Comparison between the Mean Variance optimal and the Mean Quadratic Variation optimal trading strategies *

S. T. Tse † P.A. Forsyth ‡ J.S. Kennedy § H. Windcliff ¶

November 8, 2012

Abstract

We compare optimal liquidation policies in continuous time in the presence of trading impact using numerical solutions of Hamilton Jacobi Bellman (HJB) partial differential equations (PDE). In particular, we compare the time-consistent mean-quadratic-variation strategy with the time-inconsistent (pre-commitment) mean-variance strategy. We show that the two different risk measures lead to very different strategies and liquidation profiles. In terms of the optimal trading velocities, the mean-quadratic-variation strategy is much less sensitive to changes in asset price and varies more smoothly. In terms of the liquidation profiles, the mean-variance strategy strategy is much more variable, although the mean liquidation profiles for the two strategies are surprisingly similar. On a numerical note, we show that using an interpolation scheme along a parametric curve in conjunction with the semi-Lagrangian method results in significantly better accuracy than standard axis-aligned linear interpolation. We also demonstrate how a scaled computational grid can improve solution accuracy.

Keywords: optimal trading, mean variance, pre-commitment, mean quadratic variation, time-consistent, arrival price, implementation shortfall, HJB PDE, interpolation, scaled-grid

1 Introduction

Algorithmic trade execution has become a standard technique for institutional market players in recent years, particularly in the equity market where electronic trading is most prevalent. A trade execution algorithm typically seeks to execute a trade decision optimally upon receiving inputs from a human trader.

A common form of optimality criterion seeks to strike a balance between minimizing pricing impact and minimizing timing risk. For example, in the case of selling a large number of shares, a fast liquidation will cause the share price to drop, whereas a slow liquidation will expose the seller to timing risk due to the stochastic nature of the share price.

Several approaches have been suggested in the literature to quantify the minimization of pricing impact and timing risk. The first, and perhaps the most intuitive, approach maximizes the expected revenues while minimizing a risk criterion, for example, variance [16, 15], quadratic variation [3], or value-at-risk (VaR) [19]. Another approach maximizes the expected value of a utility function of revenues, for example, a power-law

*This work was supported by the Natural Sciences and Engineering Research Council of Canada, and by a Morgan Stanley Equity Market Microstructure Research Grant. The views expressed herein are solely those of the authors, and not those of any other person or entity, including Morgan Stanley. Morgan Stanley is not responsible for any errors or omissions. Nothing in this article should be construed as a recommendation by Morgan Stanley to buy or sell any security of any kind.

†David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1 `sttse@uwaterloo.ca`

‡David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1 `paforsyt@uwaterloo.ca`

§Morgan Stanley, 1585 Broadway, New York, NY 10036, `Shannon.Kennedy@MorganStanley.com`

¶Morgan Stanley, 1585 Broadway, New York, NY 10036, `Heath.Windcliff@MorganStanley.com`

function or an exponential function [20, 30, 27]. The third approach minimizes the expected execution cost [9]. All these three approaches model the asset price process in the presence of pricing impact. Yet another approach, which is somewhat tangential to the above methodologies, minimizes the expected execution cost by modelling the dynamic distribution of bid and ask orders in a limit order book [28, 1].

In this paper we focus on maximizing revenue while minimizing a risk measure. Maximizing revenue is also know as minimizing implementation shortfall relative to the arrival (pre-trade) price, which is a popular approach in industry. More specifically, we compare the pre-commitment mean-variance strategy [16, 8, 26] with the mean-quadratic-variation strategy [17, 13]. We assume that trading takes place continuously at a finite rate, as in [5, 4]. We note that the risk-criteria in the seminal paper [3] was previously thought to be variance but is actually quadratic variation, as shown in [17]. In [25], the pre-commitment mean-variance strategy was computed in a discrete time setting. It is shown in [25] that the pre-commitment strategy outperforms the strategy in [3], when the criteria are mean and variance as seen at the initial time.

However, as discussed in [17, 13], there may be circumstances when the mean-quadratic-variation optimality is preferred. This may occur in situations where control of the risk is desired during the entire trading cycle, in contrast to only measuring *ex post* mean and variance.

Therefore, it is interesting to investigate how suboptimal the mean-quadratic-variation strategy is in terms of mean-variance efficiency, and, conversely, the question of how suboptimal the mean-variance strategy is in terms of mean-quadratic-variation efficiency.

The main contributions of this article are:

- We have improved our numerical method in [16] so that results for very challenging parametric cases can be computed using reasonable time and memory. In particular, we improve the method of interpolation at the foot of the characteristic in the semi-Lagrangian discretization of HJB PDEs. We also construct a scaled computational grid so that fewer grid nodes are needed to obtain accurate results. The improved method also guarantees convergence to the viscosity solution. We remind the reader that the method in [16] can determine the entire mean variance efficient frontier from a single solution of a nonlinear partial differential equation.

- The mean-variance formulation of the optimal liquidation problem is known to be ill-posed [16]. More specifically, many similar strategies can give rise to nearly the same efficient frontier. In this paper we analyze in detail this ill-posedness from both a mathematical and a computational perspective. In particular, we highlight the numerical challenges created by such ill-posedness and demonstrate that the choice of interpolation method can be critical.

- We then carry out a careful set of numerical tests, which show that for the same variance, the mean-quadratic-variation strategy can have a significantly suboptimal expected value compared to the mean-variance strategy. For the same quadratic-variation, the mean-variance strategy can have significantly suboptimal expected value compared to the mean-quadratic-variation strategy. By carrying out a series of grid refinement studies, we show that the differences between these strategies are significantly larger than the numerical discretization errors.

  If one wants to strike the middle ground of balancing both variance and quadratic-variation, the mean-variance strategy seems to be preferable.

- We show that the mean-variance strategy is much more sensitive to changes in the asset price than the mean-quadratic-variation strategy. Consequently, the trading profile of the mean-variance strategy is much more variable. The mean trading profiles of the two strategies, however, turn out to be surprisingly similar.

2 Optimal Execution

Let

$$\begin{aligned} P &= B + AS = \text{ Portfolio,} \\ S &= \text{ Price of the underlying risky asset,} \\ B &= \text{ Balance of risk free bank account,} \\ A &= \text{ Number of shares of underlying asset.} \end{aligned}$$

The optimal execution problem over $t \in [0, T]$ has the initial condition

$$S(0) = s_{init}, B(0) = 0, A(0) = \alpha_{init}. \tag{2.1}$$

If $\alpha_{init} > 0$, the trader is liquidating a long position (selling). If $\alpha_{init} < 0$, the trader is liquidating a short position (buying). In this article, for definiteness, we consider the selling case. At $t = T$,

$$S = S(T), B = B(T), A = A(T) = 0, \tag{2.2}$$

where $B(T)$ is the cash generated by selling shares and investing in the risk free bank account B, with a final liquidation at $t = T^-$ to ensure that $A(T) = 0$. The objective of optimal execution is to maximize $B(T)$, while at the same time minimizing a certain risk measure. The two risk measures we consider in this paper, namely variance and quadratic-variation, will be discussed in the next two sections.

In this paper, we consider Markovian trading strategies $v(\cdot)$ that specify a trading rate v as a function of the current state, i.e. $v(\cdot) : (S(t), B(t), A(t), t) \mapsto v = v(S(t), B(t), A(t), t)$. Note that in using the shorthand notations $v(\cdot)$ for the mapping, and v for the value $v = v(S(t), B(t), A(t), t)$, the dependence of v on the current state is implicit.

By definition,

$$dA(t) = v \; dt. \tag{2.3}$$

We assume that due to temporary price impact, selling shares at the rate v at the market price $S(t)$ gives an execution price $S_{exec}(v, t) \leq S(t)$. It follows that

$$dB(t) = \big(rB(t) - vS_{exec}(v, t)\big)dt \tag{2.4}$$

where r is the risk free rate.

We suppose that the market price of the risky asset S follows a Geometric Brownian Motion (GBM), where the drift term is modified due to the permanent price impact of trading [5]:

$$\begin{aligned} dS(t) &= \big(\eta + g(v)\big)S(t) \; dt + \sigma S(t) \; d\mathbb{W}(t), \\ & \quad \eta \text{ is the drift rate,} \\ & \quad g(v) \text{ is the permanent price impact function,} \\ & \quad \sigma \text{ is the volatility,} \\ & \quad \mathbb{W}(t) \text{ is a Wiener process under the real world measure.} \end{aligned} \tag{2.5}$$

2.1 Trading impact function

We assume the temporary price impact scales linearly with the asset price, i.e.

$$S_{exec}(v, t) = f(v)S(t), \tag{2.6}$$

where

$$\begin{aligned} f(v) &= (1 + \kappa_s \operatorname{sgn}(v)) \exp[\kappa_t \operatorname{sgn}(v)|v|^\beta], \\ & \quad \kappa_s = \text{ the bid-ask spread parameter,} \\ & \quad \kappa_t = \text{ the temporary price impact factor,} \\ & \quad \beta = \text{ the price impact exponent.} \end{aligned} \tag{2.7}$$

Note that we assume $\kappa_s < 1$, so that $S_{exec}(v,t) \geq 0$, regardless of the magnitude of v. For various studies which suggest the form (2.7), see [5, 24, 29].

The permanent price impact function $g(v)$ is assumed to be of the form

$$\begin{aligned} g(v) &= \kappa_p v, \\ &\kappa_p = \text{ the permanent price impact factor.} \end{aligned}$$

As explained in [17], this form of permanent price impact function eliminates the possibilities of round-trip arbitrage [5, 21].

2.2 Definition of liquidation value

Given the state $(S(T^-), B(T^-), A(T^-))$ at the instant $t = T^-$ before the end of the trading horizon, we have one final liquidation (if necessary) so that the number of shares owned at $t = T$ is $A(T) = 0$. The liquidation value $B(T)$ after this final trade is defined to be

$$\begin{aligned} B(T) &= B(T^-) + \lim_{v \to -\infty} A(T^-) S_{exec}(v, T^-) \\ &= B(T^-) \end{aligned} \tag{2.8}$$

Definition (2.8) in effect penalizes the strategy if $A(T^-) \neq 0$, so that the optimal algorithm forces the liquidation profile towards $A(T^-) = 0$. In our case, the penalty is such that the shares $A(T^-)$ are simply discarded.[1]

3 Mean-Variance Strategy

We review here the pre-commitment mean-variance strategy, as discussed in [16].

3.1 Objective functional and optimal strategy

To simplify notations, we define $x = (s, b, \alpha) = (S(t), B(t), A(t))$ for a space state. Now we specify the pre-commitment mean variance formulation as follows. For all possible states (x,t) and a fixed risk aversion parameter λ, define the family of objective functionals

$$\mathcal{F}_\lambda = \left\{ J_\lambda^{x,t}\Big(v(\cdot)\Big) : v(\cdot) \mapsto E_{v(\cdot)}^{x,t}\big[B(T)\big] - \lambda Var_{v(\cdot)}^{x,t}\big[B(T)\big] \right\}, \tag{3.1}$$

where $E_{v(\cdot)}^{x,t}[\cdot]$ is the expectation, and $Var_{v(\cdot)}^{x,t}[\cdot]$ is the variance, conditional on the state (x,t) and the control $v(\cdot) : (x,t) \mapsto v = v(x,t)$.

Note that in the notation of (3.1), the members (functionals) in the family $\mathcal{F}_\lambda$ have different initial states (x,t) but the same λ. For a given initial state (x,t), we will henceforth use the notation $v^*_{x,t,\lambda}(\cdot)$ to denote the optimal policy that maximizes the corresponding functional, i.e. $J_\lambda^{x,t}\Big(v(\cdot)\Big)$.

Let $(x_0, t = 0) = (s_{init}, 0, \alpha_{init}, 0)$ be the initial state. The optimal policy $v^*_{x_0,0,\lambda}(\cdot)$ is termed the pre-commitment mean variance optimal strategy [8, 12].

[1] In actual implementation, we would replace $\lim_{v\to-\infty}$ by a finite $v_{\min} \ll 0$ in the PDE initial condition. Also, in the case of liquidating a short position (buying), which is not considered in this paper, equation (2.8) would be defined as $B(T) = B(T^-) + \lim_{v\to\infty} A(T^-)S_{exec}(v,T^-)$, and we would replace $\lim_{v\to\infty}$ by a finite $v_{\max} \gg 0$ in implementation.

3.2 Time-inconsistency of optimal strategies

The optimal strategies in the pre-commitment mean variance formulation are time-inconsistent in the following sense. Let (x_1, t_1) be some state at time t_1 and $v^*_{x_1,t_1,\lambda}(\cdot)$ be the corresponding optimal policy. Let (x_2, t_2) be some other state at time $t_2 > t_1$ and $v^*_{x_2,t_2,\lambda}(\cdot)$ be the corresponding optimal policy.

We have time-inconsistency in the sense that

$$v^*_{x_1,t_1,\lambda}(x', t') \quad \neq \quad v^*_{x_2,t_2,\lambda}(x', t') \ ; \ t' \geq t_2 \quad . \tag{3.2}$$

The time-inconsistency (3.2) is considered as unnatural by some authors [8] and creates computational difficulties. More specifically, a dynamic programming principle cannot be directly applied to solve this problem.

Note that in [2], the author makes the case that the pre-commitment formulation optimizes trading efficiency as measured in practice.

3.3 Alternative Formulation

To solve for $v^*_{x_0,0,\lambda}(\cdot)$, we follow the method in [33, 11, 7, 18, 31]. For each fixed initial state (x, t) and risk aversion parameter λ, the optimal control $v^*_{x,t,\lambda}(\cdot)$ that maximizes (the functionals in the family) (3.1) are also the optimal controls that minimize

$$\tilde{\mathcal{F}}_\lambda = \left\{ \tilde{J}^{x,t}_\lambda\Big(v(\cdot)\Big) : v(\cdot) \mapsto E^{x,t}_{v(\cdot)}\left[\big(B(T) - \frac{\gamma(x,t;\lambda)}{2}\big)^2\right]\right\}, \tag{3.3}$$

for some

$$\gamma = \gamma(x, t; \lambda) \geq 0. \tag{3.4}$$

Note that under the alternative formulation (3.3), the optimal strategies $v^*_{x,t,\lambda}(\cdot)$ are also time-inconsistent, due to the dependence of γ on the initial state (x, t), i.e. (3.4); see [12] for more discussion on this.

Consequently, for the particular initial state $(x_0, 0)$ and an arbitrary constant $\gamma_0 \geq 0$, the optimal control $v^*_{x_0,0,\gamma_0}(\cdot)$ that minimizes the member $\hat{J}^{x_0,0}_{\gamma_0}$ in the family

$$\hat{\mathcal{F}}_\gamma = \left\{ \hat{J}^{x,t}_\gamma\Big(v(\cdot)\Big) : v(\cdot) \mapsto E^{x,t}_{v(\cdot)}\left[\big(B(T) - \frac{\gamma}{2}\big)^2\right]\right\}, \tag{3.5}$$

is also the pre-commitment optimal control $v^*_{x_0,0,\lambda_0}(\cdot)$ for a certain value of λ_0 such that $\gamma_0 = \gamma(x_0, 0, \lambda_0)$.

The benefit of reformulating (3.1) as (3.5) is that the dynamic programming principle can be applied to (3.5) to solve for $v^*_{x_0,0,\gamma_0}(\cdot)$, since γ is a constant in (3.5).

Varying γ_0 over $[0, \infty)$ gives the strategies that trace out a variance-minimizing frontier in the expected value, standard deviation plane.

4 Mean-Quadratic-Variation Strategy

Quadratic variation has been used as an approximation of variance in the algorithmic trading literature [3, 4, 17]. This approximation, however, can be poor when the trading impact is relatively large, as will be illustrated in the current paper. Instead of using quadratic variation to approximate variance, it is conceptually simpler to regard quadratic variation as an alternative risk measure.

4.1 Quadratic variation as a risk measure

Formally, the quadratic variation risk measure is defined as

$$E\left[\int_t^T \big(A(t')dS(t')\big)^2\right]. \tag{4.1}$$

Informally, the risk measure definition (4.1) can be interpreted as the quadratic variation of the portfolio value process as follows: by expanding the square of $dP(t') = dB(t') + d(A(t')S(t'))$ and ignoring higher-order terms, we have

$$\int_t^T \big(A(t')dS(t')\big)^2 = \int_t^T \big(dP(t')\big)^2, \tag{4.2}$$

when the trading velocity process $v(t)$ is bounded.

4.1.1 Static strategies

Under certain mild assumptions (including arithmetic Brownian motion), for static (asset-price-independent) strategies, quadratic variation is the same as variance [17]. In general, of course, quadratic variation is not the same as variance. In this paper, we compare (i) mean-variance optimal strategies; and (ii) mean-quadratic-variation optimal strategies, which are both dynamic (asset-price-dependent) for the geometric Brownian motion case considered in this paper.

We remind the reader that [25] compares dynamic mean variance optimal strategies and static mean variance optimal strategies (assuming arithmetic Brownian motion). However, for short trading horizons, when Geometric Brownian motion can be well approximated by arithmetic Brownian motion, mean-quadratic-variation strategies turn out to be almost static. In this case, the conclusions of the study in [25] are similar to the results of our study here. However, in order to make definitive statements about the dominance of one strategy over another, careful attention must be paid to the accuracy of the computed results. In this study, we use the provably convergent Hamilton-Jacobi-Bellman formulation and solution techniques described in [16]. Our grid refinement studies allow us to bound the discretization errors.

In addition, in contrast to [25], we also show that mean-variance strategies can be significantly worse than mean-quadratic-variation strategies, when quadratic variation is the risk measure. Mean-quadratic-variation strategies are naturally time-consistent, and control risk during the course of trading, and hence it can be argued that quadratic-variation is a sensible measure of risk [17].

4.2 Objective functional and value function

Now we specify the mean quadratic variation formulation as follows. For a fixed initial point $(s, \alpha, t) = (S(t), A(t), t)$ where $t < T$ with $B(t) = 0$, trading strategy $v(\cdot)$, and risk aversion parameter $\tilde{\lambda}$, we define the objective functional

$$J(s, \alpha, t, v(\cdot); \tilde{\lambda}) = E_{v(\cdot)}^{s,\alpha,t}\Big[B(T)\Big] - \tilde{\lambda} E_{v(\cdot)}^{s,\alpha,t}\left[\int_t^T \big(A(t')dS(t')\big)^2\right] \tag{4.3}$$

where

$$B(T) = \int_t^{T^-} e^{r(T-t')}\big(-vS_{exec}(v,t')\big)dt' + \lim_{v\to-\infty} A(T^-)S_{exec}(v,T^-) \tag{4.4}$$

and $E_{v(\cdot)}^{s,\alpha,t}[\cdot]$ is the conditional expectation at the initial point (s, α, t) using the control $v(\cdot)$.

The value function $\hat{V}^{MQV}$ is defined as

$$\hat{V}^{MQV}(s, \alpha, t; \tilde{\lambda}) \quad = \quad \sup_{v(\cdot)} J(s, \alpha, t, v(\cdot); \tilde{\lambda}). \tag{4.5}$$

For a given initial state (s, α, t), we will henceforth use the notation $v^*_{s,\alpha,t,\tilde{\lambda}}(\cdot)$ to denote the optimal policy that maximizes the corresponding functional, i.e. $J(s, \alpha, t, v(\cdot); \tilde{\lambda})$.

4.3 Time Consistency of the optimal strategies

Let (s_1, α_1, t_1) be some state at time t_1 and $v^*_{s_1,\alpha_1,t_1,\tilde{\lambda}}(\cdot)$ be the corresponding optimal strategy. Let (s_2, α_2, t_2) be some other state at time $t_2 > t_1$ and $v^*_{s_2,\alpha_2,t_2,\tilde{\lambda}}(\cdot)$ be the corresponding optimal strategy.[2]

Since the optimal controls satisfy the Bellman's principle of optimality as shown in [17], it follows that the optimal controls of (4.5) are *time consistent* in the sense that for the same state (s', α', t') at a later time $t' > t_2$,

$$v^*_{s_1,\alpha_1,t_1,\tilde{\lambda}}(s', \alpha', t') \quad = \quad v^*_{s_2,\alpha_2,t_2,\tilde{\lambda}}(s', \alpha', t') \ ; \ t' \geq t_2 \quad . \tag{4.6}$$

Hence, dynamic programming can be directly applied to this problem.

In certain special cases, it is known that strategy $v^*_{s,\alpha,t,\tilde{\lambda}}(\cdot)$ is equivalent to a time consistent mean variance strategy [12, 32]. Hence $v^*_{s,\alpha,t,\tilde{\lambda}}(\cdot)$ can be viewed as a natural time consistent strategy. In addition, as shown in [17], if arithmetic Brownian motion is assumed, the optimal strategy $v^*_{s,\alpha,t,\tilde{\lambda}}(\cdot)$ is actually identical to the optimal strategy in [3].

Let $V^{MQV}(s, \alpha, \tau; \tilde{\lambda}) = \hat{V}^{MQV}(s, \alpha, t = T - \tau; \tilde{\lambda})$. The derivation in [17] shows that V^{MQV} satisfies the HJB equation

$$V^{MQV}_\tau = \eta s V^{MQV}_s + \frac{\sigma^2 s^2}{2} V^{MQV}_{ss} - \tilde{\lambda} e^{2r\tau} \alpha^2 s^2 \sigma^2 + \sup_v \left[e^{r\tau}(-vf(v))s + g(v)sV^{MQV}_s + vV^{MQV}_\alpha \right]. \tag{4.7}$$

We note that both the value function V^{MQV} and strategy $v(\cdot)$ for the mean-quadratic-variation problem (4.5) is independent of the current bank account balance $B(t)$. In particular $v = v(s, \alpha, \tau)$.

The reader is referred to [17] for details about the numerical method used to solve equation (4.7).

4.4 Arithmetic Brownian Motion

Under the arithmetic Brownian motion approximation and the additional assumptions of zero drift, zero interest rate, unbounded control and linear price impact function detailed in [17], the optimal trading strategy has the analytic solution

$$v(\alpha, \tau) = -\alpha K \coth(K\tau), \tag{4.8}$$

where $K = \sqrt{\tilde{\lambda}\sigma^2 s_{init}/\kappa_t}$. Note that the optimal strategy is independent of the spot price s.

As noted in [17], the strategy (4.8) results in an efficient frontier that is extremely close to the true mean-quadratic-variation efficient frontier computed assuming geometric Brownian motion. We also note that (4.8) is the same strategy as used in [3].

The work in [3] is extended in [6] to consider nonlinear trading impacts similar to the form (2.7), and semi-explicit solutions are obtained.

5 Comparison between the two strategies

5.1 Risk measure

The pre-commitment mean variance strategy is optimal in the following sense [25, 2]. Suppose we carry out many thousands of trades. We then examine the post-trade data, and determine the realized mean return and the standard deviation. Assuming that the modeled dynamics very closely match the dynamics in the real world, the optimal pre-commitment strategy would result in the largest realized mean return, for a given standard deviation, compared to any other possible strategy.

From the interpretation (4.2), minimizing quadratic variation clearly corresponds to minimizing volatility in the portfolio value process. The definition (4.1) shows that quadratic variation takes into account the

[2] Note that while the initial point is changed from (s_1, α_1, t_1) to (s_2, α_2, t_2), the risk aversion level $\tilde{\lambda}$ is kept constant.

trading trajectory $A(t')$ over the whole trading horizon. This is in contrast with using variance ($Var[B(T)]$) as a risk measure, which is independent of the trading trajectory $A(t')$ given the end result $B(T)$. We note that the idea of using quadratic variation as a risk measure was first suggested in [13].

5.2 Uniqueness and smoothness

In mean-variance optimization, many similar strategies can give rise to almost the same efficient frontier (near ill-posedness). This can be advantageous as it permits more flexibility in executing the trade. On the other hand, this creates difficulties in obtaining a smoothly varying optimal strategy, as demonstrated and explained in [16] and the current paper. In our experience, these issues do not arise in mean-quadratic-variation optimization.

6 HJB Formulation: Mean Variance

6.1 Change of Variable

At first glance it seems necessary to solve the problem (3.5) for each value of γ separately. Fortunately, this can be avoided by a change of variable. Define the new variable $\mathcal{B}$ by

$$\mathcal{B}(0) = \frac{-\gamma e^{-rT}}{2} \leq 0, \quad d\mathcal{B}(t) = r\mathcal{B}(t)dt - vS_{exec}(v,t)dt \tag{6.1}$$

it is easy to see that

$$\mathcal{B}(t) = B(t) - \frac{\gamma e^{-r(T-t)}}{2}, \quad \mathcal{B}(T) = B(T) - \frac{\gamma}{2}. \tag{6.2}$$

Since equation (6.1) has the same form as equation (2.4), solving for $v^*_{x,0,\gamma}(\cdot)$ in (3.5) is equivalent to solving for the optimal control in

$$\inf_{v(\cdot)} E^{s,\mathscr{b},\alpha,t=0}_{v(\cdot)}[\mathcal{B}(T)^2]. \tag{6.3}$$

This change of variable is very convenient in the PDE context because the solutions corresponding to different values of γ can be determined by examining the PDE solutions for different values of $\mathcal{B}$ at $t = 0$. Therefore, we only need to solve the problem (6.3) once to obtain the entire variance minimizing frontier [16].

6.2 Definitions

Let $\tau = T - t$ be the backward time. Define the value function V by

$$V = V(s, \mathscr{b}, \alpha, \tau) = \inf_{v(\cdot)} E^{s,\mathscr{b},\alpha,T-\tau}_{v(\cdot)}[\mathcal{B}(T)^2]. \tag{6.4}$$

We also define the differential operator $\mathcal{L}$ by

$$\mathcal{L}V = \frac{\sigma^2 s^2}{2} V_{ss} + \eta s V_s, \tag{6.5}$$

and Lagrangian derivative $\frac{D}{D\tau}(v)$ by

$$\frac{DV}{D\tau}(v) = V_\tau - V_s g(v)s - V_{\mathscr{b}}(r\mathscr{b} - vf(v)s) - V_\alpha v, \tag{6.6}$$

which is the rate of change of V along the characteristic $s = s(\tau)$, $\mathscr{b} = \mathscr{b}(\tau)$, $\alpha = \alpha(\tau)$ defined by the trading velocity v through

$$\frac{ds}{d\tau} = -g(v)s, \quad \frac{d\mathscr{b}}{d\tau} = -(r\mathscr{b} - vf(v)s), \quad \frac{d\alpha}{d\tau} = -v. \tag{6.7}$$

6.3 PDE formulation

Following standard arguments, the optimal control is given by the solution to the nonlinear HJB equation

$$\mathcal{L}V - \max_{v \le 0} \frac{DV}{D\tau}(v) = 0. \tag{6.8}$$

in the domain $\Omega = \{s \ge 0, \mathscr{b} \in \mathbf{R}, \alpha \ge 0, \tau > 0\}$. In view of definition (2.8), the initial condition at $\tau = 0$ is

$$V(s, \mathscr{b}, \alpha, \tau = 0) = \mathcal{B}(T)^2 = (\mathscr{b} + as \lim_{v \to -\infty} f(v))^2. \tag{6.9}$$

Note that we forbid buying or holding a short position (when liquidating stock) in (6.8), i.e.

$$v(s, \mathscr{b}, \alpha, \tau) \le 0, \quad v(s, \mathscr{b}, \alpha = 0, \tau) = 0. \tag{6.10}$$

At $s = 0$, equation (6.8) reduces to

$$\max_{v \le 0} \left\{ V_\tau - r\mathscr{b}V_{\mathscr{b}} - vV_\alpha \right\} = 0. \tag{6.11}$$

Therefore, no boundary condition at $s = 0$ is needed.

At $\alpha = 0$, (6.10) causes equation (6.8) to reduce to

$$\frac{\sigma^2 s^2}{2} V_{ss} + \eta s V_s - V_\tau + r\mathscr{b}V_{\mathscr{b}} = 0. \tag{6.12}$$

Therefore, no boundary condition at $\alpha = 0$ is needed.

Solving the HJB PDE (6.4) gives the optimal control $v^*(\cdot)$. To obtain an efficient frontier, we proceed as follows. Define the value function U by

$$U = U(s, \mathscr{b}, \alpha, \tau) = E_{v^*(\cdot)}^{s,\mathscr{b},\alpha,T-\tau}[\mathcal{B}(T)], \tag{6.13}$$

which satisfies the PDE

$$\mathcal{L}U - \frac{DU}{D\tau}(v^*) = 0, \quad U(s, \mathscr{b}, \alpha, \tau = 0) = \mathcal{B}(T) = \mathscr{b} + as \lim_{v \to -\infty} f(v). \tag{6.14}$$

Since $v^*(\cdot)$ has been determined, the PDE (6.14) is linear and inexpensive to solve.

Having solved for the value functions V and U, the variance-minimizing frontier can be obtained as described in section B.1. The mean-variance frontier is then obtained by a simple sorting procedure to eliminate suboptimal points [31].

7 Limiting case

For illustration purposes, consider a limiting case with extreme parameter values $\sigma = \kappa_t = 0$, and typical parameter values $r = \kappa_p = \kappa_s = \eta = 0$. Since the asset price is constant, problem (6.3) degenerates to the deterministic control problem of minimizing $\mathcal{B}(T)^2$. Moreover, since there is no pricing impact, $\mathcal{B}(T) = \alpha s + \mathscr{b}$ with certainty. Consequently, the value function V is

$$V(s, \mathscr{b}, \alpha, \tau) = \inf_{v(\cdot)} E_{v(\cdot)}^{s,\mathscr{b},\alpha,T-\tau}[\mathcal{B}(T)^2] = \mathcal{B}(T)^2 = (\alpha s + \mathscr{b})^2, \tag{7.1}$$

which can also be verified by direct substitution into the HJB equation (6.8), (6.9) as follows. First, note that the initial condition (6.9) is satisfied because $f(v) \equiv 1$. To verify (6.8), note that the parameter values yield the simplifications

$$\mathcal{L}V = 0, \quad \frac{DV}{D\tau}(v) = V_\tau + vsV_{\mathscr{b}} - vV_\alpha. \tag{7.2}$$

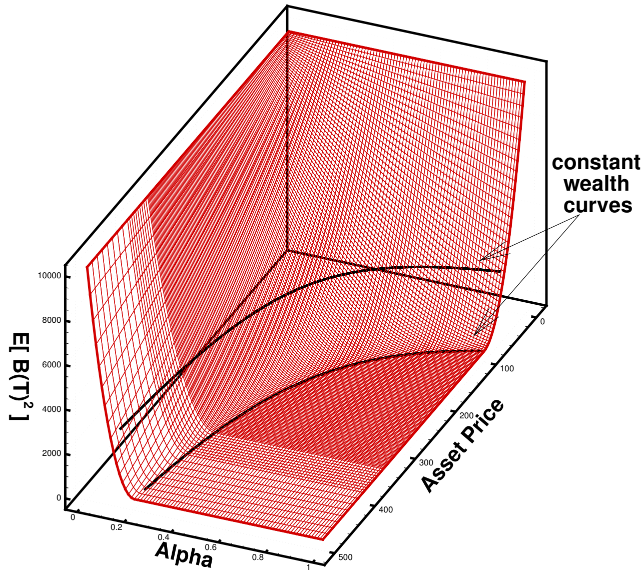

FIGURE 1: *Value function at $t = 0$ and $b = -100$ for parametric Case 1 detailed in Table 1 and Table 2. Here we show two curves of approximately constant wealth $\{\alpha s + b = 0\}$ and $\{\alpha s + b = -50\}$, which correspond to the curves $\{\alpha = 100/s\}$ and $\{\alpha = 50/s\}$, respectively, since b is fixed at -100. Note that the value function V is approximately $V \approx (\alpha s + b)^2$ and changes rapidly normal to the lines of constant wealth.*

Substituting (7.1) into (7.2) gives

$$V_\tau = 0, V_b = 2(\alpha s + b), V_\alpha = 2s(\alpha s + b) \implies \frac{DV}{D\tau}(v) \equiv 0 \text{ for all } v. \tag{7.3}$$

Since any admissible trading velocity v is optimal in this case, the problem of determining the optimal control v is completely ill-posed.

Although the above special case is degenerate, it explains what happens for realistic parametric cases. Indeed, in practical parametric cases, the values of $\sigma\sqrt{T}$ and κ_t are quite small and r has little effect. Although the actual trading velocity is highly dependent on $\sigma\sqrt{T}$ and κ_t, we expect that the actual value function will be only weakly dependent on these parameters.

7.1 Motivation for a parametric curve interpolation method

In using a semi-Lagrangian method [14, 10] to solve for the optimal velocity v, accurate interpolation at the foot of the characteristics is essential to achieving high accuracy [14, 10]. Since a monotone discretization scheme (which allows proof of convergence to the viscosity solution) is at most first-order accurate, we will deal exclusively with linear interpolation in this paper.

Let us consider again the special case in the previous section, where the value function V has the analytic solution

$$V(s, b, \alpha, \tau) = (\alpha s + b)^2. \tag{7.4}$$

It is obvious that linear interpolation along each of the three coordinate axes is not exact, since the partial derivatives V_{ss}, V_{bb} and $V_{\alpha\alpha}$ are all non-zero.

T	η	r	s_{init}	α_{init}	κ_p	κ_s	β	Action	v_{min}
1/250	0.0	0.0	100	1.0	0.0	0.0	1.0	Sell	-1000/T

TABLE 1: *Common parameters*

Case	σ	κ_t	Percentage of Daily Volume
1	1.0	2×10^{-6}	16.7%
2	0.2	2.4×10^{-6}	20.0%
3	0.2	6×10^{-7}	5.0%
4	0.2	1.2×10^{-7}	1.0%
5	0.2	2.4×10^{-8}	0.2%

TABLE 2: *Parametric cases*

Consider linear interpolation along the parametric curve of constant wealth ($\{\alpha s + \not{b} = \text{constant}\}$):

$$\frac{ds}{d\zeta} = 0, \quad \frac{d\not{b}}{d\zeta} = vs, \quad \frac{d\alpha}{d\zeta} = -v, \tag{7.5}$$

for any fixed trading velocity v. Since V is constant along this line, linear interpolation along this parametric curve is exact.

As far as the general form for the value value function is concerned, equation (7.4) suggests that the value function for realistic parameter values should be slowly varying along curves of constant wealth ($\{\alpha s + \not{b} = \text{constant}\}$). In Figure 1 we show the value function computed at $t = 0$ (for fixed $\not{b}$) for a typical set of parameters. The dark curves in Figure 1 are curves of (approximately) constant wealth. Note that the curve which passes through ($\alpha = 1, s = 100$) divides the value function into a region which is almost flat; and a region of rapidly increasing values. Interpolation normal to this curve will result in large errors, while interpolation along this curve will have relatively small errors.

To further improve accuracy, the scheme (7.5) is extended (for general parameters) in section B.4.2 in the appendix.

8 Numerical results

Our discussion on numerical results is organized as follows. First, we explain how we arrive at our parametric cases. Then, we demonstrate convergence by numerical experiments. Finally, we compare the efficient frontiers using the mean variance strategy and the mean quadratic variation strategy.

8.1 Parametric cases

The parametric cases we consider are listed in Table 1 and Table 2. Case 1 corresponds to a high volatility stock with low liquidity. Cases 2-5 correspond to a low volatility stock with various levels of liquidity. These parameters can be related to typical daily volume traded. As described in Appendix A, we estimate that $\kappa_t = 1.2 \times 10^{-7}$ corresponds to liquidating 1% of the daily volume traded of a typical large-cap liquid stock. In view of our trading model and in particular the temporary trading impact function (2.7) with the choice of $\beta = 1$, we can simulate liquidating Y% of the daily volume traded by keeping α_{init} unchanged at unity and using $\kappa_t = (1.2 \times 10^{-7})Y$.

8.2 Convergence Analysis

To demonstrate convergence numerically, we compute the mean-variance frontier from the optimal mean variance strategy using two methods. We explain both methods using the mean variance formulation as an example. The methods for the mean quadratic variation formulation is analogous.

Before discussing the two methods, we note that both methods require the following same initial step: solve for the optimal control $v^*(\cdot)$ in the HJB PDE (6.8).

8.2.1 The PDE method

In the PDE method, the mean variance efficient frontier is obtained from the value functions (6.8) and (6.14). More specifically, the PDE method consists of the following steps.

1. Note that the value function V is computed when solving for $v^*(\cdot)$.
2. Compute the value function U by solving the linear PDE (6.14).
3. Construct the variance-minimizing frontier from U and V as described in Appendix B.1.
4. Eliminate suboptimal points to obtain the mean variance efficient frontier

8.2.2 The Hybrid (PDE-Monte Carlo) method

In the Hybrid method, Monte Carlo simulations are carried out using the optimal control $v^*(\cdot)$ to estimate quantities of interest. More specifically, the Hybrid method consists of the following steps.

1. The optimal control $v^*(\cdot)$ (from solving the HJB PDE (6.8)) is an input.
2. Quantities of interest are estimated by Monte Carlo simulations, as detailed in Appendix C.

An advantage of the Hybrid method is that we can estimate quantities of interest that cannot be obtained in the PDE method. For example, it is important to understand how liquidation proceeds (in forward time) on average, which cannot be known from the PDE solutions directly. The Hybrid method also allows us to estimate both risk measures (variance and quadratic variation) of either the mean variance or the mean quadratic variation strategies, which may not be obtained directly from the value functions.

8.2.3 Computational grid

Tables 3 and 4 show the number of nodes and time steps used in the convergence study for the mean-variance strategy and the mean-quadratic-variation strategy, respectively. Note that only one node is needed in the b direction, since this variable can be eliminated using a similarity reduction (see section B.2 and [16]).[3] The v node discretization is required in order to carry out a linear search to determine the optimal control [16].

Our parametric curve interpolation scheme (see section B.4.2 in the appendix for details) suggests that the number of s nodes should be significantly more than the number of α nodes, a consideration that is also confirmed by our numerical experiments.

Note also that the same time steps are used in both PDE calculation and Monte Carlo simulations, for each refinement level. For example, the frontiers labeled with "800 time steps" in Figure 2 use the time steps as specified as Refinement Level 2 in Table 3. Similarly for the frontiers labeled with "1600 time steps" and for the frontiers in other figures in the report.

[3]For the mean-variance strategy, our numerical experiments suggest that the v grid needs to be fine near $t = 0$ (but not when t is larger) to obtain accurate estimate of the optimal v by linear search. In order to have a very fine v grid near $t = 0$ but a coarse v grid elsewhere, we perform 4 additional refinements to the v grid in the last few backward time steps in PDE solve. There is no such concern for the mean- quadratic-variation strategy because the optimal velocity is not determined by linear search (we use a one dimensional optimization method) and hence is not restricted to values in a discrete v grid.

Refinement Level	Timesteps	s nodes	$\mathfrak{b}$ node	α nodes	v nodes
0	200	369	1	11	8
1	400	737	1	21	15
2	800	1473	1	41	29
3	1600	2945	1	81	57

TABLE 3: *Grid and time step data for convergence studies for the mean variance strategy. The same time steps are used in both PDE calculation and Monte Carlo simulations. Note that there is only one $\mathfrak{b}$ node because of the use of similarity reduction (see section B.2 and [16]).*

Refinement Level	Timesteps	s nodes	α nodes	v nodes
0	800	67	41	30
1	1600	133	81	59

TABLE 4: *Grid and time step data for convergence studies for the mean quadratic variation strategy from [17]. The same time steps are used in both PDE calculation and Monte Carlo simulations. The Monte Carlo computations interpolate the optimal control from the PDE grid values.*

8.2.4 Sample size in Monte Carlo simulations

To achieve small sampling error in Monte Carlo simulations, 400,000 simulations are performed for parametric case 1 and 100,000 simulations are performed for each of the other cases. As an example, the standard error in Figure 2(a) can be estimated as follows. To be more specific, consider a point on the frontier with the maximum standard deviation, which equals 3. Since this estimate of standard deviation of $B(T)$ is calculated using 400,000 samples, its standard error is approximately $3/\sqrt{400,000} \approx 0.0047$, which is negligible in Figure 2(a). Similar calculations will show that the standard errors are negligible in other figures as well.

8.2.5 Comparison between the PDE method and the Hybrid method

For the mean variance optimal strategies, Figure 2 shows that the mean variance efficient frontiers computed by both the PDE method and the Hybrid method converge to the same frontier as the computational grid is refined.

Similarly, for the mean quadratic variation optimal strategies, Figure 3 shows that the mean quadratic variation efficient frontiers computed by both the PDE method and the Hybrid method converge to the same frontier as the computational grid is refined.

Our numerical results demonstrate that the Hybrid frontiers in general converge faster to the limit solution than the PDE frontiers. This may seem counter-intuitive as the Monte Carlo simulations use the optimal trading strategies determined by the PDE method. Nevertheless, it is plausible that Monte Carlo simulations produce a better estimate of the expected value (or standard-deviation/quadratic-variation), which is what our numerical results suggests. Given their better accuracy, we will use the Hybrid frontiers to compare the mean-variance strategy with the mean-quadratic-variation strategy. Again, note that the optimal controls are always computed by solving the HJB PDEs.

8.3 Comparisons of two risk measures

8.3.1 Making comparisons in the same units

In our figures, we plot the expected value against the risk measures in the same units. This means standard deviation is plotted instead of variance. Similarly, the square root of quadratic variation is plotted instead of quadratic variation (4.1). In our plots, the terminology $QVRisk$ stands for the square root of quadratic

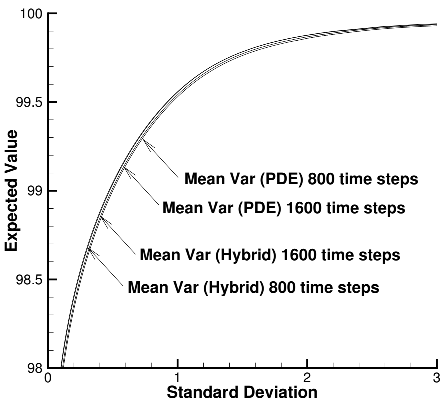

(a) σ=*1.0*, $\kappa_t = 2 \times 10^{-6}$

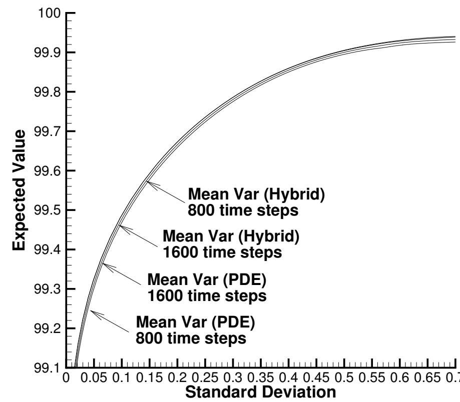

(b) σ=*0.2*, $\kappa_t = 2.4 \times 10^{-6}$

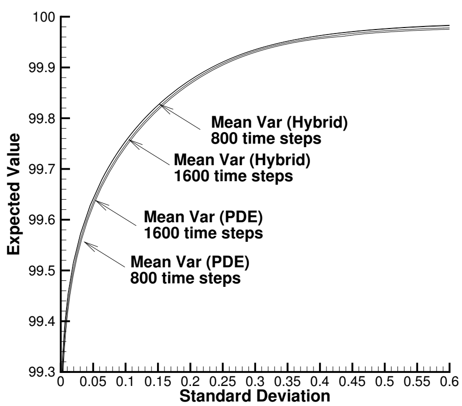

(c) σ=*0.2*, $\kappa_t = 6 \times 10^{-7}$

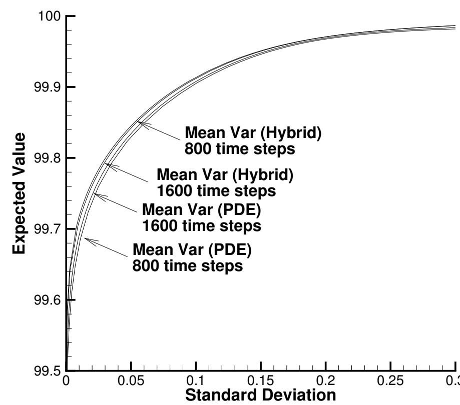

(d) σ=*0.2*, $\kappa_t = 1.2 \times 10^{-7}$

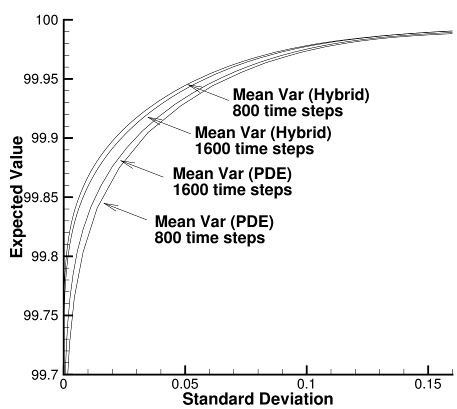

(e) σ=*0.2*, $\kappa_t = 2.4 \times 10^{-8}$

FIGURE 2: *Mean variance strategy: convergence of frontiers in the PDE method and the Hybrid method. The frontiers labeled with PDE are obtained from the PDE value functions. The frontiers labeled with Hybrid are obtained from Monte Carlo simulations which use the optimal controls determined by solving the HJB equation (6.8).*

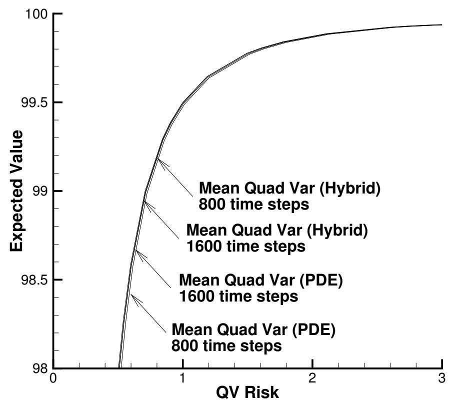

(a) σ=*1.0*, $\kappa_t = 2 \times 10^{-6}$

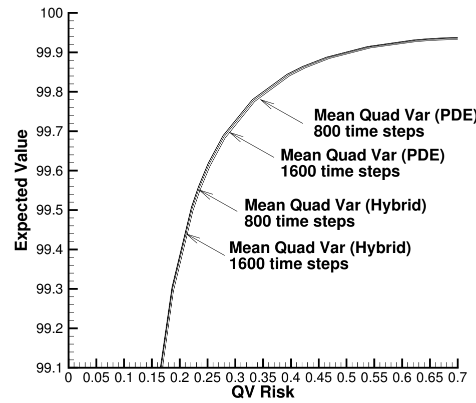

(b) σ=*0.2*, $\kappa_t = 2.4 \times 10^{-6}$

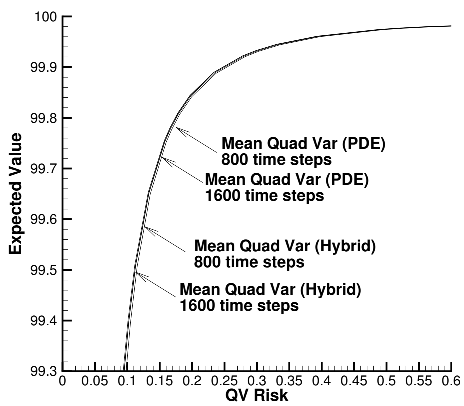

(c) σ=*0.2*, $\kappa_t = 6 \times 10^{-7}$

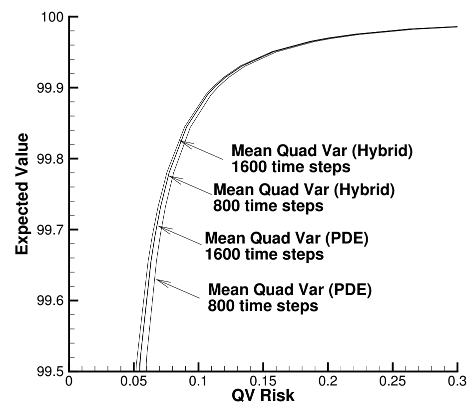

(d) σ=*0.2*, $\kappa_t = 1.2 \times 10^{-7}$

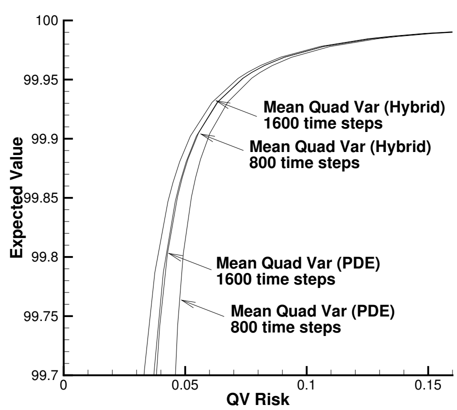

(e) σ=*0.2*, $\kappa_t = 2.4 \times 10^{-8}$

FIGURE 3: *Mean quadratic variation strategy: convergence of frontiers in the PDE method and the Hybrid method. The frontiers labeled with PDE are obtained from the PDE value functions. The frontiers labeled with Hybrid are obtained from Monte Carlo simulations which use the optimal controls determined by solving the HJB equation (4.7).*

variation, i.e.

$$QVRisk = \sqrt{E\left[\int_0^T \left(A(t')dS(t')\right)^2\right]}. \tag{8.1}$$

8.3.2 Summary of comparisons

Figures 4 to 8 compare the mean variance trade off and the mean quadratic variation trade off for both the mean variance and the mean quadratic variation strategy. For example, the left plot in Figure 4 compares the results obtained using the mean variance strategy and the mean quadratic variation strategy in terms of using standard deviation as the risk measure. Similarly, the right plot in Figure 4 compares the two strategies in terms of using QV Risk as the risk measure.

Several conclusions can be drawn from the comparisons.

- As one would expect, in terms of using standard deviation as the risk measure, the mean variance optimal strategy dominates the mean quadratic variation optimal strategy.
- Conversely, in terms of using QV Risk as the risk measure, the mean quadratic variation optimal strategy dominates the mean variance optimal strategy.
- However, it appears that the mean variance optimal strategy performs reasonably well using either risk measure. The difference between the two strategies is most pronounced at lower risk levels.

8.3.3 Remarks

Market practitioners may consider expected implementation shortfall (the relative difference between expected value and initial stock price) of 10 basis points to be significant. To achieve small implementation shortfall, liquidation must be done slowly to reduce trading impact, at the expense of increasing timing risk. Striking a good balance is important here, as it might not be wise to aim at an expected shortfall of 10 bps if the risk (as measured by either standard deviation or QV Risk) is several times larger. Our plots show that risk can be several times of a 10 bps expected shortfall in the parametric cases (a) $\sigma = 1.0$, 16.7% daily volume; (b) $\sigma = 0.2$, 20% daily volume; and (c) $\sigma = 0.2$, 5% daily volume.

The analysis above suggests that one way to choose a risk aversion level on an efficient frontier is to choose a ratio between the implementation shortfall and risk. Alternatively, a common practice among market practitioners is to pick the "corner of the frontier". Our plots show that picking the corner can result in expected implementation shortfall much larger than 10 bps.

8.4 Comparison of strategies for similar expected values

In this section we compare the mean-variance strategy with the mean-quadratic- variation strategy when they give similar expected values. In particular, we focus on the parametric case $\sigma = 1$, $\kappa_t = 2 \times 10^{-6}$ since the differences are more apparent when volatility and pricing impacts are larger.

Figures 9 to 12 correspond to comparisons across four horizontal lines in Figure 4, with four different expected values chosen to represent the more interesting part of the frontiers. For example, in Figure 9 both strategies give an expected value of around 99.29. For the mean-variance strategy, this corresponds to $\gamma = 199.82$; for the mean-quadratic-variation strategy, this corresponds to $\lambda = 1$.

8.4.1 Common observations for each level of expected value

In each of Figures 9 to 12, the subplots labeled (a) and (b) compare the optimal trading velocities at $t = 0$, where we normalized the trading velocities so that a normalized velocity of -1.0 corresponds to the constant

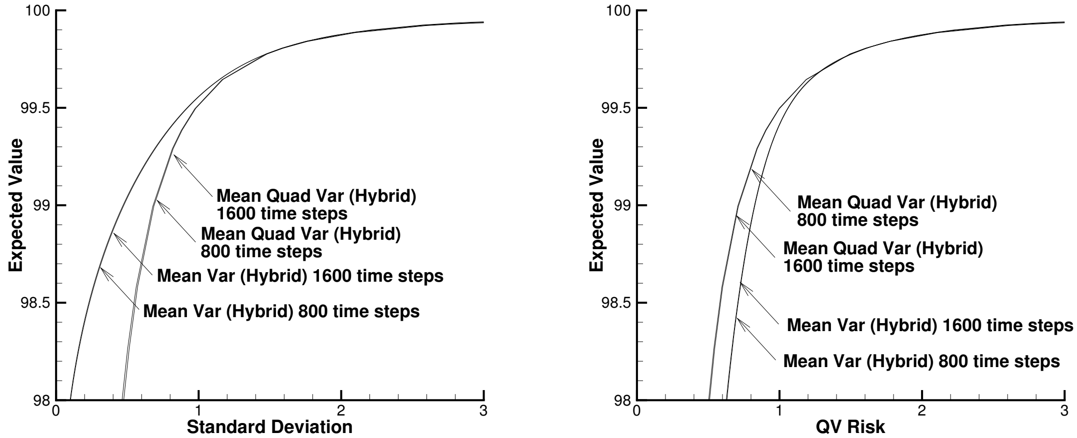

FIGURE 4: *Comparison between mean variance strategy and mean quadratic variation strategy for the case* σ*=1.0,* $\kappa_t = 2 \times 10^{-6}$.

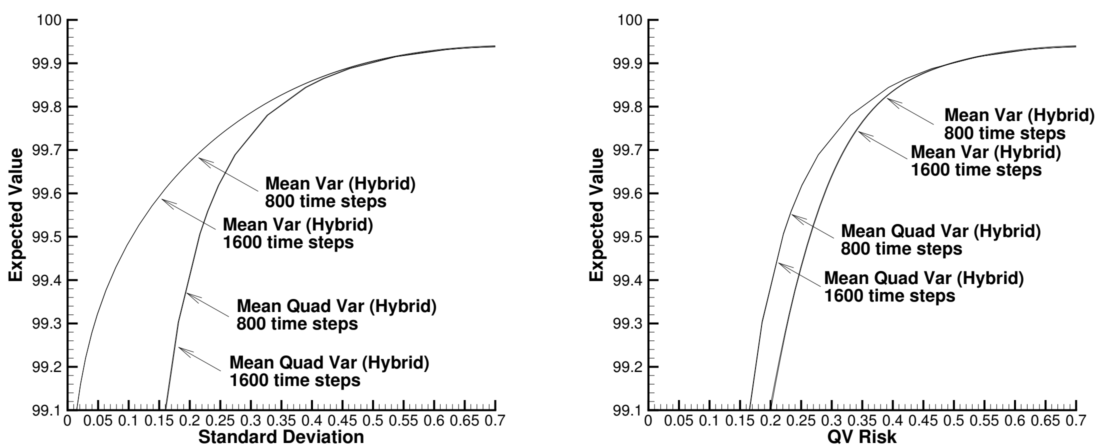

FIGURE 5: *Comparison between mean variance strategy and mean quadratic variation strategy for the case* σ*=0.2,* $\kappa_t = 2.4 \times 10^{-6}$.

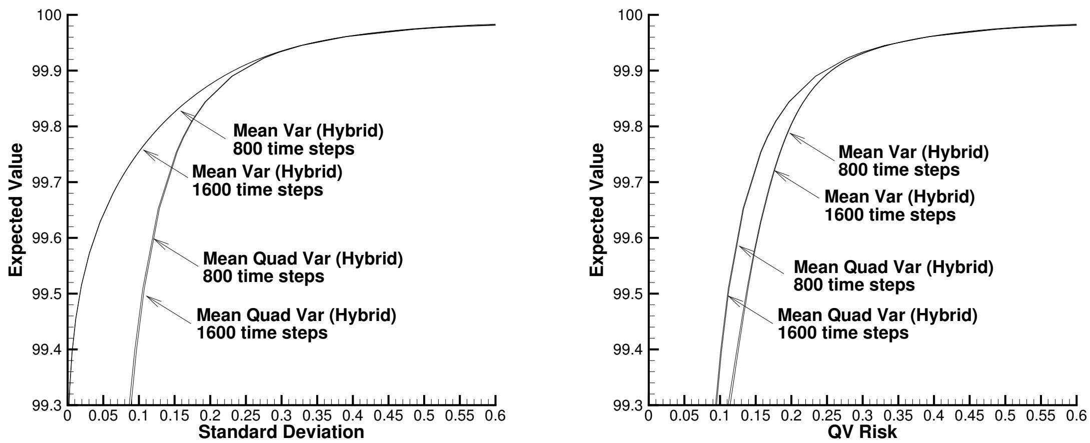

FIGURE 6: *Comparison between mean variance strategy and mean quadratic variation strategy for the case* σ*=0.2,* $\kappa_t = 6 \times 10^{-7}$.

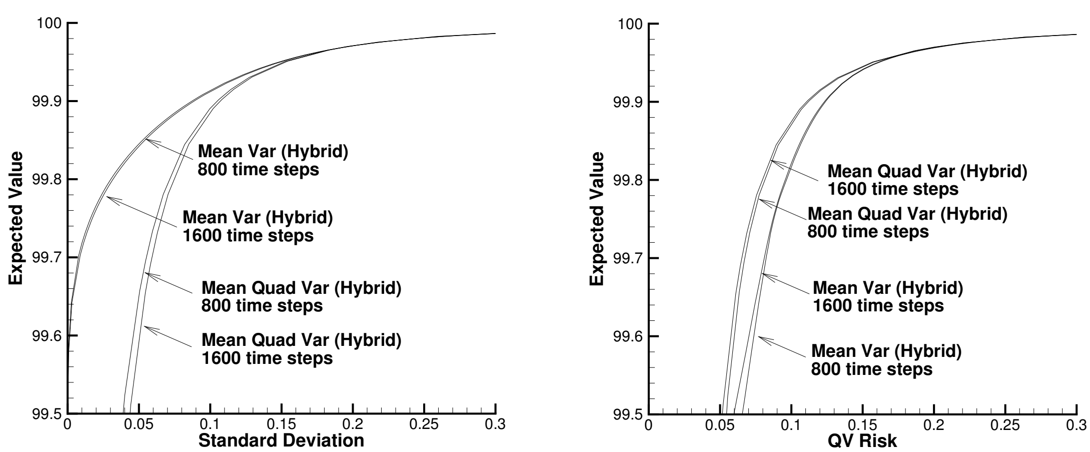

FIGURE 7: *Comparison between mean variance strategy and mean quadratic variation strategy for the case* σ*=0.2,* $\kappa_t = 1.2 \times 10^{-7}$.

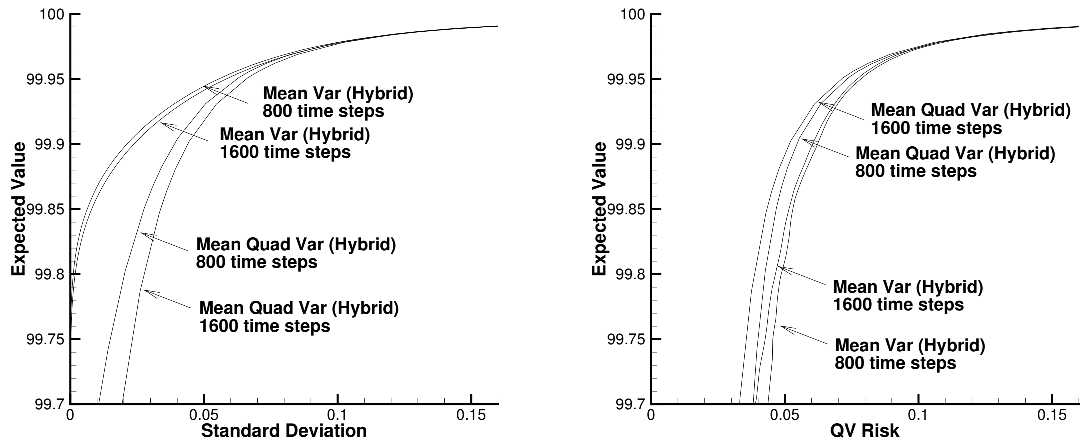

FIGURE 8: *Comparison between mean variance strategy and mean quadratic variation strategy for the case* σ*=0.2,* $\kappa_t = 2.4 \times 10^{-8}$.

liquidation rate $-\alpha_{init}/T$. It is clear that while both strategies sell faster as price becomes more favorable[4], the sensitivity in the mean-variance strategy is much more non-linear. More specifically, around the initial asset price $s_{init} = 100$, the optimal control for the mean-variance strategy is a curve with rapidly changing slope whereas that for the mean-quadratic-variation strategy is more or less a straight line. It is also worth noting that the optimal selling rates at $s_{init} = 100$ for the mean-variance strategy are close to but slightly larger than those for the mean-variation strategy in Figures 9 to 12.

Note that the trading velocity in Figure 9 (b) is nonsmooth for large values of asset price. This appears to be due to the near illposedness of the mean-variance formulation, as discussed in [16]. This can also be understood from Figure 1, which has a flat region where $V_s = V_b = V_\alpha = 0$, so that the local objective function (6.6) is independent of the control v.

In each of Figures 9 to 12, the subplots labeled (c) and (d) compare the mean and standard deviation, respectively, of the liquidation profiles $\alpha(t)$ of the mean-variance and the mean-quadratic-variation strategies over the trading horizon. It is interesting to note that while the mean profiles are very similar, the standard deviation profile of the mean-variance strategy is much larger than that of the mean-quadratic-variation strategy. This reflects the fact that the mean-variance strategy is much more sensitive to change in asset price during the liquidation, which is also suggested by the strategy subplots. We also note that the mean profiles are convex, so that the mean liquidation rate is always decreasing over time.

8.4.2 Differences among different levels of expected value

As we move from Figure 9 to 12, the expected value is increasing, and so is the standard deviation and QV Risk. By comparing subplots (a) and (b), we see that the optimal selling rates become slower as expected value increases. Recall that the mean profiles are convex, so that the mean liquidation rate is always decreasing over time. By comparing subplots (c), we observe that the convexities of the mean liquidation profiles diminish as expected value increases, and the mean liquidation profiles approach a straight line. By

[4]For the buying case, mean-variance optimal strategies would buy faster as price becomes more favorable (drops). For both selling and buying, mean-variance optimal strategies trade faster as price becomes more favorable; this property is called aggressive in the money [22]. Mean variance-optimal strategies are aggressive in the money because this introduces an anti-correlation between trading revenue and trading impacts [25]. Mean-quadratic-variation optimal strategies are different: for the buying case, mean-quadratic-variation optimal strategies would buy faster as price increases (unfavorable), to reduce quadratic variation of the remaining position; we have verified this numerically.

comparing subplots (d), we observe that the mean-variance strategy becomes less variable as expected value increases.

9 Conclusion

In this paper, we have compared the optimal trading strategies obtained using two objective functions: mean variance and mean quadratic variation. Recall that the original strategy proposed in [3] is actually a mean quadratic variation strategy [17]. The mean quadratic variation is naturally time-consistent [12, 32]. On the other hand, the pre-commitment mean variance strategy [16, 8, 26] is not time consistent. However, the pre-commitment mean variance strategy is undoubtedly optimal if performance is measured in terms of observed post-trade mean variance data.

The mean variance strategy is a highly nonlinear function of the asset price. By contrast, the mean quadratic variation strategy is approximately linear (constant) in the asset price. Consequently, the mean variance strategy has a much more variable liquidation profile than the mean quadratic variation strategy. Nevertheless, both strategies turn out to have very similar mean liquidation profiles.

In terms of using both standard deviation and QV Risk as risk measures, the mean variance strategy appears to be, overall, a good strategy. The difference between the two strategies, however, are only significant at low levels of timing risk, or equivalently, high levels of implementation shortfall.

Consequently, if a highly variable strategy is acceptable, the mean-variance strategy is perhaps the better choice. Otherwise, the mean quadratic variation strategy should be chosen if less variability in the strategy is desired.

We have improved the numerical method in [16] by using a parametric curve interpolation scheme and a scaled computational grid. The parametric curve interpolation accurately approximates the foot of the semi-Lagrangian characteristics, which is essential for obtaining accurate numerical solutions for the optimal control. The scaled computation grid concentrates computational resources on regions of interest in the state space so that sufficiently accurate results can be produced using few grid nodes.

A Example Computation for the Temporary Price Impact Factor

Here we describe a realistic scenario in which the temporary price impact factor $\kappa_t = 1.2 \times 10^{-7}$ (Case 4 in Table 2) corresponds to 1% of the daily volume of a stock.

Suppose that the initial stock price $s_{init} = 100$ dollars, buy rate = 1,000 shares/min, corresponding temporary price impact = 3 dollars/min, daily trading time = 420 minutes, and daily volume = 42,000,000 shares. For such a scenario, our trading corresponds to 1% of the daily volume, and the daily market turnover for the stock is 4.2 billion dollars, corresponding to that typical of a large-cap stock.

Assuming a constant trading rate over one day ($T = 1/250$), then the total price impact is 3×420. The ratio of total price impact to total initial value of stock is then given by

$$R = \frac{\text{total price impact}}{\text{total initial value}} = \frac{3 \times 420}{420 \times 1000 \times 100} = 3 \times 10^{-5} \tag{A.1}$$

From the trading model (2.4) and (2.7), the captured price is $s_{init} f(v) = s_{init} \exp(-\kappa_t v) \approx s_{init}(1 - \kappa_t v)$. Therefore, the ratio R is approximately $\kappa_t |v|$.

Since $\alpha_{init} = 1$ and $T = 1/250$, the constant trading rate is $v = -250$. Substituting $v = -250$ into $\kappa_t |v|$ $= 3 \times 10^{-5}$ gives $\kappa_t = 1.2 \times 10^{-7}$.

B Details of numerical method for solving PDEs (6.8) and (6.14)

The numerical method used in this paper for solving equations (6.8) and (6.14) is essentially the method used in [16] with two improvements: the use of a parametric curve linear interpolation scheme at the foot of the semi-Lagrangian characteristics and a scaled computational grid. In order to highlight the differences,

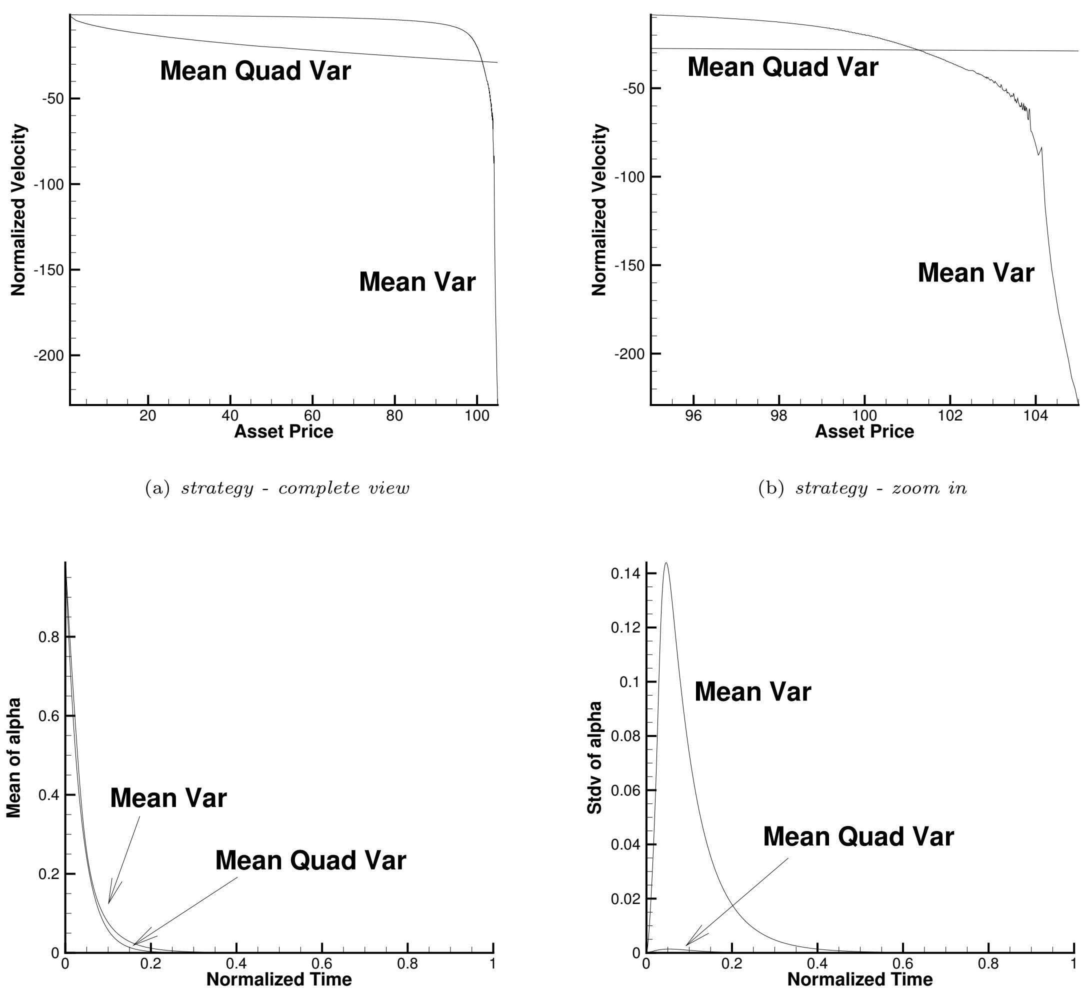

(a) *strategy - complete view*

(b) *strategy - zoom in*

(c) *mean of liquidation profile*

(d) *standard deviation of liquidation profile*

FIGURE 9: *Comparison between mean variance strategy and mean quadratic variation strategy for the case σ=1.0, $\kappa_t = 2 \times 10^{-6}$. The mean- variance strategy plotted has mean 99.29, standard deviation 0.68, QV Risk 0.93, and corresponds to γ=199.82. The mean-quadratic-variation strategy plotted has mean 99.29, standard deviation 0.82, QV Risk 0.84, and corresponds to λ=1. 1600 time steps are used to compute the results.*

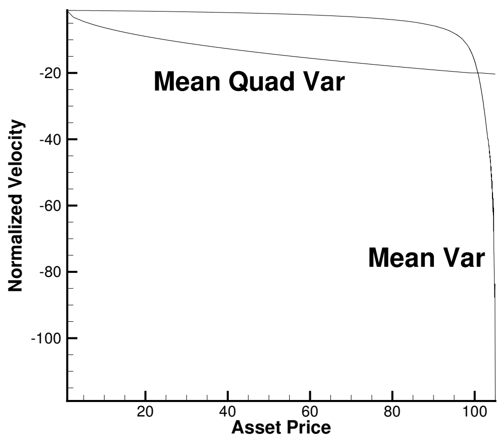

o

(a) *strategy - complete view*

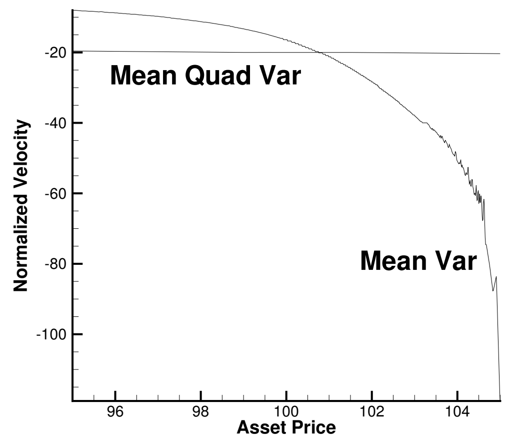

(b) *strategy - zoom in*

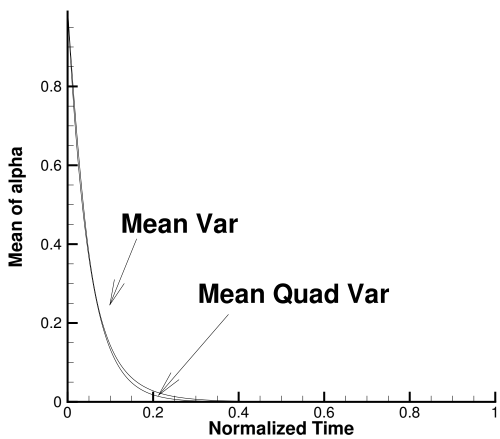

(c) *mean of liquidation profile*

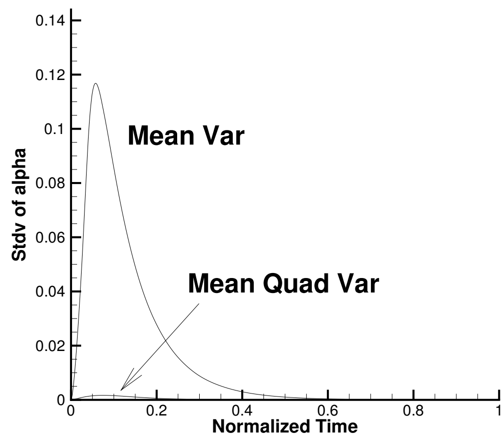

(d) *standard deviation of liquidation profile*

FIGURE 10: *Comparison between mean variance strategy and mean quadratic variation strategy for the case σ=1.0, $\kappa_t = 2 \times 10^{-6}$.The mean- variance strategy plotted has mean 99.50, standard deviation 0.90, QV Risk 1.05, and corresponds to γ=201.30. The mean-quadratic-variation strategy plotted has mean 99.50, standard deviation 0.98, QV Risk 1.00, and corresponds to λ=0.5. 1600 time steps are used to compute the results.*

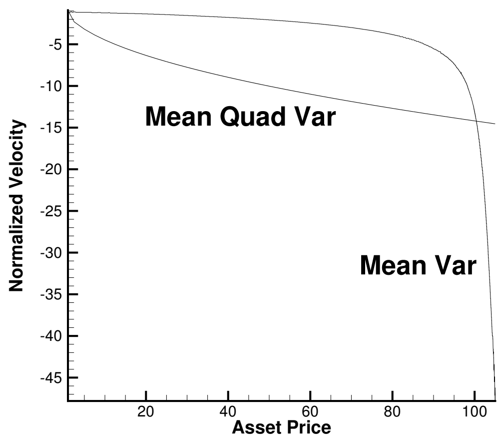

(a) *strategy - complete view*

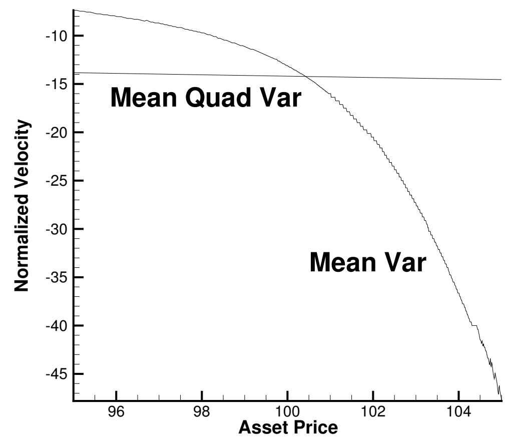

(b) *strategy - zoom in*

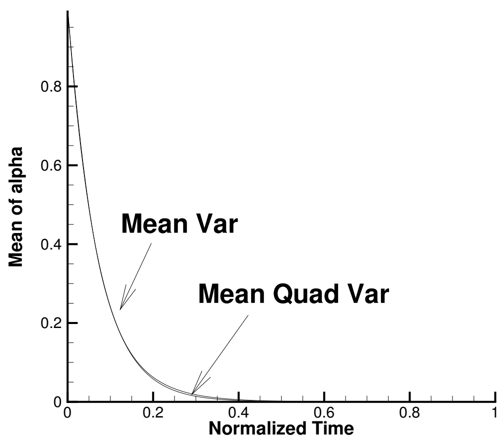

(c) *mean of liquidation profile*

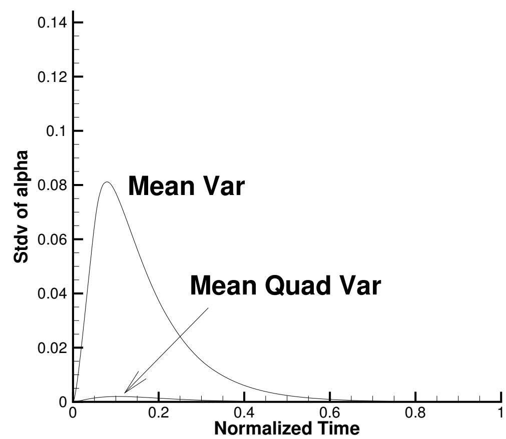

(d) *standard deviation of liquidation profile*

FIGURE 11: *Comparison between mean variance strategy and mean quadratic variation strategy for the case σ=1.0, $\kappa_t = 2 \times 10^{-6}$.The mean- variance strategy plotted has mean 99.65, standard deviation 1.13, QV Risk 1.21, and corresponds to γ=203.50. The mean-quadratic-variation strategy plotted has mean 99.65, standard deviation 1.17, QV Risk 1.19, and corresponds to λ=0.25. 1600 time steps are used to compute the results.*

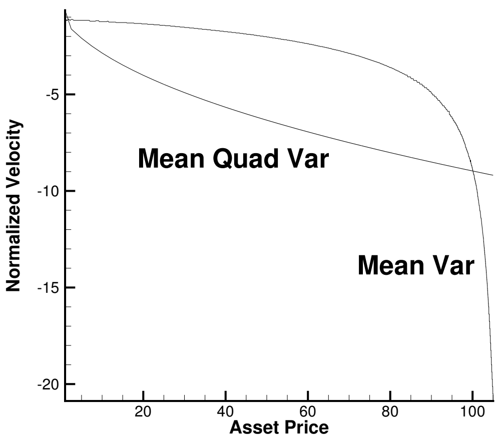

(a) *strategy - complete view*

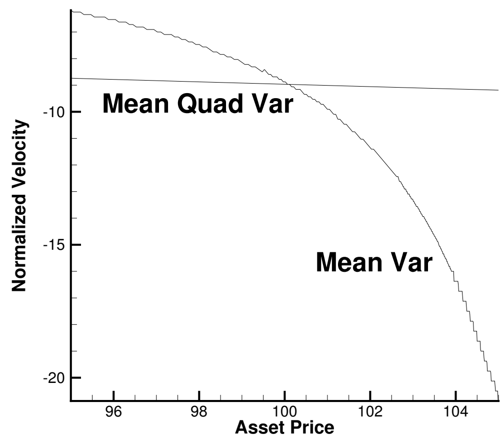

(b) *strategy - zoom in*

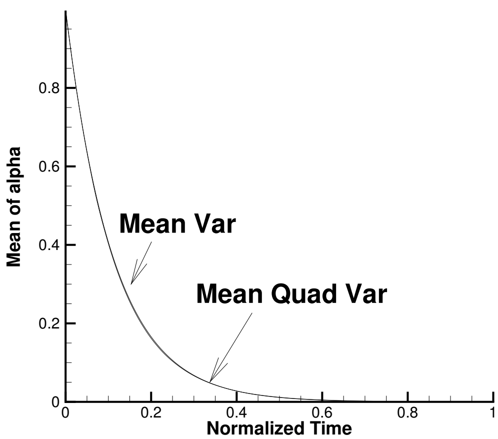

(c) *mean of liquidation profile*

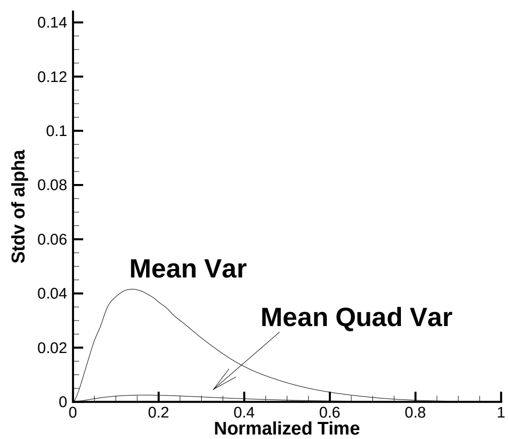

(d) *standard deviation of liquidation profile*

FIGURE 12: *Comparison between mean variance strategy and mean quadratic variation strategy for the case σ=1.0, $\kappa_t = 2 \times 10^{-6}$.The mean- variance strategy plotted has mean 99.78, standard deviation 1.46, QV Risk 1.49, and corresponds to γ=209.42. The mean-quadratic-variation strategy plotted has mean 99.78, standard deviation 1.48, QV Risk 1.49, and corresponds to λ=0.1. 1600 time steps are used to compute the results.*

we provide the discretization details only for these two improvements. Readers are referred to [16] for details on other aspects of the numerical method.

B.1 Construction of efficient frontier

Having solved (6.8) and (6.14), the variance minimizing frontier can be obtained as follows. Let $s = s_{init}$ and $\alpha = \alpha_{init}$ be the initial values of s and α in forward time. For each value of γ, the corresponding point on the variance minimizing frontier can be shown to be given by the formulae

$$\begin{aligned} E_{v^*(\cdot)}^{s,\mathscr{b},\alpha,t=0}[B(T)] &= U_0(\mathscr{b}) + \frac{\gamma}{2}, & \text{(B.1)} \\ Var_{v^*(\cdot)}^{s,\mathscr{b},\alpha,t=0}[B(T)] &= V_0(\mathscr{b}) - (U_0(\mathscr{b}))^2, & \text{(B.2)} \end{aligned}$$

where $\mathscr{b} = \mathcal{B}(0) = -\gamma e^{-rT}/2$ is computed using equation (6.1), and $U_0(\mathscr{b})$ and $V_0(\mathscr{b})$ are shorthand notations for

$$\begin{aligned} V_0(\mathscr{b}) &\equiv V(s, \mathscr{b}, \alpha, \tau = T) = E_{v^*(\cdot)}^{s,\mathscr{b},\alpha,t=0}[\mathcal{B}(T)^2], & \text{(B.3)} \\ U_0(\mathscr{b}) &\equiv U(s, \mathscr{b}, \alpha, \tau = T) = E_{v^*(\cdot)}^{s,\mathscr{b},\alpha,t=0}[\mathcal{B}(T)], & \text{(B.4)} \end{aligned}$$

which are obtained from solving the PDEs (6.8) and (6.14). The formulae (B.1) and (B.2) are obtained by solving for $E_{v^*(\cdot)}^{s,\mathscr{b},\alpha,t=0}[B(T)]$ and $E_{v^*(\cdot)}^{s,\mathscr{b},\alpha,t=0}[B(T)^2]$ from the linear system

$$\begin{aligned} E_{v^*(\cdot)}^{s,\mathscr{b},\alpha,t=0}[B(T)^2] - \gamma E_{v^*(\cdot)}^{s,\mathscr{b},\alpha,t=0}[B(T)] + \frac{\gamma^2}{4} &= E_{v^*(\cdot)}^{s,\mathscr{b},\alpha,t=0}[\mathcal{B}(T)^2] & \text{(B.5)} \\ E_{v^*(\cdot)}^{s,\mathscr{b},\alpha,t=0}[B(T)] - \frac{\gamma}{2} &= E_{v^*(\cdot)}^{s,\mathscr{b},\alpha,t=0}[\mathcal{B}(T)] & \text{(B.6)} \end{aligned}$$

The whole variance minimizing frontier is then obtained by varying γ.

B.2 Similarity Reduction

The assumption of Geometric Brownian Motion (2.5), the form of the price impact functions (2.7), (2.8), and the initial conditions (6.8), (6.14) imply the homogeneity properties

$$\begin{aligned} V(\xi s, \xi \mathscr{b}, \alpha, \tau) &= \xi^2 V(s, \mathscr{b}, \alpha, \tau), \\ U(\xi s, \xi \mathscr{b}, \alpha, \tau) &= \xi U(s, \mathscr{b}, \alpha, \tau), \\ v^*(\xi s, \xi \mathscr{b}, \alpha, \tau) &= v^*(s, \mathscr{b}, \alpha, \tau). & \text{(B.7)} \end{aligned}$$

Therefore we can use similarity reduction to reduce the original three dimensional problem to a two dimensional problem, in which we only need to solve for one fixed value of $\mathscr{b}$.

B.3 Semi-Lagrangian discretization

In this section we demonstrate how equation (6.8) can be discretized by a semi-Lagrangian method. Equation (6.14) can be discretized in a similar fashion. For more details concerning semi-Lagrangian methods for HJB equations, the reader is referred to the references in [14].

Define a set of nodes $\{s_i\}$, $\{\mathscr{b}_j\}$, $\{\alpha_k\}$ and $\{\tau^n\}$, where $0 \le i \le i_{max}$, $\mathscr{b}_j \equiv \mathscr{b}^* < 0$, $0 \le k \le k_{max}$, and $0 \le n \le n_{max}$. We order the nodes in ascending order and make $s_0 = 0$, $\alpha_0 = 0$, $\alpha_{k_{max}} = \alpha_{init}, \tau_0 = 0$, and $\tau^{n_{max}} = T$. Note that there is only one node in the $\mathscr{b}$ grid because of the use of similarity reduction. We denote the discrete approximation to V at the point $(s_i, \mathscr{b}_j, \alpha_k, \tau^n)$ by $V_{i,j,k}^n$ to distinguish it from the exact value $V(s_i, \mathscr{b}_j, \alpha_k, \tau^n)$. We also specify that the set of admissible control Z is of the form $[v_{min}, 0]$, where $v_{min} < 0$ is the fastest liquidation rate allowed.

Since the Lagrangian derivative $\frac{DV}{D\tau}(v_{i,j,k}^{n+1})$ at the node $(s_i, \mathscr{b}_j, \alpha_k, \tau^{n+1})$ is the derivative of V along the trajectory defined by (6.7). Solving equations (6.7) backwards in time from τ^{n+1} to τ^n, for a fixed $v_{i,j,k}^{n+1}$, gives the the foot of the characteristics $(s_{\hat{i}}, \mathscr{b}_{\hat{j}}, \alpha_{\hat{k}}, \tau^n)$, which in general is not on the PDE mesh. We use the notation $V_{\hat{i},\hat{j},\hat{k}}^n$ to denote an approximation of $V(s_{\hat{i}}, \mathscr{b}_{\hat{j}}, \alpha_{\hat{k}}, \tau^n)$ obtained by interpolation.

B.3.1 Local optimization

Denote the discrete form of $\mathcal{L}$ by $\mathcal{L}_h$. By using an implicit discretization of $\mathcal{L}V$ and the semi-Lagrangian discretization on equation (6.8), we obtain

$$V_{i,j,k}^{n+1} = \min_{v_{i,j,k}^{n+1} \in Z_k^{n+1}} V_{\hat{i},\hat{j},\hat{k}}^n + (\tau^{n+1} - \tau^n)(\mathcal{L}_h V)_{i,j,k}^{n+1}, \tag{B.8}$$

with the initial condition

$$V_{i,j,k}^0 = \mathscr{b}_j^2, \tag{B.9}$$

where we restrict the admissible velocities to Z_k^{n+1} so that $\alpha_{\hat{k}} \geq 0$.

Once the optimal control $(v^*)_{i,j,k}^{n+1}$ is determined, equation (6.14) can be solved by

$$U_{i,j,k}^{n+1} = U_{\hat{i},\hat{j},\hat{k}}^n\big|_{v=(v^*)_{i,j,k}^{n+1}} + (\tau^{n+1} - \tau^n)(\mathcal{L}_h V)_{i,j,k}^{n+1}, \tag{B.10}$$

with the initial condition

$$U_{i,j,k}^0 = \mathscr{b}_j. \tag{B.11}$$

Since no analytical expression is available for the local objective function, we find the optimal $v_{i,j,k}^{n+1}$ by discretizing the control space Z_k^{n+1} and look for the optimal value using a linear search. This has the advantage of not making any assumptions about the local objective function, at the expense of a higher computational cost. Numerical experiments demonstrate that accurate results can be obtained by a rather coarse discretization of the control space.

B.4 Computational challenges and solutions

B.4.1 Difficulties in determining optimal velocity numerically

Recall that in the special case considered in section 7, the Lagrangian derivative is identically zero for any admissible trading velocities. In terms of equation (B.8), this means that $V_{\hat{i},\hat{j},\hat{k}}^n(v_{i,j,k}^{n+1})$ as a function of $v_{i,j,k}^{n+1}$ is completely flat. In the parametric cases we consider, both $\sigma\sqrt{T}$ and κ_t are quite small, therefore these realistic cases are indeed similar to the completely ill-posed special case. Consequently $V_{\hat{i},\hat{j},\hat{k}}^n(v_{i,j,k}^{n+1})$ as a function of $v_{i,j,k}^{n+1}$ can be very flat, which means determining the true minimizer demands extremely high accuracy. It is also obvious that even small interpolation error can significantly alter the estimated trading velocities.

Similar computational issues also arise when ordinary finite-differencing is used instead of the semi-Lagrangian method: since the optimal velocity v will be a function that depends on the ratio of the partial derivatives in $\frac{DV}{D\tau}$. Any error in approximating the partial derivatives can also significantly alter the estimated trading velocities.

B.4.2 Parametric curve linear interpolation

The previous section has explained the importance of accurate interpolation at the foot of the semi-Lagrangian characteristics. In [16], a standard axis-aligned linear interpolation is used, which turns out to be too inaccurate. This is not surprising given the quadratic nature of the value function V and the analysis in section 7.

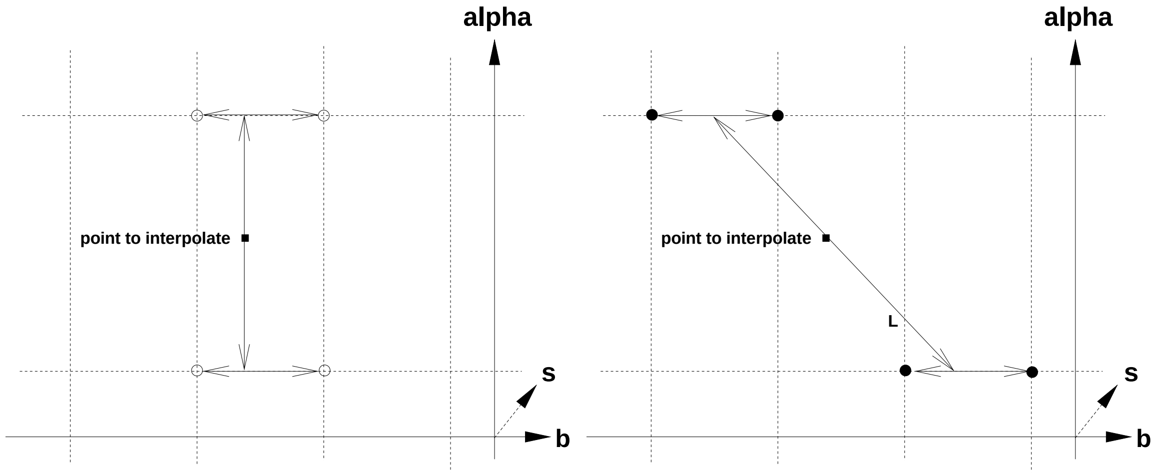

(a) *axis-aligned interpolation* (b) *Parametric curve interpolation along L*

FIGURE 13: *Comparing the two methods of interpolation at the foot of characteristics, shown as "point to interpolate" in the diagram. The dashed lines correspond to the computational grid and the dots are interpolation nodes.*

In section 7 we have shown the benefit of performing a parametric curve linear interpolation for the special case considered. Here we extend the idea to general cases. In essence, the parametric line (7.5) is generalized to the line L defined by

$$\begin{aligned} L &= (s_{\hat{i}}, \text{ƀ}_{\hat{j}}, \alpha_{\hat{k}}) + \zeta(\frac{ds}{d\zeta}, \frac{d\text{ƀ}}{d\zeta}, \frac{d\alpha}{d\zeta}), \\ \frac{ds}{d\zeta} = -g(v_{i,j,k}^{n+1})s_{\hat{i}}, \quad \frac{d\text{ƀ}}{d\zeta} &= -(r\text{ƀ}_{\hat{j}} - v_{i,j,k}^{n+1} f(v_{i,j,k}^{n+1})s_{\hat{i}}), \quad \frac{d\alpha}{d\zeta} = -v_{i,j,k}^{n+1}, \end{aligned} \tag{B.12}$$

where $v_{i,j,k}^{n+1}$ is a candidate control value.

Since equations (B.12) express how changes in α lead to changes in s and ƀ through both trading impact and pricing impact (through the terms $g(v_{i,j,k}^{n+1})$ and $f(v_{i,j,k}^{n+1})$), interpolating along L can be seen as an extension to interpolation along (7.5), which takes into account trading but not pricing impact.

Figure 13 compares the standard axis-aligned linear interpolation and the parametric curve linear interpolation along the line L. For simplicity in illustration, linear interpolation along the s coordinate axis is not shown, and there is a actual ƀ grid, i.e. no similarity reduction.

Note that the parametric curve linear interpolation as shown in Figure 13 does not require interpolation along the α coordinate axis but still requires linear interpolation along the other axes. When linear interpolation along the s direction or the ƀ direction is performed, a fine grid is still needed to reduce interpolation error. In other words, when we treat α specially as in Figure 13, we avoid the need of a fine α grid, but a fine s grid (and a fine ƀ grid when no similarity reduction is used) is still necessary.

We also note that the parametric curve linear interpolation is similar to the edge-directed interpolation method [23] in the image processing literature. Our method is similar in the sense that the direction of the line L is not fixed but adapts to the candidate control velocity $v_{i,j,k}^{n+1}$. Our method is different from that in [23] in that the parametric curve linear interpolation method does not necessarily use the neighboring grid nodes (See Figure 13(b)).

B.4.3 Remark on convergence proof

Having changed the interpolation scheme in the semi-Lagrangian method, it is important that the convergence proof in [16] is still valid. Linear interpolation is obviously consistent. To demonstrate stability, we need the following easy observation:

$V^n_{\hat{i},\hat{j},\hat{k}}$ as approximated in the new scheme takes the form

$$V^n_{\hat{i},\hat{j},\hat{k}} = \sum_p w_p V^n_p, \tag{B.13}$$

where $w_p \geq 0$, $\sum_p w_p = 1$, and V^n_p are grid node values. Therefore, we have $|V^n_{\hat{i},\hat{j},\hat{k}}| \leq ||V^n||$. This property allows the proof of stability in [16] to go through without change. In addition, since $U^n_{\hat{i},\hat{j},\hat{k}}$ as approximated in the scheme also takes the form

$$U^n_{\hat{i},\hat{j},\hat{k}} = \sum_p w_p U^n_p, \tag{B.14}$$

where the w_p's are the same as those in equation (B.13). Therefore, the proof in [16] which shows $(U^n_{i,j,k})^2 \leq V^n_{i,j,k}$ is also valid for our scheme.

B.4.4 Computational grid consideration

Consider again the analytical solution for the special case derived in section 7:

$$V(s, \mathscr{b}, \alpha, \tau) = (\alpha s + \mathscr{b})^2. \tag{B.15}$$

As can be seen in Figure 1, (B.15) is a good approximation for realistic parametric cases. Both Figure 1 and (B.15) show that V does not change a significantly along a constant line of wealth $\{(\alpha s + \mathscr{b}) = const\}$. This observation suggests constructing a computational grid with taking into account constant lines of wealth.

- For $\alpha > 0$, scale the s grid by $\{s_i\} \rightarrow \{s_i\}/\alpha$.
- For $\alpha = 0$, no scaling is performed, i.e. the original s grid $\{s_i\}$ is used.

Figure 14(a) and 14(b) illustrate the scaled computational grid, and the value function under the scaled grid, respectively. Compare the lines of constant wealth in Figure 1 and Figure 14(b).

Note that the shape of the value function is simpler in Figure 14(b) than in Figure 1. Note also that the mesh is always dense in the region $V \approx 0$ in the scaled grid in Figure 14(b), which is not the case in Figure 1. In our experience, accurate values of V in the region $V \approx 0$ are important for computing the efficient frontiers. Thus, the scaled grid is also computationally more effective.

C Details of Monte Carlo simulations

Numerically solving the mean-variance HJB equation (6.8) gives us the optimal strategy on the discrete computational mesh, i.e. $v_{MV}(s_i, \mathscr{b}^*, \alpha_k, t^n)$[5]. By using $v_{MV}(s_i, \mathscr{b}^*, \alpha_k, t^n)$ as input for Monte Carlo simulations, we can obtain information about the trading strategy which is not necessarily available in the PDE solutions. For example, we can estimate the probability distribution of $B(T)$, the quadratic variation, and the mean and the standard deviation of the liquidation profile (plot of $A(t)$ against t).

Similarly, numerically solving the mean-quadratic-variation HJB equation (4.7), for each fixed value of λ, gives us the discrete optimal strategy $v_{MQV}(s_{\tilde{i}}, \alpha_{\tilde{k}}, t^n; \lambda)$[6], which can be used as input for Monte Carlo simulations.

[5] Note the use of forward time notation.

[6] We use the notations $\tilde{i}$ and $\tilde{k}$ to emphasize that the s grid and α grid for solving the mean-quadratic-variation HJB equation (4.7) is not necessarily the same as that for solving the mean-variance HJB equation (6.8). The time grid $\{t^n\}$, however, is chosen to be the same.

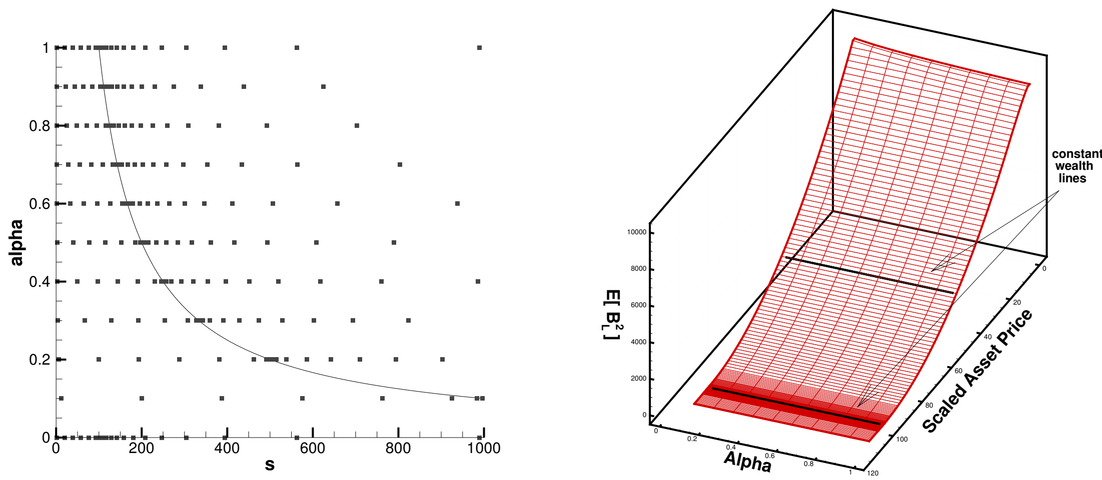

(a) *The scaled computational grid.*

(b) *Value function under the scaled grid. This is the same value function as shown in Figure 1.*

FIGURE 14: *Illustration of scaled computational grid.*

In particular, Monte Carlo simulations enable us to compute the quadratic variation of the mean-variance strategy, and conversely, the variance of the mean-quadratic-variation strategy. These allow us to compare the two strategies in terms of either variance or quadratic variation, given the same level of expected return.

The Monte Carlo simulations also provide a verification of the PDE solutions, in the sense that given the optimal control, we can obtain independent estimates of mean, variance and quadratic variation.

In the following we detail how the Monte Carlo method is conducted for the mean-variance strategy. Simulations of the mean-quadratic-variation strategy are performed in the same way, except that the mean quadratic variation strategy $v_{MQV}(s_{\tilde{i}}, \alpha_{\tilde{k}}, t^n; \lambda)$ is explicitly indexed by λ, and standard axis-aligned linear interpolation suffices.

C.1 Change of variable

Suppose that the optimal strategy $v_{MV}(s_i, \mathscr{b}^*, \alpha_k, t^n)$ is obtained from the PDE solve. For a fixed value of γ, each Monte Carlo simulation starts with the initial values $S(0), B(0), A(0)$ at time $t = 0$ and is updated at the discrete times $\{t^n\}$, i.e. the time grid nodes in the PDE solve. Below we give a full specification of the simulation procedure by detailing the simulation from time point t_{old} to the immediate next time point t_{new}.

At t_{old}, the state is $(S_{old}, B_{old}, A_{old})$. To look up the optimal trading velocity, we first need to change the variable from B to $\mathcal{B}$. For the fixed value of γ, we have $\mathcal{B}_{old} = B_{old} - \gamma e^{-r(T-t_{old})}/2$ from equation (6.2). Now the optimal trading velocity $v(S_{old}, \mathcal{B}_{old}, A_{old}, t_{old})$ needs to be interpolated from the discrete $v_{MV}(s_i, \mathscr{b}^*, \alpha_k, t^n)$.

C.1.1 Interpolation

Our numerical study shows that it is more accurate to linearly interpolate $v_{MV}(s_i, \mathscr{b}^*, \alpha_k, t^n)$ along a constant line of wealth $\{\alpha s + \mathscr{b} = \text{constant}\}$ than along the coordinate axes. Therefore, we interpolate $v_{MV}(s_i, \mathscr{b}^*, \alpha_k, t^n)$ as in Figure 13(b) with L given by the constant line of wealth $\{\alpha s + \mathscr{b} = \text{constant}\}$. This is the same as the interpolation for the limiting parametric case in section 7.1. Note that the form of the line L as defined by equation (B.12) is not applicable in the current context because there is no candidate control $v_{i,j,k}^{n+1}$.

C.1.2 Updating state variables

Let $\triangle t = t_{new} - t_{old}$, we update the state variable as follows:

$$\begin{aligned}
A_{opt} &= A_{old} + v(S_{old}, \mathcal{B}_{old}, A_{old}, t_{old})\triangle t, && \text{(C.1)}\\
A_{new} &= \max(A_{opt}, 0), && \text{(C.2)}\\
v_{opt} &= (A_{new} - A_{old})/\triangle t, && \text{(C.3)}\\
S_{new} &= S_{old}\exp\{(\eta + g(v_{opt}) - \frac{1}{2}\sigma^2)\triangle t + \sigma\sqrt{\triangle t}\mathcal{N}(0,1)\}, && \text{(C.4)}\\
B_{new} &= B_{old}\exp\{r\triangle t\} - v_{opt} f(v_{opt}) S_{old}\triangle t, && \text{(C.5)}\\
QV_{new} &= QV_{old} + (A_{old}(S_{new} - S_{old}))^2, && \text{(C.6)}
\end{aligned}$$

where $\mathcal{N}(0,1)$ is a standard normal variate and QV_{new} is an approximation of $\int_0^{t_{new}} (A(t')\ dS(t'))^2$.

References

[1] A. Alfonsi, A. Fruth, and A. Schied. Optimal execution strategies in limit order books with general shape functions. *Quantitative Finance*, 10(2):143–157, 2010.

[2] R. Almgren. Optimal trading with stochastic liquidity and volatility. *SIAM Journal on Financial Mathematics*, 3:163–181, 2012.

[3] R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–40, 2001.

[4] R. Almgren and J. Lorenz. Adaptive arrival price. *Trading*, 2007(1):59–66, 2007.

[5] R. Almgren, C. Thum, E. Hauptmann, and H. Li. Direct estimation of equity market impact. *Risk*, 18:57–62, 2005.

[6] R.F. Almgren. Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance*, 10(1):1–18, 2003.

[7] L. Bai and H. Zhang. Dynamic mean-variance problem with constrained risk control for the insurers. *Mathematical Methods of Operations Research*, 68(1):181–205, 2008.

[8] S. Basak and G. Chabakauri. Dynamic mean-variance asset allocation. *Review of Financial Studies*, 23:2970–3016, 2010.

[9] D. Bertsimas and A. Lo. Optimal control of execution costs. *Journal of Financial Markets*, 1(1):1–50, 1998.

[10] N. Besse and E. Sonnendrucker. Semi-Lagrangian schemes for the Vlasov equation on an unstructured mesh of phase space. *Journal of Computational Physics*, 191(2):341–376, 2003.

[11] T.R. Bielecki, H. Jin, S.R. Pliska, and X.Y. Zhou. continuous-time mean-variance portfolio selection with bankruptcy prohibition. *Mathematical Finance*, 15(2):213–244, 2005.

[12] Tomas Bjork and Agatha Murgoci. A General Theory of Markovian Time Inconsistent Stochastic Control Problems. *SSRN eLibrary*, 2010.

[13] P. Brugiere. Optimal portfolio and optimal trading in a dynamic continuous time framework. 6th AFIR Colloquium. *Nurenberg, Germany*, 12:89, 1996.

[14] Z. Chen and P.A. Forsyth. A semi-Lagrangian approach for natural gas storage valuation and optimal operation. *SIAM Journal on Scientific Computing*, 30(1):339–368, 2007.

[15] R.F. Engle and R. Ferstenberg. Execution Risk. *The Journal of Trading*, 2(2):10–20, 2007.

[16] P.A. Forsyth. A Hamilton Jacobi Bellman approach to optimal trade execution. *Applied Numerical Mathematics*, 61:241–265, 2011.

[17] P.A. Forsyth, J.S. Kennedy, S.T. Tse, and H. Windcliff. Optimal trade execution: a mean quadratic variation approach. *Journal of Economic Dynamics and Control*, 36:1971–1991, 2012.

[18] C. Fu, A. Lari-Lavassani, and X. Li. Dynamic mean-variance portfolio selection with borrowing constraint. *European Journal of Operational Research*, 200(1):312–319, 2010.

[19] J. Gatheral and A. Schied. Optimal Trade Execution under Geometric Brownian Motion in the Almgren and Chriss Framework. *International Journal of Theoretical and Applied Finance*, 14:353–368, 2011.

[20] H. He and H. Mamaysky. Dynamic trading policies with price impact. *Journal of Economic Dynamics and Control*, 29(5):891–930, 2005.

[21] G. Huberman and W. Stanzl. Price manipulation and quasi-arbitrage. *Econometrica*, 72(4):1247–1275, 2004.

[22] R. Kissell and Malamut R. Understanding the profit and loss of trading algorithm. *Institutional Investor*, (1):41–49, 2005.

[23] X. Li and M.T. Orchard. New edge-directed interpolation. *Image Processing, IEEE Transactions on*, 10(10):1521–1527, 2001.

[24] F. Lillo, J.D. Farmer, and R.N. Mantegna. Econophysics: Master curve for price-impact function. *Nature*, 421(6919):129–130, 2003.

[25] J. Lorenz and R. Almgren. Mean-variance optimal adaptive execution. *Applied Mathematical Finance*, 18:395–422, 2011.

[26] J.M. Lorenz. *Optimal trading algorithms: Portfolio transactions, multiperiod portfolio selection, and competitive online search*. PhD thesis, ETH Zurich, 2008.

[27] V. Ly Vath, M. Mnif, and H. Pham. A model of optimal portfolio selection under liquidity risk and price impact. *Finance and Stochastics*, 11(1):51–90, 2007.

[28] A. Obizhaeva and J. Wang. Optimal trading strategy and supply/demand dynamics. *NBER Working Paper*, 2005.

[29] M. Potters and J.P. Bouchaud. More statistical properties of order books and price impact. *Physica A: Statistical Mechanics and its Applications*, 324(1-2):133–140, 2003.

[30] A. Schied and T. Schoeneborn. Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets. *Finance and Stochastics*, 13(2):181–204, 2009.

[31] S.T. Tse, P.A. Forsyth, and Y. Li. Preservation of scalarization optimal points in the embedding technique for continuous time mean variance optimization. 2012. Working paper, Cheriton School of Computer Science, Waterloo.

[32] J. Wang and P.A. Forsyth. Numerical solution of the Hamilton-Jacobi-Bellman formulation for continuous time mean variance asset allocation. *Journal of Economic Dynamics and Control*, 34(2):207–230, 2010.

[33] X.Y Zhou and D. Li. Continuous-time mean-variance portfolio selection: A stochastic LQ framework. *Applied Mathematics and Optimization*, 42(1):19–33, 2000.