

Trading Strategies within the Edges of No-Arbitrage[☆]

Álvaro Cartea^a, Sebastian Jaimungal^b, Jason Ricci^c

^a*Department of Mathematics, University College London, UK*

^b*Department of Statistics, University of Toronto, Canada*

^c*Department of Statistics, University of Toronto, Canada*

Abstract

We develop a trading strategy which employs limit and market orders in a multi-asset economy where the assets are not only correlated, but can also be structurally dependent. To model the structural dependence, the midprice processes follow a multivariate reflected Brownian motion on the closure of a no-arbitrage region which is dictated by the assets' bid-ask spreads. We provide a formal framework for such an economy and solve for the value function and optimal control for an investor who takes positions in these assets. The optimal strategy exhibits two dominant features which depend on how far the vector of midprices is from the no-arbitrage bounds. When midprices are sufficiently far from the no-arbitrage edges, the strategy behaves as that of a market maker who posts buy and sell limit orders. And when the midprice vector is close to the edge of the no-arbitrage region, the strategy executes a combination of market orders and limit orders to profit from statistical arbitrages. Moreover, we discuss a numerical scheme to solve for the value function and optimal control, and perform a simulation study to discuss the main characteristics of the optimal strategy.

Keywords: Optimal Trading; high-frequency Trading; Algorithmic Trading; Limit Orders; Market Orders; Stochastic Control; Impulse Control; No-arbitrage bounds.

[☆]This work expresses the opinion of its authors and is in no way representing the opinion of the institutions the authors work for. AC acknowledges the research support of the Oxford-Man Institute for Quantitative Finance and the hospitality of the Finance Group at Saïd Business School. SJ would like to thank NSERC and GRI for partially funding this work. The authors would like to thank Robert Almgren for comments on an earlier version of this article. We also thank conference and seminar participants at University College London, University of Toronto.

Email addresses: a.cart@ucl.ac.uk (Álvaro Cartea), sebastian.jaimungal@utoronto.ca (Sebastian Jaimungal), jason.ricci@mail.utoronto.ca (Jason Ricci)

1. Introduction

Over the past decade, advances in technology and computing power have given rise to a new age in trading. There has been a movement from traditional pit exchanges to electronic exchanges where open quotes get amalgamated into a limit order book (LOB) and agents trade by placing limit orders (LOs) with intent to buy/sell at a particular price or by sending market orders (MOs). In this paper, we develop a strategy where the investor trades in a collection of assets which are not only correlated, but present static arbitrage opportunities if their midprices stray outside a no-arbitrage region which is dictated by the bid-ask spreads of the assets.

The investor's goal is to maximize cash profits at some terminal time T , while controlling for inventory risk by penalizing open positions. The investor may trade aggressively and take liquidity via MOs if she sees a profitable opportunity in the form of a statistical arbitrage, or may trade passively by submitting limit buy and sell orders to the LOB. The trading strategy requires aggressive or passive orders depending on where the midprices are relative to the boundaries of the no-arbitrage region. The following two examples illustrate how certain structural dependencies in a collection of assets determine the no-arbitrage region.

Suppose that an investor maximizes expected wealth by trading in an equity futures contract which matures at time T_1 and also trades in its underlying. Furthermore, assume that this underlying equity does not pay dividends and the risk-free rate is zero. Let S_t^1 denote the price of the underlying at time t and S_t^2 the price of the futures contract. In an economy where assets are traded with a zero bid-ask spread (that is, the price to sell or buy the asset is the same), we have, by no-arbitrage, $S_t^1 = S_t^2$ for all $t \leq T_1$. If this equality does not hold, then there exists an arbitrage.

Now assume that the underlying equity and its futures contract trade with bid-ask spreads $\xi_1 > 0$ and $\xi_2 > 0$, respectively. In this case, it is possible for the market to bear arbitrage-free prices of the form $S_t^1 \neq S_t^2$. To see this, assume that $S_t^1 > S_t^2$. To arbitrage these prices the investor sends a buy market order (MO) for the futures contract, for which she must cross the spread and pay $S_t^2 + \xi_2/2$, and simultaneously sends a sell MO for the underlying and obtains $S_t^1 - \xi_1/2$. This trade is an arbitrage if $(S_t^1 - \xi_1/2) - (S_t^2 + \xi_2/2) > 0$. Hence, it is simple to show that the no-arbitrage region is $|S_t^1 - S_t^2| < \frac{1}{2}(\xi_1 + \xi_2)$.

Another example is one in which an investor takes positions in the following three derivatives: i) the first nearby futures contract written on an equity having maturity T_1 , ii) the second nearby futures contract written on the *same* equity having maturity $T_2 > T_1$, and iii) the 1-2 calendar spread contract¹ having maturity T_1 . Furthermore, denote the midprices of

¹The 1-2 calendar spread is a derivative that forces one party to exchange the T_2 maturity contract for the

these three contracts by S_t^1, S_t^2, S_t^3 and their spreads ξ_1, ξ_2, ξ_3 , thus the no-arbitrage region is $|S_t^1 - S_t^2 - S_t^3| < \frac{1}{2} \sum_{j=1}^3 \xi_j$ (below we provide a formal definition of the no-arbitrage region).

Midprices outside the no-arbitrage region are immediately corrected by arbitrageurs who exert price pressure until these opportunities disappear. In our framework, this arbitrage pressure is modelled by assuming that midprices follow a multivariate reflected Brownian Motion on the closure of the no-arbitrage region. Thus, when the midprice of an asset approaches the edges of the no-arbitrage region, market forces will exert pressure so that the midprice is pulled back into the interior of the region.

We solve the investor’s optimal investment strategy and show how it depends on the proximity of the midprices to the edges of the no-arbitrage region. When midprices are ‘far’ from the no-arbitrage boundary, the investor behaves as a strategic market maker who posts two-sided LOs. However, these limit orders are strategically placed in the book to not only control for inventory exposure, but also in anticipation of midprices approaching the edges of the no-arbitrage region where the strategy employs a mix of market and limit orders to execute statistical arbitrages. And when the midprice of one asset is ‘close’ to the no-arbitrage boundary, the investor adjusts the depths of her LOs and may also execute MOs to take advantage of statistical arbitrage opportunities in anticipation of price pressures which force prices to stay within the no-arbitrage bounds, and to control exposure to inventory risk.

Our work is related to the literature that investigates investment strategies based on a collection of co-integrated assets (see for example Tourin and Yan (2013), Tourin and Lintilhac (2015), Cartea et al. (2015), Leung and Li (2015), and Cartea and Jaimungal (2015a)), but our approach is different because we focus on a collection of assets with stronger structural dependencies, and more importantly, we model the no-arbitrage region and the investor employs LOs and MOs in the strategy. Other related literature in investment strategies specific to trading in a single asset to optimally liquidate/acquire shares as in Almgren and Chriss (2000), Almgren (2003), Kharroubi and Pham (2010), Guéant et al. (2012), Bayraktar and Ludkovski (2014), Cartea and Jaimungal (2015b), or in a general profit/utility maximization setting as in Avellaneda and Stoikov (2008), Guilbaud and Pham (2013), and Cartea et al. (2014).

The rest of this paper is organized as follows. Section 2 contains definitions and some modeling assumptions. Section 3 states and discusses properties of the optimal trading strategy. Section 4 discusses the discretization scheme and convergence to the viscosity solution. Section 5 analyzes qualitative features in the interior and edges of the no-arbitrage in the 2-dimensional case. Moreover, it uses simulations to study the financial performance of the optimal strategy.

underlying asset at time T_1 . Its payoff can be decomposed into being long the T_1 contract and short the T_2 contract.

Section 6 concludes and we collect some additional details of proofs in the Appendix.

2. The Model

Here we develop a framework for investors who trade in a basket of assets. Our main focus is on assets that exhibit strong structural dependencies as in the examples above where the investor trades in futures contracts with the same underlying asset. The main question we answer is: How can an investor design an algorithmic strategy to take advantage of this structural dependency between assets?

We assume that the economy trades n risky assets with midprices at time t given by S_t^i , $i = 1, \dots, n$, and respective bid-ask spreads $\xi_i > 0$, both of which we arrange in n -dimensional vectors \mathbf{S}_t and $\boldsymbol{\xi}$ respectively. We also assume the existence of a risk-free cash asset that offers zero return and trades with zero bid-ask spread.

A well-known result in finance theory is that when one asset's payoff can be replicated as a linear combination of other traded assets, then there is a redundant asset in the economy. However, if some of these assets trade with a bid-ask spread, this redundancy ceases to exist. We classify a multiple asset economy as *over-complete* if there exists a non-zero portfolio and a finite investment horizon \bar{T} such that at time \bar{T} the portfolio's value is a constant, i.e. its variance is zero.

All random variables and stochastic processes in this paper are defined on the completed, filtered probability space $(\Omega, \mathcal{F}_{\bar{T}}, \mathbb{F}, \mathbb{P})$ with filtration $\mathbb{F} = \{\mathcal{F}_t\}_{0 \leq t \leq T}$ and where \mathbb{P} is the real-world probability measure. Note that the filtration \mathbb{F} is only defined up to time $T < \bar{T}$, whereas \mathbb{P} is defined on the larger σ -algebra $\mathcal{F}_{\bar{T}}$. Here, \mathbb{F} is the natural filtration generated by $(\mathbf{S}_t, \mathbf{M}_t^\pm, \mathbf{N}_t^\pm)$, all of these processes will be defined later.

Definition 1. *A financial system is said to be over-complete if $\exists \mathbf{b} \in \mathbb{R}^n$, $\bar{T} > 0$ with $\|\mathbf{b}\| > 0$ such that $\mathbb{V}[\mathbf{b} \cdot \mathbf{S}_{\bar{T}}] = 0$, provided $\mathbf{S}_{\bar{T}} \in L^2(\Omega, \mathcal{F}_{\bar{T}}, \mathbb{P})$, where \mathbb{V} is the variance operator, and the operation \cdot is the dot product.*

One of our motivating examples (first two nearby futures and the calendar spread) fits into this framework with portfolio $\mathbf{b} = (1, -1, -1)$ and \bar{T} being the maturity of the first nearby futures. The illustrative example that will be used throughout this paper (first nearby futures and the underlying equity) also fits into this framework with $\mathbf{b} = (1, -1)$ and \bar{T} being the maturity of the futures contract.

In Definition 1, the trivial portfolio $\mathbf{b} = \mathbf{0}$ is excluded from \mathbb{R}^n since $\mathbb{V}[\mathbf{0} \cdot \mathbf{S}_{\bar{T}}]$ is zero, making every economy over-complete and rendering this classification trivial. Notice that if the economy does *not* admit such a portfolio (i.e. there does not exist such a vector \mathbf{b}), then this problem reduces to a (correlated) multiple asset problem, see for example Tourin and Yan (2013), and Cartea and Jaimungal (2015a). The no-arbitrage region in this simplified case would be all of \mathbb{R}^n . Also observe that the set from Definition 1 consists of all the riskless portfolios when considering \bar{T} as the investment horizon. However, \bar{T} may be the only investment horizon where these portfolios are riskless.

Definition 2. *Define the space B as the set of all vectors \mathbf{b} from Definition 1 plus the origin. Observe that B is a subspace of \mathbb{R}^n .*

We require the set B in Definition 2 to include the origin so that it is, in fact, a subspace. In the simplified case where the economic system is not over-complete, the set B would only consist of the origin.

Definition 3. *The no-arbitrage region, denoted by $\mathcal{A} \subset \mathbb{R}^n$, is the set of midprices that makes this economy arbitrage free.*

It is worth pointing out that the no-arbitrage region is open. More specifically, \mathcal{A} is the set of midprices such that the initial cashflow from purchasing a portfolio \mathbf{b} plus the terminal payoff, $\mathbf{b} \cdot \mathbf{S}_{\bar{T}}$, is strictly negative, whenever this portfolio matures at a constant value (i.e. when $\mathbb{V}[\mathbf{b} \cdot \mathbf{S}_{\bar{T}}] = 0$).

This observation deserves some explanation. It is clear that the no-arbitrage set \mathcal{A} , as previously described, is open. At the boundary, these portfolios have cost equal to their maturity value. Also, it has the potential for arbitrage profits via the following strategy: Obtain portfolio \mathbf{b} now for cost $\mathbf{b} \cdot \mathbf{S}_{\bar{T}}$. There is a positive probability that prior to \bar{T} , \mathbf{S}_t will move to a value (within \mathcal{A}) such that a filled limit order (LO) on one asset, combined with liquidation of the others assets via MOs, produces a net gain. Because there is a positive probability that such an opportunity does *not* arise, having initial cost strictly greater than $\mathbf{b} \cdot \mathbf{S}_{\bar{T}}$ is enough to ensure it is not an arbitrage portfolio. This is due to the fact that if the agent can only execute MOs, the portfolio \mathbf{b} (which costs $\mathbf{b} \cdot \mathbf{S}_{\bar{T}}$ to obtain) can be liquidated for exactly $\mathbf{b} \cdot \mathbf{S}_{\bar{T}}$ only in the *best case scenario* (on the opposite boundary). In fact, this scenario is only in the closure of \mathcal{A} , and hence the midprice process can only move arbitrarily close to it. Notice that the only assumption we have made here is that the spreads are positive constants.

In the 2-dimensional illustrative example consisting of a futures contract and its underlying, we have $\mathcal{A} = \{\mathbf{s} \in \mathbb{R}^2 : |s_1 - s_2| < \frac{1}{2}(\xi_1 + \xi_2)\}$. That is, at any time $t < \bar{T}$, the midprices of

these two assets cannot differ by more than the sum of their half-spreads; if they do, then one could execute two MOs, for example one order to short the overpriced futures and the other to purchase the underpriced underlying, and hold to maturity for an arbitrage.

In general, arbitrageurs (who are willing to hold the arbitrage portfolio until \bar{T}) ensure that prices are always pulled back within the bounds of no-arbitrage. In our framework the investor assumes that for $t \in [0, T]$, the dynamics of the midprice process \mathbf{S}_t is a multivariate reflected and correlated Brownian Motion.

Assumption 1. The n -dimensional midprice process $\{\mathbf{S}_t\}_{0 \leq t \leq \bar{T}}$ satisfies

$$d\mathbf{S}_t = \boldsymbol{\sigma} d\mathbf{B}_t \quad (1)$$

for $t \in [0, T]$ and $\mathbf{S}_t \in \mathcal{A}$. Here $\boldsymbol{\sigma}$ is an n -by- n matrix with $\boldsymbol{\sigma}\boldsymbol{\sigma}'$ being positive definite, $\{\mathbf{B}_t\}_{0 \leq t \leq T}$ is a standard Brownian Motion on \mathbb{R}^n , and \mathbf{S}_t has (normal) reflection at the boundary of \mathcal{A} .

When the risky assets possess a structural dependence, as in our illustrative example, the matrix $\boldsymbol{\sigma}$ will induce a strong correlation between all or some of the n assets. For example, the case discussed above consisting of $n = 2$, a futures contract and its underlying, $\boldsymbol{\sigma}$ will induce a strong positive correlation. In general, for assets that exhibit strong structural dependence, $\boldsymbol{\sigma}$ will induce movements in the midprice processes that are generally in the same direction, on average, to some element of the subspace B^\perp (the orthogonal complement of B).

Recall that the agent will be submitting both LOs and MOs and therefore needs to model incoming MOs from other participants in addition to her own executions (which she controls). As such, she assumes that other market participants send MOs according to independent Poisson Processes. The counting processes representing the total number of MOs submitted by *other* agents up until time t are $\{N_{i,t}^- : i = 1, \dots, n\}_{0 \leq t \leq T}$ for sells and $\{N_{i,t}^+ : i = 1, \dots, n\}_{0 \leq t \leq T}$ for buys (i indexes the asset). On the other hand, the *total* number (including the agent's trades) of market sells and buys up until time t are given by the processes $\{\bar{N}_{i,t}^- : i = 1, \dots, n\}_{0 \leq t \leq T}$ and $\{\bar{N}_{i,t}^+ : i = 1, \dots, n\}_{0 \leq t \leq T}$ respectively. It is also convenient to have these processes in the following vector forms: $\mathbf{N}_t^\pm = (N_{1,t}^\pm, \dots, N_{n,t}^\pm)'$ and $\bar{\mathbf{N}}_t^\pm = (\bar{N}_{1,t}^\pm, \dots, \bar{N}_{n,t}^\pm)'$, where $'$ denotes the transpose operator. Finally, the total number of MOs sent by the high-frequency trader (HFT) up to time t is $\bar{\mathbf{N}}_t^\pm - \mathbf{N}_t^\pm$.

Thus, $N_t = \sum_{i=1}^n (N_{i,t}^- + N_{i,t}^+)$ is the total number of MOs executed by *other* agents up to time t , and $\bar{N}_t = \sum_{i=1}^n (\bar{N}_{i,t}^- + \bar{N}_{i,t}^+)$ denotes the total number of MOs executed by *all* agents up to time t . It is also convenient to have these processes in the following vector forms: $\mathbf{N}_t = (N_{1,t}, \dots, N_{n,t})'$ and $\bar{\mathbf{N}}_t = (\bar{N}_{1,t}, \dots, \bar{N}_{n,t})'$. Thus, the total number of MOs sent by the investor up to time t is $\bar{N}_t - N_t$.

The investor sends LOs to the LOB at depth δ_t^i , $i = 1, \dots, n$, and we summarize this information in the ($2n$ -dimensional) vector-valued process $\{\boldsymbol{\delta}_t\}_{0 \leq t \leq T}$. Here we use the convention that LO depth is from *across the book*: the limit buy order is measured from the best ask, and the limit sell order is measured from the best bid quote. Therefore, the investor's buy LOs are posted at $\mathbf{S}_t + \boldsymbol{\xi}/2 - \boldsymbol{\delta}_t^-$, and her sell LOs are posted at $\mathbf{S}_t - \boldsymbol{\xi}/2 + \boldsymbol{\delta}_t^+$. In addition, the running counts of the investor's filled LOs are given by the processes $\{M_{i,t}^+ : i = 1, \dots, n\}_{0 \leq t \leq T}$ and $\{M_{i,t}^- : i = 1, \dots, n\}_{0 \leq t \leq T}$ for sell and buy orders, respectively, and we summarize this information in the vector-valued processes $\{\mathbf{M}_t^+\}_{0 \leq t \leq T}$ and $\{\mathbf{M}_t^-\}_{0 \leq t \leq T}$. This allows the investor to post inside the spread.²

Assumption 2. The rate of arrival (called the fill rate functions) of MOs which hit/lift the agent's LOs for asset $i = 1, \dots, n$, as measured from "across the book", denoted $h_i^\pm(\delta)$, satisfy $h_i^\pm(0) < \infty$, $h_i^\pm(\xi_i) = \lambda_i^\pm$, and $\lim_{\delta \rightarrow \infty} \delta h_i^\pm(\delta) = 0$. Furthermore, $h_i^\pm \in C^1([0, \infty))$ with $h' < 0$.

The investor gets filled when MOs arrive and reach the level that the LO is posted at. More specifically, the deeper in the book the agent posts, the less frequently she gets filled by a matching MO; a dependence that is captured by the functions $h_i^\pm(\delta)$. This is in line with how a number of other authors have modeled optimal postings and fill probabilities as in Ho and Stoll (1981), Avellaneda and Stoikov (2008), Guéant et al. (2013), Bayraktar and Ludkovski (2014), and Cartea and Jaimungal (2015c), among others.

We may suppress the \pm and i indices on the function h and the scalar variables λ , κ (defined in Example 6), δ , and ξ , when it is clear which element is being represented.

3. The Investors's Optimization Problem

Recall that $\boldsymbol{\delta}_t$ is the depth, as measured from across the book, at which the investor sends quotes to the LOB. Let $(\boldsymbol{\tau}; \boldsymbol{\varsigma}) = \{(\tau_1, \varsigma_1), (\tau_2, \varsigma_2), \dots\}$ be an impulse control representing the investor's submission of MOs to the LOB. The ordered sequence of uniformly bounded \mathcal{F}_t stopping times, $0 \leq \tau_1 \leq \tau_2 \leq \dots \leq T$, is the sequence of times of the MOs, and the binary sequence $\varsigma_1, \varsigma_2, \dots$ indicates whether the MO is a buy or a sell order. Recall that the filtration $\{\mathcal{F}_t\}_{0 \leq t \leq T}$ is generated by the processes \mathbf{S}_t , \mathbf{M}_t^\pm , \mathbf{N}_t^\pm , $\overline{\mathbf{N}}_t^\pm$.

The cash process $X = \{X_t\}_{0 \leq t \leq T}$ of the investor satisfies the SDE

$$\begin{aligned} dX_t = & (\mathbf{S}_t - \boldsymbol{\xi}/2 + \boldsymbol{\delta}_{t-}^+) \cdot d\mathbf{M}_t^+ - (\mathbf{S}_t + \boldsymbol{\xi}/2 - \boldsymbol{\delta}_{t-}^-) \cdot d\mathbf{M}_t^- \\ & + (\mathbf{S}_t - \boldsymbol{\xi}/2) \cdot d(\overline{\mathbf{N}}_t^- - \mathbf{N}_t^-) - (\mathbf{S}_t + \boldsymbol{\xi}/2) \cdot d(\overline{\mathbf{N}}_t^+ - \mathbf{N}_t^+), \end{aligned} \quad (2)$$

²*Inside the spread* refers to LOs that improve upon the current best bid or offer.

where \mathbf{S}_t , $\boldsymbol{\delta}_t^-$, $\boldsymbol{\delta}_t^+$, and $\boldsymbol{\xi}$ are all n -dimensional row vectors, and \mathbf{M}_t^- , \mathbf{M}_t^+ , \mathbf{N}_t^- , \mathbf{N}_t^+ , $\overline{\mathbf{N}}_t^-$, $\overline{\mathbf{N}}_t^+$ are all n -dimensional column vectors.

The terms in the first line in (2) of the cash process account for the changes in cash due to filled limit sell and buy orders, and the terms in the second line account for changes in cash due to the investor's own market sell and buy orders.

The investor seeks the strategy $(\{\boldsymbol{\delta}_t \geq 0\}_{0 \leq t \leq T}, \boldsymbol{\tau}, \boldsymbol{\varsigma})$ which maximizes her expected terminal wealth while penalizing and constraining inventories, where $\boldsymbol{\delta}_t$ is a bounded \mathcal{F}_t -previsible process representing the LO quote depths, $\boldsymbol{\tau}$ is an increasing sequence of \mathcal{F}_t -stopping times (bounded above by T) representing the times of MO submissions, and $\boldsymbol{\varsigma}$ is the binary sequence of buy/sell indicators. The collection of all such strategies is called the admissible set and denoted by \mathcal{X} .

Thus, for $0 < T \ll \overline{T}$, the investor solves the control problem³

$$H(t, x, \mathbf{s}, \mathbf{q}) = \sup_{(\boldsymbol{\delta}_t, \boldsymbol{\tau}, \boldsymbol{\varsigma}) \in \mathcal{X}} \mathbb{E}_{t,x,\mathbf{s},\mathbf{q}} \left[X_T + \sum_{i=1}^n q_T^i \left(S_T^i - \text{sign}(q_T^i) \frac{\xi_i}{2} - \gamma q_T^i \right) - \phi \int_t^T \ell(\mathbf{q}_u) du \right], \quad (3)$$

where $\mathbb{E}_{t,x,\mathbf{s},\mathbf{q}}[\cdot]$ denotes the expectation conditional on $X_{t-} = x$, $\mathbf{S}_t = \mathbf{s}$, $\mathbf{q}_{t-} = \mathbf{q}$. Moreover, $\gamma \geq 0$, $\phi \geq 0$, $\mathbf{q}_t = (q_t^1, \dots, q_t^n)$ is the investor's inventory at time t , and $\ell(\cdot) \geq 0$ is an instantaneous inventory penalty.

In Equation (3), we have included two inventory penalties. The first, $q_T^i (-\text{sign}(q_T^i) \frac{\xi_i}{2} - \gamma q_T^i)$, is a cost that the investor incurs to unwind terminal inventory. To do so the investor uses an MO that crosses the spread, picks up the cost $\xi_i/2$, and walks down one side of the LOB where these costs are given by the linear "market impact" term γq_T^i . The second, $\int_t^T \ell(\mathbf{q}_u) du$, is a running inventory penalty which allows the agent to control her exposure to inventory risk through time, but is not a financial cost, see Cartea and Jaimungal (2015c) and Cartea et al. (2013). Here we assume that

$$\ell(\mathbf{q}) = \mathbb{V} \left[\mathbf{q} \cdot \widehat{\mathbf{S}}_1 \right] = \mathbf{q}' \boldsymbol{\sigma} \boldsymbol{\sigma}' \mathbf{q}, \quad (4)$$

where $\{\widehat{\mathbf{S}}_t\}_{0 \leq t \leq T}$ is the *unreflected* version of $\{\mathbf{S}_t\}_{0 \leq t \leq T}$.

A priori, it is unclear why the *unreflected* version of \mathbf{S}_t in the penalty function (4) is used. Classical mean-variance portfolio optimization problems penalize according to the variance of the terminal profit and loss (PnL). However, such a penalty fails to characterize how much inventory risk was taken on by the investor throughout the entire life of the strategy. In light

3

of this shortfall, we take the penalty to be the intra-trade quadratic variation of the investor's portfolio, similar to what was done in Forsyth et al. (2012), which is a natural characterization of the total risk taken on by the investor. The following observation gives a mathematical justification for the use of the inventory penalty in (4).

Observation 4. For $0 \leq u \leq v \leq T$, we have the following equivalence (\mathbb{P} almost surely):

$$\int_u^v \mathbb{V}[\mathbf{q}_t \cdot \widehat{\mathbf{S}}_1] dt = \langle \mathbf{q} \cdot \mathbf{S} \rangle_v - \langle \mathbf{q} \cdot \mathbf{S} \rangle_{t_{\overline{N}_v}} + \sum_{k=\overline{N}_u+1}^{\overline{N}_v} \{ \langle \mathbf{q} \cdot \mathbf{S} \rangle_{t_k} - \langle \mathbf{q} \cdot \mathbf{S} \rangle_{t_{k-1}} \}, \quad (5)$$

where t_k is the time of the k^{th} MO, t_{k-} is the left limit at t_k , and $\langle \cdot \rangle_t$ is the quadratic variation operator. Recall that $\{\widehat{\mathbf{S}}_t\}_{0 \leq t \leq T}$ is the unreflected version of $\{\mathbf{S}_t\}_{0 \leq t \leq T}$.

The relation in (5) is also highlighted in Cartea and Jaimungal (2015c) where the authors offer a similar relation for the one-dimensional case. Thus, the trader penalizes the strategy with the ‘‘variance per second’’ of her current portfolio, which is scaled by $\phi \geq 0$ according to her inventory-risk tolerance. In a single asset economy, this penalty reduces to $\int q_u^2 du$ (up to a multiplicative constant), which has been shown in Cartea et al. (2013) to stem from ambiguity aversion on the drift of the midprice process.

3.1. Solving the Value Function

By appealing to Øksendal and Sulem (2007), the value function $H(t, x, \mathbf{s}, \mathbf{q})$, as defined in (3), satisfies the following HJBQVI on the no-arbitrage region \mathcal{A} :

$$\max \left[\sup_{\delta > 0} \left\{ (\partial_t + \mathcal{L})H + \sum_i h_i^- (\delta_i^-) \Delta_i^- H + \sum_i h_i^+ (\delta_i^+) \Delta_i^+ H - \phi \ell(\mathbf{q}) \right\}; \right. \\ \left. \max_{i, \pm} \{ H(t, x \pm s_i - \xi_i/2, \mathbf{s}, \mathbf{q} \mp e_i) - H \} \right] = 0, \quad \mathbf{s} \in \mathcal{A}, \quad (6)$$

with terminal condition $H(T, \cdot) = x + \sum_{i=1}^n q_i \left(s_i - \text{sign}(q_i) \frac{\xi_i}{2} - \gamma q_i \right)$, the Neumann condition on the boundary of \mathcal{A}

$$\frac{\partial H}{\partial \mathbf{n}} = \nabla H \cdot \mathbf{n} = 0, \quad \forall \mathbf{s} \in \partial \mathcal{A}, \quad (7)$$

where \mathbf{n} is the normal vector to $\partial \mathcal{A}$, and where $\mathcal{L} = \frac{1}{2} \sum_{i,j=1}^n [\boldsymbol{\sigma} \boldsymbol{\sigma}']_{ij} \frac{\partial^2}{\partial s_i \partial s_j}$ is the infinitesimal generator of the n -dimensional correlated Brownian Motion \mathbf{B} . Finally, we use Δ_i^\pm to denote

the difference operator

$$\Delta_i^\pm \varphi(t, x, \mathbf{s}, \mathbf{q}) = \varphi(t, x \pm s_i \pm \delta_i^\pm \mp \xi_i/2, \mathbf{s}, \mathbf{q} \mp e_i) - \varphi(t, x, \mathbf{s}, \mathbf{q}),$$

where e_i is a vector with 1 in the i^{th} position and 0 elsewhere, and dependence on $\boldsymbol{\delta}$ is suppressed for convenience. For example, $\Delta_i^- \varphi$ denotes the change in φ due to the investor's limit buy order being filled: inventory for asset i increases by 1, and cash x decreases by the cost of the transaction $s_i - \delta_i^- + \xi_i/2$. The interpretation for $\Delta_i^+ \varphi$ when the investor's limit sell order is filled has a similar interpretation.

Lemma 5. *If the midprice process is in the no-arbitrage region (i.e. $\mathbf{s} \in \mathcal{A}$), then the supremum term in (6) is equal to the sum of the supremum over each component.*

PROOF. The supremum operator satisfies countable sub-additivity. Equality comes from the fact that the $2n$ variables being optimized appear exactly once in all of $2n$ distinct additive terms. \square

We apply the following ansatz for the value function

$$H(t, x, \mathbf{s}, \mathbf{q}) = x + \mathbf{q} \cdot \mathbf{s} + g(t, \mathbf{s}, \mathbf{q}), \quad (8)$$

and recalling that $h(\delta)$ is the fill rate function facing the investor for an LO placed at depth δ , the first order condition for each component of $\boldsymbol{\delta}$ in the continuation region is given by

$$h'(\delta_i^{\pm,*}) \left(\delta_i^{\pm,*} - \frac{\xi_i}{2} + \Delta_i^\pm g \right) + h_i^\pm(\delta_i^{\pm,*}) = 0. \quad (9)$$

The implicit form in Equation (9) is then used to simplify (6) and obtain a new QVI for $g(\cdot)$:

$$\max \left[(\partial_t + \mathcal{L})g - \sum_{i,\pm} \frac{(h(\delta^*))^2}{h'(\delta^*)} - \phi \ell(\mathbf{q}), \max_{i,\pm} \{ \Delta_i^\pm g - \xi_i/2 \} \right] = 0 \quad (10)$$

with terminal condition $g(T, \cdot) = -\sum q_i \left(\text{sign}(q_i) \frac{\xi_i}{2} + \gamma q_i \right)$ and boundary condition $(\mathbf{q} + \nabla g) \cdot \mathbf{n} = 0$ on $\partial \mathcal{A}$.

In order to obtain an explicit form for $\boldsymbol{\delta}^*$ (in feedback control form), we further assume that $h(\delta)$ is exponential.

Example 6. Exponential Fill Rate. Suppose that MOs arrive to the market according to a Poisson process and the probability of being filled decays exponentially. More precisely, $h_i^\pm(\delta) = \lambda_i^\pm e^{-\kappa_i^\pm(\delta - \xi_i)}$ with $0 < \kappa_i^\pm \leq \frac{1}{\xi_i}$.

Example 6 provides an additional modeling assumption⁴ on the fill rate facing the investor that will be used in all numerical demonstrations throughout this paper. The additional condition that $\kappa_i^\pm \leq \frac{1}{\xi_i}$ is imposed so that a risk-neutral investor following an optimal market making strategy does not post limit quotes inside the spread.

In this exponential fill rate setting, we have

$$\delta_i^{\pm,*} = \frac{1}{\kappa_i^\pm} + \frac{\xi_i}{2} - \Delta_i^\pm g, \quad (11)$$

and Equation (10), in the continuation region, further reduces to

$$g_t + \frac{1}{2} \sum_{i,j=1}^n [\sigma\sigma']_{ij} \frac{\partial^2 g}{\partial s_i \partial s_j} + \sum_{i,\pm} \frac{\lambda_i^\pm}{\kappa_i^\pm} \exp \left\{ -1 + \frac{\kappa_i^\pm \xi_i}{2} + \kappa_i^\pm \Delta_i^\pm g \right\} - \phi \ell(\mathbf{q}) = 0. \quad (12)$$

The optimal LO depth has two main components. The first is given by the two first terms on the right-hand side in (11) which is the optimal depth chosen by an investor who does not impose any inventory penalty, nor does she pick up any costs for unwinding terminal inventory, $\gamma = \phi = 0$. The second component, given by the third term on the right-hand of the equation, is a correction to the optimal posting to account for the adjustment to the LO depth as a result of the investor's inventory management strategy.

Proposition 7. The $\max_{i,\pm} \{\Delta_i^\pm g - \xi_i/2\}$ component in the QVI (10) give rise to the LO spread bounds $\frac{1}{\kappa_i^\pm} < \delta_i^\pm \leq \frac{1}{\kappa_i^\pm} + \xi_i$ for each $i = 1, \dots, n$.

When considering δ^- , the stated lower bound is because the agent is eager to obtain another unit of the asset. Therefore, rather than posting an LO at the strict lower bound of $1/\kappa_i^\pm$, she executes an MO. The upper bound corresponds to posting at a deep level (since the investor does not want the asset), and if the LO is filled, the agent immediately posts an MO to unwind and attempts to make a profit. Profits in this case are not guaranteed and depend on the resilience of the LOB. Note that $\delta_i^-(\mathbf{q})$ is arbitrarily close to $1/\kappa_i^\pm$ if and only if $\delta_i^+(\mathbf{q} + e_i)$ is close to $\frac{1}{\kappa_i^\pm} + \xi_i$ and similarly, $\delta_i^+(\mathbf{q})$ is arbitrarily close to $1/\kappa_i^\pm$ if and only if $\delta_i^-(\mathbf{q} - e_i)$ is

⁴It is also common for a power law function to be used as the fill rate function.

close to $\frac{1}{R_i} + \xi_i$.

4. Viscosity Solution and Numerical Scheme for Exponential Fill Rate

This section outlines the technical details regarding the existence of a unique viscosity solution to the dynamic programming equation (DPE), specifies a dimensional reduction in the spacial variables, gives explicit details regarding the discrete numerical scheme, and proves that the discretized value function converges to the unique viscosity solution.

4.1. Dimensional Reduction of Spatial Variables

We first prove the following lemma to demonstrate a dimensional reduction in the spacial variable \mathbf{s} . Due to this lemma, the relevant \mathbf{s} values live in a compact domain. For a fixed error, the total computational run time will then be proportional to $c^{\dim B}(2Q+1)^n$, where c is some constant that depends on the spreads $\xi_1, \xi_2, \dots, \xi_n$, the basis vectors that span B (or equivalently on the number of discrete steps on the dimensionally reduced \mathbf{s} -grid), and Q is the maximum inventory the investor can be long or short for each asset.

Lemma 8. *There exists an C^2 , surjective function $\Gamma : \bar{\mathcal{A}} \rightarrow \bar{\mathcal{A}}'$ such that:*

(I) For all $\mathbf{s} \in \bar{\mathcal{A}}$, $g(\mathbf{s}, \cdot) = g((\Gamma(\mathbf{s}), \mathbf{0}), \cdot)$,

(II) $\mathbf{s} \in \partial\mathcal{A} \iff \Gamma(\mathbf{s}) \in \partial\mathcal{A}'$,

(III) $\bar{\mathcal{A}}'$ is the cartesian product of closed, bounded intervals.

PROOF. Define $m := \dim B$. Then $B = \text{span}\{b_1, \dots, b_m\}$, where b_1, \dots, b_m is an orthogonal basis, and for each $i = 1, \dots, m$ we have (wlog) $\mathbb{P}[b_i \cdot \mathbf{S}_{\bar{T}} = 0] = 1$. No arbitrage (i.e. $\mathbf{s} \in \mathcal{A}$) is equivalent to the system of inequalities

$$\mathbf{y} \cdot \mathbf{s} + \sum_{j=1}^n |y_j| \frac{\xi_j}{2} > 0 \quad \text{for } \mathbf{y} = (\pm b_1, \dots, \pm b_m)',$$

which simplifies to

$$-\sum_{j=1}^n |y_j| \frac{\xi_j}{2} < \mathbf{y} \cdot \mathbf{s} < \sum_{j=1}^n |y_j| \frac{\xi_j}{2} \quad \text{for } \mathbf{y} = (b_1, \dots, b_m)'. \quad (13)$$

Because the b_i 's are linearly independent, one can see that (13) is equivalent to $\Gamma(\mathbf{s})$ being inside some rectangular subset of \mathbb{R}^m for the function $\Gamma(\mathbf{s}) = [b_j \cdot \mathbf{s}]_j$, which is also clearly C^2 . At this point, (II) should be obvious from (13) and the definition of Γ , (III) is proven, and linear independence of the basis vectors gives (IV).

To prove (I), we must appeal to some invariance properties of g . Notice that $\bar{\mathcal{A}}$ can now be reparameterized to the product space $\prod_{i=1}^m \left[-\sum_{j=1}^n |b_i \cdot e_j| \frac{\xi_j}{2}, \sum_{j=1}^n |b_i \cdot e_j| \frac{\xi_j}{2} \right] \times \mathbb{R}^{n-m}$. Looking at (10) and the Neumann boundary condition (only depends on \mathbf{s} via the direction of the normal vector, which only depends on which piece of the boundary \mathbf{s} is on; this can be determined uniquely via $\Gamma(\mathbf{s})$ by taking a limiting sequence in \mathbf{s} to the boundary), we see that under this new parameterization, g only depends on \mathbf{s} through the first m components. Hence, if $\mathbf{s}_1, \mathbf{s}_2 \in \bar{\mathcal{A}}$ are such that $\Gamma(\mathbf{s}_1) = \Gamma(\mathbf{s}_2)$, then $g(\mathbf{s}_1) = g(\mathbf{s}_2)$ and (I) is proved. \square

Lemma 8 states that it is not necessary to know \mathbf{s} to obtain the value of $g(\mathbf{s}, \cdot)$; we only require the information contained in $\Gamma(\mathbf{s})$. Furthermore, there exists such a function Γ that provides a nice, smooth mapping between the original n -dimensional no-arbitrage region $\bar{\mathcal{A}}$ and the new n' -dimensional no-arbitrage region $\bar{\mathcal{A}}'$, and that this new region is a bounded rectangle. The fact that the new region is bounded improves the accuracy of the numerical scheme discussed below in Section 4.2 since we no longer have to approximate the unbounded set $\bar{\mathcal{A}}$ by a bounded one.

There is, however, a minor detail that has been overlooked. The form $(\Gamma(\mathbf{s}), \mathbf{0})$, where $\mathbf{0}$ is an $(n - n')$ -dimensional vector of 0's, requires the assets to be ordered in a certain way. In an arbitrary ordering, the components of $(\Gamma(\mathbf{s}), \mathbf{0})$ will need to be re-ordered. We continue with the labels that produce a simplification of this form as it is more compact to write $(\Gamma(\mathbf{s}), \mathbf{0})$ instead of having the 0's in arbitrary positions (i.e. not consecutive nor all at the beginning or end).

In some sense, $\Gamma(\mathbf{s})$ is *sufficient*⁵ for g . That is, $\Gamma(\mathbf{s})$ contains the same amount of information as \mathbf{s} when computing g . Furthermore, one could say that $\Gamma(\mathbf{s})$ is *minimal sufficient* since it can be written as a function of any other function of \mathbf{s} that is sufficient for g . Stated heuristically, $\Gamma(\mathbf{s})$ contains the smallest amount of information about \mathbf{s} that still allows us to know $g(\mathbf{s})$ with certainty. The construction of Γ in the proof of Lemma 8 is used in what follows to obtain a

⁵In mathematical statistics, a statistic U is said to be *sufficient* for an unknown model parameter Θ if the conditional distribution of the data, conditional on U , is independent of Θ . Heuristically speaking, U contains the same amount of information about Θ as the data itself.

dimensional reduction, and is given by

$$\Gamma(\mathbf{s}) = [b_j \cdot \mathbf{s}]_j. \quad (14)$$

In our 2-dimensional illustrative example, we take $\Gamma(\mathbf{s}) = s_1 - s_2$. For the 3-dimensional example with two nearby futures and the 1-2 calendar spread, one could take $\Gamma(\mathbf{s}) = s_1 - s_2 - s_3$. More details for these two examples are provided in Appendix A.

The following Corollary states that the function g is invariant to movements in the direction of any element of B^\perp . This would correspond to vectors of the form $(1, 1)$ in our illustrative example (i.e. g remains unchanged when the futures and its underlying both increase or decrease by exactly the same amount).

Corollary 9. $D_v g = 0$ on $\overline{\mathcal{A}}$ whenever $v \in B^\perp$.

PROOF. Let $\mathbf{s} \in \overline{\mathcal{A}}$ and $v \in B^\perp$. Note that we necessarily have $\mathbf{s} + v \in \overline{\mathcal{A}}$ by Equation (13) since $b_i \cdot v = 0$ for $i = 1, \dots, n'$. We then have $\Gamma(\mathbf{s}) = \Gamma(\mathbf{s} + v)$ by (14). But $v \in B^\perp$ is arbitrary and B^\perp is a subspace, so one can consider a limiting sequence of the form $\{\eta_k v\}$, with the scalar sequence $\eta_k \downarrow 0$, which yields $D_v g = 0$. \square

4.2. Description of Numerical Scheme

We numerically solve (12) by applying upper and lower bounds to each component in \mathbf{q} , and solving the corresponding ODEs backwards in time using an explicit numerical scheme. For any fixed \mathbf{q} , we have an n -dimensional discrete grid to solve for various values of \mathbf{s} . Once we have this grid for each \mathbf{q} , we propagate one step backwards in t and repeat the process.

Define the domain for \mathbf{s} , $\mathcal{A}' \subset \mathbb{R}^{n'}$, on the dimensionally reduced grid. Let \mathbb{T}_{dt} be a regular partition of $[0, T]$, \mathbb{S}_{ds} be an n' -dimensional regular rectangular partition of the dimensionally reduced \mathbf{s} -grid $\overline{\mathcal{A}'}$, and \mathbb{Y}_Q be the subset of \mathbb{Z}^n that corresponds to inventory with maximum Q as previously described (i.e. $\mathbb{Y}_Q = \mathbb{Z}^n \cap [-Q, Q]^n$). Throughout the rest of this paper, $dt > 0$ and $ds > 0$ refer to the discrete step sizes in the time (t) and space ($\Gamma(\mathbf{s})$) variables, respectively.

Define, for any function $\varphi : [0, T] \times \overline{\mathcal{A}'} \times \mathbb{Z}^n \rightarrow \mathbb{R}$, the discrete dynamic evolution operator

$$\mathcal{S}^{dt, ds, Q}(t, \mathbf{s}, \mathbf{q}, \varphi) = \max \left[\mathcal{T}^{dt, ds, Q}(t, \mathbf{s}, \mathbf{q}, \varphi), \mathcal{M}^{dt, ds, Q}(t, \mathbf{s}, \mathbf{q}, \varphi) \right],$$

with the component corresponding to the continuation region given by

$$\mathcal{T}^{dt,ds,Q}(t, \mathbf{s}, \mathbf{q}, \varphi) = \varphi + dt \left[\frac{1}{2} \sum_{i,j=1}^n [\boldsymbol{\sigma}\boldsymbol{\sigma}']_{ij} \widehat{\mathcal{T}}_{ij}^{dt,ds,Q}(\cdot, \varphi) + \sum_{i,\pm} \frac{\lambda}{\kappa e} \exp \left\{ \frac{\kappa \xi}{2} + \kappa \Delta \varphi \right\} - \phi \ell(\mathbf{q}) \right],$$

$$\widehat{\mathcal{T}}_{ij}^{dt,ds,Q}(t, \mathbf{s}, \mathbf{q}, \varphi) = \frac{\varphi(\Gamma(s_i + ds, s_j + ds), \cdot) - \varphi(\Gamma(s_i + ds), \cdot) - \varphi(\Gamma(s_j + ds), \cdot) + \varphi(\cdot)}{ds^2},$$

and the component corresponding to the impulse region given by

$$\mathcal{M}^{dt,ds,Q}(t, \mathbf{s}, \mathbf{q}, \varphi) = \max_{i,\pm} \{ \varphi(t, \mathbf{s}, \mathbf{q} \pm e_i) - \xi_i/2 \}.$$

We approximate the solution g by $g^{dt,ds,Q} : \mathbb{T}_{dt} \times \mathbb{S}_{ds} \times \mathbb{Y}_Q \rightarrow \mathbb{R}$ defined by:

$$\begin{cases} g^{dt,ds,Q}(T, \mathbf{s}, \mathbf{q}) &= - \sum q_i \left(\text{sign}(q_i) \frac{\xi_i}{2} + \gamma q_i \right) \\ g^{dt,ds,Q}(dt \cdot k, \mathbf{s}, \mathbf{q}) &= \mathcal{S}^{dt,ds,Q}(dt \cdot (k+1), \mathbf{s}, \mathbf{q}, g^{dt,ds,Q}). \end{cases} \quad (15)$$

We employ an explicit scheme backwards in time as follows:

- When $t = T$, for each \mathbf{s} and \mathbf{q} on their respective grids, set

$$g^{dt,ds,Q}(T, \mathbf{s}, \mathbf{q}) = - \sum q_i \left(\text{sign}(q_i) \frac{\xi_i}{2} + \gamma q_i \right).$$

- For $t < T$ and $\mathbf{s} \in \mathcal{A}'$, we perform the following:
 - Compute $\mathcal{T}^{dt,ds,Q}(t + dt, \mathbf{s}, \mathbf{q}, g^{dt,ds,Q})$ and store the optimal control δ . Some details for handling the mixed partial term are given in Appendix A.1.
 - Compute $\mathcal{M}^{dt,ds,Q}(t + dt, \mathbf{s}, \mathbf{q}, g^{dt,ds,Q})$ and store the optimal impulse τ .
 - If $\mathcal{T}^{dt,ds,Q}(t + dt, \mathbf{s}, \mathbf{q}, g^{dt,ds,Q}) \geq \mathcal{M}^{dt,ds,Q}(t + dt, \mathbf{s}, \mathbf{q}, g^{dt,ds,Q})$, then set $g^{dt,ds,Q}(t, \mathbf{s}, \mathbf{q}) = \mathcal{T}^{dt,ds,Q}(t + dt, \mathbf{s}, \mathbf{q}, g^{dt,ds,Q})$. Otherwise, set $g^{dt,ds,Q}(t, \mathbf{s}, \mathbf{q}) = \mathcal{M}^{dt,ds,Q}(t + dt, \mathbf{s}, \mathbf{q}, g^{dt,ds,Q})$.
- For $t < T$ and $\mathbf{s} \in \partial \mathcal{A}'$, we perform the following:
 - Numerically approximate the condition $(\mathbf{q} + \nabla g) \cdot \mathbf{n} = 0$, similar to what is done in Appendix A.3 and Appendix A.4 for the $n = 2$ and $n = 3$ examples, respectively.

4.3. Convergence of the Numerical Scheme

We require the following bounds on the value function to ensure the existence of a unique viscosity solution to the DPE.

Lemma 10. *There exists constants $c_0, c_1, c_2 \in \mathbb{R}$ such that*

$$x + \sum_i q_i \left(s_i - \text{sign}(q_i) \frac{\xi_i}{2} \right) \leq H \leq x + \mathbf{q} \cdot \mathbf{s} + c_0(T - t) + c_1 \|\mathbf{q}\|_1 + c_2. \quad (16)$$

PROOF. The lower bound is attained by the strategy that immediately liquidates and then does nothing until time T . For the upper bound we have

$$\begin{aligned} H &\leq \sup_{(\boldsymbol{\delta}, \boldsymbol{\tau}) \in \mathcal{U}} \mathbb{E}_t \left[X_T + \mathbf{q}_T \cdot \mathbf{S}_T - \phi \int_t^T \ell(\mathbf{q}_u) du \right] \Big|_{\mathbf{q}=0} + \mathbf{q} \cdot \mathbf{s} + \|\mathbf{q}\|_1 \sum_i \xi_i \\ &\leq \sup_{(\boldsymbol{\delta}, \boldsymbol{\tau}) \in \mathcal{U}'} \mathbb{E}_t \left[X_T + \mathbf{q}_T \cdot \mathbf{S}_T - \phi \int_t^T \ell(\mathbf{q}_u) du \right] \Big|_{\mathbf{q}=0} + \mathbf{q} \cdot \mathbf{s} \\ &\quad + (T - t) \left(\sum_i \xi_i + \sup_{\delta, i, \pm} \delta h(\delta) \right) + \|\mathbf{q}\|_1 \sum_i \xi_i \\ &= x + \sum_i \xi_i \sup_{\omega, t, i} (\mathbf{q}_t \cdot \mathbf{e}_i) c_3(\sigma) \sqrt{T - t} \mathbb{E}[\chi_1] + (T - t) \sum_i \xi_i \sup_{\delta, i, \pm} \delta h(\delta) \\ &\quad + \mathbf{q} \cdot \mathbf{s} + (T - t) \left(\sum_i \xi_i + \sup_{\delta, i, \pm} \delta h(\delta) \right) + \|\mathbf{q}\|_1 \sum_i \xi_i, \end{aligned}$$

where χ_u is the maximum (in absolute value) of a standard n -dimensional Brownian motion up to time u and $\mathcal{U}' \subset \mathcal{U}$ is the subset of admissible controls with $\boldsymbol{\delta} = \infty$ whenever $\|\mathbf{q}\| = 0$. Finally, the result follows from Assumption 2. \square

Although this upper bound is not necessarily sharp, it is sufficient to continue. We now turn our attention to each term in the upper bound in this proof. The terms $\mathbf{q} \cdot \mathbf{s} + \|\mathbf{q}\|_1 \sum_i \xi_i$ correspond to free liquidation to $\mathbf{q} = 0$ with an added bonus of $\sum_i \xi_i$ per unit of each open position. The investor cannot do better than this (or even come close) on an expected value basis. The term $(T - t) \left(\sum_i \xi_i + \sup_{\omega, t, i, \pm} \delta h(\delta) \right)$ is an upper bound on the profit of all future open positions that use filled LOs. Lastly, we must consider all MO initiated open positions (for the case where $\mathbf{q} = 0$ is not an element of the continuation region). If these positions are aided with filled LOs, then we have already accounted for such events. Otherwise, in order to produce a profit, the Brownian motion must travel a certain distance from where the open position was initiated. The total number of times this happens is proportional to $\sqrt{T - t}$, and $\phi > 0$ guarantees that $\mathbf{q}_t(\omega)$ is uniformly bounded. We are able to bound the expected round trip profit because the dimensionally reduced no-arbitrage region is bounded.

Similar to Guilbaud and Pham (2015), we first prove monotonicity, stability, and consistency properties (Propositions 11, 12, and 13, respectively) of the discrete dynamic evolution operator $\mathcal{S}^{dt, ds, Q}(t, \mathbf{s}, \mathbf{q}, \cdot)$. These properties are critical in the proof of Theorem 14, which states that the solution to the numerical scheme, as given by (15), converges locally uniformly to the unique viscosity solution $g(t, \mathbf{s}, \mathbf{q})$.

Proposition 11. (*Monotonicity*) *For any $ds > 0$, there exists $h_t > 0$ (depends on ds) such that for any $0 < dt < h_t$, $Q \in \mathbb{N}$, and $\varphi_1, \varphi_2 \in C_b^{1,2}([0, T] \times \overline{\mathcal{A}}' \times \mathbb{Z}^n)$ with $\varphi_1 \leq \varphi_2$, we have $\mathcal{S}^{dt, ds, Q}(t, \mathbf{s}, \mathbf{q}, \varphi_1) \leq \mathcal{S}^{dt, ds, Q}(t, \mathbf{s}, \mathbf{q}, \varphi_2)$.*

PROOF. Let $\varphi_1 \leq \varphi_2$, $Q \in \mathbb{N}$, and $ds > 0$. Monotonicity of the operator \mathcal{M} is obvious. Because φ and its derivatives are continuous, and the domain of φ is compact, there exists finite constants $A_1, A_2 > 0$ such that

$$e^{\kappa\Delta\varphi_2} - e^{\kappa\Delta\varphi_1} \geq (e^{-\kappa\varphi_2} - e^{-\kappa\varphi_1})e^{\kappa\mathbb{S}\varphi_1} = (e^{-\kappa(\varphi_2 - \varphi_1)} - 1)e^{\kappa\Delta\varphi_1} \geq -A_1(\varphi_2 - \varphi_1), \quad (17)$$

where A_1 does *not* depend on φ_1 or φ_2 , \mathbb{S} is the shift operator associated with Δ , and

$$\hat{\mathcal{T}}_{ij}^{dt, ds, Q}(t, \mathbf{s}, \mathbf{q}, \varphi_2) - \hat{\mathcal{T}}_{ij}^{dt, ds, Q}(t, \mathbf{s}, \mathbf{q}, \varphi_1) \geq -A_2(\varphi_2 - \varphi_1), \quad (18)$$

where A_2 only depends on ds .

Using (17) and (18) gives

$$\begin{aligned} & \mathcal{T}^{dt, ds, Q}(t, \mathbf{s}, \mathbf{q}, \varphi_2) - \mathcal{T}^{dt, ds, Q}(t, \mathbf{s}, \mathbf{q}, \varphi_1) \\ &= \varphi_2 - \varphi_1 + dt \frac{1}{2} \sum_{i, j=1}^n [\sigma\sigma']_{ij} \left(\hat{\mathcal{T}}_{ij}^{dt, ds, Q}(\cdot, \varphi_2) - \hat{\mathcal{T}}_{ij}^{dt, ds, Q}(\cdot, \varphi_1) \right) \\ & \quad + dt \sum_{i, \pm} \frac{\lambda}{\kappa} e^{-1 + \kappa\xi/2} (e^{\kappa\Delta\varphi_2} - e^{\kappa\Delta\varphi_1}) \\ & \geq (\varphi_2 - \varphi_1) \left[1 - dt \left\{ A_1 \sum_{i, \pm} \frac{\lambda}{\kappa} e^{-1 + \kappa\xi/2} + \frac{A_2}{2} \sum_{i, j=1}^n [\sigma\sigma']_{ij} \right\} \right]. \end{aligned}$$

Taking

$$dt < \left(A_1 \sum_{i,\pm} \frac{\lambda}{\kappa} e^{-1+\kappa\xi/2} + \frac{A_2}{2} \sum_{i,j=1}^n [\boldsymbol{\sigma}\boldsymbol{\sigma}']_{ij} \right)^{-1} \quad (19)$$

gives monotonicity of \mathcal{T} . \square

Proposition 12. (*Stability*) For any $dt, ds > 0$ and $Q \in \mathbb{N}$, there exists a unique solution $g^{dt,ds,Q}(t, \mathbf{s}, \mathbf{q})$ to (15). Furthermore, the sequence $\{g^{dt,ds,Q}\}$ is uniformly bounded. More formally, there exists $c_0, c_1, c_2 \in \mathbb{R}$ such that for any $(t, \mathbf{s}, \mathbf{q}) \in \mathbb{T}_{dt} \times \mathbb{S}_{ds} \times \mathbb{Y}_Q$,

$$-\sum_i q_i \text{sign}(q_i) \frac{\xi_i}{2} \leq g^{dt,ds,Q}(t, \mathbf{s}, \mathbf{q}) \leq c_0(T-t) + c_1\|\mathbf{q}\|_1 + c_2. \quad (20)$$

PROOF. Existence and uniqueness follows immediately from the definition of the explicit discretization scheme given in (15). The upper bound is a direct consequence of Lemma 10 since $g^{dt,ds,Q}(t, \mathbf{s}, \mathbf{q}) \leq g(t, \mathbf{s}, \mathbf{q}) = H - x - \mathbf{q} \cdot \mathbf{s} \leq c_0(T-t) + c_1\|\mathbf{q}\|_1 + c_2$. The lower bound is attained by immediate liquidation of the current inventory and then doing nothing until $t = T$. \square

Proposition 13. (*Consistency*) For all $(t, \mathbf{s}, \mathbf{q}) \in [0, T] \times \overline{\mathcal{A}}' \times \mathbb{Z}^n$ and $\varphi \in C_b^{1,2}([0, T] \times \overline{\mathcal{A}}' \times \mathbb{Z}^n)$, we have

$$\begin{aligned} \lim_{\substack{(dt, ds, Q) \rightarrow (0, 0, \infty) \\ (t', \mathbf{s}') \rightarrow (t, \mathbf{s})}} \frac{1}{dt} [\mathcal{T}^{dt,ds,Q}(t' + dt, \mathbf{s}', \mathbf{q}, \varphi) - \varphi(t', \mathbf{s}', \mathbf{q})] = \\ \frac{\partial \varphi}{\partial t} + \frac{1}{2} \sum_{i,j=1}^n [\boldsymbol{\sigma}\boldsymbol{\sigma}']_{ij} \frac{\partial^2 \varphi}{\partial s_i \partial s_j} + \sum_{i,\pm} \frac{\lambda}{\kappa} \exp \left\{ -1 + \frac{\kappa\xi}{2} + \kappa\Delta\varphi \right\} - \phi \ell(\mathbf{q}) \end{aligned} \quad (21)$$

and

$$\lim_{\substack{(dt, ds, Q) \rightarrow (0, 0, \infty) \\ (t', \mathbf{s}') \rightarrow (t, \mathbf{s})}} \mathcal{M}^{dt,ds,Q}(t' + dt, \mathbf{s}', \mathbf{q}, \varphi) = \max_{i,\pm} \{ \Delta_i^\pm \varphi - \xi_i/2 \}. \quad (22)$$

PROOF. Equation (22) follows by continuity of φ . To prove (21), let $(t', \mathbf{s}', \mathbf{q}) \in [0, T] \times \overline{\mathcal{A}}' \times \mathbb{Z}^n$ and observe:

$$\begin{aligned}
\frac{1}{dt} \left[\mathcal{T}^{dt, ds, Q}(t' + dt, \mathbf{s}', \mathbf{q}, \varphi) - \varphi(t', \mathbf{s}', \mathbf{q}) \right] &= \frac{1}{dt} [\varphi(t' + dt, \mathbf{s}', \mathbf{q}) - \varphi(t', \mathbf{s}', \mathbf{q})] - \phi \ell(\mathbf{q}) \\
&+ \frac{1}{2} \sum_{i,j=1}^n [\sigma \sigma']_{ij} \hat{\mathcal{T}}_{ij}^{dt, ds, Q}(t' + dt, \mathbf{s}', \mathbf{q}, \varphi) \\
&+ \sum_{i, \pm} \frac{\lambda}{\kappa} \exp \left\{ -1 + \frac{\kappa \xi}{2} + \kappa \Delta \varphi(t' + dt, \mathbf{s}', \mathbf{q}) \right\}.
\end{aligned}$$

The only term that does not trivially converge is $\Delta \varphi(t' + dt, \mathbf{s}', \mathbf{q})$ (the other terms clearly converge since $\varphi \in C_b^{1,2}$). But $\Delta \varphi(t' + dt, \mathbf{s}', \mathbf{q})$ is also continuous in (t, \mathbf{s}) whenever $Q \geq \|\mathbf{q}\|_\infty + 1$, thus proving (21). \square

Observe that we do not need to include a sequence of inventory levels due to the fact that \mathbf{q} resides on the sparse grid \mathbb{Z}^n . More formally, one could only consider the limiting sequence $(t' + dt, \mathbf{s}', \mathbf{q}') \rightarrow (t, \mathbf{s}, \mathbf{q})$ sufficiently far in the tail when $\mathbf{q}' = \mathbf{q}$ onwards.

Theorem 14. (Convergence) *The solution $g^{dt, ds, Q}$ to the numerical scheme (15) converges locally uniformly to the unique viscosity solution $g(t, \mathbf{s}, \mathbf{q})$ on $[0, T] \times \overline{\mathcal{A}'} \times \mathbb{Z}^n$, as $(dt, ds, Q) \rightarrow (0, 0, \infty)$, provided $dt < f(ds)$ for some $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$.*

PROOF. Follows from Propositions 11, 12, 13, and Barles and Souganidis (1991). For completeness, the details are provided.

The limits

$$\begin{aligned}
g_*(t, \mathbf{s}, \mathbf{q}) &:= \liminf_{\substack{(dt, ds, Q) \rightarrow (0, 0, \infty) \\ (t', \mathbf{s}') \rightarrow (t, \mathbf{s})}} g^{dt, ds, Q}(t', \mathbf{s}', \mathbf{q}) \\
g^*(t, \mathbf{s}, \mathbf{q}) &:= \limsup_{\substack{(dt, ds, Q) \rightarrow (0, 0, \infty) \\ (t', \mathbf{s}') \rightarrow (t, \mathbf{s})}} g^{dt, ds, Q}(t', \mathbf{s}', \mathbf{q})
\end{aligned}$$

are lower/upper semi-continuous functions on $[0, T] \times \overline{\mathcal{A}'} \times \mathbb{Z}^n$, and inherit the boundedness

of $\{g^{dt, ds, Q}\}$ by the stability result in Proposition 12. It suffices to show (by symmetry and the strong comparison principle) that g_* possesses the viscosity supersolution property. To this end, let $(\bar{t}, \bar{s}, \bar{\mathbf{q}}) \in [0, T) \times \bar{\mathcal{A}}' \times \mathbb{Z}^n$ and $\varphi \in C_b^{1,2}([0, T] \times \bar{\mathcal{A}}' \times \mathbb{Z}^n)$ such that $(\bar{t}, \bar{s}, \bar{\mathbf{q}})$ attains the strict global minimum of $g_* - \varphi$. Then there exists sequences $\{(t'_k, \mathbf{s}'_k, \mathbf{q}'_k)\}_k$ in $[0, T) \times \bar{\mathcal{A}}' \times \mathbb{Z}^n$ and $\{(dt_k, ds_k, Q_k)\}_k$ such that

$$\begin{aligned} (t'_k, \mathbf{s}'_k, \mathbf{q}'_k) &\longrightarrow (\bar{t}, \bar{s}, \bar{\mathbf{q}}), \\ (dt_k, ds_k, Q_k) &\longrightarrow (0, 0, \infty), \\ g^{dt_k, ds_k, Q_k} &\longrightarrow g_*(\bar{t}, \bar{s}, \bar{\mathbf{q}}), \end{aligned}$$

and $(t'_k, \mathbf{s}'_k, \mathbf{q}'_k)$ is the global minimizer of $g^{dt_k, ds_k, Q_k} - \varphi$.

Observe that for us to apply the monotonicity result in Proposition 11, the sequence $\{(dt_k, ds_k, Q_k)\}_k$ must have dt_k sufficiently small for each ds_k , as prescribed by (19).

Denote the sequence of minimum values $\{\beta_k\}_k$ by $\beta_k := (g^{dt_k, ds_k, Q_k} - \varphi)(t'_k, \mathbf{s}'_k, \mathbf{q}'_k)$. Then we have the relation $g^{dt_k, ds_k, Q_k} \geq \varphi + \beta_k$. Applying the definition of the numerical scheme in (15) and the monotonicity result from Proposition 11 gives:

$$\begin{aligned} \beta_k + \varphi(t'_k, \mathbf{s}'_k, \mathbf{q}'_k) &= g^{dt_k, ds_k, Q_k}(t'_k, \mathbf{s}'_k, \mathbf{q}'_k) \\ &= \mathcal{S}^{dt_k, ds_k, Q_k}(t'_k + dt_k, \mathbf{s}'_k, \mathbf{q}'_k, g^{dt_k, ds_k, Q_k}) \\ &\geq \mathcal{S}^{dt_k, ds_k, Q_k}(t'_k + dt_k, \mathbf{s}'_k, \mathbf{q}'_k, \varphi + \beta_k) \\ &= \mathcal{S}^{dt_k, ds_k, Q_k}(t'_k + dt_k, \mathbf{s}'_k, \mathbf{q}'_k, \varphi) + \beta_k \\ &= \max \left[\mathcal{T}^{dt_k, ds_k, Q_k}(t'_k + dt_k, \mathbf{s}'_k, \mathbf{q}'_k, \varphi), \mathcal{M}^{dt_k, ds_k, Q_k}(t'_k + dt_k, \mathbf{s}'_k, \mathbf{q}'_k, \varphi) \right] + \beta_k. \end{aligned}$$

We then have the relation

$$\min \left\{ \frac{1}{dt_k} \left[\varphi(t'_k, \mathbf{s}'_k, \mathbf{q}'_k) - \mathcal{T}^{dt_k, ds_k, Q_k}(t'_k + dt_k, \mathbf{s}'_k, \mathbf{q}'_k, \varphi) \right], \right. \\ \left. \varphi(t'_k, \mathbf{s}'_k, \mathbf{q}'_k) - \mathcal{M}^{dt_k, ds_k, Q_k}(t'_k + dt_k, \mathbf{s}'_k, \mathbf{q}'_k, \varphi) \right\} \geq 0,$$

and applying the consistency result in Proposition 13 while taking $k \rightarrow \infty$ finally yields the

viscosity supersolution property

$$\min \left[-\frac{\partial \varphi}{\partial t} - \frac{1}{2} \sum_{i,j=1}^n [\boldsymbol{\sigma} \boldsymbol{\sigma}']_{ij} \frac{\partial^2 \varphi}{\partial s_i \partial s_j} - \sum_{i,\pm} \frac{\lambda}{\kappa} \exp \left\{ -1 + \frac{\kappa \xi}{2} + \kappa \Delta \varphi \right\} + \phi \ell(\mathbf{q}) \right],$$

$$\varphi - \max_{i,\pm} \left\{ \Delta_{i,\pm} \varphi - \frac{\xi_i}{2} \right\} \Big|_{(\bar{t}, \bar{\mathbf{s}}, \bar{\mathbf{q}})} \geq 0.$$

□

5. Optimal Trading: no-arbitrage region and performance of strategy

5.1. No-Arbitrage region

We use our 2-dimensional illustrative example to explore the behavior of the optimal strategy. Recall that the first asset S_t^1 is a traded equity and the second asset, S_t^2 , is a \bar{T} -maturity futures contract written on S_t^1 . In this economy, we have $\mathbb{V}[(1, -1)' \cdot \mathbf{S}_T] = 0$, yielding a market that is over-complete. To this end, we take $\mathbf{b} = (1, -1)'$ as the basis of B for this illustrative example, and other parameters are given in Table 1.

Parameter	Value
ϕ	1
$\boldsymbol{\sigma}$	$\begin{pmatrix} 0.005 & 0.009 \\ 0 & 0.015 \end{pmatrix}$
γ	0.001
ξ_1, ξ_2	0.02
λ_1^-, λ_1^+	1.2
λ_2^-, λ_2^+	1.0
κ_1^-, κ_1^+	40
κ_2^-, κ_2^+	25
$T - t$	60

Table 1: Model parameters used in the analysis of the function $g(t, \mathbf{s}, \mathbf{q})$, the continuation region, and the optimal inventory level. Recall that the fill rate assumption $h_i^\pm(\delta) = \lambda_i^\pm e^{-\kappa_i^\pm(\delta - \xi_i)}$ from Example 6 is being made.

For simplicity, we assume that the fill rates on both sides of the LOB are the same (see Table 1) and that the arrival of buy and sell MOs are also the same, but both the fill rate parameter and arrival of MOs differ across the two assets. For example, the underlying equity S_t^1 is modeled with a higher rate of incoming MOs than S_t^2 which is the \bar{T} -maturity futures contract. This

is because the underlying trades more actively than the first nearby futures and so we assume that there is more liquidity in the equity than the futures by taking $\kappa_1^\pm > \kappa_2^\pm$.⁶

Finally, note that because $\xi_i > 0$, we do *not* require their correlation to be 1 or even for their volatilities to be the same. However, it is financially sound for their volatilities to be similar and for their correlation to be positive. This leads to our choice of volatility matrix, σ , which induces a correlation of approximately 0.87. The total volatility of Asset 1 turns out to be less than that of Asset 2. The rationale behind this is that the sources of quadratic variation of the futures contract not only include the underlying equity (Asset 1), but also with that of interest rates and stochastic volatility (which affects the futures price via the convexity correction). So the futures contract should then have a larger total volatility than its underlying.

For a fixed time t , Figures 1(a) and 1(c) show $g(t, \mathbf{s}, \mathbf{q})$ as a function of \mathbf{q} for fixed $\mathbf{s} = (100, 100)$. It is clear that it attains its maximum at $\mathbf{q} = (0, 0)$. The asymmetry in g with respect to inventory is due to the fact that the risk associated with a single unit of Asset 1 is less than that of Asset 2. Furthermore, portfolios of the form (a, b) with $\text{sign}(a) \neq \text{sign}(b)$ are less risky than portfolios with $\text{sign}(a) = \text{sign}(b)$ as the two assets are strongly positively correlated. This is depicted in the dark blue regions in the upper right and lower left corners of Figure 1(c), which corresponds to large long positions and large short positions in both assets, respectively.

The left and right panels of Figure 1 show $g(t, \mathbf{s}, \mathbf{q})$ as a function of \mathbf{q} for two different values of $s_1 - s_2$. When $s_1 = s_2$, the maximum of $g(t, \mathbf{s}, \mathbf{q})$ is again obtained at $\mathbf{q} = (0, 0)$. However, this is no longer the case when $s_1 \neq s_2$. When $s_1 < s_2$, the investor would rather be long Asset 1 and short Asset 2, and she will post LOs in a way that reflects this desire to hold such an inventory. This is because the process $S_t^1 - S_t^2$ is bounded from below by the sum of the assets' half-spreads. Hence, when this process is near the lower boundary, there is the potential for profits as we have assumed reflection on this boundary.

Figure 2 shows how the continuation region also depends on $s_1 - s_2$. When $s_1 = s_2$, the continuation region is centered around the origin, just as one would expect. However, when one asset is cheap relative to the other, i.e. $s_1 < s_2$, then the agent prefers to long Asset 1 and short Asset 2 in an attempt to profit when Asset 1 increases relative to Asset 2. Furthermore, she will be aggressive and execute MOs if she is caught holding a portfolio that is short Asset 1 and long Asset 2 when this shift in the midprices occurs, as demonstrated by the absence of the continuation region in the upper left corner of Figure 2(b). As the price discrepancy becomes larger, the investor adjusts her strategy to accumulate inventory in a way that positions her to profit from these market conditions.

⁶Although this is true for equities, there are cases (such as commodities) when a future's underlying does not actively trade or is a non-traded asset (e.g., electricity).

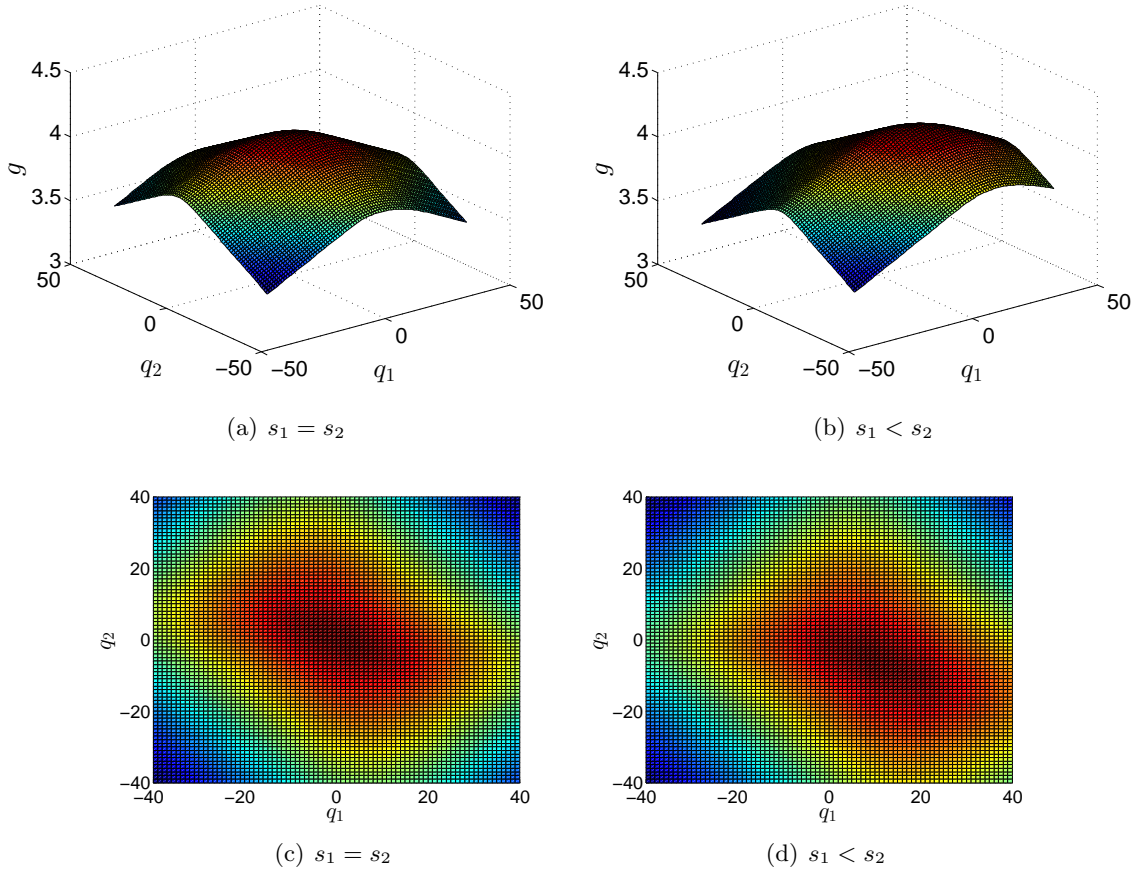


Figure 1: g surface for various \mathbf{s} . Figures 1(a) and 1(c) show the surface and heatmap, respectively, of $g(t, \mathbf{s}, \mathbf{q})$ when the assets' midprices are equal. Figures 1(b) and 1(d) show the surface and heatmap, respectively, of $g(t, \mathbf{s}, \mathbf{q})$ when Asset 1 is cheap relative to Asset 2 ($s_1 = s_2 - 0.012$).

As discussed earlier, inventory control plays a critical role in investment strategies. The investor actively tries to liquidate outstanding long positions and cover outstanding short positions. In a single asset scenario, the logistics behind this are clear. However, in a multiple asset setting with correlated assets, the investor may attempt to partially hedge a long position in one asset with a short position in another asset that is positively correlated.

Corollary 15. Mean Reversion of Inventory. *Given the optimal strategy δ^* , the expected drift of the inventory process \mathbf{q}_t inside the continuation region is given by*

$$\lim_{u \rightarrow t^+} \frac{1}{u - t} \mathbb{E}[\mathbf{q}_u - \mathbf{q}_t | \mathcal{F}_t] = \left[\lambda_i^- h_i^-(\delta_{i,t}^-) - \lambda_i^+ h_i^+(\delta_{i,t}^+) \right]_i. \quad (23)$$

PROOF. By observing that each element of δ_t is a right continuous with left limits (RCLL)




Figure 2: Continuation region for various \mathbf{s} . The different coloured regions correspond to different MO submissions. Namely, light blue: Asset 1 buy MO; orange: Asset 1 sell MO; dark blue: Asset 2 buy MO; burgundy: Asset 2 sell MO; green: continuation region; and white: unattainable. Recall that Figure 2(b) is specified by $s_1 = s_2 - 0.012$.

function of t (for fixed $\omega \in \Omega$), we can conclude that the expected drift in inventories inside the continuation region is given by the difference in the arrival rates of filled LOs at time t . \square

Figure 3 shows the inventory drift as a function of the midprice \mathbf{s} . When the trader is flat (holding zero inventory) and Asset 1 is cheap relative to Asset 2 (i.e. when $s_1 - s_2$ is small), she wants to acquire a long position in Asset 1 and short position in Asset 2. The corresponding drifts in the q_t^1 and q_t^2 processes become more pronounced as the inventory penalty parameter ϕ decreases.

Figures 4 and 5 show the drift in the process \mathbf{q}_t as a function of the current inventory \mathbf{q} for fixed values of $s_1 - s_2$ (and $\phi = 1$ as taken in previous plots where not explicitly stated). The inventory will be attracted to the curve where the reversion rate is 0 (colored green). Notice how in each figure the q_i process is attracted to a *curve* and not a *point*. More specifically, for each value of \mathbf{q} , the component q_i is attracted to a particular point that depends on the values of the other components of \mathbf{q} . When $q_1 = 0$ and $s_1 = s_2$, then the optimal amount of Asset 2 to hold is 0. However, when $q_1 > 0$, then the optimal amount of Asset 2 to hold is less than 0. This highlights the behaviour that the investor tries to partially hedge a long position in one asset by acquiring a short position in the other (as well as try to reduce the long position in Asset 1 as well).

The intersection point of the corresponding *attraction curves* for q_1 and q_2 then acts as the *attraction point* for the vector-valued inventory process \mathbf{q}_t . Figure 6 is a graphical depiction of this attraction point as a function of $s_1 - s_2$.

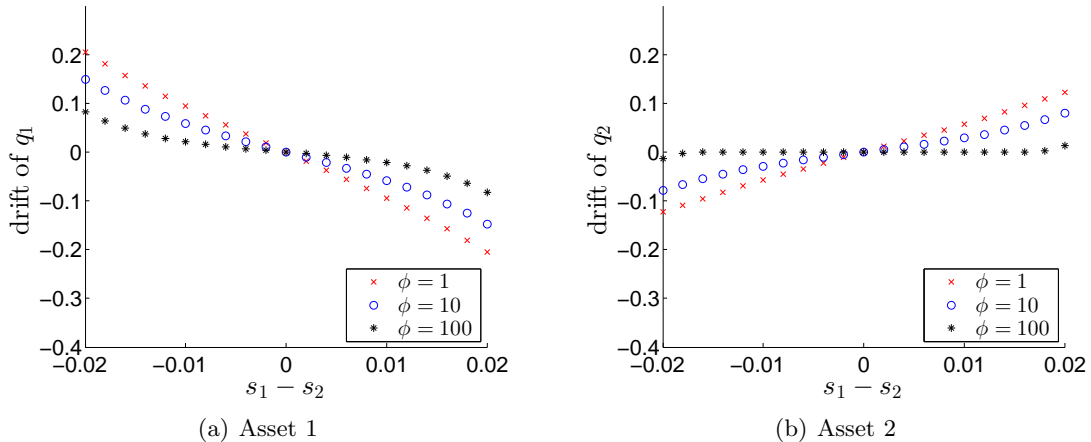


Figure 3: Inventory drift when the investor is flat (i.e. $q_1 = q_2 = 0$).

5.2. Simulation Study: PnL of the Strategy

In this section we use simulations to analyze the financial performance of an investor who follows the optimal strategy in the last section. We use the time limiting strategy for $T - t$ large since the . The investor's operating frequency is every 5 milliseconds, the maximum amount of inventory the investor can hold long or short⁷ is $Q = 40$ for each asset, and the trading horizon is $T = 300$ seconds. Other model parameters are provided in Table 2. Recall that we continue with the fill rate assumption $h_i^\pm(\delta) = \lambda_i^\pm e^{-\kappa_i^\pm(\delta - \xi_i)}$ from Example 6.

5.3. Facing Adverse Selection

Here we analyze the strategy under more realistic scenarios, in particular we assume that an important source of risk faced by the investor is adverse selection. The investor might be trading with better informed counterparties and therefore faces the risk of buying assets right before prices drop or selling assets right before prices rise. In our analysis, we include adverse selection by assuming that some MOs will affect the midprice of the asset, but the investor does not account for this in her model.

With this in mind, we suppose that an MO originates from one of two sources. It was either executed by an *informed* trader, or an *uninformed* one. We assume that the probability

⁷This bound on the investor's inventory only plays a role in the numerical solution of the value function and was not imposed in the optimal control problem above, but it is trivial to include it. Depending on the penalty parameter $\phi > 0$, the exact strategy will dictate uniform (over time and Ω) bounds on the investor's inventory since the continuation region is bounded, as argued in Section 4.3. This is visually depicted in Figure 2 where the optimal strategy keeps inventory inside the rhombus shaped region.




Figure 4: Inventory drift when the assets' midprices are equal (i.e. $s_1 = s_2$).

that any particular MO comes from an informed trader is constant, and we shall denote this probability by ρ . We describe the two possible effects by example. Suppose the current best bid-ask quotes for Assets 1 and 2 are 100.00/100.02 and 100.01/100.03, respectively. A large buy MO on Asset 1 arrives, widening the spread to 100.00/100.06. At this point, one of two things will happen immediately:

- *The MO was executed by an informed trader:* The best bid-ask quotes for Asset 1 becomes 100.04/100.06, giving an immediate impact of \$0.04. This impacts Asset 2 in the same way causing the BBO to move to 100.05/100.07. The informed trader knows that there will be upward pressure on the price, and some other informed market participants (on both assets) move their current quotes to reflect this.
- *The MO was executed by an uninformed trader:* The best bid-ask quotes for both assets remains unchanged. There is nothing special about this MO, and the depleted portion of the LOB refills with LOs.

This describes the *immediate effect* on the midprice process which now follows

$$\begin{aligned}
 d\mathbf{S}_t &= \boldsymbol{\alpha}_t dt + \boldsymbol{\sigma} d\mathbf{B}_t + \mathfrak{d}_{(N_{t-}+1)} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} d\left(\widehat{\mathbf{N}}_t^+ + \widehat{\mathbf{N}}_t^-\right), \\
 d\boldsymbol{\alpha}_t &= -\zeta \boldsymbol{\alpha}_t dt + \boldsymbol{\epsilon}_{(N_{t-}+1)} \Psi d\left(\widehat{\mathbf{N}}_t^+ - \widehat{\mathbf{N}}_t^-\right).
 \end{aligned} \tag{24}$$

Here the vector-valued processes $\{\widehat{\mathbf{N}}_t^-\}_{0 \leq t \leq T}$ and $\{\widehat{\mathbf{N}}_t^+\}_{0 \leq t \leq T}$ are the total number of *informed* market sell and buy orders up to time t , respectively, with the i^{th} component representing Asset i . The process $\boldsymbol{\alpha}_t$ denotes the persistent impact of an informed MO which affects the




Figure 5: Inventory drift when Asset 1 is cheap relative to Asset 2 (i.e. $s_1 = s_2 - 0.012$).




Figure 6: Mean reversion point of \mathbf{q} as a function of $s_1 - s_2$. The red (blue) dots correspond to small (large) values of $s_1 - s_2$.

drift of the midprice process.⁸ Moreover, \mathbf{S}_t has (normal) reflection at the boundary of \mathcal{A} , ζ and Ψ are $n \times n$ matrices with non-negative entries, $\mathfrak{d}_{(1)}, \mathfrak{d}_{(2)}, \dots$ is the sequence of random variables representing immediate market impact (will be non-negative for market buy orders and non-positive for market sell orders), and $\mathfrak{e}_{(1)}, \mathfrak{e}_{(2)}, \dots$ is an i.i.d. sequence of exponential random variables with mean 1.

Note that because the immediate market impact term adds a random multiple of $(1, 1)'$ to \mathbf{S}_t , the resulting value of the midprice vector is guaranteed to remain inside the no-arbitrage region \mathcal{A} (recall from Section 2 that $\mathcal{A} = \{\mathbf{s} \in \mathbb{R}^2 : |s_1 - s_2| < \frac{1}{2}(\xi_1 + \xi_2)\}$) and $\alpha_t^1 = \alpha_t^2$

⁸Since all of these effects occur simultaneously (or faster than the HFT can respond), this model does not admit *latency arbitrage*.

Parameter	Value
T	300
ϕ	0.3, 1, 3, 10, 30, 100, 200
σ	$\begin{pmatrix} 0.005 & 0.009 \\ 0 & 0.015 \end{pmatrix}$
γ	0.001
ξ_1, ξ_2	0.02
λ_1^-, λ_1^+	1.2
λ_2^-, λ_2^+	1.0
κ_1^-, κ_1^+	40
κ_2^-, κ_2^+	25
ζ	$\begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}$
ρ	0.40, 0.50, 0.60
ϵ	0, 0.01, 0.02, 0.04

Table 2: Model parameters used in the simulation study. Recall that the fill rate assumption $h_i^\pm(\delta) = \lambda_i^\pm e^{-\kappa_i^\pm(\delta - \xi_i)}$ from Example 6 is being made.

(which will be denoted α_t). Equation (24) is the natural multivariate extension of the drift dynamics presented in Cartea et al. (2014).

When an informed buy MO arrives on Asset i , then there will immediately be a jump in α_t equal to the i^{th} column of Ψ multiplied by a random exponential random variable. The matrix ζ serves as the decay matrix which forces each component of the process α_t to mean revert to zero. The effect of an informed sell MO is similar, except it induces a negative jump in α_t .

Observe that we are using $\widehat{\mathbf{N}}_t^\pm$ instead of \mathbf{N}_t^\pm or $\overline{\mathbf{N}}_t^\pm$ in (24). This has the interpretation that only the MOs of *informed* traders will have an impact on α_t . Furthermore, our own MOs do not move the market, unlike the informed traders' MOs, as our MOs are quite small in size and we are a small player without superior market information. For the purposes of illustration, we make the further assumption that $\Psi = \epsilon \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$, where $\epsilon \geq 0$ is a scale parameter.

Figure 7 shows sample paths for the midprice process \mathbf{S}_t and its drift process α_t . Observe the jumps in both the midprice process and in the short-term drift component due to the activity of informed traders. On this particular time window, Asset 1 is priced lower than Asset 2. Observe in Figure 7(c) how the midprice spends most of the time near the upper boundary. It will be shown in Section 4.1 that g is only a function of \mathbf{s} through the function $s_1 - s_2$. We provide a sample path of this process in Figure 7(d). Again, notice that on this particular window $S_t^{(1)} - S_t^{(2)} < 0$, but still obeying the bound $S_t^{(1)} - S_t^{(2)} > -\frac{1}{2}(\xi_1 + \xi_2) = -0.02$.




Figure 7: Sample paths of the midprice process \mathbf{S}_t and drift process α_t . The red lines in Figure 7(c) represent the boundary of the no-arbitrage region.

Notice how the investor's buy LO quote for Asset 1 (as depicted in Figure 8(a)) varies as the process $S_t^{(1)} - S_t^{(2)}$ varies over time in Figure 7(d). The discrete jumps in $\delta_{1,t}^-$ are due to the state variable $s_1 - s_2$ moving to another (discrete) grid point (recall that $g(t, \mathbf{s}, \mathbf{q})$ was solved numerically on a sparse grid for the spacial variable \mathbf{s}). In this plot, the cheaper Asset 1 is relative to Asset 2 (i.e. when $S_t^{(1)} - S_t^{(2)}$ is close to -0.02), the more willing the investor is to acquire an additional unit of Asset 1. Hence, she posts her limit buy order closer to the best bid for Asset 1 and her limit sell order closer to the best offer for Asset 2. Along this sample path, the investor manages to acquire a long position in Asset 1, as shown in 8(b), and is using her short position in Asset 2 to partially hedge away the risk associated with being long the underlying equity.

Although these adverse selection effects are present in the DGP, the trader cannot estimate the correct drift of the asset so she assumes that it is zero. Recall that the investor assumes that the dynamics of \mathbf{S}_t are given by Assumption 1, and that this was the assumption used




Figure 8: Sample paths of the investor’s LO depth and her inventory level for $\phi = 10$.

when solving the stochastic control problem above.

5.4. Results of the Simulation Study

Here we show how the investor’s strategy performs when we vary the running inventory penalty ϕ and her exposure to adverse selection by varying the scale parameter ϵ . The investor starts with zero inventory in both assets and will trade for $T = 300$ seconds. We run 2,500 simulations and record the PnL for each run. Figure 9 provides histograms depicting the distribution of the terminal PnL for two levels of running inventory penalty. In Figure 9(a) we assume that the investor is not exposed to adverse selection costs by assuming $\epsilon = 0$ and in Figure 9(b) we show the histogram of the PnLs when the investor is exposed to adverse selection effects by setting $\epsilon = 0.02$.

However, when we increase the scale parameter ϵ in the jump size of the short-term drift process, we observe that the less inventory averse investor does *not necessarily* outperform her more inventory averse counterpart. This is because the less inventory averse investor is willing to hold larger inventories for a longer period of time which further exposes her to the drift of \mathbf{S}_t , while the more inventory averse investor cuts her losses early in an attempt to reduce inventory risk. The $\epsilon = 0.04$ case in Figure 10 illustrates this result very well.

There is some debate about which risk metric better characterizes the riskiness of a strategy. Classical work in finance is typically only concerned with the terminal distribution of PnL (at time T); namely, its variance (or standard deviation) and possibly higher moments that give information about the weight of the tails. However, considering the penalty term in (3), we see that we are concerned with the optimal portfolio’s risk *throughout the entire time interval* $[0, T]$, not just at the end. To this end, we classify a strategy’s riskiness via the penalty




Figure 9: Distribution of terminal profit and loss. The probability of an informed MO is fixed at $\rho = 0.50$.

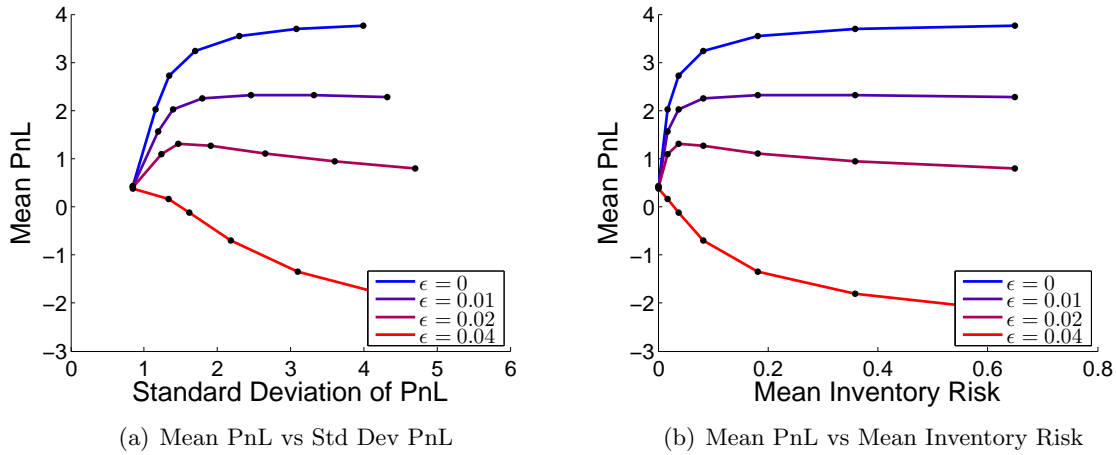


Figure 10: Performance of strategy for various ϕ with different drift impact scale parameters ϵ . The probability of an informed MO is fixed at $\rho = 0.50$.

$\int_0^T \ell(\mathbf{q}_u) du$ and consider the mean of this quantity. This is what is termed *mean inventory risk*, see Cartea and Jaimungal (2015c) in the context of algorithmic trading.

Figures 10 and 11 provide ‘efficient frontiers’ comparing expected PnL against both standard deviation of PnL and mean inventory risk for a variety of adverse selection parameters (ϵ and ρ , respectively). Decreasing ϕ (i.e. increasing inventory-risk appetite) corresponds to moving in the rightward direction of this plot. As stated in Table 2, ϕ ranges from 0.3 to 200.

Recall that the investor assumes that the drift of \mathbf{S}_t is zero, when in fact there is a short-term component in the drift that affects midprices when informed traders send MOs. As outlined in detail in Section 5.3, we have a scale parameter $\epsilon \geq 0$ that represents the magnitude of




Figure 11: Performance of strategy for various ϕ with different probabilities of informed MOs ρ . The drift impact scale parameter is fixed at $\epsilon = 0.01$.

such effects. The larger the parameter, the stronger the effects of adverse selection are in the market. It is clear from Figure 10, that adverse selection generally has a negative effect on the investor's PnL. In the extreme case of $\epsilon = 0.04$, these effects are so strong that decreasing the investor's inventory aversion actually decreases return because she is allowing her portfolio to be exposed to these effects for a longer period of time.

Up until this point, we have assumed that 50% of all incoming MOs are from informed traders and thus have an impact on the drift of the midprice. Recall that MOs of this type have two effects: an immediate effect and a persistent effect. If we increase the probability ρ that an MO comes from an informed trader, then this will directly translate into an increase in the number of MOs that immediately move the market. Because this movement is always against the investor when her LO is filled, it will *always* have a negative impact on her PnL (in contrast to increasing ϵ where this was not the case). The resulting efficient frontiers are reported in Figure 11 for three differing values of ρ .

Figure 12 shows that the larger the inventory penalty parameter ϕ is, the more MOs are executed by the investor. Larger ϕ means that the investor imposes a higher penalty on running inventory, which induces higher mean reversion in inventories. This is attained by adjusting the depth of LOs and by executing MOs to unwind positions. For example, if the inventory is long, the sell LO will be posted very close to the midprice, the buy LO far away from the midprice, and if necessary execute a sell MO. One can also see that the larger ϕ is, the smaller the continuation region, so the boundary gets breached more often inducing MOs that force the investor to cross the bid-ask spread to bring \mathbf{q} back into the continuation region.

Although there is symmetry between the buy and sell sides of the LOB, there is an intuitive

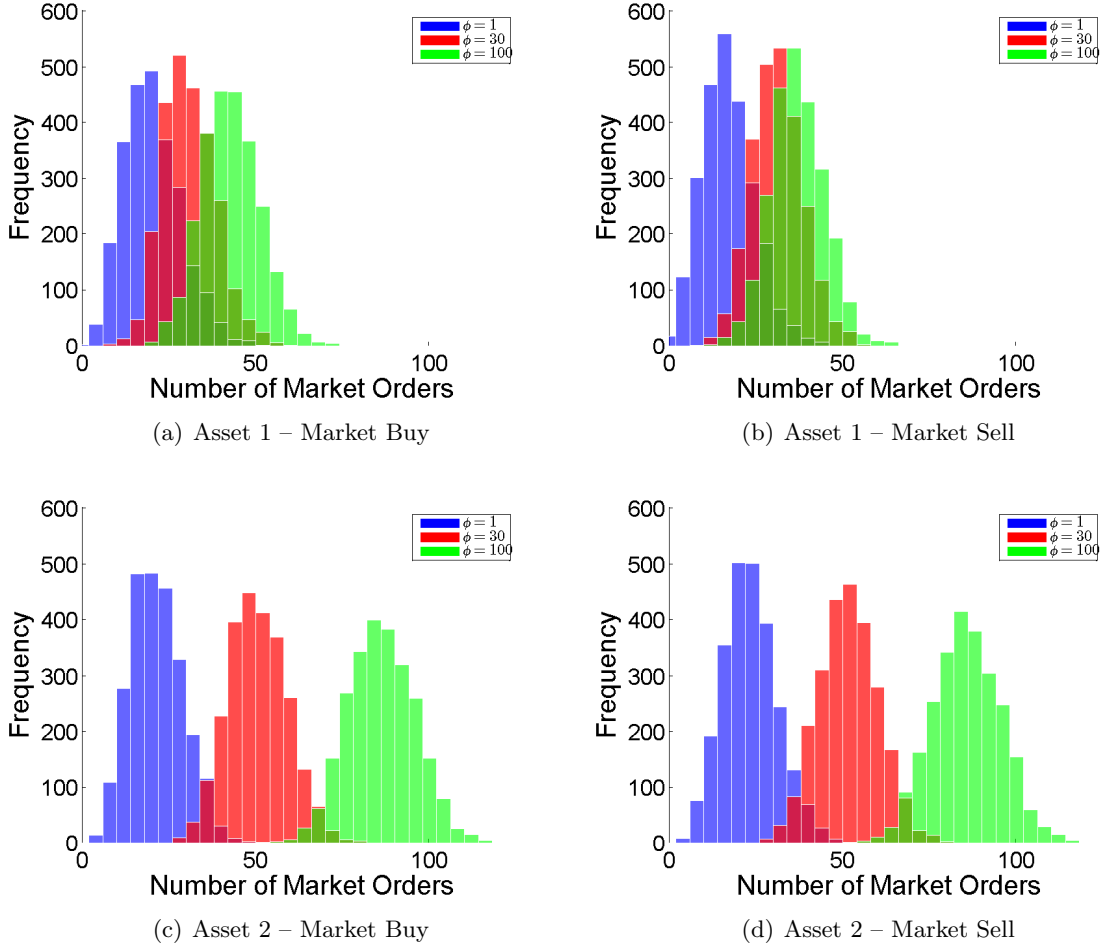


Figure 12: Distribution of the number of executed MOs for various inventory penalty levels when $\epsilon = 0.01$ and $\rho = 0.50$.

explanation for the drastic difference in number of executed MOs between the two assets. Asset 2 has a higher total volatility, which increases the penalty for holding a position in it, hence, the continuation region is smaller in the q_2 direction than in the q_1 direction (as seen in Figure 2).

6. Conclusions

We develop a trading strategy which employs limit and market orders in a multi-asset economy where assets exhibit structural dependence. For example, i) trading in a futures contract and its underlying, or ii) trading in the first and second nearby futures and the 1-2 calendar spread. When structurally dependent assets trade with positive bid-ask spread, midprices live in a no-

arbitrage region and we show how an investor can take positions in these assets to maximize expected cash over a trading horizon.

The investor's performance criteria is linear in wealth and penalizes inventory positions. Key to the strategy is the behaviour of midprices close to the edges of the no-arbitrage region. Price pressure from arbitrageurs and other traders push prices to the interior of the region, and in our framework this is modeled assuming that midprices follow a multivariate Brownian motion that is reflected on the boundaries of this region.

The optimal strategy exhibits two dominant features which depend on how far the vector of midprices is from the no-arbitrage bounds. When midprices are sufficiently far from the no-arbitrage edges, the strategy behaves as that of a market maker who posts buy and sell limit orders. However, these limit orders are strategically placed in the book to not only control for inventory exposure, but also in anticipation of midprices approaching the edges of the no-arbitrage region where the strategy employs a mix of market and limit orders to execute statistical arbitrages.

We employ simulations to show the qualitative features of the optimal trading strategy. Market orders, employed for statistical arbitrages and inventory management, play a key role in the success of the strategy. Limit orders are posted in the book to earn the spread, and to manage inventories where the instantaneous correlation between multiple assets is also used to partially hedge away the inventory risk associated with a large open position in the less liquid assets.

Finally, an explicit numerical scheme is provided to solve the derived HJBQVI and convergence results are discussed to ensure that the given discretization scheme converges to the true viscosity solution. Lastly, we derive a dimensional reduction in the spacial variable for the midprice which significantly reduces the computational time of the numerical solution, thus making it viable for practical applications.

References

- Almgren, R. (2003, January). Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance* 10(1), 1–18.
- Almgren, R. and N. Chriss (2000). Optimal execution of portfolio transactions. *Journal of Risk* 3.
- Avellaneda, M. and S. Stoikov (2008, November). High-frequency trading in a limit order book. *Quantitative Finance* 8, 217–224.
- Barles, G. and P. Souganidis (1991). Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis* 4, 271–283.

- Bayraktar, E. and M. Ludkovski (2014). Liquidation in limit order books with controlled intensity. *Mathematical Finance* 24(4), 627–650.
- Cartea, Á., S. Jaimungal, , and J. Ricci (2014). Buy low, sell high: A high-frequency trading perspective. *SIAM Journal of Financial Mathematics* 5, 415–444.
- Cartea, Á. and S. Jaimungal (2015a). Algorithmic trading of co-integrated assets. *Available at SSRN 2637883*.
- Cartea, Á. and S. Jaimungal (2015b). Optimal execution with limit and market orders. *Quantitative Finance* 15(8), 1279–1291.
- Cartea, Á. and S. Jaimungal (2015c). Risk metrics and fine tuning of high-frequency trading strategies. *Mathematical Finance* 25(3), 576–611.
- Cartea, Á., S. Jaimungal, and R. Donnelly (2013). Algorithmic trading with model uncertainty. *SSRN: <http://ssrn.com/abstract=2310645>*.
- Cartea, Á., S. Jaimungal, and J. Penalva (2015). *Algorithmic and High-Frequency Trading*. Cambridge University Press.
- Forsyth, P., S. Kennedy, J. Tse, and H. Windcliff (2012). Optimal trade execution: a mean-quadratic-variation approach. *Journal of Economic Dynamics and Control* 36, 1971–1991.
- Guéant, O., C.-A. Lehalle, and J. Fernandez Tapia (2012). Optimal portfolio liquidation with limit orders. *SIAM Journal on Financial Mathematics* 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. *Mathematics and Financial Economics* 7(4), 477–507.
- Guilbaud, F. and H. Pham (2013). Optimal high frequency trading with limit and market orders. *Quantitative Finance* 13(1), 79–94.
- Guilbaud, F. and H. Pham (2015). Optimal high-frequency trading in a pro rata microstructure with predictive information. *Mathematical Finance* 25(3), 545–575.
- Ho, T. and H. R. Stoll (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics* 9, 47–73.
- Kharroubi, I. and H. Pham (2010). Optimal portfolio liquidation with execution cost and risk. *SIAM Journal on Financial Mathematics* 1, 897–931.
- Leung, T. and X. Li (2015). Optimal mean reversion trading with transaction costs and stop-loss exit. *International Journal of Theoretical and Applied Finance*, 1550020.
- Øksendal, B. and A. Sulem (2007). *Applied Stochastic Control of Jump Diffusions*, Volume 2. Berlin, Heidelberg, New York: Universitext, Springer Verlag.

Tourin, A. and P. Lintilhac (2015). A dynamic model for pairs trading strategies. *Working paper*.

Tourin, A. and R. Yan (2013). Dynamic pairs trading using the stochastic control approach. *Journal of Economic Dynamics and Control* 37(10), 1972 – 1981.

Appendix A. Evaluating Non-Trivial Terms in the Numerical Scheme of Section 4.2

This section illustrates how to apply the dimensional reduction to terms where the result is non-trivial. All expressions are given explicitly for the $n = 2$ example. In what follows, all functions are evaluated at a fixed inventory level, \mathbf{q} , and this dependence is suppressed.

Appendix A.1. The Mixed Partial Term

For $i \neq j$, we numerically approximate the following term that appears in (12):

$$\frac{\partial^2 g}{\partial s_i \partial s_j} \approx \frac{g(s_i + d\tilde{s}, s_j + d\tilde{s}, \cdot) - g(s_i + d\tilde{s}, \cdot) - g(s_j + d\tilde{s}, \cdot) + g(\cdot)}{d\tilde{s}^2}, \quad (\text{A.1})$$

where $d\tilde{s}$ is the step size of the finite difference approximation. Note that if we require the grids to line up in our $n = 2$ example, we must have $d\tilde{s} = ds$, where ds is the step size on the dimensionally reduced grid given in Appendix A.3.

Appendix A.2. Dimensional Reduction

Recall that Corollary 9 states that there exists $v \in \mathbb{R}^n \setminus \{0\}$ such that $D_v g = 0$. In our example, $v = \mathbf{1}$, so we take $\mathbf{s} := s_1 - s_2$.

Continuing from (A.1), we then have

$$\frac{\partial^2 g}{\partial s_1 \partial s_2} \approx \frac{2g(\cdot) - g(\mathbf{s} - ds, \cdot) - g(\mathbf{s} + ds, \cdot)}{ds^2}, \quad (\text{A.2})$$

where the right hand side is defined on the dimensionally reduced grid and taken at the previous time step ($t + dt$).

For the 2-dimensional illustrative example studied in this paper, the full numerical approximation for the diffusive term in $\mathcal{L}g$ is then given by

$$\frac{g(\mathbf{s} + ds, t + dt) + g(\mathbf{s} - ds, t + dt) - 2g(\mathbf{s}, t + dt)}{2ds^2} ([\boldsymbol{\sigma}\boldsymbol{\sigma}']_{11} + [\boldsymbol{\sigma}\boldsymbol{\sigma}']_{22} - 2[\boldsymbol{\sigma}\boldsymbol{\sigma}']_{12}). \quad (\text{A.3})$$

Appendix A.3. The Neumann Condition

For $\mathbf{s} \in \partial\mathcal{A}$, numerically impose $(\mathbf{q} + \nabla g) \cdot \mathbf{n} = 0$. Take $d\bar{s} = c_0 ds$ so the grids match up, where c_0 depends only on the choice of orthogonal basis of B .

In our $n = 2$ example, we have the normal vector $\mathbf{n} = (1, -1)$ and $c_0 = 1$ resulting in $\nabla g \approx \mp \left(\frac{g(\mathbf{s}) - g(\mathbf{s} \pm ds)}{ds}, \frac{g(\mathbf{s} \pm ds) - g(\mathbf{s})}{ds} \right)$ with the upper (lower) symbol for reflection on the upper (lower) boundary when viewed graphically. The relation is only approximate because we are using a finite difference method. Finally, for $\mathbf{s} \in \partial\mathcal{A}$, we have

$$g(\mathbf{s}) \approx g(\mathbf{s} \pm ds) \pm \frac{(q_1 - q_2)ds}{2}. \quad (\text{A.4})$$

Appendix A.4. Explicit Example when $n = 3$

The analogue of the right-most coefficient in (A.3) is given by

$$[\sigma\sigma']_{11} + [\sigma\sigma']_{22} + [\sigma\sigma']_{33} - 2[\sigma\sigma']_{12} - 2[\sigma\sigma']_{13} + 2[\sigma\sigma']_{23} \quad (\text{A.5})$$

and, for $\mathbf{s} \in \partial\mathcal{A}$, we have

$$g(\mathbf{s}) \approx g(\mathbf{s} \pm ds) \pm \frac{(q_1 - q_2 - q_3)ds}{3}. \quad (\text{A.6})$$