Numerical solution of the Hamilton-Jacobi-Bellman formulation for continuous time mean variance asset allocation under stochastic volatility *

K. Ma † P. A. Forsyth ‡

May 19, 2015

Abstract

We present efficient partial differential equation (PDE) methods for continuous time mean-variance portfolio allocation problems when the underlying risky asset follows a stochastic volatility process. The standard formulation for mean variance optimal portfolio allocation problems gives rise to a two-dimensional non-linear Hamilton-Jacobi-Bellman (HJB) PDE. We use a wide stencil method based on a local coordinate rotation (Ma and Forsyth, 2014) to construct a monotone scheme. Furthermore, by using a semi-Lagrangian timestepping method to discretize the drift term and an improved linear interpolation method, accurate efficient frontiers are constructed. This scheme can be shown to be convergent to the viscosity solution of the HJB equation, and the correctness of the proposed numerical framework is verified by numerical examples. We also discuss the effects on the efficient frontier of the stochastic volatility model parameters.

Keywords: mean-variance, embedding, Pareto optimal, Hamilton-Jacobi-Bellman (HJB) equation, monotone scheme, wide stencil

JEL Codes: C63, D81, G11

1 Introduction

Consider the following prototypical asset allocation problem: an investor can choose to invest in a risk free bond, or a risky asset, and can dynamically allocate wealth between the two assets, to achieve a pre-determined criteria for the portfolio over a long time horizon. In the continuous time mean variance approach, risk is quantified by variance, so that investors aim to maximize the expected return of their portfolios, given a risk level. Alternatively, they aim to minimize the risk level, given an expected return. As a result, mean variance strategies are appealing due to their intuitive nature, since the results can be easily interpreted in terms of the trade-off between the risk and the expected return.

*This work was supported by the Bank of Nova Scotia and the Natural Sciences and Engineering Research Council of Canada

†Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1 k26ma@uwaterloo.ca

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

In the case where the asset follows a Geometric Brownian Motion (GBM), there is considerable literature on the topic (Li and Ng, 2000; Bielecki et al., 2005; Zhou and Li, 2000; Wang and Forsyth, 2010). The multi-period optimal strategy adopted in these papers is of pre-commitment type, which is not time-consistent as noted in Bjork and Murgoci (2010); Basak and Chabakauri (2010). A comparison between time-consistent and pre-commitment strategies is given in Wang and Forsyth (2012), for continuous time mean variance optimization. We note that since a time consistent strategy can be constructed from a pre-commitment policy by adding a constraint (Wang and Forsyth, 2012), the time consistent strategy is sub-optimal compared to the pre-commitment policy, i.e., it is costly to enforce time consistency. In addition, it has been shown in Vigna (2014) that pre-commitment strategies can also be viewed as a target-based optimization which involves minimizing a quadratic loss function. It is suggested in Vigna (2014) that this is intuitive, adaptable to investor preferences, and is also mean variance efficient.

Most previous literature on pre-commitment mean variance optimal asset allocation has been based on analytic techniques (Li and Ng, 2000; Zhou and Li, 2000; Bielecki et al., 2005; Zhao and Ziemba, 2000; Nguyen and Portait, 2002). These papers have primarily employed martingale methods (Bielecki et al., 2005; Zhao and Ziemba, 2000; Nguyen and Portait, 2002) or tractable auxiliary problems (Li and Ng, 2000; Zhou and Li, 2000). However, in general, if realistic constraints on portfolio selection are imposed, e.g., no trading if insolvent and a maximum leverage constraint, then a fully numerical approach is required. As shown in Wang and Forsyth (2008), in the case where the risky asset follows a GBM, realistic portfolio constraints have a significant effect on the efficient frontier.

Another modeling deficiency in previous work on pre-commitment mean variance optimal asset allocation is the common assumption that the risky asset follows a GBM. However, there is strong empirical evidence that asset return volatility is serially correlated, shocks to volatility are negatively correlated with asset returns, and the conditional variance of asset returns is not constant over time. As a result, it is highly desirable to describe the risky asset with a stochastic volatility model. In this case, the standard formulation of mean variance optimal asset allocation problems gives rise to a two-dimensional non-linear HJB PDE. The objective of this article is to develop a numerical method for the pre-commitment mean variance portfolio selection problem when the underlying risky asset follows a stochastic volatility model.

The major contributions of the paper are:

- A fully implicit, consistent, unconditionally monotone numerical scheme is developed for the HJB equation, which arises in the embedding formulation (Zhou and Li, 2000; Li and Ng, 2000) of the pre-commitment mean variance problem under our model set-up. The main difficulty in designing a discretization scheme is development of a monotone approximation of the cross derivative term in the PDE. We use the wide stencil method (Debrabant and Jakobsen, 2013; Ma and Forsyth, 2014) to deal with this difficulty.

- Accurate efficient frontiers are constructed by using a semi-Lagrangian timestepping method to handle the drift term, and an improved method of linear interpolation at the foot of the characteristic in the semi-Lagrangian discretization. In particular, the improved interpolation method uses the exact solution value at a single point, dramatically increasing the accuracy of the numerical results. Any type of constraint can be applied to the investment policy.

- We prove that the scheme developed in this paper converges to the viscosity solution of the nonlinear HJB value equation.

- In order to trace out the efficient frontier solution of our problem we use two techniques: the PDE method and the Hybrid (PDE - Monte Carlo) method (Tse et al., 2013). We also demonstrate that the Hybrid method is superior to the PDE method.

- We carry out several numerical experiments, and illustrate the convergence of the numerical scheme, as well as the effect of modeling parameters on efficient frontiers.

The remainder of this paper is organized as follows: Section 2 describes the underlying processes and the embedding framework, and gives a formulation of an associated HJB equation and a linear PDE. In Section 3, we present the discretization of the HJB equation. In Section 4, we highlight some important implementation details of the numerical method. Numerical results are presented and discussed in Section 5.

2 Mathematical formulation

Suppose there are two assets in the market: one is a risk free bond and the other is a risky equity index. The dynamics of the risk free bond B follows

$$dB(t) = rB(t)dt, \tag{2.1}$$

and an equity index S follows Heston's model (Heston, 1993) under the real probability measure

$$\frac{dS(t)}{S(t)} = (r + \xi V(t))dt + \sqrt{V(t)}dZ_1, \tag{2.2}$$

where the variance of the index, $V(t)$, follows a mean-reverting square-root process (Cox et al., 1985):

$$dV(t) = \kappa(\theta - V(t))dt + \sigma\sqrt{V(t)}dZ_2, \tag{2.3}$$

with dZ_1, dZ_2 being increments of Wiener processes. The instantaneous correlation between Z_1 and Z_2 is $dZ_1 dZ_2 = \rho dt$. The market price of volatility risk is $\xi V(t)$, which generates a risk premium proportional to $V(t)$. This assumption for the risk premium is based on Breedens's consumption-based model (Breeden, 1979), and originates from Heston (1993). Therefore, under this setup, the market is incomplete as trading in the risky asset and the bond cannot perfectly hedge the changes in the stochastic investment opportunity set.

An investor in this market is endowed at time zero with an initial wealth of w_0, and she can continuously and dynamically alter the proportion of wealth invested in each asset. In addition, let $W(t) = S(t) + B(t)$ denote the wealth at time t, let p denote the proportion of this wealth invested in the risky asset $S(t)$, consequently $(1 - p)$ then denotes the fraction of wealth invested in the risk free bond $B(t)$. The allocation strategy is a function of the current state, i.e., $p(\cdot) : (W(t), V(t), t) \to p = p(W(t), V(t), t)$. Note that in using the shorthand notations $p(\cdot)$ for the mapping, p for the value $p = p(W(t), V(t), t)$, and the dependence on the current state is implicit. From (2.1) and (2.2), we see that the investor's wealth process follows:

$$dW(t) = \left(r + p\xi V(t)\right) W(t)dt + p\sqrt{V}W(t)dZ_1. \tag{2.4}$$

2.1 Efficient frontiers and embedding methods

We assume here that the investor is guided by a pre-commitment mean variance objective based on the final wealth $W(T)$. The pre-commitment mean variance problem and its variations have been intensively studied in the literature (Li and Ng, 2000; Zhou and Li, 2000; Bielecki et al., 2005; Zhao and Ziemba, 2000; Nguyen and Portait, 2002). To best of our knowledge, there is no explicit closed-form solution for the pre-commitment mean variance problem when the risky asset follows a stochastic volatility process along with leverage constraints.

To simplify notations, we define $x = (w, v) = (W(t), V(t))$ for a state space. Let $E^{x,t}_{p(\cdot)}[W(T)]$ and $Var^{x,t}_{p(\cdot)}[W(T)]$ denote the expectation and variance of the terminal wealth conditional on the state (x, t) and the control $p(\cdot)$. Given a risk level $Var^{x,t}_{p(\cdot)}[W(T)]$, an investor desires her expected terminal wealth $E^{x,t}_{p(\cdot)}[W(T)]$ to be as large as possible. Equivalently, given an expected terminal wealth $E^{x,t}_{p(\cdot)}[W(T)]$, she wishes the risk $Var^{x,t}_{p(\cdot)}[W(T)]$ to be as small as possible. That is, she desires to find controls $p(\cdot)$ which generate Pareto optimal points. For notational simplicity, let $E^{x,t}_{p(\cdot)}[W(T)] = \mathcal{E}$ and $Var^{x,t}_{p(\cdot)}[W(T)] = \mathcal{V}$. The problem is rigorously formulated as follows.

Define the achievable mean variance objective set as

$$\mathcal{Y} = \{(\mathcal{V}, \mathcal{E}) : \ p \in \mathcal{Z}\}, \tag{2.5}$$

where $\mathcal{Z}$ is the set of admissible strategies, and denote the closure of $\mathcal{Y}$ by $\bar{\mathcal{Y}}$.

Definition 2.1. *A point $(\mathcal{V}, \mathcal{E}) \in \mathcal{Y}$ is Pareto mean variance optimal if there exists no admissible strategy $\bar{p} \in \mathcal{Z}$ such that*

$$\begin{aligned} Var^{x,t}_{\bar{p}}\{W(T)\} &\leq \mathcal{V}, \\ E^{x,t}_{\bar{p}}\{W(T)\} &\geq \mathcal{E}, \end{aligned} \tag{2.6}$$

where at least one of the inequalities in equation is strict. We denote by $\mathcal{P}$ the set of Pareto mean variance optimal points. Note that $\mathcal{P} \subseteq \bar{\mathcal{Y}}$.

Although the above definition is intuitive, determining the points in $\mathcal{P}$ requires solution of a multi-objective optimization problem, involving two conflicting criteria. A standard scalarization method can be used to combine the two criteria into an optimization problem with a single objective. In particular, for each point $(\mathcal{V}, \mathcal{E}) \in \bar{\mathcal{Y}}$, and for an arbitrary scaler $\lambda > 0$, we define the set of points $\mathcal{Y}_{P(\lambda)}$ to be

$$\mathcal{Y}_{P(\lambda)} = \left\{ (\mathcal{V}, \mathcal{E}) \in \bar{\mathcal{Y}} : \inf_{(\mathcal{V}_*, \mathcal{E}_*) \in \mathcal{Y}} (\lambda \mathcal{V}_* - \mathcal{E}_*) \right\}, \tag{2.7}$$

from which a point on the efficient frontier can be derived. The set of points on the efficient frontier are then defined as

$$\mathcal{Y}_P = \bigcup_{\lambda > 0} \mathcal{Y}_{P(\lambda)}. \tag{2.8}$$

Note that there is a difference between the set of all Pareto mean variance optimal points $\mathcal{P}$ (see Definition 2.1) and the efficient frontier $\mathcal{Y}_P$ (2.8) (Tse et al., 2014). In general,

$$\mathcal{P} \subseteq \mathcal{Y}_P,$$

but the converse may not hold if the achievable mean variance objective set $\mathcal{Y}$ (2.5) is not convex. In this paper, we restrict our attention to constructing $\mathcal{Y}_P$ (2.8).

Due to the presence of the variance term $Var^{x,t}_{p(\cdot)}[W(T)]$ in (2.7), a dynamic programming principle cannot be directly applied to solve this problem. To overcome this difficulty, we make use of the main result in (Li and Ng, 2000; Zhou and Li, 2000; Tse et al., 2014) which essentially involves the embedding technique. This result is summarized in the following Theorem.

Assumption 2.1. *We assume that $\mathcal{Y}$ is a non-empty subset of $\{(\mathcal{V},\mathcal{E}) \in \mathbb{R}^2 : \mathcal{V} > 0)\}$ and that there exists a positive scalarization parameter $\lambda_E > 0$ such that $\mathcal{Y}_{P(\lambda_E)} \neq \emptyset$.*

Theorem 2.1. *The embedded mean variance objective set $\mathcal{Y}_Q$ is defined by*

$$\mathcal{Y}_Q = \bigcup_{-\infty<\gamma<\infty} \mathcal{Y}_{Q(\gamma)}, \tag{2.9}$$

where

$$\mathcal{Y}_{Q(\gamma)} = \left\{ (\mathcal{V}_*,\mathcal{E}_*) \in \bar{\mathcal{Y}} : \mathcal{V}_* + \mathcal{E}_*^2 - \gamma\mathcal{E}_* = \inf_{(\mathcal{V},\mathcal{E})\in\mathcal{Y}} (\mathcal{V} + \mathcal{E}^2 - \gamma\mathcal{E}) \right\}. \tag{2.10}$$

If Assumption 2.1 holds and $\lambda > \lambda_E$, then $\mathcal{Y}_{P(\lambda)} \neq \emptyset$. Assume $(\mathcal{V}_0,\mathcal{E}_0) \in \mathcal{Y}_{P(\lambda)}$. Then if

$$\lambda\mathcal{V}_0 - \mathcal{E}_0 = \inf_{(\mathcal{V},\mathcal{E})\in\mathcal{Y}} (\lambda\mathcal{V} - \mathcal{E}), \tag{2.11}$$

then

$$\mathcal{V}_0 + \mathcal{E}_0^2 - \gamma\mathcal{E}_0 = \inf_{(\mathcal{V},\mathcal{E})\in\mathcal{Y}} (\mathcal{V} + \mathcal{E}^2 - \gamma\mathcal{E}), \quad i.e.\ (\mathcal{V}_0,\mathcal{E}_0) \in \mathcal{Y}_{Q(\gamma)}, \tag{2.12}$$

where $\gamma = \frac{1}{\lambda} + 2\mathcal{E}_0$. Consequently, $\mathcal{Y}_P \subseteq \mathcal{Y}_Q$.

Proof. See details in (Li and Ng, 2000; Zhou and Li, 2000; Dang et al., 2015). □

Theorem 2.1 states that the mean and variance $(\mathcal{V},\mathcal{E})$ of $W(T)$ are embedded in a scalarization optimization problem with the objective function being $\mathcal{V} + \mathcal{E}^2 - \gamma\mathcal{E}$. Noting that

$$\begin{aligned} \mathcal{V} + \mathcal{E}^2 - \gamma\mathcal{E} &= E^{x,t}_{p(\cdot)}[W^2(T)] - (E^{x,t}_{p(\cdot)}[W(T)])^2 + (E^{x,t}_{p(\cdot)}[W(T)])^2 - \gamma E^{x,t}_{p(\cdot)}[W(T)] \\ &= E^{x,t}_{p(\cdot)}[W^2(T) - \gamma W(T)] \\ &= E^{x,t}_{p(\cdot)}[(W(T) - \frac{\gamma}{2})^2] + \frac{\gamma^2}{4}, \end{aligned} \tag{2.13}$$

and that we can ignore the constant $\frac{\gamma^2}{4}$ term for the purposes of minimization, we then define the value function

$$\mathcal{U}(x,t) = \inf_{p(\cdot)\in\mathcal{Z}} E^{x,t}_{p(\cdot)}[(W(T) - \frac{\gamma}{2})^2]. \tag{2.14}$$

Theorem 2.1 implies that there exists a γ, such that, for a given positive λ, a control p^* which maximizes (2.7) also minimizes equation (2.14). Dynamic programming can then be directly applied to equation (2.14) to determine the optimal control $p^*(\cdot)$.

The procedure for determining the points on the efficient frontier is as follows. For a given value of γ, the optimal strategy p^* is determined by solving for the value function problem (2.14). Once this optimal policy $p^*(\cdot)$ is known, it is then straightforward to determine a point

$(Var^{x,t}_{p^*(\cdot)}[W(T)], E^{x,t}_{p^*(\cdot)}[W(T)])$ on the frontier. Varying γ traces out a curve in the $(\mathcal{V}, \mathcal{E})$ plane (see details in Section 4.2). Consequently, the numerical challenge is to solve for the value function (2.14). More precisely, the above procedure for constructing the efficient frontier generates points that are in the set $\mathcal{Y}_Q$. As pointed out in Tse et al. (2014), the set $\mathcal{Y}_Q$ may contain spurious points, i.e., points which are not in $\mathcal{Y}_P$. For example, when the original problem is nonconvex, spurious points can be generated. An algorithm for removing spurious points is discussed in Tse et al. (2014). The set of points in $\mathcal{Y}_Q$ with the spurious points removed generates all points in $\mathcal{Y}_P$. Reference (Dang et al., 2015) also discusses the convergence of finitely sampled γ to the efficient frontier.

2.2 The value function problem

Following standard arguments, the value function $\mathcal{U}(w, v, \tau),\ \tau = T - t$ (2.14) is the viscosity solution of the HJB equation

$$\mathcal{U}_\tau = \inf_{p \in \mathcal{Z}} \left\{ (r + p\xi v) w \mathcal{U}_w + \kappa(\theta - v)\mathcal{U}_v + \frac{1}{2}(p\sqrt{v}w)^2 \mathcal{U}_{ww} + p\rho\sigma\sqrt{v}w\mathcal{U}_{wv} + \frac{1}{2}\sigma^2 v \mathcal{U}_{vv} \right\}, \qquad (2.15)$$

on the domain $(w, v, \tau) \in [0, +\infty] \times [0, +\times] \times [0, T]$, and with the terminal condition

$$\mathcal{U}(w, v, 0) = \left(w - \frac{\gamma}{2} \right)^2. \qquad (2.16)$$

Remark 2.1. *In one of our numerical tests, we allow p to become unbounded, which may occur when $w \to 0$ (Wang and Forsyth, 2010). However, although $p \to \infty$ as $w \to 0$, we must have $(pw) \to 0$ as $w \to 0$, i.e., the amount invested in the risky asset converges to zero as $w \to 0$. This is required in order to ensure that the no-bankruptcy boundary condition is satisfied (Wang and Forsyth, 2010). As a result, we can then formally eliminate the problem with unbounded control by using $q = pw$ as the control, and assume q remains bounded. See details in (Wang and Forsyth, 2010).*

2.3 The expected wealth problem

2.3.1 The PDE formulation

Given the solution for the value function (2.14), with the optimal control $p^*(\cdot)$. We then need to determine the expected value $E^{x,t}_{p^*(\cdot)}[W(T)]$, denoted as

$$\mathcal{E}(w, v, t) = E^{x,t}_{p^*(\cdot)}[W(T)], \qquad (2.17)$$

Then, $\mathcal{E}(w, v, \tau),\ \tau = T - t$ is given from the solution to the following linear PDE

$$\mathcal{E}_\tau = (r + p^*\xi v) w \mathcal{E}_w + \kappa(\theta - v)\mathcal{E}_v + \frac{1}{2}(p^*\sqrt{v}w)^2 \mathcal{E}_{ww} + p^*\rho\sigma\sqrt{v}w\mathcal{E}_{wv} + \frac{1}{2}\sigma^2 v \mathcal{E}_{vv} \qquad (2.18)$$

with the initial condition $\mathcal{E}(w, v, 0) = w$, where p^* is obtained from the solution of the HJB equation (2.15).

2.3.2 The Hybrid (PDE - Monte Carlo) method

Alternatively, given the stored control $p^*(\cdot)$ determined from the solution of equation (2.15), we can directly estimate $(Var^{x,t}_{p^*(\cdot)}[W(T)], E^{x,t}_{p^*(\cdot)}[W(T)])$ by using a Monte Carlo method, based on solving the SDEs (2.4-2.3). The details of the SDE discretization are given in Section 4.2. This hybrid(PDE - Monte Carlo) method was originally proposed in (Tse et al., 2013).

2.4 Allowable portfolios

In order to obtain analytical solutions, many previous papers typically make assumptions which allow for the possibility of unbounded borrowing and bankruptcy. Moreover, these models assume a bankrupt investor can still keep on trading. The ability to continue trading even though the value of an investor's wealth is negative is highly unrealistic. In this paper, we enforce the condition that the wealth value remains in the solvency regions by applying certain boundary conditions to the HJB equation (Wang and Forsyth, 2008). Thus, bankruptcy is prohibited, i.e.,

$$w \in [0, +\infty).$$

We will also assume that there is a leverage constraint, i.e., the investor must select an asset allocation satisfying

$$p = \frac{\text{The risky asset value}}{\text{The total wealth}} = \frac{pW(t)}{W(t)} < p_{\max},$$

which can be interpreted as the maximum leverage condition, and $p_{\max}$ is a known positive constant with typical value in $[1.0, 2.0]$. Thus, the control set

$$p \in \mathcal{Z} = [0, p_{\max}].$$

Note that when the risk premium ξ (2.2) is positive, it is not optimal to short the risky asset, since we have only a single risky asset in our portfolio.

3 Numerical Discretization of the HJB equation

3.1 Localization

We will assume that the discretization is posed on a bounded domain for computational purposes. The discretization is applied to the localized finite region $(w, v) \in [0, w_{\max}] \times [0, v_{\max}]$. Asymptotic boundary conditions will be imposed at $w = w_{\max}$ and $v = v_{\max}$ which are compatible with a monotone numerical scheme.

3.1.1 The localization of V

The proper boundary on $v = 0$ needs to be specified to be compatible with the corresponding SDE (2.3), which has a unique solution (Feller, 1951). If $2\kappa\theta \geq \sigma^2$, the so-called Feller condition holds, and $v = 0$ is unattainable. If the Feller condition is violated, $2\kappa\theta < \sigma^2$, then $v = 0$ is an attainable boundary but is strongly reflecting (Feller, 1951). The appropriate boundary condition can be obtained by setting $v = 0$ into the equation (Ekström and Tysk, 2010). That is,

$$\mathcal{U}_\tau = rw\mathcal{U}_w + \kappa\theta\mathcal{U}_v, \tag{3.1}$$

and the equation degenerates to a linear PDE. On the lower boundary $v = 0$, the variance and the risk premium vanishes, according to (2.4), so that the wealth return is always the risk free rate r. The control value p vanishes in the degenerate equation (3.1), and we can simply define $p^*(w, v = 0, t) \equiv 0$ which we need in the estimation of $(Var^{x,t}_{p^*(\cdot)}[W(T)], E^{x,t}_{p^*(\cdot)}[W(T)])$ using the Monte Carlo simulation. In this case, since the *risky* asset is riskless, the distinction between risky and risk free asset is meaningless.

The validity of this boundary condition is intuitively justified by the fact that the solution to the SDE for v is unique, such that the behavior of v at the boundary $v = 0$ is determined by the SDE itself, and hence the boundary condition is determined by setting $v = 0$ in equation (2.15). A formal proof that this boundary condition is correct is given in (Ekström and Tysk, 2010). If the boundary at $v = 0$ is attainable, then this boundary behaviour serves as a boundary condition and guarantees uniqueness in the appropriate function spaces. On the other hand, if the boundary is non-attainable, then the boundary behaviour is not needed to guarantee uniqueness, but is nevertheless very useful in a numerical scheme.

On the upper boundary $v = v_{\max}$, $\mathcal{U}_v$ is set to zero. Thus, the boundary condition on $v_{\max}$ is set to

$$\mathcal{U}_\tau = \inf_{p \in \mathcal{Z}} \left\{ (r + p\xi v) w \mathcal{U}_w + \frac{1}{2}(p\sqrt{v}w)^2 \mathcal{U}_{ww} \right\}. \tag{3.2}$$

The optimal control p^* at $v = v_{\max}$ is determined by solving the equation (3.2). This boundary condition can be justified by noting that as $v \to \infty$, then the diffusion term in the w direction in equation (2.15) becomes large. As well, the initial condition (2.16) is independent of v. As a result, we expect that

$$\mathcal{U} \approx C'w + C'', \quad v \to \infty,$$

where C' and C'' are constants, and hence $\mathcal{U}_v \approx 0$ at $v = v_{\max}$.

3.1.2 The localization for W

We prohibit the possibility of bankruptcy ($W(t) < 0$) by requiring that $\lim_{w\to 0}(pw) = 0$ (Wang and Forsyth, 2010), so, on $w = 0$, the equation (2.15) reduces to

$$\mathcal{U}_\tau = \kappa(\theta - v)\mathcal{U}_v + \sigma^2 v \mathcal{U}_{vv}. \tag{3.3}$$

When $w \to +\infty$, we assume that asymptotic form of the exact solution is

$$\mathcal{U}(w \to +\infty, v, \tau) = \bar{\mathcal{U}}(w) = H_2(\tau)w^2 + H_1(\tau)w + H_0(\tau), \tag{3.4}$$

and make the assumption that $p^*(w_{\max}, v, 0)$ at $w = w_{\max}$ is set to zero. That is, once the investor's wealth is very large, she prefers the risk free asset. This can be justified from the arguments in (Cui et al., 2012; Dang and Forsyth, 2014a).

3.1.3 Alternative localization for w

$\mathcal{U}(w, v, \tau)$ is the viscosity solution of the HJB equation (2.15). Recall that the initial condition for problem (2.14) is

$$\mathcal{U}(w, v, 0) = \left(W(T) - \frac{\gamma}{2} \right)^2.$$

For a fixed gamma, we define the optimal embedded terminal wealth at time t, denoted by $W_{opt}(t)$, as

$$W_{opt}(t) = \frac{\gamma}{2} e^{-r(T-t)}. \tag{3.5}$$

It is easy to verify that $W_{opt}(t)$ is a globally minimum state of the value function $\mathcal{U}(w, v, t)$. Consider the state $(W_{opt}(t), v), t \in [0, T]$, and the optimal strategy $p^*(\cdot)$ such that $p^*(w, v, \mathcal{T}) \equiv 0,\ \mathcal{T} > t$. Under $p^*(\cdot)$, the wealth is all invested in the risk free bond without further re-balancing from time t. As a result, the wealth will accumulate to $W(T) = \frac{\gamma}{2}$ with certainty, i.e., the optimal embedded terminal wealth $\frac{\gamma}{2}$ is achievable. By definition (2.14), we have,

$$\mathcal{U}(W_{opt}(t), v, t) = \inf_{p(\cdot)\in\mathcal{Z}} \left\{ E^{x,t}_{p(\cdot)}[(W(T) - \frac{\gamma}{2})^2] \right\} = E^{x,t}_{p^*(\cdot)}[(W(T) - \frac{\gamma}{2})^2] = 0. \tag{3.6}$$

Since the value function is the expectation of a non-negative quantity, it can never be less than zero. Then, the exact solution for the value function problem at the special point $W_{opt}(t)$ must be zero. This result holds for both the discrete and continuous re-balancing case. For the formal proof, we refer the reader to (Dang and Forsyth, 2014a).

Consequently, the point $w = \frac{\gamma}{2} e^{-r\tau}$ is a Dirichlet boundary $\mathcal{U}(\frac{\gamma}{2} e^{-r\tau}, v, \tau) = 0$, and information for $w > \frac{\gamma}{2} e^{-r\tau}$ is not needed. We can then restrict the size the computational domain to be $0 \le w \le \frac{\gamma}{2}$. Note that the optimal control will ensure that $\mathcal{U}(\frac{\gamma}{2} e^{-r\tau}, v, \tau) = 0$ without any need to enforce this boundary condition. This will occur since we assume continuous rebalancing. This effect that $W(t) \le W_{opt}(t)$ is also discussed in Vigna (2014). It is interesting to note that in the case of discrete rebalancing that it is optimal to withdraw cash from the portfolio if it is ever observed that $W(t) > W_{opt}(t)$. This is discussed in Cui et al. (2012); Dang and Forsyth (2014a).

We have verified, experimentally, that restricting the computational domain to $w \in [0, \gamma/2]$ gives the same results as the domain $w \in [0, w_{\max}],\ w_{\max} \gg \frac{\gamma}{2}$, with asymptotic boundary condition (3.4).

Remark 3.1 (Significance of $W(t) \le W_{opt}(t)$)**.** *If we assume that initially $W(0) < W_{opt}(0)$ (otherwise the problem is trivial if we allow cash withdrawals), then the optimal control will ensure that $W(t) \le W_{opt}(t), \forall t$. Hence continuous time mean variance optimization is* time consistent in efficiency *(Cui et al., 2012). Another interpretation is that continuous time mean variance optimization is equivalent to minimizing the quadratic loss with respect to the wealth target $W_{opt}(T)$ (Vigna, 2014).*

Remark 3.2 (Significance of $W(T) \le \gamma/2$)**.** *From Remark 3.1 we have trivially that $W(T) \le \gamma/2$, hence from equation (2.14), the investor is never penalized for large gains, i.e. the quadratic utility function (2.14) is always well behaved. Consequently, continuous time mean variance optimization is fundamentally different from the single period counterpart.*

3.2 Discretization

In the following section, we discretize equation (2.15) over a finite grid $N = N_1 \times N_2$ in the space (w, v). Define a set of nodes $\{w_1, w_2, \ldots, w_{N_1}\}$ in w direction and $\{v_1, v_2, \ldots, v_{N_2}\}$ in the v direction. Denote the n^{th} time step by $\tau^n = n\Delta\tau,\ n = 0, \ldots, N_\tau$, with $N_\tau = \frac{T}{\Delta\tau}$. Let $\mathcal{U}^n_{i,j}$ be the approximate solution of the equation (2.15) at (w_i, v_j, τ^n).

It will be convenient to define

$$\begin{aligned} \Delta w_{\max} &= \max_i (w_{i+1} - w_i), \quad \Delta w_{\min} = \min_i (w_{i+1} - w_i), \\ \Delta v_{\max} &= \max_i (v_{i+1} - v_i), \quad \Delta v_{\min} = \min_i (v_{i+1} - v_i). \end{aligned} \tag{3.7}$$

We assume that there is a mesh discretization parameter h such that

$$\Delta w_{\max} = C_1 h,\ \Delta w_{\min} = C_2 h,\ \Delta v_{\max} = C_1^{'} h,\ \Delta v_{\min} = C_2^{'} h,\ \Delta \tau = C_3 h, \tag{3.8}$$

where $C_1, C_2, C_1^{'}, C_2^{'}, C_3$ are constants independent of h.

In the following sections, we will give the details of the discretization for a reference node $(w_i, v_j),\ 1 < i < N_1,\ 1 < j < N_2$.

3.2.1 The wide stencil

We need a monotone discretization scheme in order to guarantee convergence to the desired viscosity solution (Barles and Souganidis, 1991). We remind the reader that seemingly reasonable non-monotone discretizations can converge to the incorrect solution (Pooley et al., 2003). Due to the cross derivative term in (2.15), however, a classic finite difference method can not produce such a monotone scheme. Since the control appears in the cross derivative term, it will not be possible (in general) to determine a grid spacing or global coordinate transformation which eliminates this term. We will adopt the wide stencil method developed in Ma and Forsyth (2014) to discretize the second derivative terms. Suppose we discretize equation (2.15) at grid node (i, j) for a fixed control. For a fixed p, consider a virtual rotation of the local coordinate system clockwise by the angle $\eta_{i,j}$

$$\eta_{i,j} = \frac{1}{2}\tan^{-1}\left(\frac{2\rho p \sigma w_i v_j}{(p\sqrt{v_j}w_i)^2 - (\sigma\sqrt{v_j})^2}\right). \tag{3.9}$$

That is, (y_1, y_2) in the transformed coordinate system is obtained by using the following matrix multiplication

$$\begin{pmatrix} w \\ v \end{pmatrix} = \begin{pmatrix} \cos\eta_{i,j} & -\sin\eta_{i,j} \\ \sin\eta_{i,j} & \cos\eta_{i,j} \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}. \tag{3.10}$$

We denote the rotation matrix in (3.10) as $\mathbf{R}_{i,j}$. This rotation operation will result in a zero correlation in the diffusion tensor of the rotated system. Under this grid rotation, the second order terms in equation (2.18) are, in the transformed coordinate system (y_1, y_2),

$$a_{i,j}\frac{\partial^2 \mathcal{W}}{\partial y_1^2} + b_{i,j}\frac{\partial^2 \mathcal{W}}{\partial y_2^2}, \tag{3.11}$$

where $\mathcal{W}$ is the value function $\mathcal{W}(y_1, y_2, \tau)$ in the transformed coordinate system, and

$$\begin{aligned} a_{i,j} &= \left(\frac{1}{2}(p\sqrt{v_j}w_i)^2\cos(\eta_{i,j})^2 + \rho p\sigma w_i v_j \sin(\eta_{i,j})\cos(\eta_{i,j}) + \frac{1}{2}(\sigma\sqrt{v_j})^2\sin(\eta_{i,j})^2\right), \\ b_{i,j} &= \left(\frac{1}{2}(p\sqrt{v_j}w_i)^2\sin(\eta_{i,j})^2 - \rho p\sigma w_i v_j \sin(\eta_{i,j})\cos(\eta_{i,j}) + \frac{1}{2}(\sigma\sqrt{v_j})^2\cos(\eta_{i,j})^2\right). \end{aligned} \tag{3.12}$$

The diffusion tensor in (3.11) is diagonally dominant with no off-diagonal terms, and consequently a standard finite difference discretization for the second partial derivatives results in a monotone scheme. The rotation angle $\eta_{i,j}$ depends on the grid node and the control, therefore it is impossible to rotate the global coordinate system by a constant angle and build a grid over the entire space (y_1, y_2). The local coordinate system rotation is only used to construct a virtual grid

which overlays the original mesh. We have to approximate the values of $\mathcal{W}$ on our virtual local grid using an interpolant $\mathcal{J}_h\mathcal{U}$ on the original mesh. To keep the numerical scheme monotone, $\mathcal{J}_h$ must be a linear interpolation operator. Moreover, to keep the numerical scheme consistent, we need to use the points on our virtual grid whose Euclidean distances are $O(\sqrt{h})$ from the central node, where h is the mesh discretization parameter (3.8). This results in a wide stencil method since the relative stencil length increases as the grid is refined ($\frac{\sqrt{h}}{h} \to +\infty$ as $h \to 0$). For more details, we refer the reader to Ma and Forsyth (2014).

Let us rewrite the HJB equation (2.15) as

$$\sup_{p \in \mathcal{Z}} \{\mathcal{U}_\tau - (r + p\xi v) w \mathcal{U}_w - \mathcal{L}^p \mathcal{U}\} = 0, \tag{3.13}$$

where the linear operator $\mathcal{L}^p$ is defined as

$$\mathcal{L}^p \mathcal{U} = \kappa(\theta - v)\mathcal{U}_v + \frac{1}{2}(p\sqrt{v}w)^2 \mathcal{U}_{ww} + p\rho\sigma\sqrt{v}w\mathcal{U}_{wv} + \frac{1}{2}\sigma^2 v \mathcal{U}_{vv}. \tag{3.14}$$

The drift term $\kappa(\theta - v)\mathcal{U}_v$ in equation (3.14) is discretized by a standard backward or forward finite differencing discretization, depending on the sign of $\kappa(\theta - v)$. Overall, the discretized form of the linear operator $\mathcal{L}^p$ is then denoted by L_h^p

$$\begin{aligned} L_h^p \mathcal{U}_{i,j}^{n+1} = {} & 1_{\kappa(\theta - v_j) \geq 0} \frac{\kappa(\theta - v_j)}{h} \mathcal{U}_{i,j+1}^{n+1} - 1_{\kappa(\theta - v_j) < 0} \frac{\kappa(\theta - v_j)}{h} \mathcal{U}_{i,j-1}^{n+1} \\ & + \frac{a_{i,j}}{h} \mathcal{J}_h \mathcal{U}^{n+1} \left(x_{i,j} + \sqrt{h}(\mathbf{R}_{i,j})_1 \right) + \frac{a_{i,j}}{h} \mathcal{J}_h \mathcal{U}^{n+1} \left(x_{i,j} - \sqrt{h}(\mathbf{R}_{i,j})_1 \right) \\ & + \frac{b_{i,j}}{h} \mathcal{J}_h \mathcal{U}^{n+1} \left(x_{i,j} + \sqrt{h}(\mathbf{R}_{i,j})_2 \right) + \frac{b_{i,j}}{h} \mathcal{J}_h \mathcal{U}^{n+1} \left(x_{i,j} - \sqrt{h}(\mathbf{R}_{i,j})_2 \right) \\ & - \left(1_{\kappa(\theta - v_j) \geq 0} \frac{\kappa(\theta - v_j)}{h} - 1_{\kappa(\theta - v_j) < 0} \frac{\kappa(\theta - v_j)}{h} + \frac{2a_{i,j}}{h} + \frac{2b_{i,j}}{h} \right) \mathcal{U}_{i,j}^{n+1}, \end{aligned} \tag{3.15}$$

where h is the discretization parameter, and the superscript p in L_h^p indicates that the discretization depends on the control p. $x_{i,j} = \begin{pmatrix} w_i \\ v_j \end{pmatrix}$, $a_{i,j}$ and $b_{i,j}$ are given in (3.12), and the presence of $\mathcal{J}_h \mathcal{U}^{n+1} \left(x_{i,j} \pm \sqrt{h}(\mathbf{R}_{i,j})_k \right)$, $k = 1, 2$ is due to the discretization of the second derivative terms. $(\mathbf{R}_{i,j})_k$ is k-th column of the rotation matrix.

3.2.2 Semi-Lagrangian timestepping scheme

When $p \to 0$, equation (2.15) degenerates, with no diffusion in the w direction. As a result, we will discretize the drift term $(r + p\xi v) w \mathcal{U}_w$ in equation (2.15) by a semi-Lagrangian timestepping scheme in this section. Initially introduced by Douglas and Russell (1982); Pironneau (1982) for atmospheric and weather numerical prediction problems, semi-Lagrangian schemes can effectively reduce the numerical problems arising from convection dominated equations.

Firstly, we define the Lagrangian derivative $\frac{D\mathcal{U}}{D\tau}(p)$ by

$$\frac{D\mathcal{U}}{D\tau}(p) = \mathcal{U}_\tau - (r + p\xi v) w \mathcal{U}_w, \tag{3.16}$$

which is the rate of change of $\mathcal{U}$ along the characteristic $w = w(\tau)$ defined by the risky asset fraction p through

$$\frac{dw}{d\tau} = -(r + p\xi v)w. \tag{3.17}$$

We can then rewrite equation (3.13) as

$$\sup_{p\in\mathcal{Z}} \left\{ \frac{D\mathcal{U}}{D\tau} - \mathcal{L}^p\mathcal{U} \right\} = 0. \tag{3.18}$$

Solving equation (3.17) backwards in time from τ^{n+1} and τ^n, for a fixed $\mathcal{U}_{i,j}^{n+1}$ gives the point at the foot of the characteristic

$$(w_{i^*}, v_j) = (w_i e^{(r+p\xi v_j)\Delta\tau^n}, v_j), \tag{3.19}$$

which in general is not on the PDE grid. We use the notation $\mathcal{U}_{i^*,j}^n$ to denote an approximation of the value $\mathcal{U}(w_{i^*}, v_j, \tau^n)$, which is obtained by linear interpolation to preserve monotonicity. The Lagrangian derivative at a reference node (i, j) is then approximated by

$$\frac{D\mathcal{U}}{D\tau}(p) \approx \frac{\mathcal{U}_{i,j}^n - \mathcal{U}_{i^*,j}^n(p)}{\Delta\tau^n}, \tag{3.20}$$

where $\mathcal{U}_{i^*,j}^n(p)$ denotes that w_{i^*} depends on the control p through equation (3.19). For the details of the semi-Lagrangian timestepping scheme, we refer the reader to (Chen and Forsyth, 2007).

Finally, by using the implicit timestepping method, combining the expressions (3.15) and (3.20), the HJB equation (3.18) at a reference point (w_i, v_j, τ^{n+1}) is then discretized as

$$\sup_{p\in\mathcal{Z}_h} \left\{ \frac{\mathcal{U}_{i,j}^{n+1}}{\Delta\tau^n} - \frac{\mathcal{U}_{i^*,j}^n(p)}{\Delta\tau^n} - L_h^p\mathcal{U}_{i,j}^n \right\} = 0, \tag{3.21}$$

where $\mathcal{Z}_h$ is the discrete control set. Since there is no simple analytic expression which can be used to minimize the discrete equations (3.21), and we need to discretize the admissible control set $\mathcal{Z}$ and perform linear search. This guarantees that we find the global maximum of equation (3.21), since the objective function has no known convexity properties. If the discretization step for the controls is also $O(h)$, where h is the discretization parameter, then this is a consistent approximation (Wang and Forsyth, 2008).

3.3 Matrix form of the discrete equation

Our discretization is summarized as follows. The domains are defined in Table 3.1. For the case $(w_i, v_j) \in \Omega_{in}$, we need to use a wide stencil based on a local coordinate rotation to discretize the second derivative terms, and use the semi-Lagrangian timestepping scheme to handle the drift term $(r + p\xi v)w\mathcal{U}_w$. The HJB equation is discretized as (3.21), and the optimal p^* in this case is determined by solving (3.21). For the case $\Omega_{v_{\max}}$, the HJB equation degenerates to (3.2). In this case, the drift term is also handled by the semi-Lagrangian timestepping scheme. With vanishing cross-derivative term, the degenerate linear operator $\mathcal{L}^p$ can be discretized by a standard finite difference method. The corresponding discretized form D_h^p is given in Appendix A. The value for case $\Omega_{w_{\max}}$ is obtained by the asymptotic solution (3.4), and the optimal p^* is set to zero. At the lower boundaries $\Omega_{w_{\min}}$ and $\Omega_{v_{\min}}$, the HJB equation degenerates to a linear equation. The

wide stencil and the semi-Lagrangian timestepping scheme may require the value of the solution at a point outside the computational domain, denoted as Ω_{out}. Details on how to handle this case are given in Section 4.3. From the discretization (3.21), we can see that the measure of Ω_{out} convergences to zero as $h \to 0$. Lastly, fully implicit time-stepping is used to ensure unconditional monotonicity of our numerical scheme. Fully implicit timestepping requires solution of highly nonlinear algebraic equations at each timestep. For the applications addressed in (Forsyth and Labahn, 2007) an efficient method for solving the associated nonlinear algebraic systems makes use of a policy iteration scheme. We refer the reader to (Huang et al., 2012; Forsyth and Labahn, 2007) for the details of the policy iteration algorithm.

Notation	The domain
Ω	$[0, w_{\max}] \times [0, v_{\max}]$
Ω_{in}	$(0, w_{\max}) \times (0, v_{\max})$
$\Omega_{w_{\max}}$	The upper boundary $w = w_{\max}$
$\Omega_{v_{\max}}$	The upper boundary $v = v_{\max}$
$\Omega_{w_{\min}}$	The lower boundary $w = 0$
$\Omega_{v_{\min}}$	The lower boundary $v = 0$
Ω_{out}	$(w_{\max}, +\infty) \times (0, +\infty) \cup (0, +\infty) \times (v_{\max}, +\infty)$

Table 3.1: The domain definitions.

It is convenient to use a matrix form to represent the discretized equations for computational purposes. Let $\mathcal{U}^n_{i,j}$ be the approximate solution of the equation (2.15) at (w_i, v_j, τ^n), $1 \le i \le N_1$, $1 \le j \le N_2$ and $0 \le \tau^n \le N_\tau$, and form the solution vector

$$\mathbf{U}^n = \left(\mathcal{U}^n_{1,1}, \mathcal{U}^n_{2,1}, \ldots, \mathcal{U}^n_{N_1,1}, \ldots, \mathcal{U}^n_{1,N_2}, \ldots, \mathcal{U}^n_{N_1,N_2}\right). \tag{3.22}$$

It will sometimes be convenient to use a single index when referring to an entry of the solution vector

$$\mathcal{U}^n_\ell = \mathcal{U}^n_{i,j}, \quad \ell = i + (j-1)N_1.$$

Let $N = N_1 \times N_2$, and we define the $N \times N$ matrix $\mathbf{L}^{n+1}(\mathcal{P})$, where

$$\mathcal{P} = \{p_1, \ldots, p_N\} \tag{3.23}$$

is an indexed set of N controls, and each p_ℓ is in the set of admissible controls. $\mathbf{L}^{n+1}_{\ell,k}(\mathcal{P})$ is the entry on the ℓ-th row and k-th column of the discretized matrix $\mathbf{L}^{n+1}(\mathcal{P})$. We also define a vector of boundary conditions $\mathbf{F}^{n+1}(\mathcal{P})$.

For the case $(w_i, v_j) \in \Omega_{w_{\max}}$ where the Dirichlet boundary condition (3.4) is imposed, we then have

$$\mathbf{F}^{n+1}_\ell(\mathcal{P}) = \bar{\mathcal{U}}(w_{\max}), \tag{3.24}$$

and

$$\mathbf{L}^{n+1}_{\ell,k}(\mathcal{P}) = 0, \quad k = 1, \ldots, N. \tag{3.25}$$

For the case $(w_i, v_j) \in \Omega_{v_{\min}} \cup \Omega_{w_{\min}} \cup \Omega_{v_{\max}}$, the differential operator degenerates, and the entries $\mathbf{L}^{n+1}_{\ell,k}(\mathcal{P})$ are constructed from the discrete linear operator D^p_h (see the Appendix, equation (A.1)). That is,

$$[\mathbf{L}^{n+1}(\mathcal{P})\mathbf{U}^{n+1}]_\ell = D^p_h \mathcal{U}^{n+1}_{i,j}. \tag{3.26}$$

For the case $(w_i, v_j) \in \Omega_{in}$, we need to use the values at the following four off-grid points $x_{i,j} \pm \sqrt{h}(\mathbf{R}_{i,j})_k,\ k = 1, 2$ in (3.15), and we denote those values by $\Psi^m_{i,j},\ m = 1, 2, 3, 4$, respectively. When $\Psi^m_{i,j} \in \Omega$, using linear interpolation, values at these four points are approximated as follows

$$\mathcal{J}_h\mathcal{U}^{n+1}(\Psi^m_{i,j}) = \begin{cases} \sum_{\substack{d=0,1\\e=0,1}} \omega^{f_m+d,g_m+e}_{i,j}\mathcal{U}^{n+1}_{f_m+d,g_m+e}, & \Psi^m_{i,j} \in \Omega \\ 0, & \text{otherwise} \end{cases}. \tag{3.27}$$

For linear interpolation, we have that $\omega^{f_m+d,g_m+e}_{i,j} \geq 0$ and $\sum_{\substack{d=0,1\\e=0,1}} \omega^{f_m+d,g_m+e}_{i,j} = 1$. Then, inserting (3.27) in (3.15), the entries $\mathbf{L}^{n+1}_{\ell,k}(\mathcal{P})$ on ℓ-th row are specified. When we use $\Psi^m_{i,j} \in \Omega_{out}$, we directly use its asymptotic solution $\bar{\mathcal{U}}(\Psi^m_{i,j})$ (3.4). Thus, we need to define the vector $\mathbf{G}^{n+1}(\mathcal{P})$ to facilitate the construction of the matrix form in this situation when we use a point in the domain Ω_{out}.

$$\mathbf{G}^{n+1}_\ell(\mathcal{P}) = \begin{cases} 1_{\Psi^1_{i,j}\in\Omega_{out}}\frac{a_{i,j}}{h}\bar{\mathcal{U}}(\Psi^1_{i,j}) + 1_{P^2_{i,j}\in\Omega_{out}}\frac{a_{i,j}}{h}\bar{\mathcal{U}}(\Psi^2_{i,j}) \\ \qquad + 1_{\Psi^3_{i,j}\in\Omega_{out}}\frac{b_{i,j}}{h}\bar{\mathcal{U}}(\Psi^3_{i,j}) + 1_{\Psi^4_{i,j}\in\Omega_{out}}\frac{b_{i,j}}{h}\bar{\mathcal{U}}(\Psi^4_{i,j}), & (w_i, v_j) \in \Omega_{in}, \\ 0, & \text{otherwise} \end{cases} \tag{3.28}$$

where $a_{i,j}$ and $b_{i,j}$ are defined in equation (3.12). As a result, for the case $(w_i, v_j) \in \Omega_{in}$,

$$[\mathbf{L}^{n+1}(\mathcal{P})\mathbf{U}^{n+1}]_\ell + \mathbf{G}^{n+1}_\ell(\mathcal{P}) = L^p_h\mathcal{U}^{n+1}_{i,j}, \tag{3.29}$$

where L^p_h is defined in equation (3.15).

Let $\Phi^{n+1}(\mathcal{P})$ be a linear Lagrange interpolation operator such that

$$[\Phi^{n+1}(\mathcal{P})\mathbf{U}]_l = \begin{cases} \mathcal{J}_h\mathcal{U}^n_{i^*,j}, & (w_{i^*}, v_j) \in \Omega \\ \bar{\mathcal{U}}(w_{i^*})\ (3.4), & (w_{i^*}, v_j) \in \Omega_{out} \end{cases}, \tag{3.30}$$

where $\mathcal{U}^n_{i^*,j}$ is defined in (3.19).

The final matrix form of the discretized equations is then

$$\begin{aligned} \left[\mathbf{I} - \Delta\tau^n\mathbf{L}^{n+1}(\hat{\mathcal{P}})\right]\mathbf{U}^{n+1} &= \Phi^{n+1}(\mathcal{P})\mathbf{U}^n + \Delta\tau^n\mathbf{G}^{n+1}(\mathcal{P}) + \mathbf{F}^{n+1} - \mathbf{F}^n, \\ \hat{p}_\ell \in \arg\min_{p\in\mathcal{Z}_h} &\left[\Phi^{n+1}(\mathcal{P})\mathbf{U}^n + \Delta\tau^n\left(\mathbf{L}^{n+1}(\mathcal{P})\mathbf{U}^{n+1} + \mathbf{G}^{n+1}(\mathcal{P})\right)\right]_\ell, \\ \ell = i + (j-1)N_1,&\quad i = 2, \ldots, N_1 - 1,\quad j = 2, \ldots, N_2, \end{aligned} \tag{3.31}$$

where $\mathcal{Z}_h$ is the discretized control set $\mathcal{Z}$.

Remark 3.3. *Note that* $\left[\mathbf{I} - \Delta\tau^n\mathbf{L}^{n+1}(\mathcal{P})\right]_{\ell,k}$, $\left[\Phi^{n+1}(\mathcal{P})\right]_\ell$ *and* $\left[\mathbf{G}^{n+1}(\mathcal{P})\right]_\ell$ *depend only on* p_ℓ.

3.4 Convergence to the viscosity solution

Assumption 3.1. *If the control p is bounded, Equation (2.15) satisfies the strong comparison property, hence a unique continuous viscosity solution to equation (2.15) exists (Debrabant and Jakobsen, 2013).*

Provided that the original HJB satisfies Assumption 3.1, we can show that the numerical scheme (3.31) is ℓ_∞ stable, consistent and monotone, and then the scheme converges to the unique and continuous viscosity solution (Barles and Souganidis, 1991). We give a brief overview of the proof as follows.

- **Stability:** From the formation of matrix $\mathbf{L}$ in (3.25), (3.26) and (3.29), it is easily seen that $[\mathbf{I} - \Delta\tau\mathbf{L}^{n+1}(\mathcal{P})]$ (3.31) has positive diagonals, non-positive offdiagonals, and the ℓ-th row sum for the matrix is

$$\sum_k \left[\mathbf{I} - \Delta\tau\mathbf{L}^{n+1}(\mathcal{P})\right]_{\ell,k} > 0, \quad i = 1, \ldots, N_1,\ j = 1, \ldots, N_2, \tag{3.32}$$

  where $\ell = i + (j-1)N_1$, hence the matrix $[\mathbf{I} - \Delta\tau\mathbf{L}^{n+1}(\mathcal{P})]$ is diagonally dominant, and thus it is an M-matrix (Varga, 2009). We can then easily show that the numerical scheme is l_∞ stable by a straightforward maximum analysis as in (d'Halluin et al., 2004).

- **Monotonicity:** To guarantee monotonicity, we use a wide stencil to discretize the second derivative terms in the discrete linear operator L_h^p (3.15) (see proof in (Ma and Forsyth, 2014)). Note that using linear interpolation to compute $\mathcal{U}^n_{i^*,j}$ (3.20) in the semi-Lagrangian timestepping scheme also ensures monotonicity.

- **Consistency:** A simple Taylor series verifies consistency. As noted in Section 4.3, we may shrink the wide stencil length to avoid using points below the lower boundaries. We can use the same proof in Ma and Forsyth (2014) to show this treatment retains local consistency. Since we have either simple Dirichlet boundary conditions, or the PDE at the boundary is the limit from the interior, the we need only use the classical definition of consistency here (See proof in Ma and Forsyth (2014)). The only case where the point $\mathcal{U}^n_{i^*,j}$ (3.20) in the semi-Lagrangian timestepping scheme is outside computational domain is through the upper boundary $w = w_{max}$, where the asymptotic solution (3.4) is used. Thus, unlike the semi-Lagrangian timestepping scheme in Chen and Forsyth (2007), we do not need the more general definition of consistency (Barles and Souganidis, 1991) to handle the boundary data.

3.5 Policy iteration

Our numerical scheme requires the solution of highly nonlinear algebraic equations (3.31) at each timestep. We use the policy iteration algorithm (Forsyth and Labahn, 2007) to solve the associated algebraic systems. For the details of the algorithm we refer the reader to Forsyth and Labahn (2007); Huang et al. (2012). Regarding the convergence of the policy iteration, since the matrix $[\mathbf{I} - \Delta\tau\mathbf{L}^{n+1}(\mathcal{P})]$ (3.31) is an M-matrix and the control set $\mathcal{Z}_h$ is a finite set, it is easy to show that the policy iteration is guaranteed to converge (Forsyth and Labahn, 2007).

4 Implementation Details

4.1 Complexity

Examination of the algorithm for solving discrete equations (3.31) reveals that each timestep requires

- In order to solve the local optimization problems at each node, we perform a linear search to find the minimum for $p \in \mathcal{Z}_h$. Thus, with total $O(1/h^2)$ nodes, this gives a complexity $O(1/h^3)$ for solving the local optimization problems at each time step.

- We use a preconditioned Bi-CGSTAB iterative method for solving the sparse matrix at each policy iteration. The time complexity of solving the sparse M-matrix is $O((1/h^2)^{\frac{5}{4}})$ (Saad, 2003). Note that in general, we need to reconstruct the data structure of the sparse matrix for each iteration.

Assuming that the number of policy iterations is bounded, as the mesh size tends to zero, which is in fact observed in our experiments, the complexity of the time advance is thus dominated by the solution of the local optimization problems. Finally, the total complexity is $O(1/h^4)$.

4.2 The efficient frontier

In order to trace out the efficient frontier solution of problem (2.7), we proceed in the following way. Pick an arbitrary value of γ and solve problem (2.14), which determines the optimal control $p^*(\cdot)$. There are then two methods to determine the quantities of interest $(Var_{p^*}^{x_0,0}[W(T)], E_{p^*}^{x_0,0}[W(T)])$, namely the PDE method and the Hybrid (PDE - Monte Carlo) method. We will compare the performance of these methods in the numerical experiments.

4.2.1 The PDE Method

For a fixed γ, given $\mathcal{U}(w_0, v_0, 0)$ and $\mathcal{E}(w_0, v_0, 0)$ obtained solving the corresponding equations (2.15) and (2.18) at the initial time with $W_0 = w_0$ and $V_0 = v_0$, we can then compute the corresponding pair $(Var_{p^*(\cdot)}^{x_0,0}[W(T)], E_{p^*(\cdot)}^{x_0,0}[W(T)])$, where $x_0 = (w_0, v_0)$. That is,

$$\begin{aligned} E_{p^*(\cdot)}^{x_0,0}[W(T)] &= \mathcal{E}(w_0, x_0, 0), \\ Var_{p^*(\cdot)}^{x_0,0}[W(T)] &= \mathcal{U}(w_0, v_0, 0) - \gamma\mathcal{E}(w_0, x_0, 0) - \frac{\gamma^2}{4} - \mathcal{E}(w_0, v_0, 0)^2, \end{aligned} \tag{4.1}$$

which gives us a single candidate point $\mathcal{Y}_{Q(\gamma)}$. Repeating this for many values of γ gives us a set of candidate points. Finally, the efficient frontier is constructed from the upper left convex hull of $\mathcal{Y}_Q$ (Tse et al., 2014) to remove spurious points. In our case, however, it turns out that all the points are on the efficient frontier, and there are no spurious points.

We are effectively using the parameter γ to trace out the efficient frontier. From Theorem 2.1, we have that $\gamma = \frac{1}{\lambda} + 2\mathcal{E}_0$. If $\lambda \to \infty$, the investor is infinitely risk averse, and invests only the risk free bond, hence in this case, we must have smallest possible value of γ

$$\gamma_{\min} = 2w_0 \exp(rT). \tag{4.2}$$

In practice, the interesting part of the efficient frontier is in the range $\gamma \in [\gamma_{\min}, 10\gamma_{\min}]$.

4.2.2 The Hybrid (PDE - Monte Carlo) discretization

In the hybrid method, given the stored optimal control $p^*(\cdot)$ from solving the HJB PDE (2.15), $(Var_{p^*(\cdot)}^{x_0,0}[W(T)], Var_{p^*(\cdot)}^{x_0,0}[W(T)])$ are then estimated by Monte Carlo simulations. We use the Euler scheme to generate the Monte Carlo simulation paths of the wealth (2.4), and an implicit Milstein

scheme to generate the Monte Carlo simulation paths of the variance process (2.3). Starting with $W_0 = w_0$ and $V_0 = v_0$, the Euler scheme for the wealth process (2.4) is

$$W_{t+\Delta t} = W_t \exp\left(\left(r + p^*\xi V_t - 0.5(p^*\sqrt{V_t})^2\right)\Delta t + p^*\sqrt{V_t \Delta t}\phi_1\right), \tag{4.3}$$

and the implicit Milstein scheme of the variance process (2.3) (Kahl and Jäckel, 2006) is

$$V_{t+\Delta t} = \frac{V_t + \kappa\theta\Delta t + \sigma\sqrt{V_t\Delta t}\phi_2 + \sigma^2\Delta t(\phi_2^2 - 1)/4}{1 + \kappa\Delta t}, \tag{4.4}$$

where ϕ_1 and ϕ_2 are standard normal variables with correlation ρ. Note that this discretization scheme will result in strictly positive paths for the variance process if $4\kappa\theta > \sigma^2$ (Kahl and Jäckel, 2006). For the cases where this bound does not hold, it will be necessary to modify (4.4) to prevent problems with the computation of $\sqrt{V_t}$. For instance, whenever V_t drops below zero, we could use the Euler discretization

$$V_{t+\Delta t} = V_t + \kappa(\theta - V_t^+)\Delta t + \sigma\sqrt{V_t^+}\sqrt{\Delta t}\phi_2, \tag{4.5}$$

where $V_t^+ = \max(0, V_t)$. (Lord et al., 2010) reviews a number of similar remedies to get around the problem when V_t becomes negative and concludes that the simple fix (4.5) works best.

4.3 Outside the computational domain

To make the numerical scheme consistent in a wide stencil method (Section 3.2.1), the stencil length needs to be increased to use the points beyond the nearest neighbors of the original grid. Therefore, when solving the PDE in a bounded region, the numerical discretization may require points outside the computational domain. When a candidate point we use is outside the computational region at the upper boundaries, we can directly use its asymptotic solution (3.4). For a point outside the upper boundary $w = w_{\max}$, the asymptotic solution is specified by the equation (3.4). For a point outside the upper boundary $v = v_{\max}$, by the implication of the boundary condition $\mathcal{U}_v = 0$ on $v = v_{\max}$, we have,

$$\mathcal{U}(w, v, \tau) = \mathcal{U}(w, v_{\max}, \tau), \quad v > v_{\max}. \tag{4.6}$$

However, we have to take special care when we may use a point below the lower boundaries $w = 0$ or $v = 0$, because the equation (2.15) is defined over $[0, \infty] \times [0, \infty]$. The possibility of using points below the lower boundaries only occurs when the node (i, j) falls in a possible region close to the lower boundaries

$$[h, \sqrt{h}] \times (0, w_{\max}] \cup (0, v_{\max}] \times [h, \sqrt{h}],$$

as discussed in Ma and Forsyth (2014). We can use the algorithm proposed in Ma and Forsyth (2014) to avoid this problem, and which retains consistency. That is, when one of the four candidate points $x_{i,j} \pm \sqrt{h}(\mathbf{R}_{i,j})_k,\ k = 1, 2$ (3.15) is below the lower boundaries, we then shrink its corresponding distance (from the reference node (i, j)) to h, instead of the original distance $\sqrt{h}$. This simple treatment ensures that all data required is within the domain of the HJB equation. It is straightforward to show that this discretization is consistent (Ma and Forsyth, 2014).

In addition, due to the semi-Lagrangian timestepping (Section 3.2.2) , we may need to evaluate the value of an off-grid point $(w_{i^*} = w_i e^{(r - p\xi v_j)\Delta\tau^n}, v_j)$ (3.19). This point maybe outside computational domain through the upper boundary $w = w_{\max}$ (the only possibility). When this situation occurs, the asymptotic solution (3.4) is used.

4.4 An improved linear interpolation scheme

When solving the value function problem (2.15) or the expected value problem (2.18) on a computational grid, it is required to evaluate $\mathcal{U}(\cdot)$ and $\mathcal{E}(\cdot)$, respectively, at points other than a node of the computational grid. This is especially important when using semi-Lagrangian timestepping. Hence, interpolation must be used. As mentioned earlier, to preserve the monotonicity of the numerical schemes, linear interpolation for an off-grid node is used in our implementation. Dang and Forsyth (2014b) introduces a special linear interpolation scheme applied along the w-direction to significantly improve the accuracy of the interpolation in a 2-D impulse control problem. We modify this algorithm in our problem set-up.

We then take advantage of the results in Section 3.1.3 to improve the accuracy of the linear interpolation. Assume that we want to proceed from timestep τ^n to τ^{n+1}, and that we want to compute $\mathcal{U}(\bar{w}, v_j, \tau^n)$ where $\bar{w}$ is neither a grid point in the w-direction nor the special value $W_{opt}(T-\tau^n)$, where W_{opt} is defined in equation (3.5). Furthermore, assume that $w_k < \bar{w} < w_{k+1}$ for some grid points w_k and w_{k+1}. For presentation purposes, let $w_{special} = W_{opt}(T-\tau^n)$ and $\mathcal{U}_{special} = 0$. An improved linear interpolation scheme along the w-direction for computing $\mathcal{U}(\bar{w}, v_j, \tau^n)$ is shown in Algorithm 4.1. Note that the interpolation along v-direction is a plain linear interpolation, thus we only illustrate the interpolation algorithm in w-direction.

Algorithm 4.1 Improved linear interpolation scheme along the w-direction for the function value problem

if $w_{special} < w_k$ or $w_{special} > w_{k+1}$ **then**
  set $w_{left} = w_k$, $\mathcal{U}_{left} = \mathcal{U}^n_{k,j}$, $w_{right} = w_{k+1}$, and $\mathcal{U}^n_{right} = \mathcal{U}^n_{k+1,j}$
else
  if $w_{special} < \bar{w}$ **then**
    set $w_{left} = w_{special}$, $\mathcal{U}_{left} = \mathcal{U}_{special}$, $w_{right} = w_{k+1}$, and $\mathcal{U}^n_{right} = \mathcal{U}^n_{k+1,j}$
  else
    set $w_{left} = w_k$, $\mathcal{U}_{left} = \mathcal{U}^n_{k,j}$, $w_{right} = w_{special}$, and $\mathcal{U}^n_{right} = \mathcal{U}_{special}$
  end if
end if
Apply linear interpolation to $(w_{left}, \mathcal{U}_{left})$ and $(w_{right}, \mathcal{U}_{right})$ to compute $\mathcal{U}(\bar{w}, v_j, \tau^n)$

Following the same line of reasoning used for the function value problem, we have that

$$\mathcal{E}(v, W_{opt}(t), t) = \frac{\gamma}{2}.$$

By using this result, a similar method as Algorithm 4.1 can be used to improve the accuracy of linear interpolation when computing the expected value $\mathcal{E}(\bar{w}, v_j, \tau^n)$.

Remark 4.1. *For the discretization of the expected value problem (2.18), we still use the semi-Lagrangian timestepping to handle the drift term $(r + p^*\xi v)w\mathcal{E}_w$. Since it may be necessary to evaluate $\mathcal{E}^n_{i^*,j}$ at points other than a node of the computational grid, we need to use linear interpolation.*

5 Numerical Experiments

In this section, we present numerical results of solution of equation (2.15) applied to the continuous time mean variance portfolio allocation problem. In our problem, the risky asset (2.2) follows the Heston model. The parameter values of the Heston model used in our numerical experiments are taken from (Aıt-Sahalia and Kimmel, 2007) based on empirical calibration from S&P 500 index and VIX index dataset during 1990 to 2004 (under the real probability measure). Table 5.1 lists the Heston model parameters, and Table 5.2 lists the parameters of the mean variance portfolio allocation problem.

κ	θ	σ	ρ	ξ
5.07	0.0457	0.48	-0.767	1.605

Table 5.1: Parameter values in the Heston model

Investment Horizon T	10
The risk free rate r	0.03
Leverage constraint $p_{\max}$	2
Initial wealth w_0	100
Initial variance v_0	0.0457

Table 5.2: Input parameters for the mean variance portfolio allocation problem.

For all the experiments, unless otherwise noted, the details of the grid, the control set, and timestep refinement levels used are given in Table 5.3.

Refinement	Timesteps	W Nodes	V Nodes	$\mathcal{Z}_h$ Nodes
0	160	112	57	8
1	320	223	113	16
2	640	445	225	32
3	1280	889	449	64

Table 5.3: Grid and timestep refinement levels used during numerical tests. On each refinement, a new grid point is placed halfway between all old grid points and the number of timesteps is doubled. A constant timestep size is used. $w_{\max} = 6 \times 10^6$ and $v_{\max} = 3.0$. The number of finite sampled γ is 50. Note that increasing $w_{\max}$ by an order of magnitude and doubling $v_{\max}$ results in no change to the points on the efficient frontier to five digits. Increasing the number of γ points did not result in any appreciable change to efficient frontier (no spurious points in this case).

5.1 Effects of the improved interpolation scheme for the PDE method

In this subsection, we discuss the effects on numerical results of the linear interpolation scheme described in Section 4.4. We plot expected values against standard deviation, since both variables have the same units. Figure 5.1a illustrates the numerical efficient frontiers obtained using standard

linear interpolation. It is clear that the results are very inaccurate for small standard deviations. It appears that the numerical methods were not able to construct the known point on the exact efficient frontier

$$(Var^{x,t}_{p^*(\cdot)}[W(T)], E^{x,t}_{p^*(\cdot)}[W(T)]) = (0, w_0 e^{rT}) \approx (0, 134.9859).$$

This trivial case corresponds to the case where $\gamma = \gamma_{\min}$ (4.2), and the investor invests only in the risk free bond and not in the risky asset. However, as shown in Figure 5.1a, in this special case, the standard deviation obtained by the numerical scheme using standard linear interpolation is far from the exact solution.

Figure 5.1b shows the numerical efficient frontiers obtained with the improved linear interpolation scheme, where Algorithm 4.1 is utilized. It is obvious that the numerical efficient frontiers obtained with the improved linear interpolation scheme are more reasonable, especially for the small standard deviation region. In particular, the special point where the variance is zero is now approximated accurately. This result illustrates the importance of using the optimal embedded terminal wealth $W_{opt}(t)$ and its function value for linear interpolation in constructing accurate numerical efficient frontiers. In all our numerical experiments in the following, the improved linear interpolation scheme is used.

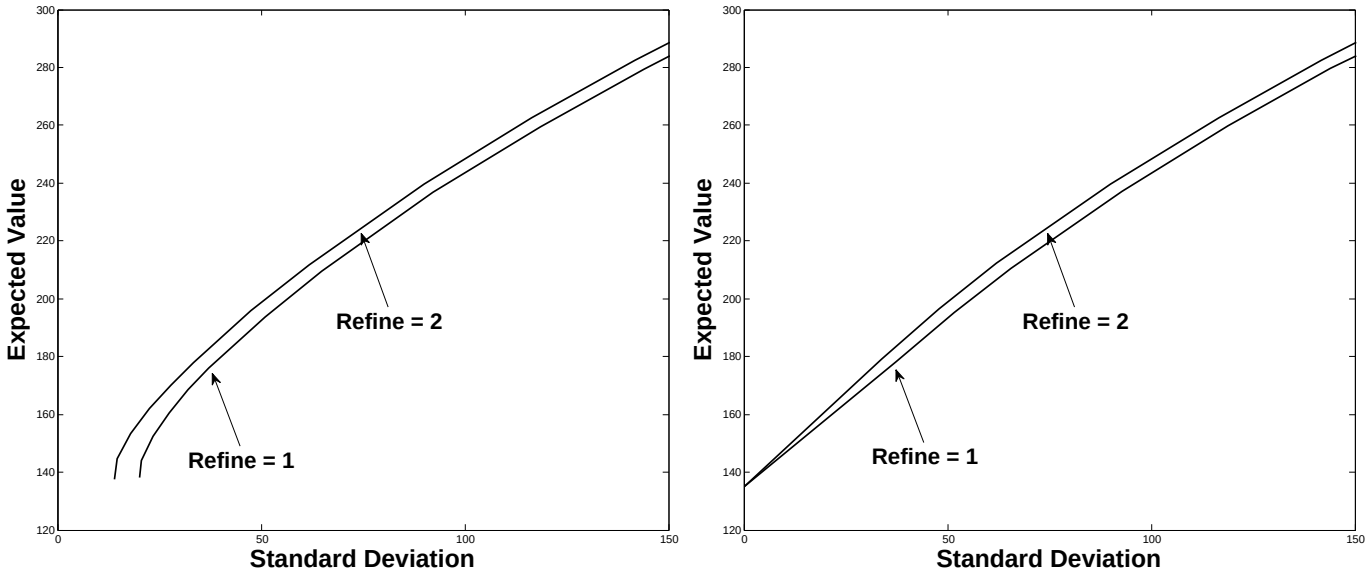

(a) No Special Interpolation

(b) Special Interpolation

Figure 5.1: Close-up of efficient frontier for small standard deviations. (a) No Special Interpolation. (b) Special interpolation.

5.2 Convergence analysis

In this section, we illustrate the convergence of our numerical scheme, and compare the performance of two methods, namely the PDE method (Section 4.2.1) and the Hybrid method (4.2.2), for constructing the mean variance frontier under our model set-up.

Figure 5.2 shows that the mean standard deviation 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 results than the pure PDE solutions. This same phenomenon was observed in (Tse et al., 2013). As shown in Figure 5.2, the frontiers obtained by the Hybrid method are almost identical for refinement level 1 and 2. Note that for both methods, the optimal control is always computed by solving the HJB PDEs.

The same timesteps are used in both PDE method and Monte Carlo simulations, for each refinement level. For example, the frontiers labeled with "Refine = 1" for both methods in Figure 5.2 use the time steps as specified as Refinement level 1 in Table 5.3. To achieve small sampling error in Monte Carlo simulations, 10^6 simulations are performed for the numerical experiments. The standard error in Figure 5.2 can then be estimated. For example, consider a point on the frontier with the large standard deviation value which is about 350. For the expected value of $W(T)$, the sample error is approximately $350/\sqrt{10^6} \approx 0.35$, which could be negligible in Figure 5.2.

We will verify our conclusion by examining several specific points on these efficient frontiers in Figure 5.2. Table 5.4 and Table 5.5 show computed means and standard deviations for different refinement levels when $\gamma = 540$. The numerical results indicate first order convergence is achieved for both the PDE method and the Hybrid method. In this case, our numerical results demonstrate that the Hybrid frontiers converge faster to the limit results than the PDE solutions. Table 5.6 and Table 5.7 show computed means and standard deviations for different refinement levels when $\gamma = 1350$. The numerical results indicate first order convergence is achieved for the PDE method. In this case, our numerical results also demonstrate that the Hybrid frontiers converge faster to the limit results than the PDE solutions. However, the convergence ratio for the Hybrid method is erratic. As we noted before, in this case, the sample error for the estimate of the mean value is about 0.2 ($200/\sqrt{10^6}$). The sample error may cause the phenomenon of the erratic convergence ratio in the Hybrid method results. To decrease the sample error to, for example, 0.01, the number of simulation paths would have to increase to 100×10^6, which is unaffordable in terms of the computational cost. Note that in the case $\gamma = 540$, with the small standard deviation, the sample error decreases to about 0.01.

Remark 5.1 (Efficiency of the Hybrid method.)**.** *We remind that reader that for both the Hybrid and PDE methods, the same (computed) control used. The more rapid convergence of the Hybrid method is simply due to a more accurate estimate of the expected quantities (with a known control). This result is somewhat counter-intuitive, since it suggests that a low accuracy control can be used to generate high accuracy expected values. We also observe this from the fact that a fairly coarse discretization of the admissible set $\mathcal{Z}_h$ generates fairly accurate solutions.*

Refine	Mean	Change	Ratio	Standard Deviation	Change	Ratio
0	207.1434			71.3924		
1	210.4694	3.3260		65.5090	−5.88336	
2	212.1957	1.7263	1.92	62.0862	−3.42288	1.72
3	213.1481	0.95238	1.81	60.4738	−1.61237	2.12

Table 5.4: The convergence table for the PDE method. Small standard deviation case with $\gamma = 540$.

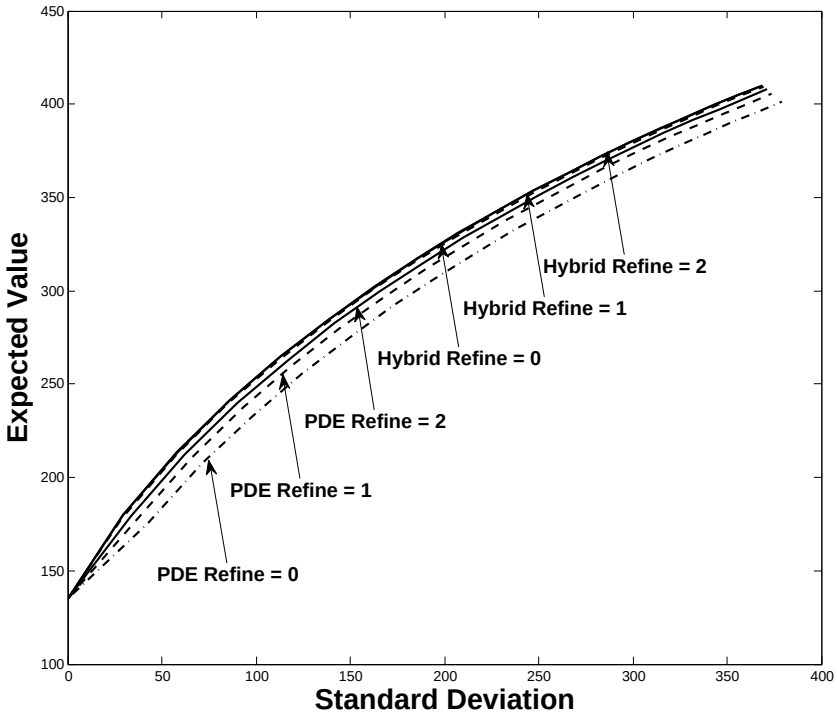

Figure 5.2: convergence of frontiers in the PDE method and the Hybrid method. The frontiers labeled with "PDE" are obtained from the PDE method (Section 4.2.1). The frontiers labeled with "Hybrid" (Section 4.2.2) are obtained from a Monte Carlo simulation which uses the optimal controls determined by solving the HJB equation (2.15).

Refine	Mean	Change	Ratio	Standard Deviation	Change	Ratio
0	212.2993			56.6128		
1	213.2077	0.908		57.7652	1.152	
2	213.7573	0.550	1.65	58.2987	0.534	2.16
3	213.9903	0.233	2.36	58.5253	0.227	2.35

Table 5.5: The convergence table for the Hybrid method. Small standard deviation case with $\gamma = 540$.

Refine	Mean	Change	Ratio	Standard Deviation	Change	Ratio
0	320.5139			217.0009		
1	325.5443	5.030		212.1886	−4.812	
2	328.2670	2.723	1.85	209.8434	−2.345	2.05
3	329.8172	1.550	1.76	208.9045	−0.939	2.50

Table 5.6: The convergence table for the PDE method. Large standard deviation case with $\gamma = 1350$.

Refine	Mean	Change	Ratio	Standard Deviation	Change	Ratio
0	329.4411			206.0875		
1	330.5172	1.076		206.8351	0.748	
2	330.7066	0.189	5.68	207.1958	0.361	2.07
3	331.2820	0.575	0.33	207.3707	0.175	2.06

Table 5.7: The convergence table for the Hybrid method. Large standard deviation case with $\gamma = 1350$.

5.3 Sensitivity of Efficient Frontiers

In this subsection, we show some numerical sensitivity analysis for the major market parameters, namely the leverage constraints $p_{\max}$, the market risk ξ, the mean reversion level for the variance θ, the volatility of the variance σ, the correlation ρ between the risky asset and the variance, and the mean reversion speed κ. In our numerical tests, the corresponding frontiers are generated as the market parameter of interest changes, and the values of the remaining parameters are fixed and are listed in Table 5.1 and Table 5.2. We use the Hybrid method with the discretization level 2.

As observed in Figure 5.3, with $p_{\max} = \{1, 1.5, 2, +\infty\}$, larger values of the leverage constraints $p_{\max}$ result in much more dominant efficient frontiers. From Figure 5.4, with $\xi = \{0.5, 1.605, 2.5\}$, we can see that larger values of ξ result in much more dominant efficient frontiers. The maximal standard deviation point ($\gamma = +\infty$) on the efficient frontier with $\xi = 0.5$ is only about 191, which is much smaller than those with larger ξ values. From Figure 5.5, $\theta = \{0.01, 0.0457, 0.36\}$, we can see that larger values of the mean reversion level θ for the variance, result in much more dominant efficient frontiers. The maximal standard deviation point ($\gamma = +\infty$) on the efficient frontier with $\theta = 0.01$ is only about 108, which is much smaller than those with larger θ values. From Figure 5.6, $\sigma = \{0.7, 0.0457, 0.2\}$, we can see that larger values of the volatility of the variance σ result in a slightly more dominant efficient frontiers in general. In particular, these efficient frontiers in large standard deviation region with different σ values values are almost identical.

On the other hand, from Figure 5.7, with $\rho = \{-0.767, -0.3, 0\}$, we can see that an increase in the correlation ρ produces frontiers with a sightly smaller expected value for a given standard deviation. These efficient frontiers in the large standard deviation region with different ρ values are almost identical. The effect of the κ values on the efficient frontier is more complex. From Figure 5.8, $\kappa = \{1, 5.07, 20\}$, in the small standard deviation region, an increase in κ produces frontiers with a smaller expected value for a given standard deviation. However, when the standard deviation increases to about 230, the larger values of κ gradually result in more significant dominant efficient frontiers.

5.4 Comparison between constant volatility and stochastic volatility cases

In this paper, the risky asset follows the stochastic volatility model (2.2-2.3). In this Section, we will compare the constant volatility and stochastic volatility cases in terms of mean variance efficiency for the continuous time pre-commitment mean variance problem. With a constant volatility, the risky asset is the governed by the following geometric Brownian Motion (GBM) process:

$$\frac{dS}{S} = (r + \mu)dt + \sigma_S dZ_s. \tag{5.1}$$

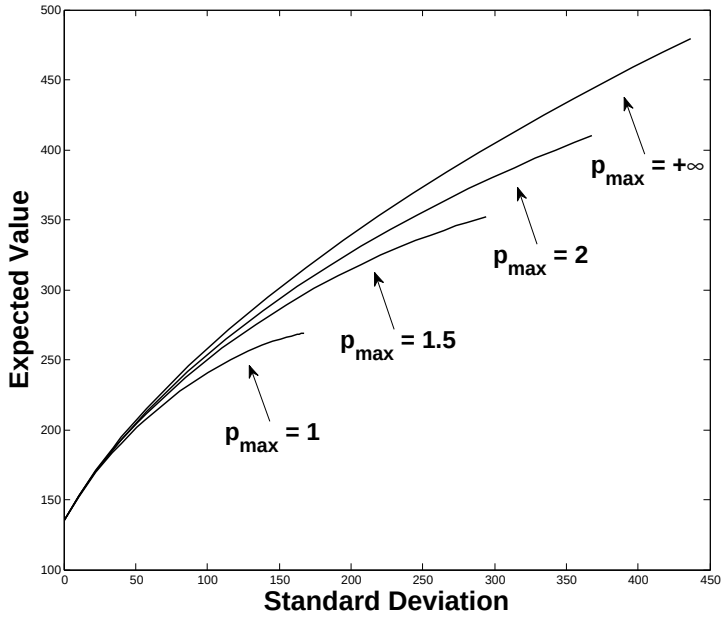

Figure 5.3: Sensitivity analysis of the efficient frontiers with respect to different leverage constraints $p_{\max}$. The Heston parameters and the remaining model parameters are given in Table 5.1 and Table 5.2. The Hybrid method with discretization level 2 is used.

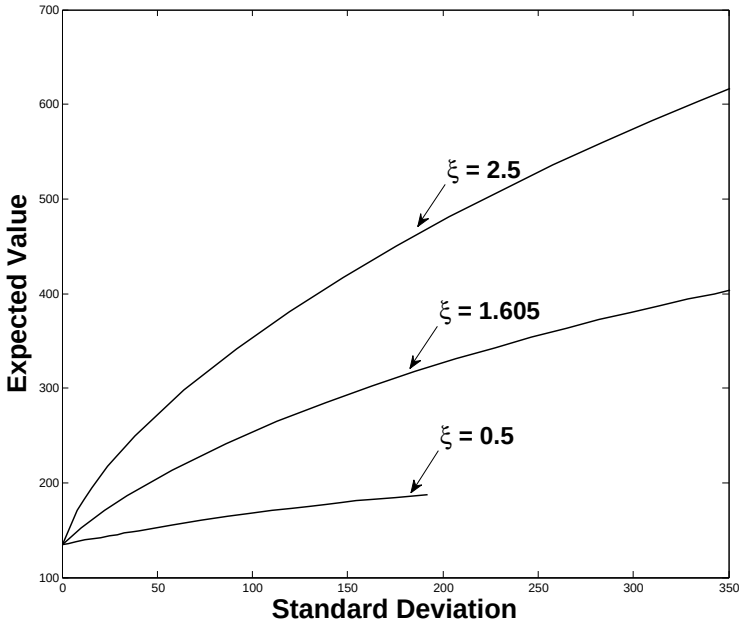

Figure 5.4: Sensitivity analysis of the efficient frontiers with respect to different risk premium factor ξ values. The remaining Heston parameters and the model parameters are given in Table 5.1 and Table 5.2. The Hybrid method with discretization level 2 is used.

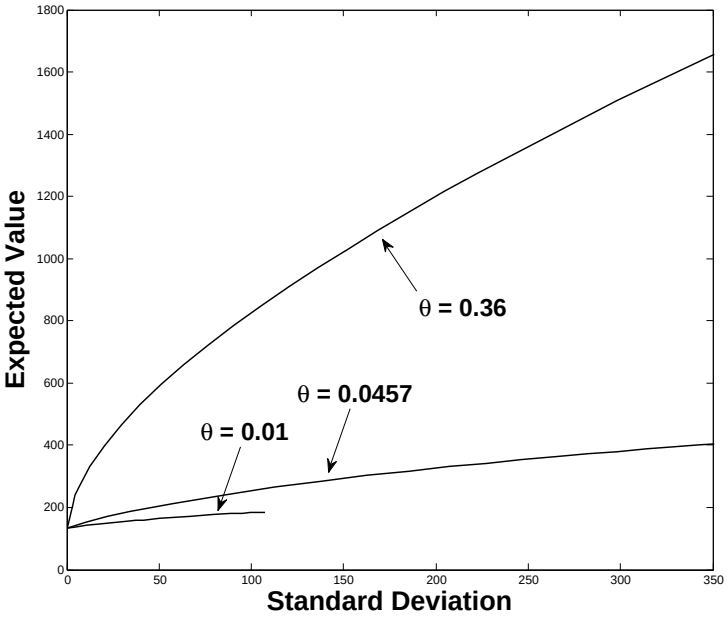

Figure 5.5: Sensitivity analysis of the efficient frontiers with respect to different mean reversion level θ values. The remaining Heston parameters and the model parameters are given in Table 5.1 and Table 5.2. The Hybrid method with discretization level 2 is used.

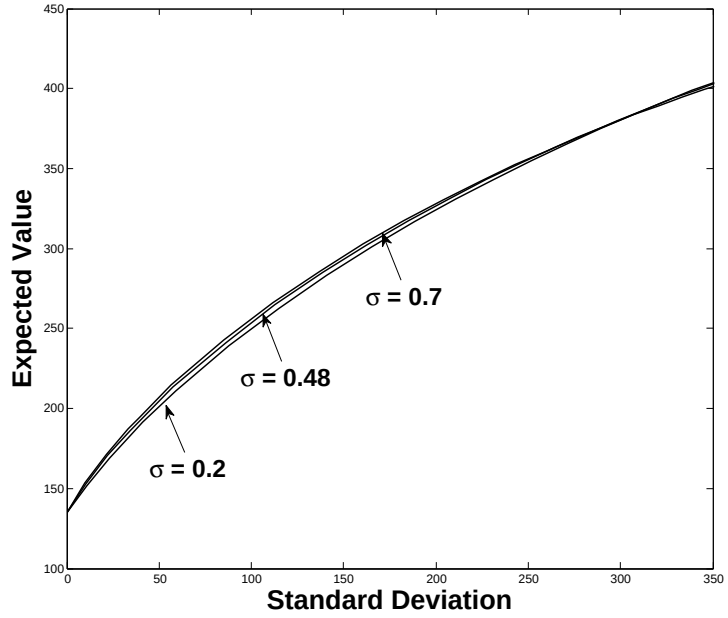

Figure 5.6: Sensitivity analysis of the efficient frontiers with respect to different σ values. The remaining Heston parameters and the model parameters are given in Table 5.1 and Table 5.2. The Hybrid method with discretization level 2 is used.

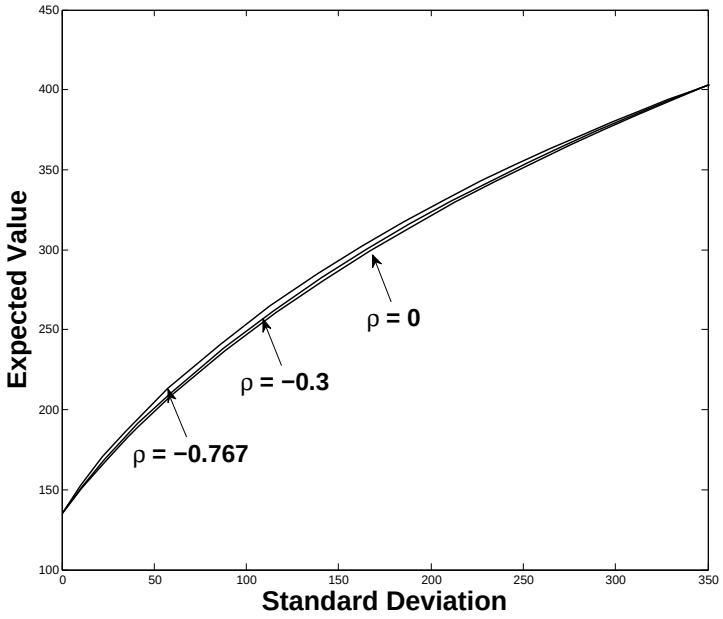

Figure 5.7: Sensitivity analysis of the efficient frontiers with respect to different ρ values. The remaining Heston parameters and the model parameters are given in Table 5.1 and Table 5.2. The Hybrid method with discretization level 2 is used.

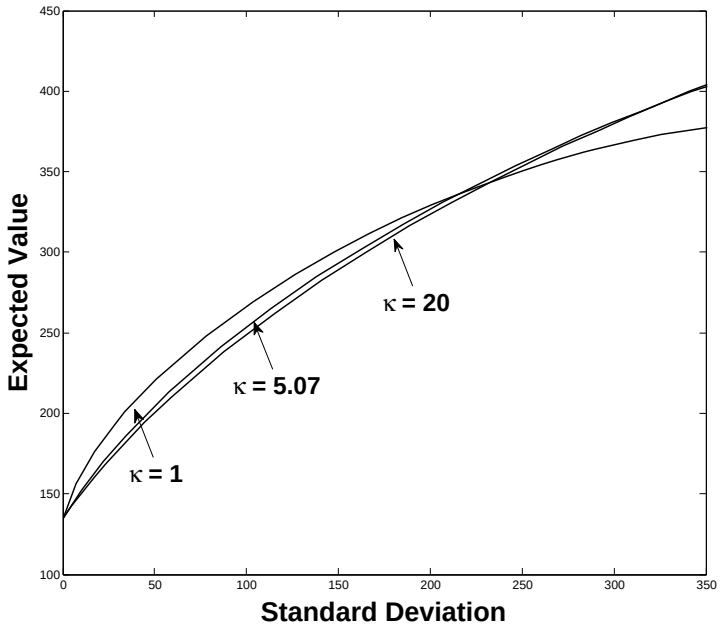

Figure 5.8: Sensitivity analysis of the efficient frontiers with respect to different κ values. The Heston parameters and the remaining model parameters are given in Table 5.1 and Table 5.2. The Hybrid method with discretization level 2 is used.

		$\gamma = 540$		$\gamma = 1350$	
Control Process	Price Process	Mean	Stnd Dev	Mean	Stnd Dev
GBM	GBM	209.50	59.68	330.09	213.01
GBM	Stoch. Vol.	212.68	58.42	329.13	207.23
Stoch. Vol.	Stoch. Vol	213.99	58.53	331.28	207.37

Table 5.8: Given a γ, the optimal portfolio allocation strategy is computed and stored assuming a control process, which is either GBM or stochastic volatility. The mean variance pairs are then estimated by Monte Carlo Simulation, using the stored controls, assuming that the actual price process follows either GBM or stochastic volatility. For the stochastic volatility case, the parameters are given in Table 5.1. For the GBM case, the variance is fixed to the mean value of the stochastic volatility case.

To compare with the stochastic volatility case in Table 5.1, the constant volatility σ_S is set to $\sqrt{\theta} \approx 0.2138$, and the risky return over the risk free rate μ is set to $\xi\sigma_S^2 = 0.0733485$, which has the same mean premium of the volatility risk as the stochastic volatility model (2.2). This then corresponds to the case where the variance $V(t)$ in (2.2) is fixed to the mean reversion level θ. The remaining mean variance problem parameters are the same as listed in Table 5.2.

Figure 5.9 illustrates the fact that the efficient frontiers produced by using the stochastic volatility sightly dominates the curve produced by the constant volatility model. With the Heston model's parameters in Table 5.1, we may conclude that the efficient frontier produced by the constant volatility is a good approximation of the frontier generated by the stochastic volatility model. From Figure 5.9, however, we see that if the mean reversion speed κ is set to a small value, e.g. one, in the stochastic volatility case, the efficient frontiers computed using a constant volatility model will be considerably different from those computed using the stochastic volatility model. The quantity $1/\kappa$ is measured in years and is related to the time over which a volatility shock dissipates. Specially, the half-life of a volatility shock is $\frac{\ln 2}{\kappa}$.

Finally, using the portfolio allocation strategy that is precomputed and stored from the constant volatility case, we then carry out a Monte Carlo simulation where the risky asset follows the stochastic volatility model. We then compare the results using this approximate control, with the optimal control computed using the full stochastic volatility model. From Table 5.8, we can see that the mean variance pairs computed using the optimal strategy are very close to the strategy computed using the GBM approximation. Based on several tests, a good heuristic guideline is that if $\kappa T > 40$, then the GBM control is a good approximation to the true (optimal control).

6 Conclusion

In this paper, we develop an efficient fully numerical PDE approach for the pre-commitment continuous time mean variance asset allocation problem when the risky asset follows a stochastic volatility model. We use the wide stencil method (Ma and Forsyth, 2014) to overcome the main difficulty in designing a monotone approximation. We show that our numerical scheme is monotone, consistent, and ℓ_∞-stable. Hence, the numerical solution is guaranteed to converge to the unique viscosity solution of the corresponding HJB PDE, assuming that the HJB PDE satisfies a strong comparison property. Furthermore, using semi-Lagrangian timestepping to handle the drift term and an improved method of linear interpolation, allows us to compute accurate efficient

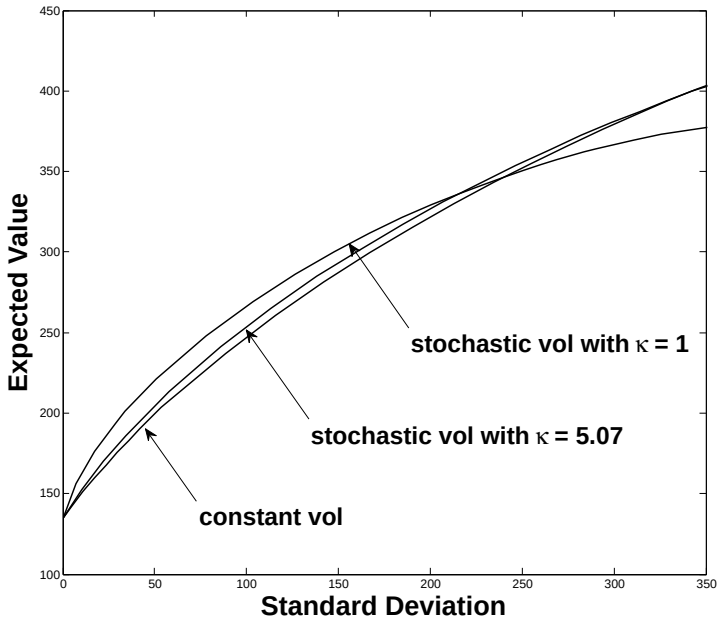

Figure 5.9: Efficient Frontier Comparison between constant volatility and stochastic volatility cases. For the stochastic volatility cases, $\kappa = 1,\ 5.07$, and the remaining stochastic volatility parameters are given in Table 5.1. The GBM parameters are given in Section 5.4.

frontiers. When tracing out the efficient frontier solution of our problem, we demonstrate that the Hybrid (PDE - Monte Carlo) method (Tse et al., 2013) converges faster than the pure PDE method. Similar results are observed in Tse et al. (2013). Finally, if the mean reversion time $\frac{1}{\kappa}$ is small compared to the investment horizon T, then a constant volatility GBM approximation to the stochastic volatility process gives a very good approximation to the optimal strategy.

Appendix

A The discrete linear operator D_h^p

With vanishing cross-derivative term, the degenerate linear operator $\mathcal{L}^p$ (3.14) can be discretized by a standard finite difference method. The degenerate linear operators $\mathcal{L}^p$ in (3.1), (3.2), and (3.3) are approximated as the discrete form

$$D_h^p \mathcal{U}_{i,j}^n = \alpha_{i,j}^w \mathcal{U}_{i-1,j}^n + \beta_{i,j}^w \mathcal{U}_{i+1,j}^n + \alpha_{i,j}^v \mathcal{U}_{i,j-1}^n + \beta_{i,j}^v \mathcal{U}_{i,j+1}^n - (\alpha_{i,j}^w + \beta_{i,j}^w + \alpha_{i,j}^v + \beta_{i,j}^v)\mathcal{U}_{i,j}, \qquad \text{(A.1)}$$

where $\alpha^w_{i,j}, \beta^w_{i,j}, \alpha^v_{i,j}$ and $\beta^v_{i,j}$ are defined as follows

$$
\begin{aligned}
\alpha^w_{i,j} &= \frac{(\sqrt{v}pw_i)^2}{(w_i - w_{i-1})(w_{i+1} - w_{i-1})}, \\
\beta^w_{i,j} &= \frac{(\sqrt{v}pw_i)^2}{(w_{i+1} - w_i)(w_{i+1} - w_{i-1})}, \\
\alpha^v_{i,j} &= \left[\frac{(\sigma\sqrt{v_j})^2}{v_j - v_{j-1})(v_{j+1} - v_{j-1})} + \max\left(0, -\frac{\kappa(\theta - v_j)}{v_j - v_{j-1}}\right) \right], \\
\beta^v_{i,j} &= \left[\frac{(\sigma\sqrt{v_j})^2}{(v_{j+1} - v_j)(v_{j+1} - v_{j-1})} + \max\left(0, \frac{\kappa(\theta - v_j)}{v_{j+1} - v_j}\right) \right].
\end{aligned}
\tag{A.2}
$$

The coefficients $\alpha^w_{i,j}, \beta^w_{i,j}, \alpha^v_{i,j}$ and $\beta^v_{i,j}$ are all non-negative, and is compatible with a monotone scheme. On the upper boundary $v = v_{\max}$, the coefficients α^v_{i,N_2} and $\beta^v_{i,N_2} = 0$ degenerate to zero, and On the lower boundary $w = 0$, $\alpha^w_{1,j}$ and $\beta^w_{1,j}$ are set to 0. On the lower boundary $v = 0$, $\alpha^w_{i,1} = 0$, $\beta^w_{i,1} = 0$, $\alpha^v_{i,1} = 0$, and $\beta^w_{i,1} = \frac{\kappa\theta}{v_{j+1} - v_j}$, $j = 1$.

References

Aıt-Sahalia, Y. and R. Kimmel (2007). Maximum likelihood estimation of stochastic volatility models. *Journal of Financial Economics 83*(2), 413–452.

Barles, G. and P. E. Souganidis (1991). Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic analysis 4*(3), 271–283.

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

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

Bjork, T. and A. Murgoci (2010). A general theory of Markovian time inconsistent stochastic control problems. *Available at SSRN 1694759*.

Breeden, D. T. (1979). An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of financial Economics 7*(3), 265–296.

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

Cox, J. C., J. Ingersoll, E. Jonathan, and S. A. Ross (1985). A theory of the term structure of interest rates. *Econometrica: Journal of the Econometric Society 53*(2), 385–407.

Cui, X., D. Li, S. Wang, and S. Zhu (2012). Better than dynamic mean-variance: Time inconsistency and free cash flow stream. *Mathematical Finance 22*(2), 346–378.

Dang, D. M. and P. A. Forsyth (2014a). Better than pre-commitment mean-variance portfolio allocation strategies: a semi-self-financing Hamilton-Jacobi-Bellman equation approach. *Working paper, Cheriton School of Computer Science, University of Waterloo*.

Dang, D. M. and P. A. Forsyth (2014b). Continuous time mean-variance optimal portfolio allocation under jump diffusion: An numerical impulse control approach. *Numerical Methods for Partial Differential Equations 30*(2), 664–698.

Dang, D. M., P. A. Forsyth, and Y. Li (2015). Convergence of the embedded mean-variance optimal points with discrete sampling. To appear in Numerische Mathematik.

Debrabant, K. and E. Jakobsen (2013). Semi-Lagrangian schemes for linear and fully non-linear diffusion equations. *Mathematics of Computation 82*(283), 1433–1462.

d'Halluin, Y., P. A. Forsyth, and G. Labahn (2004). A penalty method for American options with jump diffusion processes. *Numerische Mathematik 97*(2), 321–352.

Douglas, J. J. and T. F. Russell (1982). Numerical methods for convection-dominated diffusion problems based on combining the method of characteristics with finite element or finite difference procedures. *SIAM Journal on Numerical Analysis 19*(5), 871–885.

Ekström, E. and J. Tysk (2010). The Black–Scholes equation in stochastic volatility models. *Journal of Mathematical Analysis and Applications 368*(2), 498–507.

Feller, W. (1951). Two singular diffusion problems. *Annals of mathematics 54*(1), 173–182.

Forsyth, P. A. and G. Labahn (2007). Numerical methods for controlled Hamilton-Jacobi-Bellman PDEs in finance. *Journal of Computational Finance 11*(2), 1–44.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of financial studies 6*(2), 327–343.

Huang, Y., P. A. Forsyth, and G. Labahn (2012). Combined fixed point and policy iteration for HJB equations in finance. *SIAM Journal on Numerical Analysis 50*(4), 1849–1860.

Kahl, C. and P. Jäckel (2006). Fast strong approximation Monte Carlo schemes for stochastic volatility models. *Quantitative Finance 6*(6), 513–536.

Li, D. and W.-L. Ng (2000). Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical Finance 10*(3), 387–406.

Lord, R., R. Koekkoek, and D. V. Dijk (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance 10*(2), 177–194.

Ma, K. and P. A. Forsyth (2014). An unconditionally monotone numerical scheme for the two factor uncertain volatility model. *Working paper, Cheriton School of Computer Science, University of Waterloo*.

Nguyen, P. and R. Portait (2002). Dynamic asset allocation with mean variance preferences and a solvency constraint. *Journal of Economic Dynamics and Control 26*(1), 11–32.

Pironneau, O. (1982). On the transport-diffusion algorithm and its applications to the Navier-Stokes equations. *Numerische Mathematik 38*(3), 309–332.

Pooley, D. M., P. A. Forsyth, and K. R. Vetzal (2003). Numerical convergence properties of option pricing PDEs with uncertain volatility. *IMA Journal of Numerical Analysis 23*(2), 241–267.

Saad, Y. (2003). *Iterative methods for sparse linear systems*. SIAM.

Tse, S. T., P. A. Forsyth, J. S. Kennedy, and H. Windcliff (2013). Comparison between the mean-variance optimal and the mean-quadratic-variation optimal trading strategies. *Applied Mathematical Finance 20*(5), 415–449.

Tse, S. T., P. A. Forsyth, and Y. Li (2014). Preservation of scalarization optimal points in the embedding technique for continuous time mean variance optimization. *SIAM Journal on Control and Optimization 52*(3), 1527–1546.

Varga, R. S. (2009). *Matrix Iterative Analysis*, Volume 27. Springer-Verlag, Berlin.

Vigna, E. (2014). On efficiency of mean–variance based portfolio selection in defined contribution pension schemes. *Quantitative finance 14*(2), 237–258.

Wang, J. and P. A. Forsyth (2008). Maximal use of central differencing for Hamilton-Jacobi-Bellman PDEs in finance. *SIAM Journal on Numerical Analysis 46*(3), 1580–1601.

Wang, J. and P. A. Forsyth (2010). 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.

Wang, J. and P. A. Forsyth (2012). Comparison of mean variance like strategies for optimal asset allocation problems. *International Journal of Theoretical and Applied Finance 15*(2).

Zhao, Y. and W. T. Ziemba (2000). Mean-variance versus expected utility in dynamic investment analysis. *Working paper, University of British Columbia*.

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