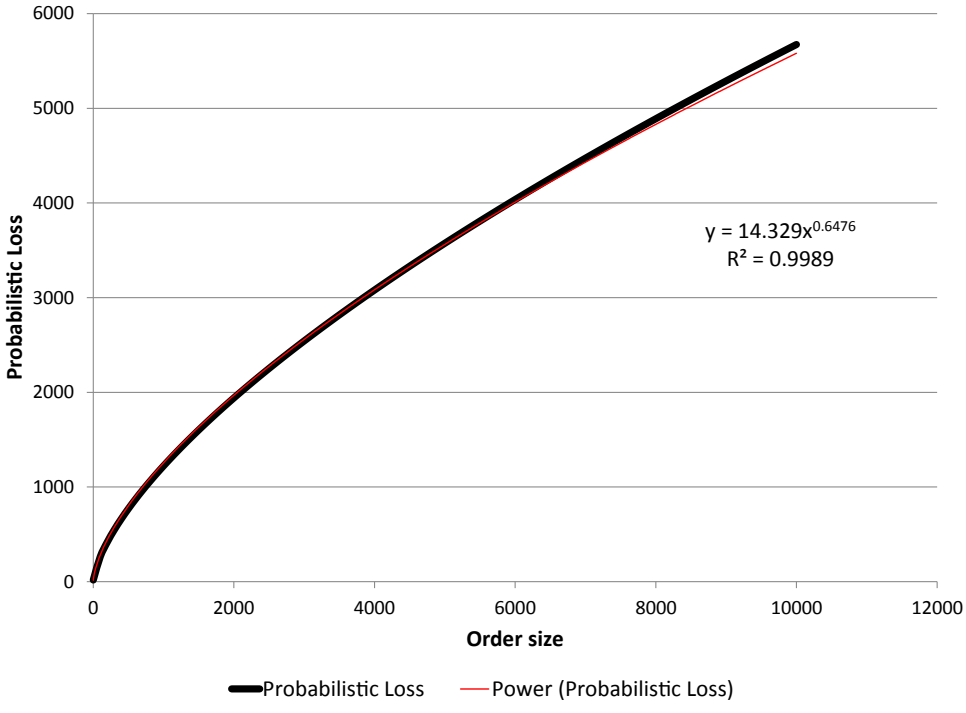

FIGURE 9.1. Probabilistic loss under $v^B = \frac{1}{2}$ and $\varphi[|m|]$ linear. When order flow is balanced, the probabilistic loss follows a functional form close to the square root.

component can be estimated as $|\widetilde{OI}_{OEH,\tau}|[\bar{S} - \underline{S}]|m|$, following equations (4) and (7). Thus, the total profit during volume bucket τ is

$$PL_{OEH,\tau} = \underbrace{-|\widetilde{OI}_{OEH,\tau}|[\bar{S} - \underline{S}]|m|}_{PL^L_{OEH,\tau}} + \underbrace{(S\tau - \bar{F}_{OEH,\tau})m}_{PL^T_{OEH,\tau}} \tag{10.1}$$

where we have expressed the profit in terms of its liquidity ($PL^L_{OEH,\tau}$) and timing ($PL^T_{OEH,\tau}$) components. Similarly, if the trader executes through a VWAP with fixed horizon, the total profit during volume bucket τ is

$$PL_{VWAP,\tau} = \underbrace{-|\widetilde{OI}_{VWAP,\tau}|[\bar{S} - \underline{S}]|m|}_{PL^L_{VWAP,\tau}} + \underbrace{(S_\tau - \bar{F}_{VWAP,\tau})m}_{PL^T_{VWAP,\tau}}. \tag{10.2}$$

Based on the previous equations, we can compute the relative outperformance of OEH over VWAP in terms of its information ratio,

$$IR = \frac{E[PL_{OEH,\tau} - PL_{VWAP,\tau}]}{\sigma[PL_{OEH,\tau} - PL_{VWAP,\tau}]}\sqrt{n} \tag{10.3}$$

where $\sqrt{n}$ is the annualization factor, and n the number of independent trades per year.

Our goal is to evaluate the performance of OEH relative to a VWAP benchmark, for a trader that needs to execute a large order on a daily basis ($n = 260$). To compute

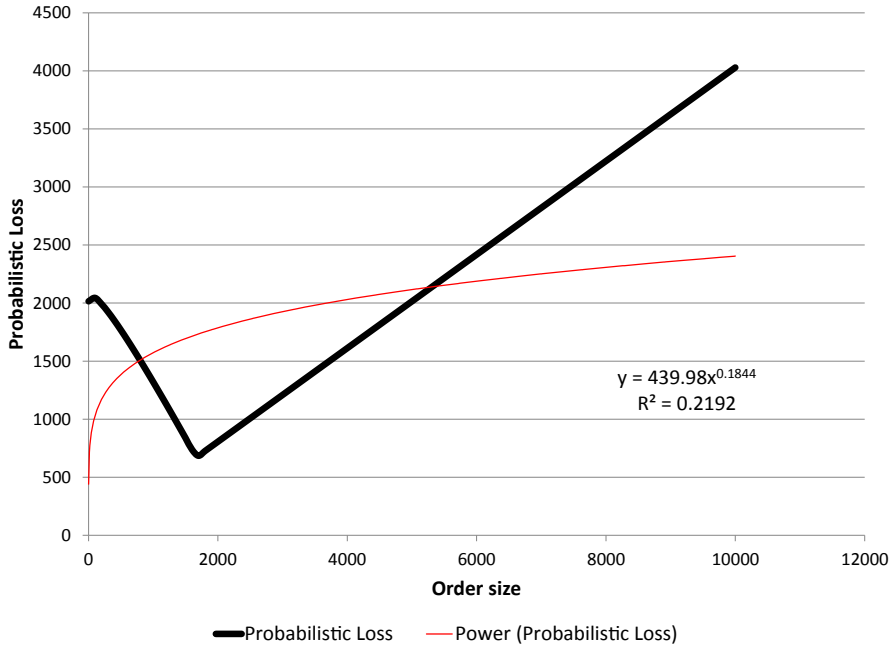

FIGURE 9.2. Probabilistic loss under $v^B = 0.4$ and $\varphi[|m|]$ linear. When order flow is leaning against the market, the probabilistic loss has a piecewise linear functional form.

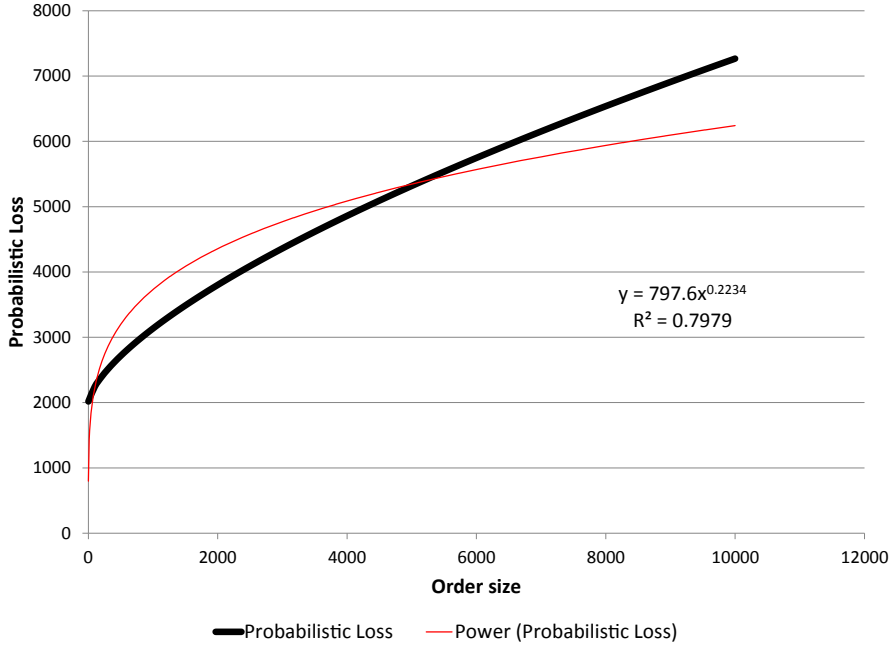

FIGURE 9.3. Probabilistic loss under $v^B = 0.6$ and $\varphi[|m|]$ linear. When the order is large and competing for liquidity, the probabilistic loss is greater than predicted by the square root.

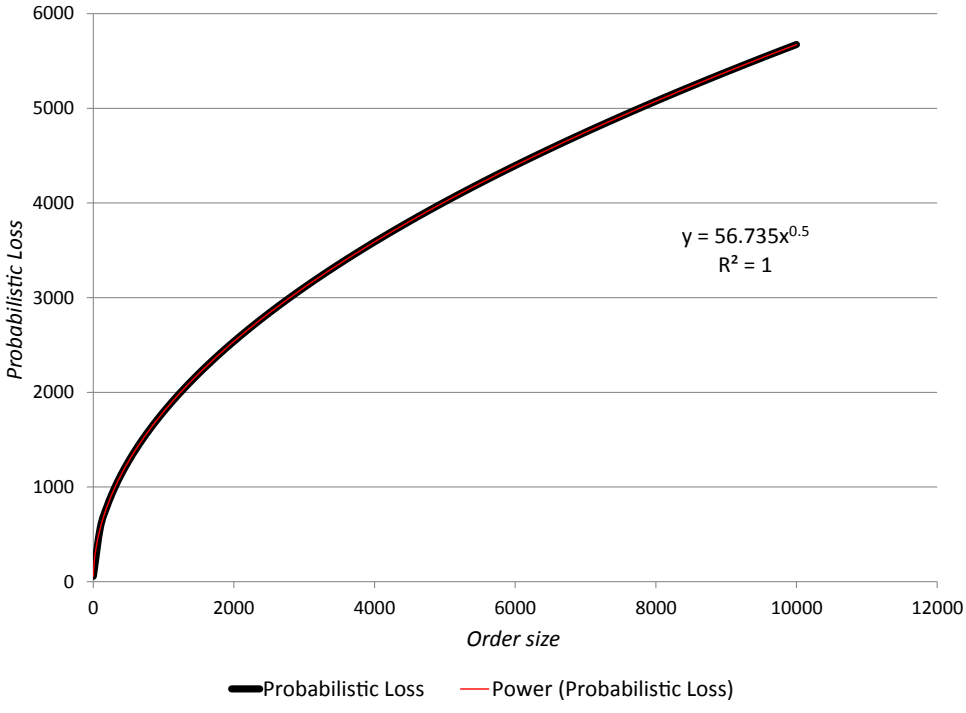

FIGURE 9.4. Probabilistic loss under $v^B = \frac{1}{2}$ and $\varphi[|m|] \propto \sqrt{m}$. The probabilistic loss exactly fits the square root when $v^B = \frac{1}{2}$ and $\varphi[|m|] \propto \sqrt{m}$.

this performance, we estimate v^B using the methodology discussed in Easley, López de Prado, and O'Hara (2012b), on one volume bucket per day, where each volume bucket is composed of 25 volume bars. σ_τ is estimated for each bucket looking back enough trades to have a combined volume of 5 times the ADV (an average of one week of trading activity). We then compute the OEHs for trades of size $q = \{\frac{ADV}{100}, \frac{ADV}{20}, \frac{ADV}{10}\}$, with a risk aversion $\lambda = 0.05$ and an informational leakage function $\varphi[|m|] = min(\sqrt{\frac{|m|}{ADV}}, .999)$, where ADV is the average daily volume.[16]

Table 10.1 summarizes the data used in our calculations. We have selected these products because they encompass a wide variety of asset classes and liquidity conditions. For example, E-Mini S&P500 Futures are traded in the Chicago Mercantile Exchange, it is an equity index product, our sample contains 476,676,009 transactions recorded between January 1, 2007 and July 26, 2012, rolls occur 12 days prior to expiration date, and the ADV for that period has been 1,964,844.89 contracts.

Tables 10.2, 10.3, and 10.4 report the outperformance of OEH over VWAP for trading sizes equivalent to 1%, 5%, and 10% of ADV respectively, when the trader has information about the *sign* of the price move over the next volume bucket (not the size of the move). To interpret these results, suppose that a trader has information that the price of the E-Mini S&P500 futures will increase over the next volume bucket, and for that reason she wishes

[16]These are rather arbitrary values, and alternative ones could be adopted, depending on the user's specific objective. For example, λ could have been calibrated in order to maximize OEH's risk-adjusted performance over VWAP.

TABLE 10.1
Description of the Data Series Used in the Numerical Examples

Futures contract	Exchange	Group	Start	End	Roll	Records	ADV
E-Mini S&P500	CME	Equity	1/1/2007	7/26/2012	12	476,676,009	1,964,844.89
T-Note	CBOT	Rates	1/1/2007	7/26/2012	28	95,091,010	921,056.33
EUR/USD	CME	FX	1/1/2007	7/26/2012	10	188,197,121	233,201.17
WTI Crude Oil	NYMEX	Energy	1/1/2007	7/26/2012	19	164,619,912	194,902.36
Gold	COMEX	Metals	1/1/2007	7/26/2012	27	62,672,073	81,854.96
Corn	CBOT	Softs	1/1/2007	7/26/2012	20	41,833,299	73,860.53
Natural Gas	NYMEX	Energy	1/1/2007	7/26/2012	Volume	50,575,494	61,685.78
Lean Hogs	CME	Meat	1/1/2007	7/26/2012	24	5,499,602	6,544.67
Cotton#2	ICE	Softs	1/1/2007	7/26/2012	20	4,494,294	6,171.32

TABLE 10.2
OEH's Outperformance over VWAP for Trades Equivalent to 1% of ADV and Information Regarding the Side of the Price Move over the Next Bucket

Futures contract	Information	Trade size	Max profit	OEH profit (pts)	Outperf.(pts)	Outperf.(%)	IR
E-Mini S&P500	Sign	0.01*ADV	12.1428	10.5104	4.3262	35.63%	10.04
T-Note	Sign	0.01*ADV	0.3966	0.3441	0.1322	33.33%	9.18
EUR/USD	Sign	0.01*ADV	0.0074	0.0064	0.0028	37.28%	10.62
WTI Crude Oil	Sign	0.01*ADV	1.3913	1.1949	0.4582	32.93%	10.02
Gold	Sign	0.01*ADV	9.4932	8.1780	3.2875	34.63%	9.68
Corn	Sign	0.01*ADV	8.4173	7.2806	3.1640	37.59%	9.67
Natural Gas	Sign	0.01*ADV	0.1098	0.0945	0.0409	37.26%	9.50
Lean Hogs	Sign	0.01*ADV	0.7451	0.6334	0.2613	35.07%	10.51
Cotton#2	Sign	0.01*ADV	1.3211	1.1358	0.4675	35.38%	7.66

TABLE 10.3
OEH's Outperformance over VWAP for Trades Equivalent to 5% of ADV and Information Regarding the Side of the Price Move over the Next Bucket

Futures contract	Information	Trade size	Max profit	OEH profit (pts)	Outperf.(pts)	Outperf.(%)	IR
E-Mini S&P500	Sign	0.05*ADV	12.1428	8.2047	2.2770	18.75%	5.63
T-Note	Sign	0.05*ADV	0.3966	0.2682	0.0649	16.37%	4.91
EUR/USD	Sign	0.05*ADV	0.0074	0.0051	0.0015	20.78%	6.51
WTI Crude Oil	Sign	0.05*ADV	1.3913	0.9275	0.2217	15.94%	5.33
Gold	Sign	0.05*ADV	9.4932	6.3202	1.6253	17.12%	5.15
Corn	Sign	0.05*ADV	8.4173	5.6471	1.7222	20.46%	5.73
Natural Gas	Sign	0.05*ADV	0.1098	0.0731	0.0221	20.14%	5.68
Lean Hogs	Sign	0.05*ADV	0.7451	0.4820	0.1230	16.51%	5.32
Cotton#2	Sign	0.05*ADV	1.3211	0.8776	0.2372	17.96%	4.24

to acquire a position equivalent to 1% of E-Mini S&P500 futures' ADV, i.e., buy 19,648 contracts. We compute the liquidity and timing components for that trade using VWAP, and evaluate by how much OEH beats VWAP on that same trade in dollar terms. After repeating that calculation for each of the 1,450 volume buckets, we can estimate OEH's performance over VWAP. For instance, Table 10.2 reports that the maximum profit

TABLE 10.4
OEH's Outperformance over VWAP for Trades Equivalent to 10% of ADV and Information Regarding the Side of the Price Move over the Next Bucket

Futures contract	Information	Trade size	Max profit	OEH profit (Pts)	Outperf.(pts)	Outperf.(%)	IR
E-Mini S&P500	Sign	0.1*ADV	12.1428	6.3551	0.9577	7.89%	2.59
T-Note	Sign	0.1*ADV	0.3966	0.2065	0.0212	5.36%	1.82
EUR/USD	Sign	0.1*ADV	0.0074	0.0039	0.0007	9.18%	3.17
WTI Crude Oil	Sign	0.1*ADV	1.3913	0.7037	0.0600	4.31%	1.58
Gold	Sign	0.1*ADV	9.4932	4.7746	0.5125	5.40%	1.82
Corn	Sign	0.1*ADV	8.4173	4.2081	0.6715	7.98%	2.63
Natural Gas	Sign	0.1*ADV	0.1098	0.0549	0.0091	8.30%	2.57
Lean Hogs	Sign	0.1*ADV	0.7451	0.3580	0.0314	4.22%	1.46
Cotton#2	Sign	0.1*ADV	1.3211	0.6582	0.0869	6.58%	1.72

TABLE 10.5
OEH's Outperformance over VWAP for Trades Equivalent to 1% of ADV and Information Regarding the Side and Size of the Price Move over the Next Bucket

Futures contract	Information	Trade size	Max profit	OEH profit (pts)	Outperf.(pts)	Outperf.(%)	IR
E-Mini S&P500	Sign, Size	0.01*ADV	15.8723	14.1671	6.4076	40.37%	8.52
T-Note	Sign, Size	0.01*ADV	0.5291	0.4721	0.1959	37.03%	6.74
EUR/USD	Sign, Size	0.01*ADV	0.0098	0.0087	0.0039	39.98%	8.74
WTI Crude Oil	Sign, Size	0.01*ADV	1.8682	1.6672	0.6830	36.56%	8.39
Gold	Sign, Size	0.01*ADV	12.5753	11.4060	4.7222	37.55%	6.96
Corn	Sign, Size	0.01*ADV	12.3966	11.0999	5.1200	41.30%	5.77
Natural Gas	Sign, Size	0.01*ADV	0.1380	0.1230	0.0566	40.98%	7.97
Lean Hogs	Sign, Size	0.01*ADV	0.8442	0.7552	0.3348	39.66%	7.92
Cotton#2	Sign, Size	0.01*ADV	1.7879	1.6020	0.7070	39.54%	6.52

attainable with this information is 12.1428 index points on average, if execution were instantaneous and costless.[17] Out of that, OEH was able to capture 10.5104 index points on average, or 86.56% of the average maximum profit. This represents an outperformance of 4.3262 index points over VWAP's performance on the same trade, or 35.63% of the average maximum profit, with an information ratio of 10.04. As expected, OEH's IR edge over VWAP decays as we approach very large daily trades. For example, the information ratio associated with an OEH trade for 196,484 E-Mini S&P500 futures contracts (10% of its ADV) is 2.59. This occurs because for trades of that extremely large size, it becomes increasingly difficult to conceal the trader's intentions. Although the information ratio is smaller, the average dollar amount of savings is greater (0.9577 points per contract), because the trade is for a size 10 times larger.

Tables 10.5, 10.6, and 10.7 present the results of applying our methodology when the trader has information about the *side and size* of the price change, for an amount equivalent to 1%, 5%, and 10% of ADV, respectively. OEH's edge over VWAP also decays as we approach very large daily trades, however the extent of this performance decay is much less pronounced when the trader holds size as well as sign information with regards to price changes.

[17]The value of an index point is USD50 in the case of the E-Mini S&P500 futures.

TABLE 10.6
OEH's Outperformance over VWAP for Trades Equivalent to 5% of ADV and Information Regarding the Side and Size of the Price Move over the Next Bucket

Futures contract	Information	Trade size	Max profit	OEH Profit (pts)	Outperf.(Pts)	Outperf.(%)	IR
E-Mini S&P500	Sign, Size	0.05*ADV	15.8723	11.4107	4.0384	25.44%	7.34
T-Note	Sign, Size	0.05*ADV	0.5291	0.3744	0.1120	21.17%	6.30
EUR/USD	Sign, Size	0.05*ADV	0.0098	0.0071	0.0025	25.40%	8.26
WTI Crude Oil	Sign, Size	0.05*ADV	1.8682	1.3453	0.4091	21.90%	7.25
Gold	Sign, Size	0.05*ADV	12.5753	9.3511	2.9978	23.84%	5.94
Corn	Sign, Size	0.05*ADV	12.3966	8.6294	3.0115	24.29%	6.28
Natural Gas	Sign, Size	0.05*ADV	0.1380	0.0988	0.0358	25.98%	6.91
Lean Hogs	Sign, Size	0.05*ADV	0.8442	0.6141	0.2119	25.11%	7.50
Cotton#2	Sign, Size	0.05*ADV	1.7879	1.2726	0.4329	24.21%	5.46

TABLE 10.7
OEH's Outperformance over VWAP for Trades Equivalent to 10% of ADV and Information Regarding the Side and Size of the Price Move over the Next Bucket

Futures contract	Information	Trade size	Max profit	OEH profit (Pts)	Outperf.(Pts)	Outperf.(%)	IR
E-Mini S&P500	Sign, Size	0.1*ADV	15.8723	8.8113	2.2197	13.98%	5.98
T-Note	Sign, Size	0.1*ADV	0.5291	0.2891	0.0546	10.32%	4.69
EUR/USD	Sign, Size	0.1*ADV	0.0098	0.0055	0.0014	14.79%	7.10
WTI Crude Oil	Sign, Size	0.1*ADV	1.8682	1.0389	0.2004	10.73%	5.16
Gold	Sign, Size	0.1*ADV	12.5753	6.7359	1.4143	11.25%	6.36
Corn	Sign, Size	0.1*ADV	12.3966	6.4158	1.5333	12.37%	5.39
Natural Gas	Sign, Size	0.1*ADV	0.1380	0.0768	0.0208	15.11%	5.26
Lean Hogs	Sign, Size	0.1*ADV	0.8442	0.4827	0.1199	14.21%	6.57
Cotton#2	Sign, Size	0.1*ADV	1.7879	0.9576	0.2466	13.79%	4.26

11. CONCLUSIONS

The choice of execution horizon is a critical input required by many optimal execution strategies. These strategies attempt to minimize the trading costs associated with a particular order. They do not typically address the footprint that those actions leave in the liquidity provision process. In particular, most execution models do not incorporate information regarding the order's side, without which it is not possible to understand the asymmetric impact that the order will have on the liquidity provision process.

In this paper, we introduce the OEH model, which builds on asymmetric information market microstructure theory to determine the OEH. OEH allows existing optimal execution models to minimize both an order's trading costs and its footprint in the market. In a high frequency world, this latter ability takes on increased importance. OEH is shown to perform better than schemes that target a participation rate. Our model also provides an explanation of the apparent disagreement in the literature regarding the functional form of the market impact function as a result of not controlling for the order side and its relation to the order imbalance. Overall, our analysis suggests a new way to trade in the high frequency environment that characterizes current market structure.

Our empirical study shows that OEH allows traders to achieve greater profits on their information, as compared to VWAP, which we have theoretically justified by OEH's use

of the trader's private information. If the trader's information is correct, OEH will allow her to capture greater profits on that trade. If her information is inaccurate, OEH will deliver smaller losses than VWAP. OEH is not an investment strategy on its own, but it delivers substantial "execution alpha" by boosting the performance of "investment alpha."

A.1 APPENDIX: COMPUTATION OF THE OPTIMAL EXECUTION HORIZON

We need to solve the optimization problem

$$\min_V \Pi[V, v^B, m, [\bar{S} - \underline{S}], \lambda, \hat{\sigma}, V_\sigma] \equiv \left|\varphi[|m|]\left[(2v^B - 1)\left(1 - \frac{|m|}{V}\right) + \frac{m}{V}\right]\right.$$

$$\left. + (1 - \varphi[|m|])(2v^B - 1)\right|[\bar{S} - \underline{S}] - Z\lambda\sqrt{\frac{V}{V_\sigma}}\hat{\sigma} \tag{A.1}$$

subject to $V \geq |m|$.
Let's denote

$$\widetilde{OI} \equiv \varphi[|m|]\left[(2v^B - 1)\left(1 - \frac{|m|}{V}\right) + \frac{m}{V}\right] + \underbrace{(1 - \varphi[|m|])(2v^B - 1)}_{c}$$

$$= \varphi[|m|]\left[\frac{m - (2v^B - 1)|m|}{V} + 2v^B - 1\right] + c. \tag{A.2}$$

Observe that, if $V \geq |m|$, then $-1 \leq \widetilde{OI} \leq 1$. The objective function contains the absolute value of $\widetilde{OI}$ which depends on the choice variable V and it is not continuously differentiable around $\widetilde{OI} = 0$. Accordingly, we solve the problem in cases based on the sign of $\widetilde{OI}$ at an optimum.

A.1.1. CASE 1: Suppose that $\widetilde{OI} > 0$. Then

$$\frac{\partial\left|\widetilde{OI}\right|}{\partial V} = \varphi[|m|]\frac{(2v^B - 1)|m| - m}{V^2}, \tag{A.3}$$

and

$$\frac{\partial\Pi(V)}{\partial V} = \varphi[|m|]\frac{(2v^B - 1)|m| - m}{V^2}[\bar{S} - \underline{S}] - \frac{Z_\lambda\hat{\sigma}}{2\sqrt{V_\sigma V}}. \tag{A.4}$$

Note that the second term in (A.4) is positive, and so an interior solution can occur only if the first term in (A.4) is negative. We first suppose that the solution to the problem does not occur at the constraint $V \geq |m|$. Using the first order necessary condition for an interior solution that $\frac{\partial\Pi(V)}{\partial V} = 0$ and multiplying by $\sqrt{V}$, we obtain

$$\varphi[|m|][(2v^B - 1)|m| - m][\bar{S} - \underline{S}]V^{-3/2} = \frac{Z_\lambda\hat{\sigma}}{2\sqrt{V_\sigma}}. \tag{A.5}$$

Thus,

$$V^* = \left(2\varphi[|m|][(2v^B - 1)|m| - m][\bar{S} - \underline{S}]\frac{\sqrt{V_\sigma}}{Z_\lambda \hat{\sigma}}\right)^{2/3}. \tag{A.6}$$

The V^* given by equation (A.3) is a candidate for an interior solution if $\widetilde{OI} > 0$ evaluated at V^* and $V^* \geq |m|$. It is straightforward to show that $V^* \geq |m|$ if $v^B \leq v^{B+}$ where

$$v^{B+} = \frac{Z_\lambda \hat{\sigma}\sqrt{|m|}}{4\varphi[|m|][\bar{S} - \underline{S}]\sqrt{V_\sigma}} + \frac{\frac{|m|}{m} + 1}{2}. \tag{A.7}$$

Alternatively, if $v^B > v^{B+}$, then the candidate for a solution in this region is $|m|$. Note that the condition $v^B \leq v^{B+}$ implies that m is negative which in turn implies that the first term in (A.4) is positive and so the V^* given by equation (A.3) is well defined.
In this region the second derivative of $\Pi(V)$ is

$$\frac{\partial^2 \Pi(V)}{\partial V^2} = -2\varphi[|m|]\frac{(2v^B - 1)|m| - m}{V^3}[\bar{S} - \underline{S}] + \frac{Z_\lambda \hat{\sigma}}{4V\sqrt{V_\sigma V}}. \tag{A.8}$$

It can be shown that $\frac{\partial^2 \Pi(V)}{\partial V^2} > 0$ if $v^B \leq v^{B+}$.

A.1.2. CASE 2: Suppose that $\widetilde{OI} < 0$. Then

$$\frac{\partial |\widetilde{OI}|}{\partial V} = \varphi[|m|]\frac{m - (2v^B - 1)|m|}{V^2}. \tag{A.9}$$

Using an argument similar to that in Case 1, we see that in this region the candidate for an interior solution is

$$V^* = \left(2\varphi[|m|][m - (2v^B - 1)|m|][\bar{S} - \underline{S}]\frac{\sqrt{V_\sigma}}{Z_\lambda \hat{\sigma}}\right)^{2/3}. \tag{A.10}$$

The V^* given by equation (A.10) is a candidate for a solution if $\widetilde{OI} < 0$ evaluated at V^* and $V^* \geq |m|$. Note that $V^* \geq |m|$ if $v^B \geq v^{B-}$ where

$$v^{B-} = \frac{\frac{|m|}{m} + 1}{2} - \frac{Z_\lambda \hat{\sigma}\sqrt{|m|}}{4\varphi[|m|][\bar{S} - \underline{S}]\sqrt{V_\sigma}}. \tag{A.11}$$

Note that the condition $v^B \geq v^{B-}$ implies that m is negative which in turn implies that the V^* given by equation (A.10) is well defined.

Alternatively, if $v^B < v^{B-}$, then the candidate for a solution in this region is $|m|$.
In this region the second derivative of $\Pi(V)$ is

$$\frac{\partial^2 \Pi(V)}{\partial V^2} = -2\varphi[|m|]\frac{m - (2v^B - 1)|m|}{V^3}[\bar{S} - \underline{S}] + \frac{Z_\lambda \hat{\sigma}}{4V\sqrt{V_\sigma V}}. \tag{A.12}$$

It can be shown that $\frac{\partial^2 \Pi(V)}{\partial V^2} > 0$ if $v^B \geq v^{B-}$.

A.1.3. CASE 3: Suppose that $\widetilde{OI} = 0$. In this case $\frac{\partial \Pi(V)}{\partial V}$ does not exist**,** and the optimum cannot be computed using calculus. However, we can still compute V^*, because

in this case

$$\varphi[|m|]\left[\frac{m-(2v^B-1)|m|}{V}+(2v^B-1)\right]+(1-\varphi[|m|])(2v^B-1)=0, \tag{A.13}$$

and thus

$$V^*=\varphi[|m|]\left(|m|-\frac{m}{2v^B-1}\right). \tag{A.14}$$

It is straightforward to show that $V^* \geq |m|$ if $v^B \geq v^{B=}$ where

$$v^{B=}=\frac{1}{2}\left(\frac{|m|\varphi[|m|]}{m(\varphi[|m|]-1)}+1\right). \tag{A.15}$$

A.2 INFORMATIONAL LEAKAGE ON THE MID-PRICE

The literature devoted to optimal execution typically differentiates between the temporary and the permanent impact that an order m has on mid-prices (e.g., Stoll 1989). The temporary impact arises from the actual order slicing (also called *trajectory*) within the execution horizon, while the permanent impact is a result of the information leaked by the order flow. Because our model is not concerned with the trajectory computation, we do not take into consideration the temporary impact of the actual order slicing. With regard to the permanent impact, Almgren and Chriss (2000) and Almgren (2003) propose a linear permanent impact function that depends on $|m|$.

Let us suppose that, as in the above referred models, the permanent impact function depends on $|m|$, but that unlike those models it is also a function of our volume participation rate, $\frac{|m|}{V^*}$, and whether we are leaning with or against the market. As we have noted, this is not the approach chosen by most execution strategies, but we believe it is worth considering because for large enough trades our orders will have a noticeable impact. For example, in the extreme case that $\frac{|m|}{V^*} \approx 1$, and a sufficiently large $|m|$, our parent trade will not go unnoticed, no matter how well we slice it (especially if we are leaning with the market). We would expect leakage to be greatest when informed traders dominate market activity over a considerable horizon. Under this conjecture, we could conceive a permanent impact function that indeed depends on factors other than $|m|$.

In fact, the specification presented earlier in the paper already modeled an informational leakage process. In the text, we only considered the effect on the trading range, but there is no reason why we could not assume that the same way our parent trade leaves a footprint in the liquidity component, it also leaves a footprint in the mid-price. The probabilistic loss function could then take the form

$$\begin{aligned}\Pi \equiv &\left|\varphi[|m|]\left[(2v^B-1)\left(1-\frac{|m|}{V}\right)+\frac{m}{V}\right]\right.\\ &\left.+\,(1-\varphi[|m|])\,(2v^B-1)\right|([\bar{S}-\underline{S}]+|m|k)-Z_\lambda\hat{\sigma}\sqrt{\frac{V}{V_\sigma}}\end{aligned} \tag{A.16}$$

for $k \geq 0$. If $k=0$, the probabilistic loss reduces to our original specification. If $k>0$, there is a permanent impact on prices that is linear in the informational leakage and the parent order's size. In this way we are taking into account not only the size of the order, but also its side, as we should expect a greater permanent impact when we compete with

the market for scarce liquidity. The optimal execution horizon can then be computed with the following algorithm:

(1) If $(2v^B - 1)|m| < m$, try $V_1 = (2\varphi[|m|][(2v^B - 1)|m| - m][\bar{S} - \underline{S} + |m|k]\frac{\sqrt{V_\sigma}}{Z_\lambda \hat{\sigma}})^{2/3}$ and compute the value of $\widetilde{OI}$ associated with V_1, $\widetilde{OI}[V_1]$.

   (a) If $\widetilde{OI}[V_1] > 0$ and $v^B \leq v^{B+}$, then $V^* = V_1$ is the solution.
   (b) If $\widetilde{OI}[V_1] > 0$ and $v^B > v^{B+}$, then $V^* = |m|$ is the solution.

(2) If $(2v^B - 1)|m| > m$, try $V_2 = (2\varphi[|m|][m - (2v^B - 1)|m|][\bar{S} - \underline{S} + |m|k]\frac{\sqrt{V_\sigma}}{Z_\lambda \hat{\sigma}})^{2/3}$ and compute the value of $\widetilde{OI}$ associated with V_2, $\widetilde{OI}[V_2]$.

   (a) If $\widetilde{OI}[V_2] < 0$ and $v^B \geq v^{B-}$, then $V^* = V_2$ is the solution.
   (b) If $\widetilde{OI}[V_2] < 0$ and $v^B < v^{B-}$, then $V^* = |m|$ is the solution.

(3) If $(2v^B - 1)|m| = m$, then $V_3 = |m|$ is the solution.
(4) Else, try $V_4 = \varphi[|m|](|m| - \frac{m}{2v^B - 1})$.

   (a) If $v^B \geq v^{B=}$, then $V^* = V_4$ is the solution.
   (b) If $v^B < v^{B=}$, then $V^* = |m|$ is the solution.

with

$$\widetilde{OI}[V^*] \equiv \varphi[|m|]\left[\frac{m - (2v^B - 1)|m|}{V^*} + 2v^B - 1\right] + (1 - \varphi[|m|])(2v^B - 1) \tag{A.17}$$

$$v^{B+} = \frac{\frac{|m|}{m} + 1}{2} + \frac{Z_\lambda \hat{\sigma}\sqrt{|m|}}{4\varphi[|m|][\bar{S} - \underline{S} + |m|k]\sqrt{V_\sigma}} \tag{A.18}$$

$$v^{B-} = \frac{\frac{|m|}{m} + 1}{2} - \frac{Z_\lambda \hat{\sigma}\sqrt{|m|}}{4\varphi[|m|][\bar{S} - \underline{S} + |m|k]\sqrt{V_\sigma}} \tag{A.19}$$

$$v^{B=} = \frac{1}{2}\left(\frac{|m|\varphi[|m|]}{m(\varphi[|m|] - 1)} + 1\right). \tag{A.20}$$

A.3 ALGORITHM IMPLEMENTATION

The following is an implementation in Python of the algorithm described in this paper. Set the values given in the "Parameters" section according to your particular problem. The k parameter is optional. If k is not provided (or given a value $k = 0$), the procedure described in Appendix 1 is followed. Otherwise, the specification discussed in Appendix 2 is applied. If you run this code with the parameters indicated below, you should get $V^* = 6,000$ (the solution in Section 6.1).

```
#!/usr/bin/env python
# By MLdP on 20130427 <mlopezdeprado@hetco.com>
# It computes the Optimal Execution Horizon
#-------------------------------------------------
# PARAMETERS
sigma = 1000
volSigma = 10000
S_S = 10000
zLambda = -1.644853627 #CDF(0.05) from the Std Normal dist
vB = 0.5
phi = 1
m = 1000
k = 0
#-------------------------------------------------
def signum(int):
   if(int < 0):return -1
   elif(int > 0):return 1
   else:return 0
#-------------------------------------------------
def getOI(v,m,phi,vB,sigma,volSigma):
   return phi*(float(m-(2*vB-1)*abs(m))/v+2*vB-1)+ \
      (1-phi)*(2*vB-1)
#-------------------------------------------------
def getBounds(m,phi,vB,sigma,volSigma,S_S,zLambda,k = 0):
   vB_l = float(signum(m)+1)/2-zLambda*sigma*abs(m)**0.5/ \
      float(4*phi*(S_S+abs(m)*k)*volSigma**0.5)
   vB_u = float(signum(m)+1)/2+zLambda*sigma*abs(m)**0.5/ \
      float(4*phi*(S_S+abs(m)*k)*volSigma**0.5)
   vB_z = (signum(m)*phi/float(phi-1)+1)/2.
   return vB_l,vB_u,vB_z
#-------------------------------------------------
def minFoot(m,phi,vB,sigma,volSigma,S_S,zLambda,k = 0):
   # compute vB boundaries:
   if phi< = 0:phi+ = 10**-12
   if phi> = 1:phi- = 10**-12
   vB_l,vB_u,vB_z = getBounds(m,phi,vB,sigma,volSigma,S_S, \
      zLambda,k)
   # try alternatives
   if (2*vB-1)*abs(m)<m:
      v1 = (2*phi*((2*vB-1)*abs(m)-m)*(S_S+abs(m)*k)*volSigma**0.5/ \
         float(zLambda*sigma))**(2./3)
      oi = getOI(v1,m,phi,vB,sigma,volSigma)
      if oi>0:
         if vB< = vB_u: return v1
         if vB>vB_u: return abs(m)
```

```
  elif (2*vB-1)*abs(m)>m:
    v2 = (2*phi*(m-(2*vB-1)*abs(m))*(S_S+abs(m)*k)*volSigma**0.5/ \
      float(zLambda*sigma))**(2./3)
    oi = getOI(v2,m,phi,vB,sigma,volSigma)
    if oi<0:
      if vB> = vB_l: return v2
      if vB<vB_l: return abs(m)
  elif (2*vB-1)*abs(m) = = m: return abs(m)
  if m<0:
    if vB<vB_z: return phi*(abs(m)-m/float(2*vB-1))
    if vB> = vB_z: return abs(m)
  else:
    if vB> = vB_z: return phi*(abs(m)-m/float(2*vB-1))
    if vB<vB_z: return abs(m)
#-------------------------------------------------
def main():
    print minFoot(m,phi,vB,sigma,volSigma,S_S,zLambda,k)
#-------------------------------------------------
if __name__ = = '__main__': main()
```

A.4 APPENDIX: EXPECTED ORDER IMBALANCE

Given a sample of L buckets, we fit a forecasting regression model of the form

$$Ln\left(v_{\tau+1}^{B}\right) = \beta_0 + \beta_1 Ln\left(v_{\tau}^{B}\right) + \xi_\tau \tag{A.21}$$

and because $v_\tau^B \in [0, 1]$ we must obtain that $|\beta_1| < 1$.[18] The expected value of the order imbalance at τ over $\tau + 1$ is

$$E_\tau[OI_{\tau+1}] = 2e^{E_\tau\left[Ln\left(v_{\tau+1}^{B}\right)\right]} - 1 \tag{A.23}$$

where

$$E_\tau\left[Ln\left(v_{\tau+1}^{B}\right)\right] = \hat{\beta}_0 + \hat{\beta}_1 Ln\left(v_{\tau}^{B}\right). \tag{A.24}$$

In general we would expect roughly balanced markets, and so $v_\tau^B = \frac{1}{2}E_\tau[v_{\tau+1}^B] = \frac{1}{2}$. This means that the regression should cross through the equilibrium point $(Ln(v_\tau^B), E_\tau[Ln(v_{\tau+1}^B)]) = (Ln(\frac{1}{2}), Ln(\frac{1}{2}))$, rather than having an equilibrium is systemically disrupted by some arbitrary intercept value $\hat{\beta}_0$. We impose that condition as

$$Ln\left(v_{\tau+1}^{B}\right) - Ln\left(\frac{1}{2}\right) = \beta_1^*\left[Ln\left(v_{\tau}^{B}\right) - Ln\left(\frac{1}{2}\right)\right] + \xi_\tau \tag{A.24}$$

and we use

$$Ln\left(v_{\tau+1}^{B}\right) = \beta_0^* + \beta_1^* Ln\left(v_{\tau}^{B}\right) + \xi_\tau \tag{A.25}$$

[18] When fitting this regression, observations pairs $(v_{\tau+1}^B, v_\tau^B)$ where either $v_{\tau+1}^B = 0$ or $v_\tau^B = 0$ are removed.

where $\hat{\beta}_1^* = \dfrac{\sum_{l=1}^{L}(Ln(v_{\tau+1}^B)Ln(v_\tau^B) + [Ln(\frac{1}{2})]^2 - Ln(\frac{1}{2})(Ln(v_{\tau+1}^B) + Ln(v_\tau^B)))}{\sum_{l=1}^{L}([Ln(v_\tau^B)]^2 + [Ln(\frac{1}{2})]^2 - 2Ln(\frac{1}{2})Ln(v_\tau^B))}$ and $\hat{\beta}_0^* = Ln(\frac{1}{2})(1 - \hat{\beta}_1^*)$.

REFERENCES

ABAD, D., and J. YAGÜE (2012): From PIN to VPIN: An Introduction to Order Flow Toxicity, *Spanish Rev. Financial Eco.* 10(2), 74–83.

ALMGREN, R. (2009): Optimal Trading in a Dynamic Market, NYU Mathematics In Finance Working Paper 2009-2.

ALMGREN, R. (2003): Optimal Execution with Nonlinear Impact Functions and Trading-Enhanced Risk, *Appl. Math. Finance* 10, 1–18.

ALMGREN, R., and N. CHRISS (2000): Optimal Execution of Portfolio Transactions, *J. Risk* Winter, 5–39.

ALMGREN, R., C. THUM, E. HAUPTMANN, and H. LI (2005): Direct Estimation of Equity Market Impact, Working Paper, University of Toronto.

BARRA (1997): *Market Impact Model Handbook*, Berkeley, CA: BARRA.

BAYRAKTAR, E., and M. LUDKOVSKI (2011): Optimal Trade Execution in Illiquid Markets, *Math. Finance* 21(4), 681–701.

BERKOWITZ, S., D. LOGUE, and E. NOSER (1988): The Total Cost of Transactions on the NYSE, *J. Finance* 41, 97–112.

BERTSIMAS, D., and A. LO (1998): Optimal Control of Execution Costs, *J. Fin. Markets* 1, 1–50.

BETHEL, W., D. LEINWEBER, O. RUEBEL, K. WU (2011): Federal Market Information Technology in the Post Flash Crash Era: Roles for Supercomputing, *J. Trading* 7(2), 9–25. http://ssrn.com/abstract=1939522

BREEN, W., L. HODRICK, and R. KORAJCZYK (2002): Predicting Equity Liquidity, *Manag. Sci.* 48(4), 470–483.

CFTC-SEC (2010): Findings Regarding the Market Events of May 6, 2010. Washington, D.C.: Securities and Exchange Commission.

DAI, M., and Y. ZHANG (2012): Optimal Selling/Buying Strategy with Reference to the Ultimate Average, *Math. Finance* 22(1), 165–184.

DONGES, J., and F. HEINEMANN (2006): Crossing Network versus Dealer Market: Unique Equilibrium in the Allocation of Order Flow, Working Paper, Humboldt University.

EASLEY, D., and M. O'HARA (1992b): Time and the Process of Security Price Adjustment, *J. Finance* 47, 576–605.

EASLEY, D., N. KIEFER, M. O'HARA, and J. PAPERMAN (1996): Liquidity, Information, and Infrequently Traded Stocks, *J. Finance* 51(4), 1405–1436.

EASLEY, D., R. F. ENGLE, M. O'HARA, and L. WU (2008): Time-Varying Arrival Rates of Informed and Uninformed Traders, *J. Fin. Econ* 6(2), 171–207.

EASLEY, D., M. LÓPEZ de PRADO, and M. O'HARA (2011a): The Microstructure of the Flash Crash: Flow Toxicity, Liquidity Crashes and the Probability of Informed Trading, *J.Portfolio Manag.* 37(2), 118–128. http://ssrn.com/abstract=1695041.

EASLEY, D., M. LÓPEZ de PRADO, and M. O'HARA (2011b): The Exchange of Flow Toxicity, *J. Trading* 6(2), 8–14. http://ssrn.com/abstract=1748633

EASLEY, D., M. LÓPEZ de PRADO, and M. O'HARA (2012a): Flow Toxicity and Liquidity in a High Frequency World, *Rev. Fin. Studies* 25(5), 1457–1493. http://ssrn.com/abstract=1695596.

EASLEY, D., M. LÓPEZ de PRADO, and M. O'HARA (2012b): Bulk Volume Classification, SSRN, Working Paper. http://ssrn.com/abstract=1989555

EASLEY, D., M. LÓPEZ de PRADO, and M. O'HARA (2012c): The Volume Clock: Insights into the High Frequency Paradigm, *J. Portfolio Manag.* 39(1), 19–29. http://ssrn.com/abstract=2034858

FORSYTH, P. (2011): A Hamilton Jacobi Bellman Approach to Optimal Trade Execution, *Appl. Num. Math.* 61(2), 241–265.

GATHERAL, J., and A. SCHIED (2011): Optimal Trade Execution under Geometric Brownian Motion in the Almgren and Chriss Framework, *Int. J. Theor. Appl. Finance* 14(3), 353–368.

GRINOLD, R., and R. KAHN (1999): *Active Portfolio Management*, New York: McGraw-Hill, 473–475.

HASBROUCK, J., and R. SCHWARTZ (1988): Liquidity and Execution Costs in Equity Markets, *J. Portfolio Manag.* 14, 10–16.

HENDERSHOTT, T., and H. MENDELSON (2000): Crossing Networks and Dealer Markets: Competition and Performance, *J. Finance* 55(5), 2071–2115.

KISSELL, R., and M. GLANTZ (2003): *Optimal Trading Strategies*, New York: American Management Association.

KONISHI, H., and N. MAKIMOTO (2001): Optimal Slice of a Block Trade, *J. Risk* 3(4), 33–51.

KRATZ, P., and T. SCHÖNEBORN (2014): Optimal Liquidation in Dark Pools, *Quant. Finance* 14(9), 1519–1539.

LILLO, F., J. FARMER, and R. MANTEGNA (2003): Master Curve for Price Impact Function, *Nature* 421, 129–130.

MADHAVAN, A. (2002): VWAP Strategies, Trading Performance: The Changing Face of Trading, *Trading* 2002(1), 32–39.

MENKVELD, A., and B YUESHEN (2013): Anatomy of the Flash Crash, http://ssrn.com/abstract=2243520

OBIZHAEVA, A., and J. WANG (2013): Optimal Trading Strategy and Supply/Demand Dynamics, *J. Fin.Markets* 16(1), 1–32.

PEROLD, A. F. (1988): The Implementation Shortfall: Paper versus Reality, *J. Portfolio Manag.* 14, 4–9.

SCHIED, A., and T. SCHÖNEBORN (2009): Risk Aversion and the Dynamics of Optimal Liquidation Strategies in Illiquid Markets, *Finance Stoch.* 13(2), 181–204.

STOLL, H. (1989): Inferring the Components of the Bid-Ask Spread: Theory and Empirical Tests, *J. Finance* 44, 115–134.

YE, M. (2011): A Glimpse into the Dark: Price Formation, Transaction Costs, and Market Share in the Crossing Network, Working Paper, University of Illinois at Urbana-Champaign.