Time-consistent mean-variance portfolio optimization: a numerical impulse control approach

Pieter Van Staden* Duy-Minh Dang† Peter A. Forsyth‡

Abstract

We investigate the time-consistent mean-variance (MV) portfolio optimization problem, popular in investment-reinsurance and investment-only applications, under a realistic context that involves the simultaneous application of different types of investment constraints and modelling assumptions, for which a closed-form solution is not known to exist. We develop an efficient numerical partial differential equation method for determining the optimal control for this problem. Central to our method is a combination of (i) an impulse control formulation of the MV investment problem, and (ii) a discretized version of the dynamic programming principle enforcing a time-consistency constraint. We impose realistic investment constraints, such as no trading if insolvent, leverage restrictions and different interest rates for borrowing/lending. Our method requires solution of linear partial integro-differential equations between intervention times, which is numerically simple and computationally effective. The proposed method can handle both continuous and discrete rebalancings. We study the substantial effect and economic implications of realistic investment constraints and modelling assumptions on the MV efficient frontier and the resulting investment strategies. This includes (i) a comprehensive comparison study of the pre-commitment and time-consistent optimal strategies, and (ii) an investigation on the significant impact of a wealth-dependent risk aversion parameter on the optimal controls.

Keywords: Asset allocation, constrained optimal control, time-consistent, pre-commitment, impulse control

JEL Subject Classification: G11, C61

1 Introduction

Originating with Markowitz (1952), the standard criterion in modern portfolio theory has been maximizing the (terminal) expected return of a portfolio, given an acceptable level of risk, where risk is quantified by the (terminal) variance of the portfolio returns. This is referred to as mean-variance (MV) portfolio optimization. Mean-variance strategies are appealing due to their intuitive nature, since the results can be easily interpreted in terms of the trade-off between risk (variance) and reward (expected return).

Broadly speaking, there are three main approaches, or strategies, to perform MV optimization, namely (i) the pre-commitment approach, (ii) the dynamically optimal approach, and (iii) the time-consistent (or game theoretical) approach. The pre-commitment approach (see, for example, Dang and Forsyth (2014); Li and Ng (2000); Zhou and Li (2000)) is time-inconsistent in the sense that, for

*School of Mathematics and Physics, The University of Queensland, St Lucia, Brisbane 4072, Australia, email: pieter.vanstaden@uq.edu.au

†School of Mathematics and Physics, The University of Queensland, St Lucia, Brisbane 4072, Australia, email: duyminh.dang@uq.edu.au

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

$\Delta t > 0$, the MV optimal strategy for time $t + 2\Delta t$, computed at time t, may not agree with the MV optimal strategy for time $t+2\Delta t$, computed at time $t+\Delta t$. This time-inconsistency is addressed in the dynamically optimal approach of Pedersen and Peskir (2017) by recomputing the pre-commitment MV optimal strategy at each time instant, thereby achieving instantaneous optimality by performing an infinite number of optimization problems. The time-consistent approach (Basak and Chabakauri, 2010; Bjork and Murgoci, 2010; Cong and Oosterlee, 2016; Wang and Forsyth, 2011) addresses the time-inconsistency by imposing a time-consistency constraint, i.e. constraining the original MV optimization problem to ensure that the optimal strategy for time $t + 2\Delta t$, computed at time t, necessarily agrees with the optimal strategy for time $t + 2\Delta t$, computed at time $t + \Delta t$.

The pre-commitment MV perspective typically requires that an investor/insurer make a credible commitment that the MV-optimal strategy, computed at time zero, will be followed over the entire investment time horizon, which could be substantially long (e.g. pension fund investments). This may not be feasible in institutional settings. The time-consistent MV approach has received considerable attention in recent literature; see, for example, Alia et al. (2016); Bensoussan et al. (2014); Cui et al. (2015); Li et al. (2015c); Liang and Song (2015); Sun et al. (2016); Zhang and Liang (2017), among many other publications. In particular, as evidenced by these publications, this approach has been very popular in institutional settings - especially in insurance-related applications, where MV-utility insurers are typically concerned with investment-reinsurance or investment-only optimization problems.

With the notable exception of Wang and Forsyth (2011) and Cong and Oosterlee (2016), virtually all of the available literature on time-consistent MV optimization is based on solving the resulting equations using closed-form (analytical) techniques, which necessarily requires very restrictive, and hence unrealistic, modelling and investment assumptions. These assumptions include continuous rebalancing, zero transaction costs, allowing insolvency and infinite leverage. Formulating problems without realistic investment constraints usually results in conclusions that are difficult to justify, and/or are potentially infeasible to implement in practice.

Specifically, in the time-consistent MV literature, the effect of the commonly encountered assumption, namely trading continues even if the investor is insolvent, is rarely considered. A few exceptions include Zhou et al. (2016), where the bankruptcy implications from multi-period time-consistent MV and pre-commitment MV optimization problems are compared; however, a bankruptcy constraint is not explicitly enforced in this work. A conclusion in Zhou et al. (2016) is that the time-consistent strategy "can diversify bankruptcy risk efficiently", since the resulting probability of insolvency over the investment time horizon is lower, and therefore, the time-consistent strategy might be preferred by a rational investor over the pre-commitment strategy. However, in practice, real portfolios have bankruptcy constraints. Hence, such conclusions are questionable. In the case of other time-consistent MV applications, such as asset-liability management, the explicit incorporation of insolvency considerations is critical to ensure that the results are of any practical use. The analytical solutions in, for example, Wei et al. (2013) and Wei and Wang (2017), while useful, necessarily assume trading continues in the case of insolvency.

Moreover, in the time-consistent MV literature, it is typical for analytical techniques to allow for a leverage ratio, i.e. the ratio of the investment in the risky asset to the total wealth, substantially larger than a ratio that brokers would typically allow retail investors or financial regulators would likely allow institutions to undertake in practice. More specifically, while a leverage ratio of around 1.5 times is typically allowed in practice (for retail investors), some of the analytical techniques illustrated in the available literature call for much larger leverage ratios, for example 2.4 times in Li et al. (2012), 3 times in Zeng et al. (2013), 2.6 times in Liang and Song (2015), 2.5 times in Li et al. (2015c), and as high as 14 times in Li et al. (2015a), none of which are practically feasible, and which only further increases the probability of insolvency. In a number of publications, a leverage constraint is completely ignored,

such as Lioui (2013), and this potentially leads to misplaced economic conclusions. For example, it is concluded in Lioui (2013) that the time-consistent strategy is preferred over the pre-commitment strategy, since the latter requires "huge and unrealistic positions in risky assets; in some cases, the pre-commitment strategy is more than 60 times the time consistent strategy". However, such a conclusion appears unconvincing, since the pre-commitment MV strategy's positions in the risky asset would have been significantly smaller, if a realistic leverage constraint had been incorporated into the problem formulation.

In addition, failing to incorporate transaction costs may also lead to strategies which are not economically viable. For example, a numerical example provided in Li et al. (2015b), where no transaction costs are considered, shows the risky asset price undergoing reasonable changes over the course of a month, but the resulting time-consistent MV-optimal analytical solution calls for an almost three-fold increase in the risky asset holdings as the risky asset price declines, only to unwind the entire position again as the risky asset price recovers at the end of the month.

Also, any strategy which allows leverage, even if limited, should take into account that borrowing rates will be larger than lending rates, which will clearly affect any conclusions drawn regarding trading strategies.

Furthermore, the use of a wealth-dependent risk-aversion parameter has been popular in time-consistent MV literature, especially in insurance-related applications, such as Zeng and Li (2011), Wei et al. (2013), Li and Li (2013), as well as Liang and Song (2015)). While arguments in favour of, for example, a risk aversion parameter inversely proportional to wealth appear to be reasonable when considered in the absence of investment constraints (see for example Bjork et al. (2014) and Li and Li (2013)), in the presence of realistic constraints this formulation may have some unintended and undesirable economic consequences from both a risk and a return perspective, as will become evident below.

As a result, in order to ensure that economically viable strategies can be developed and economically reasonable conclusions can be drawn, a number of realistic investment constraints need to be incorporated simultaneously as part of the formulation of the MV optimization problem. Such a comprehensive treatment with realistic investment constraints cannot be expected to yield analytical solutions, and hence a fully numerical solution approach must be used in this case. This is the main focus of this work.

The literature on numerical methods for time-consistent MV portfolio optimization is virtually limited to the case of diffusion dynamics, i.e. Geometric Brownian Motion, for the risky asset, including notable works of Cong and Oosterlee (2016); Wang and Forsyth (2011). However, it is well-documented in the finance literature that jumps are often present in the price processes of risky assets (see, for example, Cont and Tankov (2004); Ramezani and Zeng (2007)). Jump processes permit modelling of non-normal asset returns and fat tails. We focus on jump-diffusions in this work, since previous studies indicate that mean-reverting stochastic volatility processes have a very small effect on the efficient frontier for long term (> 10 years) investors (Ma and Forsyth, 2016). Using a Monte Carlo approach, Cong and Oosterlee (2016) compare pre-commitment and time-consistent policies with leverage and bankruptcy constraints in the case of diffusion dynamics.[1] In the present work, we go a step forward by considering both the continuous and discrete rebalancing versions of the time-consistent MV portfolio optimization problem with jump-diffusion dynamics for the risky asset and realistic investment constraints, such as transaction costs and different borrowing and lending interest rates. Moreover, we also provide a comprehensive comparison between the time-consistency and pre-commitment approaches, not only in terms of the resulting efficient frontiers, but also in terms of the optimal investment policies over time under the above-mentioned realistic context. Furthermore, our use of partial integro-differential equation (PIDE) methods for solution of the optimal control problem

[1] The bankruptcy constraint in (Cong and Oosterlee, 2016) is not quite the same as considered in this work.

allows us to illustrate the strategies in terms of easy-to-interpret heat maps.

Generally speaking, the impulse control approach is suitable for many complex situations in stochastic optimal control (Oksendal and Sulem, 2005). In particular, in the context of pre-commitment MV portfolio optimization under jump diffusion, it has been demonstrated in Dang and Forsyth (2014) that an impulse control formulation of the investment problem is very computationally advantageous. This is because an impulse control formulation can avoid the presence of the control in the integrand of the jump terms, which, in turn, facilitates the use of a fast computational method, such as the FFT, for the evaluation of the integral. In addition, an impulse control formulation also allows for efficient handling of realistic modelling assumptions, such as transaction costs.

For time-consistent MV portfolio optimization with jump-diffusion dynamics, an impulse control approach can also be utilized to potentially achieve similar computational advantages. In the realistic context considered in this work, applying the popular method of Bjork et al. (2016); Bjork and Murgoci (2014), together with relevant results from Oksendal and Sulem (2005), the value function under an impulse control formulation can be shown to satisfy a strongly coupled, nonlinear system of equations, the so-called an extended Hamilton-Jacobi-Bellman (HJB) quasi-integro-variational inequality. This system of equations must be solved numerically, since a closed-form solution for it is not known to exist, except in special cases. However, it is not clear how such a very complex system of equations can be solved effectively numerically. As a result, in this case, the method of Bjork et al. (2016); Bjork and Murgoci (2014) does not appear to result in equations amenable for computational purposes. Hence, for numerical purposes, an alternative formulation of this problem is desirable.

The objective of this paper is two-fold. Firstly, we develop a numerically a computationally efficient partial differential equation (PDE) method for the solution of the time-consistent MV portfolio optimization problem under different types of investment constraints and realistic modelling assumptions. We formulate this problem in such a way as to avoid some of the numerical difficulties resulting from the approach of Bjork et al. (2016); Bjork and Murgoci (2014). Secondly, using actual long-term data, we present a comprehensive study of the impact of simultaneously imposing those investment constraints on the efficient frontier, as well as on the optimal investment strategies, for both the time-consistent and pre-commitment approaches.

The main contributions of this paper are as follows.

- We formulate the time-consistent MV portfolio optimization problem as a system of two-dimensional impulse control problems, with a time-consistency constraint enforced via a discretized version of the dynamic programming principle.

  This approach results in only linear partial integro-differential equations (PIDEs) to solve between intervention times, which is not only numerically simpler than the approach of Bjork et al. (2016); Bjork and Murgoci (2014), but also computationally efficient.

- We study the simultaneous application of realistic investment constraints, including (i) discrete (infrequent) rebalancing of the portfolio, (ii) liquidation in the event of insolvency, (iii) leverage constraints, (iv) different interest rates for borrowing and lending, and (v) transaction costs.

- Since the viscosity solution theory (Crandall et al. (1992)) does not apply in this case, we have no formal proof of convergence of our numerical PDE method. However, we (i) show that our method converges to analytical solutions, where available, and (ii) validate the results from our method using Monte Carlo simulations, where analytical solutions are unavailable.

- Extensive numerical experiments are conducted with model parameters calibrated to real (i.e. inflation adjusted) long-term US market data (89 years), enabling realistic conclusions to be drawn from the results. Through these experiments, the (significant) impact of various modelling assumptions and investment constraints on the MV efficient frontiers are investigated.

We also present a comprehensive comparison study of the time-consistent and pre-commitment MV optimal strategies.

- For the popular case of a wealth-dependent risk aversion parameter in the time-consistent MV literature, our results show that a seemingly reasonable definition of a wealth-dependent risk-aversion parameter, when used in combination with investment and bankruptcy constraints, can result in conclusions that are not economically reasonable. Not only does this finding pose questions about the use of such wealth-dependent risk aversion parameters in existing time-consistent MV literature, but it also highlights the importance of incorporating realistic constraints in investment models.

The remainder of the paper is organized as follows. Section 2 describes the underlying processes and the impulse control approach, and introduce the pre-commitment and time-consistent MV optimization approaches. A numerical algorithm for solving the time-consistency MV portfolio optimization problem is discussed in detail in Section 3. In Section 4, we discuss the localization and numerical techniques, including discrete rebalancing case. Numerical results are presented and discussed in Section 5. Section 6 concludes the paper and outlines possible future work.

2 Formulation

2.1 Underlying processes

We consider the investment-only problem[2] from the perspective of a mean-variance investor/insurer investing in portfolios consisting of just two assets, namely a risky asset and a risk-free asset. The lack of allowance for investment in multiple risky assets may initially appear to be overly restrictive, but we argue that this is not the case, due to the following reasons. Firstly, in the applying the approach presented in this paper, we use a diversified index, rather than a single stock (see Section 5). Secondly, in the available analytical solutions for multi-asset time-consistent MV problems, the composition of the risky asset basket remains relatively stable over time (see for example Zeng and Li (2011)). Finally, investment problems with long time horizons has a strong strategic component - the investor/insurer may be more interested in overall global portfolio shifts from stocks to bonds and vice versa[3], rather than the more secondary questions relating to risky asset basket compositions.

Let $S(t)$ and $B(t)$ respectively denote the amounts (i.e. total dollars) invested in the risky and risk-free asset, at time $t \in [0, T]$, where $T > 0$ is the fixed horizon investment. Define $t^- = \lim_{\epsilon \downarrow 0} (t - \epsilon)$, $t^+ = \lim_{\epsilon \downarrow 0} (t + \epsilon)$, i.e. t^- (resp. t^+) as the instant of time before (resp. after) the (forward) time t. First, consider the risky asset. Let ξ be a random number representing a jump multiplier, with probability density function (pdf) $p(\xi)$. When a jump occurs, $S(t) = \xi S(t^-)$. As a specific example, we consider two jump distributions for ξ, namely the log-normal distribution (Merton, 1976) and the log-double-exponential distribution (Kou, 2002). Specifically, in the former case, $\log \xi$ is normally distributed, so that

$$p(\xi) = \frac{1}{\xi \sqrt{2\pi \widetilde{\gamma}^2}} \exp\left\{ -\frac{(\log \xi - \widetilde{m})^2}{2\widetilde{\gamma}^2} \right\}, \tag{2.1}$$

[2] As noted in the conclusion to this paper, we leave the investment-reinsurance problem for future work.

[3] It is natural for institutions, answerable to their stockholders regarding their chosen investment strategies, to be sensitive to these global trends. As a typical example of an article discussing these trends, see "Global stock optimism drives rotation from bonds into equities", by Kate Allen, which appeared in the Financial Times (FT) on January 16, 2018.

with mean $\widetilde{m}$ and standard deviation $\widetilde{\gamma}$, and $E[\xi] = \exp(\widetilde{m}+\widetilde{\gamma}^2/2)$, where $E[\cdot]$ denotes the expectation operator. In the latter case, $\log \xi$ has an asymmetric double-exponential distribution, so that

$$p(\xi) = \nu\zeta_1 \xi^{-(\zeta_1+1)} \mathbb{I}_{[\xi \geq 1]} + (1-\nu)\,\zeta_2 \xi^{\zeta_2 - 1} \mathbb{I}_{[0 \leq \xi < 1]}. \tag{2.2}$$

Here, $\nu \in [0,1]$, $\zeta_1 > 1$ and $\zeta_2 > 0$, and $\mathbb{I}_{[A]}$ denotes the indicator function of the event A. Given that a jump occurs, ν is the probability of an upward jump, and $(1-\nu)$ is the probability of a downward jump. Furthermore, in this case, we have $E[\xi] = \dfrac{\nu\zeta_1}{\zeta_1 - 1} + \dfrac{(1-\nu)\zeta_2}{\zeta_2 + 1}$.

In the context of pre-commitment MV analysis, the results in (Ma and Forsyth, 2016) indicate that the effects of mean-reverting stochastic volatility are unimportant for long-term (i.e. greater than 10 years) investors. Hence we focus here on the effect of jump processes, as a major source of risk. In the absence of control, i.e. if we do not adjust the amount invested according to our control strategy, the amount S invested in the risky asset is assumed to follow the process

$$\frac{dS(t)}{S(t^-)} = (\mu - \lambda\kappa)\,dt + \sigma dZ + d\left(\sum_{i=1}^{\pi(t)} (\xi_i - 1)\right). \tag{2.3}$$

Here, $\kappa = \mathbb{E}[\xi - 1]$; Z denotes a standard Brownian motion; μ and σ are the real world drift and volatility, respectively; $\pi(t)$ a Poisson process with intensity $\lambda \geq 0$; and ξ_i are i.i.d. random variables having the same distribution as ξ. Moreover, ξ_i, π_t and Z are assumed to all be mutually independent. For later use in the paper, we also define $\kappa_2 = \mathbb{E}\left[(\xi-1)^2\right]$.

It is assumed that the investor can earn a (continuously compounded) rate r_ℓ on cash deposits, and borrow at a rate of $r_b > 0$, with $r_\ell < r_b$. In the absence of control, the dynamics of the amount $B(t)$ invested in the risk-free asset are given by

$$dB(t) = \mathcal{R}(B(t))\,B(t)\,dt, \quad \text{where } \mathcal{R}(B(t)) = r_\ell + (r_b - r_\ell)\,\mathbb{I}_{[B(t)<0]}. \tag{2.4}$$

We make the standard assumption that the real world drift rate μ of S is strictly greater than r_ℓ. Since there is only one risky asset, it is never optimal in an MV setting to short stock, i.e. $S(t) \geq 0$, $t \in [0,T]$. However, we do allow short positions in the risk-free asset, i.e. it is possible that $B(t) < 0$, $t \in [0,T]$.

In some of the examples considered in this paper, we assume that, in the absence of the control, the dynamics for $S(t)$ follows GBM. This is implemented by suppressing any possible jumps in (2.3), i.e. by setting the intensity parameter λ to zero.

2.2 Dynamics of the controlled system

We denote by $X(t) = (S(t), B(t))$, $t \in [0,T]$, the multi-dimensinal controlled underlying process, and by $x = (s,b)$ the state of the system. Furthermore, the liquidation value of the (controlled) wealth, denoted by $W(t)$. We note that $W(t)$ may include liquidation costs (see (2.8)).

Let $(\mathcal{F}_t)_{t\geq 0}$ be the natural filtration associated with the wealth process $\{W(t) : t \in [0,T]\}$. We use $C_t(\cdot)$ to denote the control, representing a strategy as a function of the underlying state, computed at time $t \in [0,T]$, i.e. $C_t(\cdot) : (X(t), t) \mapsto C_t = C(X(t), t)$, for the time interval $[t,T]$. Following Dang and Forsyth (2014), we make use of impulse controls, which allows for efficient handling of jumps, as well as other realistic modelling assumptions, such as transaction costs. A generic impulse control $\mathcal{C}_t$ is defined as a double, possibly finite, sequence (Oksendal and Sulem, 2005)

$$\mathcal{C}_t \;=\; \{t_1, , t_2, \ldots, t_n \;;\; \eta_1, \eta_2, \ldots, \eta_n, \ldots\}_{n \leq n_{\max}} = \{\{t_n, \eta_n\}\}_{n \leq n_{\max}}, \quad n_{\max} \leq \infty. \tag{2.5}$$

Here, intervention times $t \leq t_1 < \ldots < t_{n_{\max}} < T$ are any sequence of $(\mathcal{F}_t)$-stopping times, associated with a corresponding sequence of random variables $(\eta_n)_{n \leq n_{\max}}$ denoting the impulse values, with each

η_n being $\mathcal{F}_{t_n}$-measurable, for all t_n. We denote by $\mathcal{Z}$ the set of admissible impulse values, and by $\mathcal{A}$ the set of admissible impulse controls. For use later in the paper, we denote by $\mathcal{C}_t^* = (\{t_n, \eta_n^*\})_{n \leq n_{\max}}$, $n_{\max} \leq \infty$, the optimal impulse control.

In our context, the intervention time t_n correspond to the re-balancing times of the portfolio, and the impulse η_n corresponds to readjusting the amounts of the stock and bond in the investor's portfolio at time t_n. Recalling definition (2.5), t_n can formally be *any* $(\mathcal{F}_t)$-stopping time. However, in any numerical implementation, we are of course limited to a finite set of pre-specified potential intervention[4] times (see for example equation (3.7) below). In what follows, we will consider both "continuous rebalancing" - see Section 5.2 (where, as $\max_n (t_n - t_{n-1}) \to 0$, we recover the ability to intervene as per definition (2.5)), as well as "discrete rebalancing", where the set of potential intervention times remain fixed - see Section 4.4.

The dynamics of portfolio rebalancing is as follows. Assume that the system is in state $x = (s, b)$ at time t_n^-. We denote by $(S^+(t_n), B^+(t_n)) \equiv (S^+(s, b, \eta_n), B^+(s, b, \eta_n))$ the state of the system immediately after application of the impulse η_n at time t_n. More specifically, we assume that fixed and proportional transaction costs, respectively denoted by $c_1 > 0$ and c_2, where $c_2 \in [0, 1)$, may be imposed on each rebalancing of the portfolio. Applying the impulse η_n at time t_n results in

$$\begin{aligned} B^+(t_n) \equiv B^+(s, b, \eta_n) &= \eta_n, \\ S^+(t_n) \equiv S^+(s, b, \eta_n) &= (s + b) - \eta_n - c_1 - c_2 \left| S^+(s, b, \eta_n) - s \right|, \end{aligned} \tag{2.6}$$

where the transaction costs have been taken into account.

Between intervention times, for $t \in \left[t_n^+, t_{n+1}^-\right]$, the amounts S and B evolve according to the dynamics specified in (2.4) and (2.3), respectively. Specifically,

$$\begin{aligned} \frac{dS(t)}{S(t^-)} &= (\mu - \lambda\kappa)\, dt + \sigma dZ + d\left(\sum_{i=1}^{\pi\left[t_n^+, t_{n+1}^-\right]} (\xi_i - 1) \right), \\ dB(t) &= \mathcal{R}(B(t))\, B(t)\, dt, \qquad t \in \left[t_n^+, t_{n+1}^-\right], \quad n = 0, 1, 2, \ldots, n_{\max} - 1, \end{aligned} \tag{2.7}$$

where $\pi\left[t_n^+, t_{n+1}^-\right]$ denotes the number of jumps in the Poisson process $\pi(t)$ in the time interval $\left[t_n^+, t_{n+1}^-\right]$.

2.3 Admissible portfolios

To include transaction costs, the liquidation value $W(t)$ of the portfolio is defined to be

$$W(t) = W(s, b) = b + \max\left[(1 - c_2)\, s - c_1,\, 0\right], \quad t \in [0, T]. \tag{2.8}$$

We strictly enforce two investment constraints on the *joint* values of S and B, namely a solvency condition and a maximum leverage condition. The solvency condition takes the following form: if insolvent, defined to be the case when $W(s, b) \leq 0$, we require that the position in the risky asset be liquidated, the total remaining wealth be placed in the risk-free asset, and the ceasing of all subsequent trading activities. More formally, we define a solvency region $\mathcal{N}$ and an insolvency or bankruptcy region $\mathcal{B}$ as follows:

$$\mathcal{N} = \{(s, b) \in \Omega^\infty :\ W(s, b) > 0\}, \tag{2.9}$$

$$\mathcal{B} = \{(s, b) \in \Omega^\infty :\ W(s, b) \leq 0\}, \tag{2.10}$$

[4] As is evident from Algorithm 3.1, the investor is not forced to rebalance the portfolio at a potential intervention time t_n, but can retain existing investments unchanged if it is optimal to do so, which is equivalent to "non-intervention".

where

$$\Omega^{\infty} = [0, \infty) \times (-\infty, \infty) \,. \tag{2.11}$$

The solvency condition can then be stated as

$$\text{If } (s,b) \in \mathcal{B} \text{ at } t_n^- \quad \Rightarrow \quad \begin{cases} \text{we require } \ (S^+(t_n) = 0, B^+(t_n) = W(s,b)) \,, \\ \text{and remains so for } \forall t \in [t_n, T]. \end{cases} \tag{2.12}$$

The investors net debt then accumulates at the borrowing rate. It is noted that due to the S-dynamics (2.3), the wealth can jump into the bankruptcy region (regardless of whether we trade continuously or not).

We also constrain the leverage ratio, i.e. at each intervention time t_n, the investor must select an allocation satisfying

$$\frac{S^+(t_n)}{S^+(t_n) + B^+(t_n)} < q_{\max} \tag{2.13}$$

for some positive constant $q_{\max}$, typically in the range $[1.0, 2.0]$.

2.4 Mean-variance (MV) optimization

Let $E_{\mathcal{C}_t}^{x,t}[W(T)]$ and $Var_{\mathcal{C}_t}^{x,t}[W(T)]$ denote the mean and variance of the liquidation value of the terminal wealth, respectively, given the state $x = (s, b)$ at time t and using impulse control $\mathcal{C}_t \in \mathcal{A}$ over $[t, T]$.

2.4.1 Pre-commitment

Using the standard linear scalarization method for multi-criteria optimization problems (Yu, 1971), we define the (time-t) *pre-commitment MV* (PCMV) problem by

$$(PCMV_t(\rho)): \qquad \sup_{\mathcal{C}_t \in \mathcal{A}} \left(E_{\mathcal{C}_t}^{x,t}[W(T)] - \rho Var_{\mathcal{C}_t}^{x,t}[W(T)] \right), \quad \rho > 0. \tag{2.14}$$

Here, the scalarization parameter ρ reflects the investor's level of risk aversion. The MV "efficient frontier" is defined as the following set of points in $\mathbb{R}^2$:

$$\left\{ \left(\sqrt{Var_{\mathcal{C}_0^*}^{x_0,0}[W(T)]}, \ E_{\mathcal{C}_0^*}^{x_0,0}[W(T)] \right) : \ \rho > 0 \right\}, \tag{2.15}$$

traced out by solving (2.14) for each $\rho > 0$. In other words, given a fixed level of risk aversion, an "efficient" portfolio, i.e. any point in the set (2.15), cannot be improved upon in the MV sense, using any other admissible strategy in $\mathcal{A}$.

There are two important issues related to the pre-commitment MV problem (2.14). First, since variance does not satisfy the smoothing property of conditional expectation, dynamic programming cannot be used directly to solve (2.14). To overcome this challenge, an embedding approach for solving (2.14) is proposed in Li and Ng (2000); Zhou and Li (2000) - see Dang and Forsyth (2014); Dang et al. (2016); Wang and Forsyth (2010) for the numerical treatment of the problem as well as a discussion of technical issues.

The second issue is the "time-inconsistency" of the resulting optimal control (see Bjork et al. (2016); Bjork and Murgoci (2014)). To this end, with a slight abuse of notation, we denote by $\mathcal{C}_{t,u}^*$ the optimal control for problem $PCMV_t(\rho)$ computed at time t for a fixed time $u \in [t, T]$. For the pre-commitment approach, the "time-inconsistency" phenomenon means that, in general,

$$\mathcal{C}_{t,u}^* \neq \mathcal{C}_{t',u}^*, \quad t' > t, \quad u \in \left[t', T\right]. \tag{2.16}$$

However, we note that the issue of time-inconsistency is, to some extent, a matter of perspective. This is because it can be shown that problem (2.14) is equivalent to target-based optimization, where a quadratic loss function is minimized - see Vigna (2014). From this perspective, for a fixed target, the optimal control *is* time consistent, but it remains time-inconsistent when viewed from the perspective of the original problem (2.14).

2.4.2 Time-consistent approach

In contrast, the *time-consistent MV* (TCMV) optimization problem - first studied in Basak and Chabakauri (2010) and later substantially expanded and generalized in Bjork et al. (2016); Bjork and Murgoci (2014); Hu et al. (2012) - enforces equality, instead of inequality, in (2.16). Specifically, a "time-consistency constraint" is imposed on (2.14), and this gives the time-consistent MV (TCMV)

$$\left(TCMV_t\left(\rho\right)\right): \qquad V\left(s,b,t\right) = \sup_{\mathcal{C}_t\in\mathcal{A}}\left(E_{\mathcal{C}_t}^{x,t}\left[W\left(T\right)\right]-\rho Var_{\mathcal{C}_t}^{x,t}\left[W\left(T\right)\right]\right), \tag{2.17}$$

$$\text{s. t.} \quad \mathcal{C}_{t,u}^{*} = \mathcal{C}_{t',u}^{*}, \quad \text{for all } t'\geq t \text{ and } u\geq t'. \tag{2.18}$$

Dynamic programming can be applied directly to (2.17)-(2.18) to compute optimal controls and the TCMV efficient frontier. See, for example Wang and Forsyth (2011), for the pure-diffusion case.

Since the constrained optimization problem (2.17)-(2.18) always leads to MV outcomes inferior to, or at most, the same as, those of the unconstrained optimization problem (2.14), a natural question is: what makes time-consistent MV optimization potentially attractive? As discussed in the introduction, the pre-commitment approach may not be feasible in institutional settings, while, on the contrary, the time-consistent approach is typically popular in these settings. However, it should be noted that neither the pre-commitment nor the time-consistent approach is "better" in some objective sense - see Vigna (2016, 2017) for a discussion of a number of subtle issues involved.

Remark 2.1. (Game-theoretic perspective; notion of optimality). In Bjork and Murgoci (2014), the terminology "equilibrium" control is used as opposed to "optimal" control, since the time-consistent optimal control $\mathcal{C}_t^*$ satisfies the conditions of a subgame perfect Nash equilibrium control. We will follow the example of Basak and Chabakauri (2010); Cong and Oosterlee (2016); Li and Li (2013); Wang and Forsyth (2011) and retain the terminology "optimal" (time-consistent) control for simplicity.

3 Algorithm development

For subsequent use, we write the value function $V(s,b,t)$ of the time-consistent problem (2.17)-(2.18) in terms of two auxiliary functions $U\left(s,b,t\right)$ and $Q\left(s,b,t\right)$ as follows

$$V(s,b,t) = U\left(s,b,t\right) - \rho Q\left(s,b,t\right) + \rho(U\left(s,b,t\right))^2, \tag{3.1}$$

where

$$U\left(s,b,t\right) = E_{\mathcal{C}_t^*}^{x,t}\left[W\left(T\right)\right], \tag{3.2}$$

$$Q\left(s,b,t\right) = E_{\mathcal{C}_t^*}^{x,t}\left[\left(W\left(T\right)\right)^2\right], \tag{3.3}$$

where, it is implicitly understood hereafter that $\mathcal{C}_t^*$ is the optimal control for the $TCMV_t\left(\rho\right)$ problem. We also define the following operators, applied to an appropriate test function f:

$$\mathcal{L}f\left(s,b,t\right) = \left(\mu-\lambda\kappa\right)sf_s + \mathcal{R}\left(b\right)bf_b + \frac{1}{2}\sigma^2s^2f_{ss} - \lambda f, \tag{3.4}$$

$$\mathcal{J}f\left(s,b,t\right) = \lambda\int_0^\infty f\left(\xi s,b,t\right)p\left(\xi\right)d\xi. \tag{3.5}$$

We now primarily focus on the continuous re-balancing case. The discrete rebalancing case is discussed in Subsection (4.4).

Fix an arbitrary point in time $t \in [0, T)$, and assume we are in state $x = (s, b)$ at time t^-. We define the intervention operator, a fundamental object in impulse control problems (Oksendal and Sulem, 2005), applied to the value function V of the time-consistent problem (2.17)-(2.18) as

$$\mathcal{M}V(s, b, t) = \sup_{\eta \in \mathcal{Z}} \left[V\left(S^+(s, b, \eta), B^+(s, b, \eta), t\right) \right], \tag{3.6}$$

where $S^+(\cdot)$ and $B^+(\cdot)$ are defined in (2.6).

In analogy to the case of continuous controls, where an extended HJB system of equations is obtained (see Bjork et al. (2016)), as discussed in the Introduction, in our case, the techniques of Bjork et al. (2016); Bjork and Murgoci (2014) results in an extended HJB quasi-integrovariational inequality - a strongly coupled, nonlinear system of equations that needs to solve simultaneously to obtain the value function. Under realistic modelling assumptions and investment constraints, a closed-form solution for this highly complex system of equations is not known to exist, except for very special cases, and hence a numerical method must be used. However, it is not clear how such a highly complex system of equations can be solved effectively numerically for practical purposes.

To overcome the above-mentioned hurdle, we choose to enforce the dynamic programming principle on the discretized time variable, i.e. the time-consistency constraint (2.18) is enforced on a set of discrete intervention times obtained from discretizing the time variable. The intervention operator $\mathcal{M}$, defined in (3.6), is applied across each of these times As shown later, this approach results in only linear partial integro-differential equations to solve between intervention times. Furthermore, when combined with a semi-Lagrangian timestepping scheme, we just have a set of one-dimensional PIDE in the s-variable to solve between intervention times. As a result, our approach is not only numerically simpler than the approach of Bjork et al. (2016); Bjork and Murgoci (2014), but also computationally effective.

3.1 Recursive relationships

We consider the following uniform partition of the time interval $[0, T]$

$$\mathcal{T}_{n_{\max}} = \{t_n \mid t_n = n\Delta t\}, \quad \Delta t = T/n_{\max}, \quad \Delta t = C_1 h, \tag{3.7}$$

where C_1 is positive and independent of the discretization parameter $h > 0$. In the limit as $h \to 0$, we shall demonstrate via numerical experiments that, at least for some known cases, the numerical solution of the time-discretized formulation converges to the closed-form solution of the continuous time formulation.

To avoid heavy notation, we now introduce the following notational convention: any admissible impulse control $\mathcal{C} \in \mathcal{A}$ will be written as the set of impulses

$$\mathcal{C} = \{\eta_n \in \mathcal{Z} : \ n = 0, \ldots, n_{\max}\}, \tag{3.8}$$

where the corresponding set of (discretized) intervention times is implicitly understood to be $\{t_n\}_{n=0}^{n_{\max}}$. Given an impulse control $\mathcal{C}$ as in (3.8), we also define the control $\mathcal{C}_n \equiv \mathcal{C}_{t_n} \subseteq \mathcal{C}$, $n = 0, \ldots, n_{\max}$, as the subset of impulses (and, implicitly, corresponding intervention times) of $\mathcal{C}$ applicable to the time interval $[t_n, T]$:

$$\mathcal{C}_n = \{\eta_n, \ldots, \eta_{n_{\max}}\} \subseteq \mathcal{C} = \{\eta_0, \ldots, \eta_{n_{\max}}\}. \tag{3.9}$$

Subsequently, we use

$$\mathcal{C}_n^* = \left\{\eta_n^*, \ldots, \eta_{n_{\max}}^*\right\} \tag{3.10}$$

to denote the optimal impulse control to the problem $(TCMV_{t_n}(\rho))$ defined in (2.17)-(2.18).

With this time discretization and notational conventions, for a given scalarization parameter $\rho > 0$ and an intervention time t_n, we define the scalarized time-consistent MV problem $(TCMV_{t_n}(\rho))$ as follows:

$$
\begin{aligned}
(TCMV_{t_n}(\rho)): \qquad V(s,b,t_n) &= \sup_{\mathcal{C}_n \in \mathcal{A}} \left(E_{\mathcal{C}_n}^{x,t_n}[W(T)] - \rho Var_{\mathcal{C}_n}^{x,t_n}[W(T)] \right) && (3.11)\\
\text{s.t. } \mathcal{C}_n &= \left\{ \eta_n, \mathcal{C}_{n+1}^* \right\} \coloneqq \left\{ \eta_n, \eta_{n+1}^*, \ldots, \eta_{n_{\max}-1}^*, \eta_{n_{\max}}^* \right\} && (3.12)\\
&\quad \text{where } \mathcal{C}_{n+1}^* \text{ is optimal for problem } \left(TCMV_{t_{n+1}}(\rho) \right).
\end{aligned}
$$

We note that the definition of (3.11)-(3.12) agrees conceptually with the continuous-time definition given by (2.17)-(2.18), but is more convenient from a computational perspective. The particular form of the time-consistency constraint in (3.12) is a discretized equivalent of the constraint in (2.18), since, given the optimal impulse control $\mathcal{C}_{n+1}^* = \left\{ \eta_{n+1}^*, \ldots, \eta_{n_{\max}}^* \right\}$ of problem $\left(TCMV_{t_{n+1}}(\rho) \right)$ applicable to the time period $[t_{n+1}, T]$, any *arbitrary* admissible impulse control $\mathcal{C}_n \in \mathcal{A}$ will necessarily be of the form

$$
\mathcal{C}_n = \left\{ \eta, \eta_{n+1}^*, \ldots, \eta_{n_{\max}}^* \right\} = \left\{ \eta, \mathcal{C}_{n+1}^* \right\} \tag{3.13}
$$

for some admissible impulse value $\eta \in \mathcal{Z}$ applied at time t_n.

We use the notation $E_\eta^{x,t_n}[\cdot]$ to indicate that the expectation is evaluated using an (arbitrary) impulse value $\eta \in \mathcal{Z}$ at time t_n, with the implied application of $\mathcal{C}_{n+1}^*$ over the time interval $[t_{n+1}, T]$. We note that, given $X\left(t_{n+1}^-\right) = \left(S\left(t_{n+1}^-\right), B\left(t_{n+1}^-\right)\right)$ at time t_{n+1}^-, we have the following recursive relationships for $U(s,b,t_n)$ and $Q(s,b,t_n)$:

$$
\begin{aligned}
U(s,b,t_n) &= E_{\eta_n^*}^{x,t_n}\left[U\left(S\left(t_{n+1}^-\right), B\left(t_{n+1}^-\right), t_{n+1}\right)\right], && (3.14)\\
Q(s,b,t_n) &= E_{\eta_n^*}^{x,t_n}\left[Q\left(S\left(t_{n+1}^-\right), B\left(t_{n+1}^-\right), t_{n+1}\right)\right], && (3.15)
\end{aligned}
$$

where, as defined previously in (3.10), η_n^* is the optimal impulse value for time t_n. For the special case of $t_{n_{\max}} = T$, we have

$$
\begin{aligned}
U(s,b,T) &= U(s,b,t_{n_{\max}}) = W(s,b), && (3.16)\\
Q(s,b,T) &= Q(s,b,t_{n_{\max}}) = (W(s,b))^2. && (3.17)
\end{aligned}
$$

We similarly obtain a recursive relationship for the value function (3.11)

$$
\begin{aligned}
V(s,b,t_n) = \sup_{\eta \in \mathcal{Z}} \Big\{ & E_\eta^{x,t_n}\left[U\left(S\left(t_{n+1}^-\right), B\left(t_{n+1}^-\right), t_{n+1}\right)\right] - \rho E_\eta^{x,t_n}\left[Q\left(S\left(t_{n+1}^-\right), B\left(t_{n+1}^-\right), t_{n+1}\right)\right] \\
& + \rho\left(E_\eta^{x,t_n}\left[U\left(S\left(t_{n+1}^-\right), B\left(t_{n+1}^-\right), t_{n+1}\right)\right]\right)^2 \Big\}, \qquad (3.18)
\end{aligned}
$$

where, for the special case of $t_{n_{\max}}$, we have $V(s,b,t_{n_{\max}}) = V(s,b,T) = W(s,b)$. This is effectively the discretized version of the intervention operator $\mathcal{M}$, defined in (3.6).

Assume that $E_\eta^{x,t_n}[\cdot]$ is a bounded, upper semi-continuous function of the admissible impulse value η. If we can determine $U\left(S\left(t_{n+1}^-\right), B\left(t_{n+1}^-\right), t_{n+1}\right)$ and $Q\left(S\left(t_{n+1}^-\right), B\left(t_{n+1}^-\right), t_{n+1}\right)$, then

$$
\begin{aligned}
\eta_n^* \in \arg\max_{\eta \in \mathcal{Z}} \Big\{ & E_{\eta_n}^{x,t_n}\left[U\left(S\left(t_{n+1}^-\right), B\left(t_{n+1}^-\right), t_{n+1}\right)\right] - \rho E_{\eta_n}^{x,t_n}\left[Q\left(S\left(t_{n+1}^-\right), B\left(t_{n+1}^-\right), t_{n+1}\right)\right] \\
& + \rho\left(E_{\eta_n}^{x,t_n}\left[U\left(S\left(t_{n+1}^-\right), B\left(t_{n+1}^-\right), t_{n+1}\right)\right]\right)^2 \Big\}. \qquad (3.19)
\end{aligned}
$$

Relations (3.14)-(3.19) form the basis for a recursive algorithm to determined the value function and the optimal impulse value.

3.2 Computation of expectations

We now introduce the change of variable $\tau = T - t$, and let

$$\bar{U}(s,b,\tau) = U(s,b,T-t), \quad \bar{Q}(s,b,\tau) = Q(s,b,T-t), \quad \bar{V}(s,b,\tau) = V(s,b,T-t), \tag{3.20}$$

and hence (3.1) becomes

$$\bar{V}(s,b,\tau) = \bar{U}(s,b,\tau) - \rho\bar{Q}(s,b,\tau) + \rho\left(\bar{U}(s,b,\tau)\right)^2 \tag{3.21}$$

In terms of τ, time grid (3.7) now becomes

$$\{\tau_n = T - t_{n_{\max}-n} : \ n = 0,1,\ldots,n_{\max}\}. \tag{3.22}$$

Next, we define the following "candidate" expectation values at the rebalancing time τ_n under an arbitrary impulse $\eta \in \mathcal{Z}$:

$$\hat{U}_\eta^n(s,b) = E_\eta^{x,\tau_n}\left[\bar{U}\left(S\left(\tau_{n-1}^+\right), B\left(\tau_{n-1}^+\right), \tau_{n-1}^+\right)\right], \tag{3.23}$$

$$\hat{Q}_\eta^n(s,b) = E_\eta^{x,\tau_n}\left[\bar{Q}\left(S\left(\tau_{n-1}^+\right), B\left(\tau_{n-1}^+\right), \tau_{n-1}^+\right)\right]. \tag{3.24}$$

To handle the computation of expectations in (3.23) and (3.24), we proceed as follows. For solvent portfolios, i.e. $(s,b) \in \mathcal{N}$, we first solve the following associated two PIDEs from τ_{n-1}^+ to τ_n^- (Oksendal and Sulem, 2005)

$$\Psi_\tau(s,b,\tau) - \mathcal{L}\Psi(s,b,\tau) - \mathcal{J}\Psi(s,b,\tau) = 0 \qquad (s,b,\tau) \in \mathcal{N} \times \left(\tau_{n-1}^+, \tau_n^-\right] \tag{3.25}$$

$$\text{with initial condition } \Psi\left(s,b,\tau_{n-1}^+\right) = \bar{U}(s,b,\tau_{n-1}) \tag{3.26}$$

and

$$\Phi_\tau(s,b,\tau) - \mathcal{L}\Phi(s,b,\tau) - \mathcal{J}\Phi(s,b,\tau) = 0 \qquad (s,b,\tau) \in \mathcal{N} \times \left(\tau_{n-1}^+, \tau_n^-\right] \tag{3.27}$$

$$\text{with initial condition } \Phi\left(s,b,\tau_{n-1}^+\right) = \bar{Q}(s,b,\tau_{n-1}) \tag{3.28}$$

where, for the special case of $\tau_0 = 0$, we have

$$\bar{U}(s,b,0) = W(s,b), \quad \bar{Q}(s,b,0) = (W(s,b))^2. \tag{3.29}$$

Here, the operators $\mathcal{L}$ and $\mathcal{J}$ in the PDEs (3.25) and (3.27) are defined in (3.4) and (3.5), respectively. Then, for a given arbitrary impulse $\eta \in \mathcal{Z}$, we obtain the "candidate" expectation values $\hat{U}_\eta^n(s,b)$ and $\hat{Q}_\eta^n(s,b)$ by

$$\hat{U}_\eta^n(s,b) = \Psi\left(S\left(\tau_n^+\right), B\left(\tau_n^+\right), \tau_n^-\right), \tag{3.30}$$

$$\hat{Q}_\eta^n(s,b) = \Phi\left(S\left(\tau_n^+\right), B\left(\tau_n^+\right), \tau_n^-\right), \tag{3.31}$$

where $B(\tau_n^+) = \eta$ and $S(\tau_n^+) = (s+b) - \eta - c_1 - c_2 \cdot |S(\tau_n^+) - s|$, as per (2.6), subject to the leverage constraint (2.13). Finally, using (3.30)-(3.31), we can find the optimal impulse value η_n^* via $\eta_n^* \in \arg\max_{\eta\in\mathcal{Z}} \left\{\hat{U}_\eta^n(s,b) - \rho\hat{Q}_\eta^n(s,b) + \rho\left(\hat{U}_\eta^n(s,b)\right)^2\right\}$.

For insolvent portfolios, i.e. $(s,b) \in \mathcal{B}$, the solvency constraint (2.12) results in enforced liquidation. This is captured by a Dirichlet condition

$$\bar{U}\left(s,b,\tau_n^-\right) = \bar{U}\left(0, W(s,b)e^{\mathcal{R}(s+b)\tau_n}, 0\right),$$

$$\bar{Q}\left(s,b,\tau_n^-\right) = \bar{Q}\left(0, W(s,b)e^{\mathcal{R}(s+b)\tau_n}, 0\right), \qquad (s,b) \in \mathcal{B}. \tag{3.32}$$

In Algorithm 3.1, we present a recursive algorithm for the time-consistent MV ($TCMV_n(\rho)$) for a fixed $\rho > 0$.

Algorithm 3.1 Recursive algorithm to solve $(TCMV_n(\rho))$ for a fixed $\rho > 0$.

1: set $\bar{U}(s,b,0) = W(s,b)$ and $\bar{Q}(s,b,0) = (W(s,b))^2$;
2: **for** $n = 1, \ldots, n_{\max}$ **do**
3: **if** $(s,b) \in \mathcal{B}$ **then**
4: enforce the solvency constraint (2.12) via (3.32) to obtain $\bar{U}(s,b,\tau_n)$ and $\bar{Q}(s,b,\tau_n)$;
5: **else**
6: solve (3.25)-(3.26) and (3.27)-(3.28) from τ_{n-1}^+ to τ_n^- to obtain $\Psi(s,b,\tau_n^-)$ and $\Phi(s,b,\tau_n^-)$;
7: **for** each $\eta \in \mathcal{Z}$ **do**
8: set $B^+ = \eta$ and $S^+ = s + b - \eta - c_1 - c_2 \cdot |S^+ - s|$ as per (2.6), subject to the leverage constraint (2.13);
9: compute $\hat{U}_\eta^n(s,b) = \Psi(S^+, B^+, \tau_n^-)$ and $\hat{Q}_\eta^n(s,b) = \Phi(S^+, B^+, \tau_n^-)$;
10: **end for**
11: find $\eta_n^* \in \underset{\eta \in \mathcal{Z}}{\arg\max} \left\{ \hat{U}_\eta^n(s,b) - \rho \hat{Q}_\eta^n(s,b) + \rho \left(\hat{U}_\eta^n(s,b) \right)^2 \right\}$;
12: set $\bar{U}(s,b,\tau_n) = \hat{U}_{\eta_n^*}^n(s,b)$ and $\bar{Q}(s,b,\tau_n) = \hat{Q}_{\eta_n^*}^n(s,b)$;
13: **end if**
14: **end for**
15: return $V(s,b,\tau_{n_{\max}}) = \bar{U}(s,b,\tau_{n_{\max}}) - \rho \bar{Q}(s,b,\tau_{n_{\max}}) + \rho(\bar{U}(s,b,\tau_{n_{\max}}))^2$;

Remark 3.1. (Convergence of numerical solution). Since the viscosity solution theory (Crandall et al. (1992)) does not apply in this case, we have no proof that Algorithm 3.1 converges to an appropriately defined (weak) solution of the corresponding extended HJB quasi-integrovariational inequality in the limit as $\Delta\tau \to 0$. However, we can show, as in Cong and Oosterlee (2016); Wang and Forsyth (2011), that our numerical solution converges to known analytical solutions available in special cases. Where no analytical solutions are available, the numerical PDE results are validated using Monte Carlo simulation.

4 Localization

4.1 Semi-Lagrangian timestepping scheme

Recall the definition of the operator $\mathcal{L}$, defined in (3.4). We observe that the PIDEs (3.25) and (3.27) for $\Psi(s,b,\tau)$ and $\Phi(s,b,\tau)$, respectively, that need to be solved in Step 6 in Algorithm 3.1. involves partial derivatives with respect to both s and b. Direct implementation would be therefore computationally expensive.

With this in mind, we introduce the semi-Lagrangian timestepping scheme proposed in Dang and Forsyth (2014). The intuition behind the the semi-Lagrangian timestepping scheme is that, instead of obtaining the PIDEs by modelling the change (via Ito's lemma) in a test function $f(S(\tau), B(\tau), \tau)$ with both S and B varying, we consider the Lagrangian derivative along the trajectory where B is held fixed over the length of the timestep. Specifically, we model the change in $f(S(\tau), B(\tau), \tau)$ with $(S(\tau), B(\tau) = b)$ for $\tau \in \left[\tau_{n-1}^+, \tau_n^-\right]$, with interest paid only at the end of the timestep, i.e. at time τ_n, at which time the amount in the risk-free asset would jump to $b \cdot \exp\{\mathcal{R}(b)\Delta\tau\}$, reflecting the settlement (payment or receipt) of interest due for the time interval $[\tau_{n-1}, \tau_n]$. Along this trajectory, the partial derivative of the test function $f(s,b,\tau)$ with respect to the b-variable is zero, resulting in a decoupling of the PIDE for every value of the b-variable.

We emphasize that the above argument is an intuitive explanation of the semi-Lagrangian scheme. In fact, we can prove rigorously that in the limit as $\Delta\tau \to 0$, this treatment converges to the case

where interest is paid continuously.[5] Moreover, this approach is also valid for discrete rebalancing, regardless of whether the interest is paid continuously or discretely.

Applying this reasoning to the two PIDEs (3.25) and (3.27), we have

$$\Psi_b(s,b,\tau) = \Phi_b(s,b,\tau) = 0, \qquad (s,b,\tau) \in \mathcal{N} \times \left(\tau_{n-1}^+, \tau_n^-\right],$$

and we can replace the operator $\mathcal{L}$ in the PDEs (3.25) and (3.27) by the operator $\mathcal{P}$ defined as

$$\mathcal{P}f(s,b,t) = \quad (\mu - \lambda\kappa) s f_s + \tfrac{1}{2}\sigma^2 s^2 f_{ss} - \lambda f. \tag{4.1}$$

Therefore, instead of solving a *two*-dimensional PDE in space variables (s,b) for both Ψ and Φ, we now solve, for each discrete value of b, two *one*-dimensional PIDEs (in a single space variable s):

$$\begin{aligned} \Psi_\tau(s,b,\tau) - \mathcal{P}\Psi(s,b,\tau) - \mathcal{J}\Psi(s,b,\tau) &= 0, \qquad (s,b,\tau) \in \mathcal{N} \times \left(\tau_{n-1}^+, \tau_n^-\right] \\ \text{with initial condition } \Psi\left(s,b,\tau_{n-1}^+\right) &= \bar{U}(s,b,\tau_{n-1}), \end{aligned} \tag{4.2}$$

and

$$\begin{aligned} \Phi_\tau(s,b,\tau) - \mathcal{P}\Phi(s,b,\tau) - \mathcal{J}\Phi(s,b,\tau) &= 0, \qquad (s,b,\tau) \in \mathcal{N} \times \left(\tau_{n-1}^+, \tau_n^-\right] \\ \text{with initial condition } \Phi\left(s,b,\tau_{n-1}^+\right) &= \bar{Q}(s,b,\tau_{n-1}). \end{aligned} \tag{4.3}$$

The second consequence of semi-Lagrangian timestepping is that the calculation of the value of $S(\tau_n^-)$, used in computing $\hat{U}_\eta^n(s,b)$ and $\hat{Q}_\eta^n(s,b)$ as per (3.30) and (3.31), has to be adjusted to reflect the payment of interest at time τ_n:

$$S\left(\tau_n^+\right) = \left(s + be^{\mathcal{R}(b)\Delta\tau}\right) - \eta - c_1 - c_2 \cdot \left|S\left(\tau_n^+\right) - s\right|. \tag{4.4}$$

4.2 Localization

Each set of PIDEs (4.2) - (4.3), together with the Dirichlet conditions (3.32), are to be solved in the domain $(s,b,\tau) \in \Omega^\infty \equiv [0,\infty) \times (-\infty,+\infty) \times [\tau_{n-1}^+, \tau_n^-]$. For computational purposes, we localize this domain to the set of points

$$(s,b,\tau) \in \Omega \times [\tau_{n-1}^+, \tau_n^-] = [0, s_{\max}) \times [-b_{\max}, b_{\max}] \times [\tau_{n-1}^+, \tau_n^-],$$

where $s_{\max}$ and $b_{\max}$ are sufficiently large positive numbers. Let $s^* < s_{\max}$. Following Dang and Forsyth (2014), we define the following sub-computational domains

$$\Omega_{s^*} = (s^*, s_{\max}] \times [-b_{\max}, b_{\max}], \tag{4.5}$$

$$\Omega_{s_0} = \{0\} \times [-b_{\max}, b_{\max}], \tag{4.6}$$

$$\Omega_{\mathcal{B}} = \{(s,b) \in \Omega \setminus \Omega_{s^*} \setminus \Omega_{s_0} : W(s,b) \le 0\}, \tag{4.7}$$

$$\Omega_{in} = \Omega \setminus \Omega_{s^*} \setminus \Omega_{s_0} \setminus \Omega_{\mathcal{B}}, \tag{4.8}$$

$$\Omega_{b_{\max}} = (0, s^*] \times \left[-b_{\max}e^{r_{\max}T}, -b_{\max}\right) \cup \left(b_{\max}, b_{\max}e^{r_{\max}T}\right], \tag{4.9}$$

where $r_{\max} = \max(r_b, r_\ell)$. Following the steps in Dang and Forsyth (2014), we have the following localized problem for Ψ:

$$\begin{aligned} \Psi_\tau(s,b,\tau) - \mathcal{P}\Psi(s,b,\tau) - \mathcal{J}_\ell\Psi(s,b,\tau) &= 0, \quad (s,b,\tau) \in \Omega_{in} \times \left[\tau_{n-1}^+, \tau_n^-\right], \\ \Psi_\tau(s,b,\tau) - \mu\Psi(s,b,\tau) &= 0, \quad (s,b,\tau) \in \Omega_{s^*} \times \left[\tau_{n-1}^+, \tau_n^-\right], \\ \Psi(s,b,\tau) - U(0,b,\tau_{n-1}) &= 0, \quad (s,b,\tau) \in \Omega_{s_0} \times \left[\tau_{n-1}^+, \tau_n^-\right], \\ \Psi(s,|b| > |b_{\max}|,\tau) - \frac{|b|}{b_{\max}}\Psi(s,\operatorname{sgn}(b) b_{\max},\tau) &= 0, \quad (s,b,\tau) \in \Omega_{b_{\max}} \times \left[\tau_{n-1}^+, \tau_n^-\right], \\ \text{with } \Psi(s,b,\tau = \tau_{n-1}) - U(s,b,\tau_{n-1}) &= 0, \quad (s,b) \in \Omega. \end{aligned} \tag{4.10}$$

[5] See Dang and Forsyth (2014) for the consistency proof in the context of the pre-commitment mean-variance problem.

while the localized problem for Φ is as follows:

$$
\begin{aligned}
\Phi_\tau(s,b,\tau) - \mathcal{P}\Phi(s,b,\tau) - \mathcal{J}_\ell\Phi(s,b,\tau) &= 0, \quad (s,b,\tau) \in \Omega_{in} \times \left[\tau_{n-1}^+, \tau_n^-\right] \\
\Phi_\tau(s,b,\tau) - \left[2\mu + \sigma^2 + \lambda\kappa_2\right]\Phi(s,b,\tau) &= 0, \quad (s,b,\tau) \in \Omega_{s^*} \times \left[\tau_{n-1}^+, \tau_n^-\right], \\
\Phi(s,b,\tau) - Q(0,b,\tau_{n-1}) &= 0, \quad (s,b,\tau) \in \Omega_{s_0} \times \left[\tau_{n-1}^+, \tau_n^-\right], \\
\Phi(s,|b| > |b_{\max}|,\tau) - \left(\frac{b}{b_{\max}}\right)^2 \Phi(s,\operatorname{sgn}(b)\, b_{\max},\tau) &= 0, \quad (s,b,\tau) \in \Omega_{b_{\max}} \times \left[\tau_{n-1}^+, \tau_n^-\right] \\
\text{with } \Phi(s,b,\tau = \tau_{n-1}) - Q(s,b,\tau_{n-1}) &= 0, \quad (s,b) \in \Omega. \qquad (4.11)
\end{aligned}
$$

Here,

$$
\mathcal{J}_\ell f(s,b,\tau) = \lambda \int_0^{s_{\max}/s} f(\xi s, b, \tau)\, p(\xi)\, d\xi, \qquad (4.12)
$$

We note that the problem is only solved in Ω. Should any values in $\Omega_{b_{\max}}$ be required, they are deduced from the computed solution in Ω. Some guidelines for choosing $s^*, s_{\max}$ which minimize the effect of the localization error for the jump terms can be found in d'Halluin et al. (2005). We refer the reader to Dang and Forsyth (2014) for relevant details regarding a derivation of the localized problems (4.10)-(4.11).

We solve the localized problems (4.10)-(4.11) using finite differences as described in Dang and Forsyth (2014). Specifically, in addition to the time grid in (3.22), we also introduce nodes, not necessarily equally spaced, in the s-direction $\{s_i : i = 1, \ldots, i_{\max}\}$ and b-direction $\{b_j : j = 1, \ldots, j_{\max}\}$, with $\Delta s_{\max} = \max_i(s_{i+1} - s_i) = C_3 h$ and $\Delta b_{\max} = \max_j(b_{j+1} - b_j) = C_4 h$, where C_3 and C_4 are positive and independent of h. Using the nodes in the b-direction, we define $\mathcal{Z}_h = \{b_j : j = 1, \ldots, j_{\max}\} \cap \mathcal{Z}$ to be the discretization of the admissible impulse space. We use linear interpolation onto the computational grid if the spatial point $\left(s_i, b_j e^{\mathcal{R}(b_j)\Delta\tau}\right)$, arising from the implementation of the semi-Lagrangian timestepping scheme (see Section 4.1), does not correspond to any available grid point.

Central differencing is used as much as possible for the discrete approximation to the operator $\mathcal{P}$ in (4.1), but we require that the scheme be a positive coefficient method (Wang and Forsyth, 2008). The operator $\mathcal{J}_\ell$ in (4.12) is handled using the method described in d'Halluin et al. (2005), which avoids a dense matrix solve (due to the presence of the jump term) by using a fixed-point iteration to solve the discrete equations arising at each b-grid node and timestep.

4.3 Construction of efficient frontier

We assume that the given initial wealth, denoted by $W(t=0) = W_{init}$, is invested in the risk-free asset, so that the time $t = 0$ portfolio is given by $(S(0), B(0)) = (0, W_{init})$. For initial wealth W_{init}, and given the positive discretization parameter h, the goal is the tracing out of the efficient frontier using the scalarization parameter ρ:

$$
\mathcal{Y}_h = \bigcup_{\rho \geq 0} \left(\sqrt{\left(Var_{\mathcal{C}_0^*}^{t=0}\left[W(T)\right]\right)_h}, \left(E_{\mathcal{C}_0^*}^{t=0}\left[W(T)\right]\right)_h \right)_\rho, \qquad (4.13)
$$

where $(\cdot)_h$ refers to a discretization approximation to the expression in the brackets.

This can be achieved as follows. For a fixed value $\rho \geq 0$ in $\{\rho_{\min}, \ldots, \rho_{\max}\} \subset [0, \infty)$, executing Algorithm 3.1 gives us the following quantities:

$$
U_0(W_{init}) \simeq \left(E_{\mathcal{C}_0^*}^{(s=0,b=W_{init}),t=0}\left[W(T)\right]\right)_h, \quad Q_0(W_{init}) \simeq \left(E_{\mathcal{C}_0^*}^{(s=0,b=W_{init}),t=0}\left[(W(T))^2\right]\right)_h,
$$

Using these, we compute the corresponding single point on the efficient frontier $\mathcal{Y}_h$ (4.13):

$$
\left(Var_{\mathcal{C}_0^*}^{t=0}\left[W(T)\right]\right)_h = Q_0(W_{init}) - (U_0(W_{init}))^2, \quad \left(E_{\mathcal{C}_0^*}^{t=0}\left[W(T)\right]\right)_h = U_0(W_{init}). \qquad (4.14)
$$

Remark 4.1. (Complexity) For each timestep, we have to perform i) a local optimization problem to search for the optimal impulse η_n^* at each node, and ii) a time advance step for the two PIDEs (4.10) and (4.11). From the perspective of a complexity analysis, this is similar to the case encountered in Dang and Forsyth (2014), with the exception that there are two PIDEs to be solved for each value of b, instead of one. As a result, the complexity analysis of Dang and Forsyth (2014) holds for the algorithm described here as well. Recalling the positive discretization parameter h in (3.7), we conclude that the total complexity of constructing an efficient frontier is $\mathcal{O}\left(1/h^5\right)$.

4.4 Discrete rebalancing

The formulation of the problem up to this point assumes continuous rebalancing of the portfolio - equivalently, in the discretized setting, the portfolio is rebalanced at every timestep. While the continuous rebalancing treatment is crucial for numerical tests showing convergence to the known closed form solutions (see Section 5.2 below), it is not realistic - and in the presence of transaction costs, it is also not practically feasible.

For the construction of efficient frontiers (see Section 5), we therefore assume discrete rebalancing. That is, the portfolio is only rebalanced at a set of pre-determined intervention times $0 = \tilde{t}_0 \leq \tilde{t}_1 < \ldots < \tilde{t}_{m_{\max}} < T$, where t_0 is the inception of the investment. With the change of variable $\tau = T - t$, the set of intervention times become

$$0 = \tilde{\tau}_0 < \tilde{\tau}_1 < \ldots < \tilde{\tau}_{m_{\max}} = T, \quad m_{\max} < \infty. \tag{4.15}$$

Algorithm 3.1 can easily be modified to handle discrete rebalancing. Specifically, in Step 6, the PIDEs (3.25)-(3.26) and (3.27)-(3.28) are solved from from $\tilde{\tau}_{m-1}^+$ to $\tilde{\tau}_m^-$, $m = 1, \ldots, m_{\max}$, possibly using multiple timesteps for the solution of the corresponding PIDE, to obtain $\Psi\left(s, b, \tilde{\tau}_m^-\right)$ and $\Phi\left(s, b, \tilde{\tau}_m^-\right)$. Other steps of the algorithm remain unchanged. In this case, the complexity of the algorithm for constructing the entire efficient frontier is $\mathcal{O}(1/h^4|\log h|)$.

5 Numerical results

5.1 Empirical data and calibration

In order to obtain the required process parameters, the same data and calibration technique is used as in Dang and Forsyth (2016); Forsyth and Vetzal (2017). The empirical data sources are as follows:

- Risky asset data: Daily total return data covering the period *1926:1 - 2014:12* - which includes dividends and other distributions - from the Center for Research in Security Prices (CRSP), in the form of the VWD index has been used.[6] This is a capitalization-weighted index of all domestic stocks on major US exchanges, with data used dating back to 1926. For calibration purposes, the index is adjusted for inflation prior to the calculation of returns.

- Risk-free rate: The risk-free rate is based on 3-month US T-bill rates for the period *1934:1-2014:12*,[7] augmented by National Bureau of Economic Research (NBER) short-term government bond yields for *1926:1 - 1933:12* [8] to incorporate the effect of the 1929 crash. More specifically, a T-bill index is created, inflation-adjusted, then a sample average of the monthly returns is calculated, and annualized to obtain the constant risk-free rate estimate r.

[6] More specifically, results presented here were calculated based on data from Historical Indexes, ©2015 Center for Research in Security Prices (CRSP), The University of Chicago Booth School of Business. Wharton Research Data Services was used in preparing this article. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers.

[7] See `http://research.stlouisfed.org/fred2/series/TB3MS`.

[8] See `http://www.nber.org/databases/macrohistory/contents/chapter13.html`.

- Inflation: In order to adjust the time series for inflation, the annual average CPI-U index (inflation for urban consumers) from the US Bureau of Labor Statistics has been used.[9]

In order to avoid problems, such as multiple local maxima, ill-posedness, associated with the use of maximum likelihood estimation to calibrate the jump models, the thresholding technique of Cont and Mancini (2011); Mancini (2009) has been used, as applied in Dang and Forsyth (2016); Forsyth and Vetzal (2017), for the calibration. Specifically, if $\Delta\hat{X}_i$ denotes the ith inflation-adjusted, detrended log return in the historical risky asset index time series, we identify a jump in period i if

$$\left|\Delta\hat{X}_i\right| \quad > \quad \alpha\hat{\sigma}\sqrt{\Delta t}, \tag{5.1}$$

where $\hat{\sigma}$ is the estimate of the diffusive volatility, Δt is the time period over which the log return has been calculated, and α is the "threshold parameter" for identifying a jump. Distinguishing between "up" and "down" jumps for the Kou model is achieved using upward and downward jump indicators - see Forsyth and Vetzal (2017) for further details, including the simultaneous estimation of the diffusive volatility. We will use $\alpha = 3$ in what follows - in other words, we would only detect a jump in the historical time series if the (absolute, inflation-adjusted, and detrended) log return in that period exceeds 3 standard deviations of the "geometric Brownian motion change", which is a very unlikely event. In the case of GBM, we use standard maximum likelihood techniques. The resulting calibrated parameters are provided in Table 5.1.

Table 5.1: Calibrated risky and risk-free asset process parameters ($\alpha = 3$ used in (5.1) for the Merton and Kou models).

	Models		
Parameters	GBM	Merton	Kou
μ (drift)	0.0816	0.0817	0.0874
σ (diffusive volatility)	0.1863	0.1453	0.1452
λ (jump intensity)	n/a	0.3483	0.3483
$\widetilde{m}$ (log jump multiplier mean)	n/a	-0.0700	n/a
$\widetilde{\gamma}$ (log jump multiplier stdev)	n/a	0.1924	n/a
ν (probability of up-jump)	n/a	n/a	0.2903
ζ_1 (exponential parameter up-jump)	n/a	n/a	4.7941
ζ_2 (exponential parameter down-jump)	n/a	n/a	5.4349
r (Risk-free rate)	0.00623	0.00623	0.00623

5.2 Convergence analysis

In this subsection, we demonstrate that the numerical PDE solution converges to known analytical solutions available in special cases where such solutions are available, and rely on Monte Carlo simulation to verify results in the cases where analytical solutions are not available.

5.2.1 Analytical solutions

Analytical solutions for the time-consistent problem are available if the risky asset follows GBM (see Basak and Chabakauri (2010)) or any of the commonly-encountered jump models, including the Merton and Kou models (see Bjork and Murgoci (2010) and Zeng et al. (2013)), under the following

[9] CPI data from the U.S. Bureau of Labor Statistics.In particular, we use the annual average of the all urban consumers (CPI-U) index. See `http://www.bls.gov/cpi`.

assumptions: (i) continuous rebalancing of the portfolio, (ii) trading continues in the event of insolvency, (iii) no investment constraints or transaction costs, and (iv) same lending and borrowing rate ($= r$). Under these assumptions, the efficient frontier solution is given by

$$\begin{aligned} E_{\mathcal{C}_0^*}^{t=0}\left[W\left(T\right)\right] &= W\left(0\right)e^{rT} + \frac{1}{2\rho}\left[\frac{\left(\mu - r\right)^2}{\sigma^2 + \lambda\kappa_2}\right]T, \\ Stdev_{\mathcal{C}_0^*}^{t=0}\left[W\left(T\right)\right] &= \frac{1}{2\rho}\left(\frac{\mu - r}{\sqrt{\sigma^2 + \lambda\kappa_2}}\right)\sqrt{T}, \end{aligned} \tag{5.2}$$

where we set $\lambda = 0$ to obtain the special solution in the case where the risky asset follows GBM.

Table 5.2 provides the timestep and grid information for testing convergence to the analytical solution (5.2). While equal timesteps are used, the grids in the s- and b-directions are not uniform.

Table 5.2: Grid and timestep refinement levels for convergence analysis to the analytical solution (5.2)

Refinement level	Timesteps	s-grid nodes	b-grid nodes
0	30	70	147
1	60	139	293
2	120	277	585
3	240	553	1089

Table 5.3 illustrates the numerical convergence analysis for an initial wealth of $W(0) = 100$, maturity $T = 10$ years, and scalarization parameter $\rho = 0.005$. For illustrative purposes, we assume the risky asset follows the Merton model - qualitatively similar results are obtained if the Kou or GBM models are assumed. The "Error" column shows the difference between the analytical solution and the PDE solution, while the "Ratio" column shows the ratio of successive errors for each increase in the refinement level. We observe first-order convergence of the numerical PDE efficient frontier values to the analytical values obtained from (5.2) as the mesh is refined, which is expected.

Table 5.3: Convergence to analytical solution - Merton model

Refinement level	Expected value (Analytical solution: 274.5)			Standard deviation (Analytical solution: 129.7)		
	PDE solution	Error	Ratio	PDE solution	Error	Ratio
0	250.7	23.8	-	120.2	9.5	-
1	263.1	11.4	2.08	125.2	4.6	2.08
2	269.2	5.3	2.16	127.7	2.1	2.22
3	272.0	2.5	2.13	128.7	1.0	2.01

5.2.2 Monte Carlo validation

Consider now the following case where analytical solutions are *not* available: we assume discrete periodic rebalancing of the portfolio at the end of each year, with liquidation in the event of insolvency, and a maximum allowable leverage ratio of $q_{\max} = 1.5$. Additionally, we assume the risky asset follows the Kou model, with initial wealth of $W\left(0\right) = 100$, maturity $T = 20$ years, and scalarization parameter $\rho = 0.0014$. For the numerical PDE solution, using 7,280 equal timesteps, and 1,121 and 2,209 s-grid and b-grid nodes, respectively, we obtain the following approximations to the expectation and standard deviation:

$$\left(E_{\mathcal{C}_0^*}^{t=0}\left[W\left(T\right)\right], Stdev_{\mathcal{C}_0^*}^{t=0}\left[W\left(T\right)\right]\right) = \left(544.58, 400.20\right). \tag{5.3}$$

At each timestep of our numerical PDE procedure, we output and store the computed optimal strategy for each discrete state value. We then carry out Monte Carlo simulations for the portfolio (using the specified parameters) from $t = 0$ to $t = T$, rebalancing the portfolio in accordance with the stored PDE-computed optimal strategy at each discrete rebalancing time. If necessary, we use interpolation to determine the optimal strategy for a given state value. We then compare the Monte Carlo computed means and standard deviations of the terminal wealth with the corresponding values computed by the numerical PDE method, given in (5.3). The results are shown in Table 5.4. Note that, for the MC method, due to the possibility of insolvency, it is not possible to take finite timesteps between rebalancing times without incurring timestepping errors.

Table 5.4: Convergence analysis to numerical PDE solution using Monte Carlo simulation - Kou model.

Nr of simulations	Nr of timesteps / year	Expectation (PDE solution: 544.58)		Standard deviation (PDE solution: 400.20)	
		Value	Relative error	Value	Relative error
4,000	728	537.03	-1.39%	388.69	-2.88%
16,000	1,456	540.28	-0.79%	391.48	-2.18%
64,000	2,912	540.92	-0.67%	396.80	-0.85%
256,000	5,824	542.60	-0.36%	398.38	-0.46%
1,024,000	11,648	544.33	-0.05%	399.08	-0.28%

We observe that, as the number of Monte Carlo simulations and timesteps increase, the Monte Carlo computed means and standard deviations converge to the corresponding values computed by the numerical PDE method, given in (5.3).

5.3 Time-consistent MV efficient frontiers

In this subsection, we study time-consistent MV efficient frontiers. All efficient frontier results in this section are based on an initial wealth of $W(0) = 100$ and a maturity $T = 20$ years, along with annual (discrete) rebalancing, and approximately daily interest payments (364 payments per year) on the amount in the risk-free asset. To construct a point on the efficient frontier via the PDE scheme, we use 7,280 equal timesteps, and 561 and 1,105 s-grid and b-grid nodes, respectively.

In order to consider the impact of investment constraints and other assumptions, including transaction costs, we construct five experiments as outlined in Table 5.5. We highlight the following:

Table 5.5: Details of experiments

Experiment	Lending/ borrowing rates		If insolvent	Leverage constraint	Transaction costs	
	r_ℓ	r_b			Fixed (c_1)	Prop. (c_2)
Experiment 1	0.00623	0.00623	Continue trading	None	0	0
Experiment 2	0.00623	0.00623	Liquidate	None	0	0
Experiment 3	0.00623	0.00623	Liquidate	$q_{\max} = 1.5$	0	0
Experiment 4	0.00400	0.06100	Liquidate	$q_{\max} = 1.5$	0	0
Experiment 5	0.00400	0.06100	Liquidate	$q_{\max} = 1.5$	0.001	0.005

- The interest rates for Experiments 4 and 5 were obtained by assuming that the approximate relationship between current interest rates paid on margin accounts in relation to current 3-month US T-bill rates[10], also holds in relation to the historically observed 3-month US T-bill rates used to obtain the constant rate of 0.00623 (see Table 5.1).

- The transaction costs in the case of Experiment 5 are perhaps somewhat extreme. As in the case of Dang and Forsyth (2014), the costs were chosen to emphasize the effect of transaction costs in particular when compared to an Experiment 4 (which has the same borrowing/lending rates as Experiment 5, but with zero transaction costs).

5.3.1 Model choice

We consider the efficient frontiers obtained for the time-consistent MV problem using the numerical PDE scheme as outlined above, starting with the impact of model choice, namely GBM, Merton, or Kou dynamics, on the efficient frontiers. In Figure 5.1, we present the time-consistent MV efficient frontiers for Experiments 1 and 2, with the risky asset dynamics following GBM, Merton and Kou models. We observe that the Kou model results in a lower efficient frontier relative to the GBM and Merton models, whose efficient frontiers are basically indistinguishable.

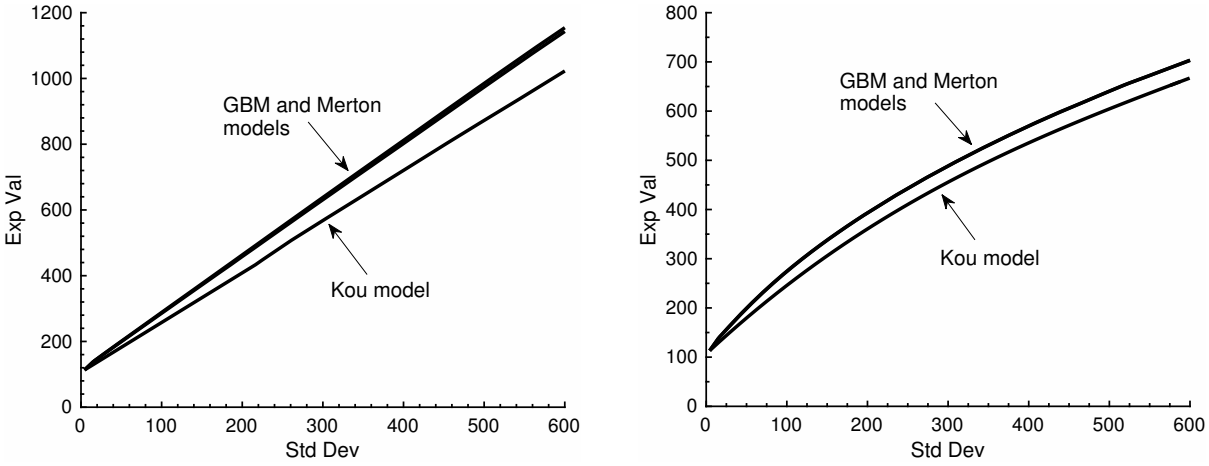

(a) Experiment 1 - No constraints

(b) Experiment 3 - Solvency and leverage constraints

Figure 5.1: Time-consistent MV efficient frontiers - Effect of model choice (GBM, Merton, Kou)

Since these results are obtained using discrete (annual) rebalancing of the portfolio, no analytical solution exists, even in the case of the Experiment 1 frontiers seen in Figure 5.1(a). However, if we assume continuous rebalancing of the portfolio and no constraints, we can use the analytical solution in (5.2) to guide our intuition. Note that (5.2) can be re-arranged to give the expected value in terms of the standard deviation,

$$E_{\mathcal{C}_0^*}^{t=0}\left[W\left(T\right)\right] \quad = \quad W\left(0\right)e^{rT} + \left(\frac{\mu - r}{\sqrt{\sigma^2 + \lambda\kappa_2}}\right)\sqrt{T}\cdot\left(Stdev_{\mathcal{C}_0^*}^{t=0}\left[W\left(T\right)\right]\right). \tag{5.4}$$

[10]The interest paid/charged currently on margin accounts at major stockbrokers can be obtained with relative ease. For these experiments, the information was obtained as follows. On 15 March 2017, Merrill Edge (an online brokerage service of the Bank of America Merrill Lynch) charged roughly 5.75% on negative balances in margin accounts - the exact rate can depend on a number of factors. At that time, the short-term deposit rates of 0.03% paid by Bank of America was used as the interest rate paid on positive balances. These figures were then inflation-adjusted and scaled with the difference between current and historical real returns on T-bills, so that we assume in effect that the observed spread (difference between borrowing and lending rates) remained the same historically as they were in early 2017. This resulted in the rates of 6.10% and 0.40% shown in Table 5.5.

Fixing a standard deviation value on the efficient frontier, we observe that the effect of model choice on the associated expected value on the efficient frontier is entirely due to the multiplier $(\mu - r)/\sqrt{\sigma^2 + \lambda\kappa_2}$ in (5.4). With calibrated process parameters as given in Table 5.1, we have combinations of parameters as given in Table 5.6. In particular, we conclude that the multiplier $(\mu - r)/\sqrt{\sigma^2 + \lambda\kappa_2}$ is lower for the Kou model, due to the higher variance of the log-double exponential distribution of the jump multipliers (resulting in a higher value of $\kappa_2 = \mathbb{E}\left[(\xi - 1)^2\right] = Var(\xi) + \kappa^2$) compared to the that of the lognormal distribution in the case of the Merton model. We also note that, as observed from Table 5.6, both the GBM and Merton models have almost the same value of the multiplier $(\mu - r)/\sqrt{\sigma^2 + \lambda\kappa_2}$.

Table 5.6: Combinations of parameters ($\alpha = 3$ used in (5.1) for the Merton and Kou models)

Combinations of parameters	GBM	Merton	Kou
$\kappa = \mathbb{E}[(\xi - 1)]$	0.0000	-0.0502	-0.0338
$\kappa_2 = \mathbb{E}\left[(\xi - 1)^2\right]$	0.0000	0.0365	0.0844
$(\mu - r)/\sqrt{\sigma^2 + \lambda\kappa_2}$	0.4046	0.4103	0.3612

Returning to the results shown in Figure 5.1 where no analytical solutions are available, we conclude the following. With the exception of parameters affecting the jump distribution, the other model parameters (drift, diffusive volatility, jump intensity) of the Kou and Merton models in Table 5.1 are very similar. Since the jump multipliers have a higher variance in the Kou model compared to the Merton model (both calibrated to the same data), then for a given level of expected terminal wealth, the Kou model results in a larger standard deviation of the terminal wealth. Consequently, the efficient frontier is lower for the Kou model than for the Merton model. Furthermore, similar multiplier values for the GBM and Merton models (observed above) imply that the relatively higher diffusive volatility of the GBM model has a similar effect as the incorporation of jumps using the Merton model over this long investment time horizon, resulting in similar efficient frontiers for the GBM and Merton models.

5.3.2 Investment constraints

The effect of investment constraints on the time-consistent MV efficient frontiers are shown in Figure 5.2 for the Kou model only, since the results for other models are qualitatively similar.

Figure 5.2(a) illustrates the significant impact of requiring liquidation in the event of insolvency (Experiment 1 vs. Experiment 2). Furthermore, it is observed that, once liquidation in the event of insolvency is a requirement, the impact of the leverage constraint is comparatively much smaller (Experiment 2 vs. Experiment 3).

If we additionally incorporate more realistic interest rates, i.e. different lending and borrowing rates, (Experiment 4), then Figure 5.2(b) shows a substantial reduction in the expected terminal wealth that can be achieved, especially for high levels of risk. (Compare Experiments 3 and 4 on Figure 5.2(b).) The reason for this is that, in order to achieve a high standard deviation of terminal wealth, a comparatively large amount needs to be invested in the risky asset, which is achieved by borrowing to invest. If the cost of borrowing is substantially increased (Experiment 4 vs. Experiment 3), the achievable expected terminal wealth reduces, reflecting the increased effective cost of executing such a strategy. By comparison, the effect of additionally introducing transaction costs (Experiment 5) is relatively negligible.

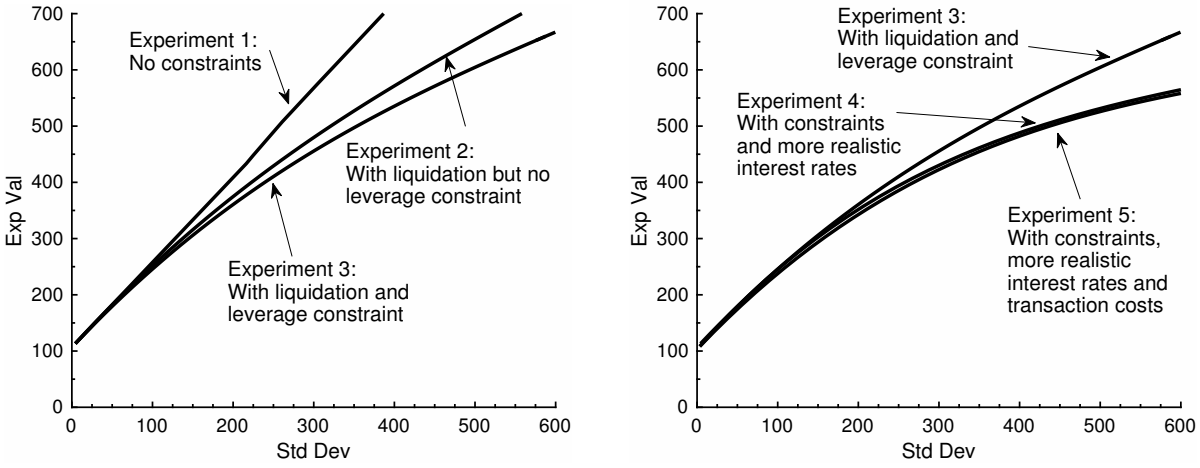

(a) Effect of liquidation and leverage constraints

(b) Effect of interest rates and transaction costs

Figure 5.2: Time-consistent MV efficient frontiers - Kou model: Effect of investment constraints

5.4 Time-consistent MV vs. Pre-commitment MV strategies

In this section, we compare the time-consistent and the pre-commitment strategies, not only in terms of the resulting efficient frontiers, but also in terms of the optimal investment policies over time. We focus on the Kou model, since the other models yield qualitatively similar results. Process parameters are as in Table 5.1, investment parameters are as outlined at the beginning of Subsection 5.3, and details of the experiments are as in Table 5.5. The pre-commitment MV problem is formulated using impulse controls and solved according to the techniques outlined in Dang and Forsyth (2014). In order to provide a *fair* comparison with the standard time-consistent formulation, we do not optimally withdraw cash for the pre-commitment MV case (Cui et al., 2012; Dang and Forsyth, 2016). Allowing optimal cash withdrawals will move the efficient upward for the pre-commitment MV strategy.

5.4.1 Combined investment constraints

Figure 5.3 compares the efficient frontiers associated with the pre-commitment and time-consistent problems in Experiments 1 and 3. As expected, the pre-commitment strategy is more MV efficient in the sense that the associated efficient frontier lies above that of the time-consistent strategy. This follows since the time-consistent problem carries the additional time-consistency constraint. However, under both the solvency and leverage constraints (Figure 5.3(b)), the difference between the two efficient frontiers is substantially reduced. A similar effect has also been observed in Wang and Forsyth (2011) for the case of continuous trading and no jumps in the risky asset process.

In Figures 5.3a and 5.3b, points on the efficient frontiers corresponding to a standard deviation of terminal wealth equal to 400 have been highlighted. The resulting MV-optimal strategies corresponding to these points will be investigated in more detail below (see Subsection 5.4.3).

5.4.2 Leverage constraint

Next, we focus on the impact of the leverage constraint. Figure 5.4 illustrates the effect of different maximum leverage constraint $q_{\max}$ assumptions on the efficient frontiers associated with the pre-commitment and time-consistent MV problems. (In these tests, the solvency constraint is also imposed.) Since leverage may not be allowed for pension fund investments, we also consider the effect of setting $q_{\max} = 1$ (so that the fraction of total wealth invested in the risky asset may not exceed one) in Experiment 3.

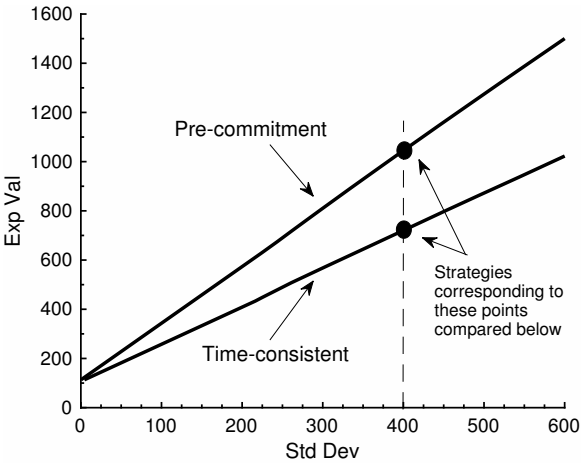

(a) Experiment 1 - No constraints

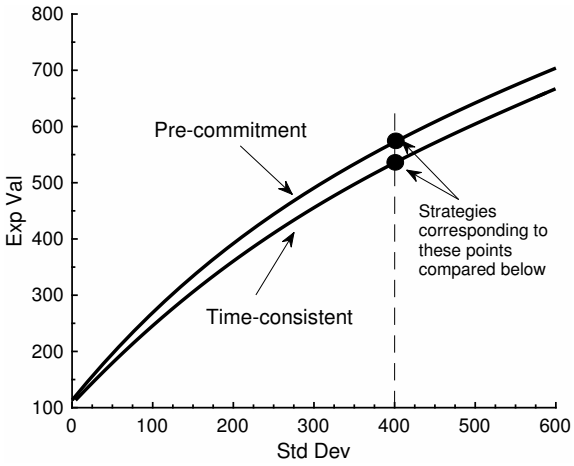

(b) Experiment 3 - Solvency and leverage constraints

Figure 5.3: Pre-commitment MV vs. Time-consistent MV efficient frontiers - Kou model

It is observed that the effect on the efficient frontiers of not allowing leverage is quite dramatic. Interestingly, especially for high standard deviation of terminal wealth, the effect of setting $q_{\max} = 1$ on the pre-commitment efficient frontier (Figure 5.4(a)) is comparatively larger than the effect on the time-consistent efficient frontier (Figure 5.4(b)).

The above observation is not entirely unexpected. As shown below (subsection 5.4.3), the pre-commitment MV optimal strategy generally favors much higher investment in the risky asset during the early years of the investment period, compared to the time-consistent MV optimal strategy. (See Figures 5.7 and 5.6 and the relevant discussion). Not allowing any leverage, therefore, has a larger relative impact on the pre-commitment MV efficient frontier.

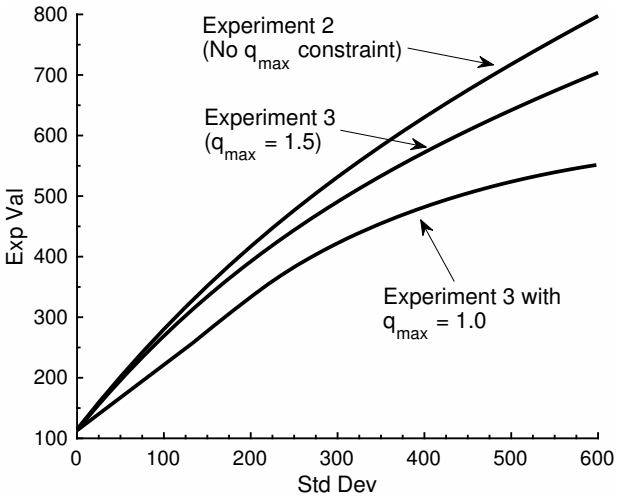

(a) Pre-commitment strategy

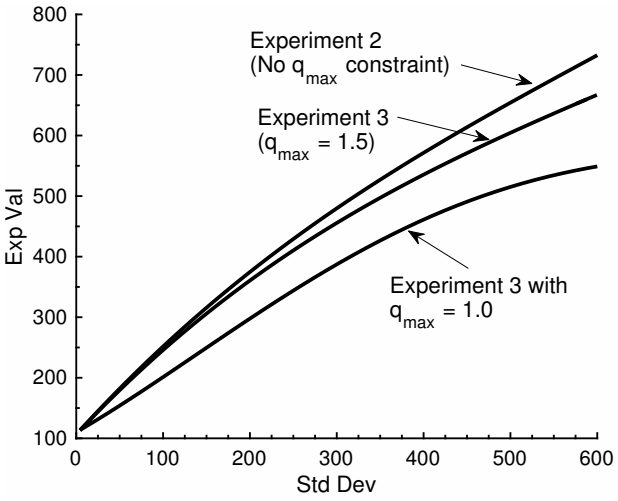

(b) Time-consistent strategy

Figure 5.4: Pre-commitment MV vs. Time-consistent MV - Kou model: Effect of maximum leverage constraint $q_{\max}$.

5.4.3 Comparison of optimal controls

To gain further insight into the optimal control strategy of the time-consistency and pre-commitment approaches, we perform additional Monte Carlo simulations, using the same steps outlined in Subsection 5.2.2, to Experiments 1 and 3 previously reported in Figure 5.3 (a)-(b).

Specifically, we first fix the standard deviation of the terminal wealth at a value of 400, as shown in Figure 5.3 (a)-(b). When solving the pre-commitment and time-consistent problems corresponding to these points on the efficient frontiers, at each timestep of our numerical PDE procedure, we output and store the computed optimal strategy for each discrete state value. We then carry out Monte Carlo simulations for the portfolio, using the specified parameters, from $t = 0$ to $t = T$, rebalancing the portfolio in accordance with the stored PDE-computed optimal strategy at each discrete rebalancing time. We compute, for each path and for each point in time, the fraction of wealth invested in the risky asset.

The results of this study are summarized in Figure 5.5 and Figure 5.6, where we show the median (50th percentile), as well as the 25th and 75th percentiles, of the distribution of the MV-optimal fraction of wealth invested in the risky asset over time.

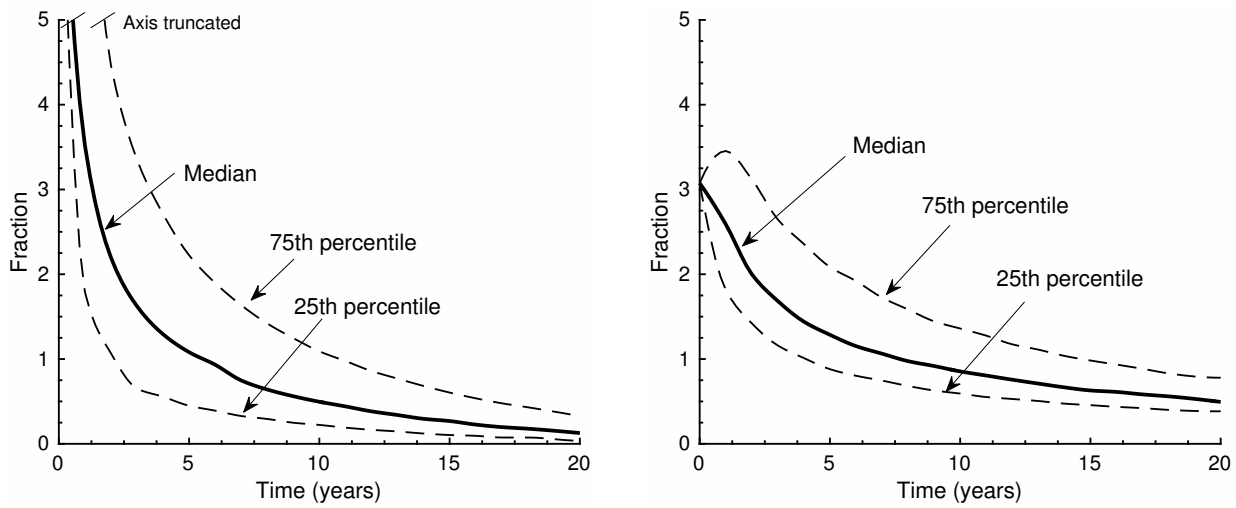

(a) Pre-commitment strategy

(b) Time-consistent strategy

Figure 5.5: MV-optimal fraction of wealth in the risky asset: Kou model, Experiment 1, standard deviation of terminal wealth equal to 400.

Figure 5.5 compares the fraction of wealth in the risky asset for Experiment 1 (no investment constraints). In the case of the pre-commitment strategy (Figure 5.5(a)), the investment in the risky asset is initially much higher than in the case of the time-consistent strategy (Figure 5.5(b)). This changes as time progresses, with the fraction of wealth invested in the risky asset decreasing substantially for the pre-commitment strategy. While a decrease can also be observed for the time-consistent strategy, it is much more gradual. Furthermore, at about $t = 3$ (years) in this case, the median fraction of wealth in the risky asset for the time-consistent strategy exceeds that of the pre-commitment strategy.

The above observation can be explained by recalling from Vigna (2014) that the pre-commitment problem can also be viewed as a target-based optimization problem, where a quadratic loss function is minimized. This means that once the portfolio wealth is sufficiently large, so that the (implicitly) targeted terminal wealth becomes more achievable, the pre-committed investor will reduce the risk by reducing the investment in the risky asset. In contrast, the time-consistent investor has no investment target, and instead, acts consistently with the mean-variance risk preferences throughout the investment time horizon (see for example Cong and Oosterlee (2016) for a relevant discussion).

If we impose liquidation in the event of insolvency, as well as a maximum leverage ratio of $q_{\max} = 1.5$, i.e. Experiment 3, Figure 5.6 shows that the resulting MV-optimal fraction of wealth invested in the risky asset changes substantially compared to Figure 5.5. In particular, we observe that the fraction invested in the risky asset for the pre-commitment strategy (Figure 5.6(a)) is more strongly

affected by the maximum leverage constraint than the fraction for the time-consistent strategy (Figure 5.6(b)). While this only considers only one point on the efficient frontier, where the standard deviation of terminal wealth is equal to 400, we have observed the higher sensitivity of the pre-commitment strategy to the maximum leverage constraint across the efficient frontier in Figure 5.4. This is due to the very large pre-commitment MV-optimal investment in the risky asset required during the early stages of the investment time period in order to achieve the implicit wealth target. On the other hand, it is interesting to observe that the pre-commitment strategy at the 25th percentile shows a very rapid de-risking compared to the time-consistent strategy.

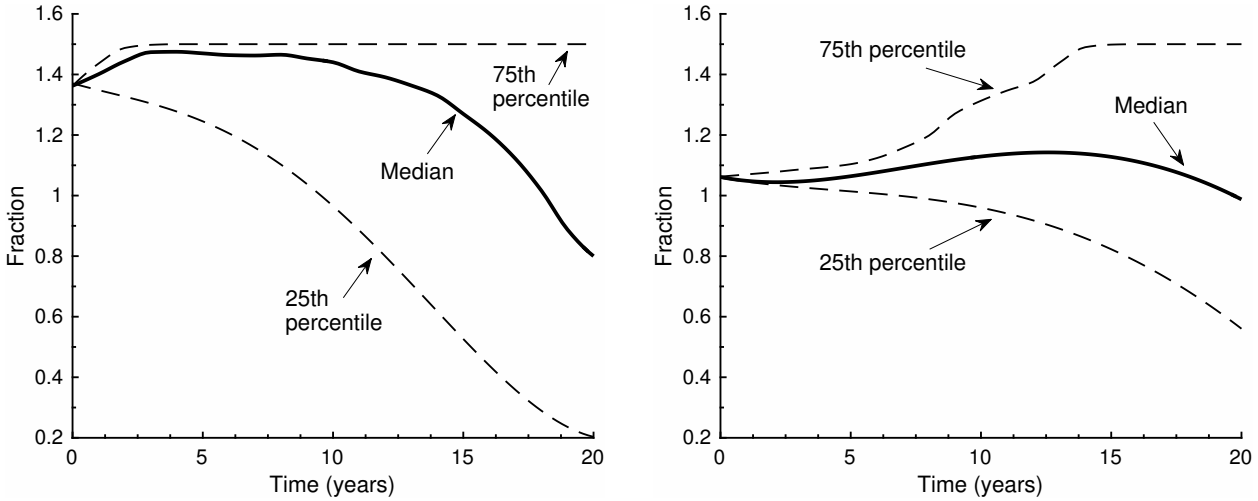

(a) Pre-commitment strategy (b) Time-consistent strategy

Figure 5.6: MV-optimal fraction of wealth in the risky asset: Kou model, Experiment 3, standard deviation of terminal wealth equal to 400.

To further investigate the differences between the pre-commitment and time-consistency optimal strategies, in Figure 5.7, we present the heatmaps of the MV-optimal control (as the fraction of wealth invested in the risky asset) as a function of time and wealth, which is used in the Monte Carlo simulation to generate the results in Figure 5.6.

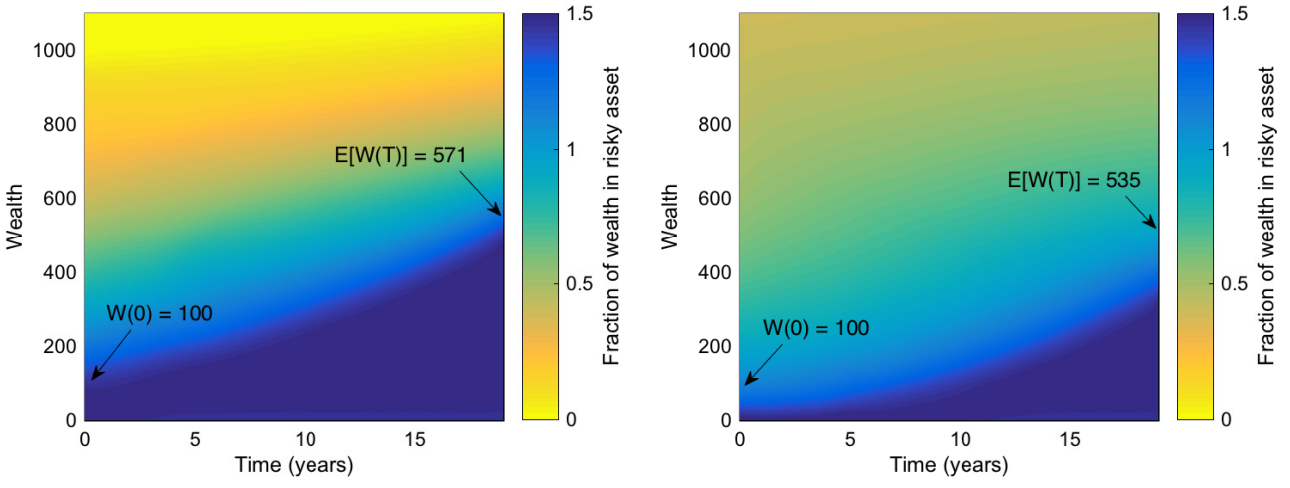

(a) Pre-commitment strategy (b) Time-consistent strategy

Figure 5.7: Optimal control as a fraction of wealth in risky asset: Kou model, Experiment 3, standard deviation of terminal wealth equal to 400.

We observe that, in the case of the pre-commitment optimal control (Figure 5.7(a)), for initial wealth of $W(0) = 100$ the optimal control requires a very large investment (very close to the maximum leverage of 1.5) in the risky asset. If returns are favourable - and therefore if wealth becomes sufficiently large over time - the optimal control specifies a reduction in the investment in the risky asset, possibly even to zero. If returns are unfavourable - so that wealth remains relatively small over time - the optimal strategy requires a very large fraction of wealth (again very close, if not equal to, the maximum leverage allowed) to remain invested in the risky asset. This is consistent with the interpretation of the pre-commitment strategy as a target-based strategy. If it becomes likely that the target will be achieved (past returns have been favourable), risk exposure is reduced; in contrast, if returns have been unfavourable in the past, risk is increased in order to make the achievement of the target more likely.

In contrast, in the case of the time-consistent optimal control (Figure 5.7(b)), there are a number of qualitative similarities to the pre-commitment optimal control (Figure 5.7(a)), but also key differences. Both of the strategies are contrarian, in the sense that all else being equal, investment in the risky asset is increased if its returns in the past have been unfavourable. However, compared to the pre-commitment optimal control, the time-consistent optimal control requires generally higher investment in the risky asset if past returns have been favourable (resulting higher wealth), and lower investment in the risky asset if past returns have been unfavourable (resulting in lower wealth). Even if the risky asset performs extremely well, the time-consistent strategy never calls for zero exposure to the risky asset. Figure 5.7 also shows why the pre-commitment strategy would be more heavily impacted if the maximum leverage ratio is reduced; the time-consistent strategy calls for generally lower leverage, and would therefore be less sensitive to the maximum leverage constraint.

5.5 Effect of a wealth-dependent scalarization parameter

Under the assumptions listed in Subsection 5.2.1 (in particular, under no investment constraints and where trading continues in the event of bankruptcy), the time-consistent MV-optimal control leading to the analytical efficient frontier solution in equation (5.2) does not depend on the investor's wealth at any point in time - see Basak and Chabakauri (2010) and Zeng et al. (2013). In other words, an investor following the resulting investment strategy is required to invest a particular *amount* in the risky asset at each point in time, entirely independent of their available wealth, which is not an economically reasonable conclusion. We emphasize that this is only true for the time-consistent MV optimal control in the absence of any investment constraints.

To remedy this situation, Bjork et al. (2014) proposes the use of a state-dependent scalarization (or risk aversion) parameter. Applied in our setting, we obtain a time-consistent MV problem otherwise identical to equations (2.17) - (2.18), with the difference being that the risk aversion parameter at each point in time is explicitly modelled by a deterministic function of the wealth $W(t)$, i.e. $\rho = \rho(W(t))$. That is (2.17) now becomes

$$\sup_{\mathcal{C}_t \in \mathcal{A}} \left(E_{\mathcal{C}_t}^{x,t} [W(T)] - \rho(W(t)) \, Var_{\mathcal{C}_t}^{x,t} [W(T)] \right)$$

In Bjork et al. (2014), it is argued that a natural choice for the function $\rho(W(t))$ is of the form

$$\rho(W(t)) = \frac{\theta}{W(t)}, \qquad \theta > 0 \tag{5.5}$$

where for each θ, we obtain a point on the resulting efficient frontier. The use of a wealth-dependent scalarization parameter has been popular in time-consistent MV literature within the non-constraint setting, especially in insurance-related applications (see for example Zeng and Li (2011), Wei et al. (2013), Li and Li (2013), as well as Liang and Song (2015)).

While some alternatives to (5.5) has also been suggested - see for example Bensoussan et al. (2014), we explore the effect of using the definition (5.5) in our setting, specifically because this simple case reveals how a seemingly reasonable definition of a wealth-dependent scalarization parameter, when used in combination with investment constraints and liquidation in the event of bankruptcy, can result in conclusions that are not economically reasonable.

Given Algorithm 3.1, implementing a wealth-dependent scalarization parameter such as (5.5) is straightforward, since we simply replace ρ in the algorithm with $\rho(W(s,b)) = \theta/W(s,b)$, where $W(s,b)$ is given by equation (2.8), without any further changes required. Varying $\theta > 0$ in this case traces out the efficient frontier.

We consider Experiment 3 in Table 5.5 (in other words we impose both liquidation in bankruptcy and a leverage constraint), since - as pointed out in Wang and Forsyth (2011) - allowing for negative wealth in equation (5.5) would lead to inappropriate risk aversion coefficients. In Figure 5.8, the efficient frontier obtained with a constant scalarization parameter ρ is compared with the efficient frontier obtained with wealth-dependent scalarization parameter of the form (5.5). We observe a similar result as in Wang and Forsyth (2011), where the case of continuous controls and no jumps was investigated: the resulting time-consistent MV efficient frontier with a wealth-dependent scalarization parameter is significantly lower than that obtained using a constant scalarization parameter. In other words, given an acceptable level of risk as measured by variance, a strategy based on the wealth-dependent scalarization parameter given by (5.5) would result in much lower expected terminal wealth, and is therefore less efficient from a MV-optimization perspective.

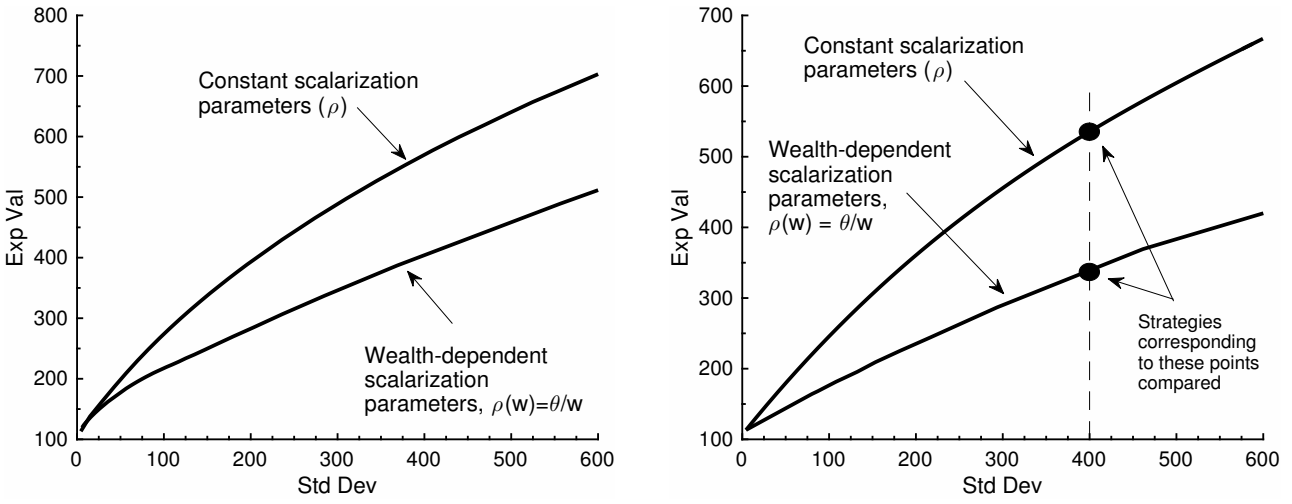

(a) GBM model

(b) Kou model

Figure 5.8: Time-consistent MV efficient frontiers - Experiment 3 (solvency and leverage constraints): Effect of using a constant scalarization parameter vs. using a wealth-dependent scalarization parameter of the form $\rho(w) = \theta/w$.

We now further compare the optimal trading strategies for the Kou model in both scenarios, namely a constant scalarization parameter and a wealth-dependent scalarization parameter of the form (5.5). In this case, we pick two points on the efficient frontiers corresponding to a standard deviation of terminal wealth equal to 400, as highlighted in Figure 5.8(b). In Figure 5.9, we now compare the resulting MV-optimal strategies corresponding to these points. Specifically, proceeding as in Subsection 5.4.3, using Monte Carlo simulations and rebalancing the portfolio in accordance with the stored PDE-computed optimal strategy at each discrete rebalancing time, we consider the resulting MV-optimal fraction of wealth invested in the risky asset over time.

Figure 5.9 (a) compares the median of the time-consistent MV-optimal fraction of wealth in the

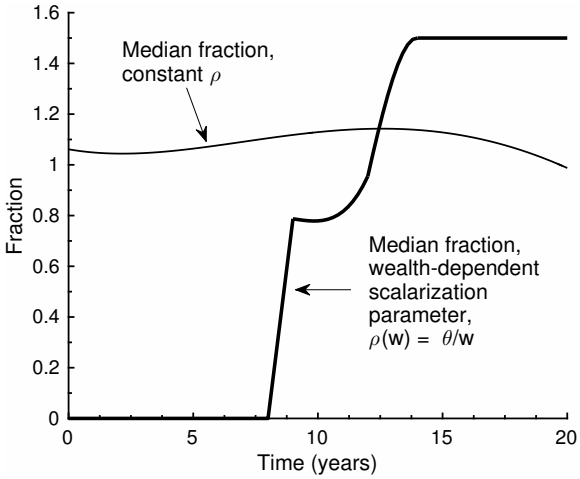

(a) Median MV-optimal fraction of wealth in the risky asset

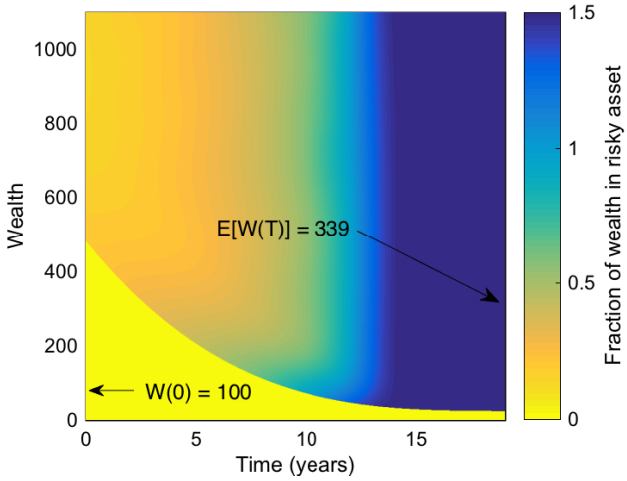

(b) Optimal control as fraction of wealth in risky asset, wealth-dependent scalarization parameter $\rho(w) = \theta/w$

Figure 5.9: Effect of using a using a wealth-dependent scalarization parameter of the form $\rho(w) = \theta/w$ on the median time-consistent MV-optimal fraction of wealth in the risky asset and on the resulting optimal controls. Kou model - Experiment 3 (solvency and leverage constraints), standard deviation of terminal wealth equal to 400.

risky asset in both scenarios.[11] Figure 5.9 (b) illustrates the heatmap of the time-consistent MV-optimal control (as the fraction of wealth invested in the risky asset) as a function of time and wealth in the case of a wealth-dependent scalarization parameter of the form (5.5). The heatmap for the time-consistent MV-optimal control in the case of a constant scalarization parameter (also for the Kou model, Experiment 3, and a standard deviation of terminal wealth equal to 400) is provided in Figure 5.7(b).

We make the following interesting observations. While the increase in exposure to the risky asset over time has been observed in the case of the wealth-dependent risk aversion parameter in the setting of no jumps, constraints or bankruptcy (see, for example, Bjork et al. (2014)), in the case of realistic investment constraints this is even more dramatic. Such observed dramatic impact can be explained as follows. The form of the wealth-dependent risk aversion in (5.5) implies that the risk aversion is inversely related to wealth. As such, it is possible (and indeed observed in Figure 5.9 (a)) that the investment in the risky asset can be zero until wealth has grown sufficiently and/or the time to maturity has decreased sufficiently, to make an investment in the risky asset MV-optimal. The level of risk aversion then steadily decreases, ensuring that the maximum exposure to the risky asset (only limited by the leverage constraint in this case) is reached as the investment maturity is approached.

While the economic merits of such a strategy depends on the particular application, it is unlikely to be economically reasonable in institution-related applications of MV optimization (such as in the case of pension funds or insurance). Specifically, relatively low investments in the risky asset during early years (due to high risk aversion resulting from relatively lower wealth levels) might result in lower terminal wealth - indeed, the expectation of terminal wealth is substantially lower with wealth-dependent scalarization parameter of the form (5.5) - which in turn might make it harder to fund liabilities, while the increase in risky asset exposure over time does not actually reduce the variance of terminal wealth (compared to the case of a constant ρ).

Therefore, in contrast to, for example Li and Li (2013), we conclude that a wealth-dependent scalarization parameter defined by (5.5) does not appear well-suited for obtaining realistic time-consistent

[11] For the constant scalarization scenario, this corresponds to the median line in Figure 5.6(b).

MV optimal strategies in the presence of investment constraints, since the resulting terminal wealth is less MV-efficient (as compared with the results obtained using a constant scalarization parameter), while the steady increase in risk exposure over time might be undesirable in many applications of time-consistent MV optimization.

6 Conclusions

In this paper, we develop a fully numerical PDE approach to solve the investment-only time-consistent MV portfolio optimization problem when the underlying risky asset follows a jump-diffusion process. The algorithm developed allows for the application of multiple simultaneous realistic investment constraints, including discrete rebalancing of the portfolio, the requirement of liquidation in the event of insolvency, leverage constraints, different interest rates for borrowing and lending, and transaction costs. The semi-Lagrangian timestepping scheme of Dang and Forsyth (2014) is extended to the system of equations for the time-consistent problem, resulting in a set of only one-dimensional PIDEs to be solved at each timestep. While no formal proof of convergence is given, numerical tests, including a numerical convergence analysis where analytical solutions are available, as well as the validation of results using Monte Carlo simulation, indicate that the algorithm provides reliable and accurate results.

The economic implications of investment constraints on the efficient frontiers and on the resulting optimal controls have been explored in detail. The numerical results illustrate that these realistic considerations can have a substantial impact on the efficient frontiers and associated optimal controls, resulting in economically plausible conclusions. In addition, the results from the time-consistent problem are compared to those of the pre-commitment problem, leading to the conclusion that the time-consistent problem is less sensitive to the maximum leverage constraint than the pre-commitment problem. In addition, we explored the consequences of implementing a popular form of a wealth-dependent risk aversion parameter (where risk aversion is inversely related to wealth), and find that the resulting optimal investment strategy has both undesirable terminal wealth outcomes and an undesirable evolution of risk characteristics over time. Not only does this finding pose questions about the use of such wealth-dependent risk aversion parameters in existing time-consistent MV literature, but it also highlights the importance of incorporating realistic constraints in investment models.

As a result of the popularity of the application of time-consistent MV optimization to investment-reinsurance problems (see for example Alia et al. (2016); Li et al. (2015c); Liang and Song (2015)), we leave the extension of the algorithm from the investment-only case to the investment-reinsurance problem for our future work.

References

Alia, I., F. Chighoub, and A. Sohail (2016). A characterization of equilibrium strategies in continuous-time mean–variance problems for insurers. *Insurance: Mathematics and Economics* (68), 212–223.

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

Bensoussan, A., K. C. Wong, S. C. P. Yam, and S. P. Yung (2014). Time-consistent portfolio selection under short-selling prohibition: From discrete to continuous setting. *SIAM Journal on Financial Mathematics* 5, 153–190.

Bjork, T., M. Khapko, and A. Murgoci (2016). A theory of Markovian time-inconsistent stochastic control in continuous time. *Working paper* .

Bjork, T. and A. Murgoci (2010). A general theory of Markovian time inconsistent stochastic control problems. *Working paper* Available at http://ssrn.com/abstract=1694759.

Bjork, T. and A. Murgoci (2014). A theory of Markovian time-inconsistent stochastic control in discrete time. *Finance and Stochastics* (18), 545–592.

Bjork, T., A. Murgoci, and X. Zhou (2014). Mean-variance portfolio optimization with state-dependent risk aversion. *Mathematical Finance* (1), 1–24.

Cong, F. and C. Oosterlee (2016). On pre-commitment aspects of a time-consistent strategy for a mean-variance investor. *Journal of Economic Dynamics and Control* 70, 178–193.

Cont, R. and C. Mancini (2011). Nonparametric tests for pathwise properties of semi-martingales. *Bernoulli* (17), 781–813.

Cont, R. and P. Tankov (2004). *Financial modelling with jump processes*. Chapman and Hall / CRC Press.

Crandall, M., H. Ishii, and P. Lions (1992). User's guide to viscosity solutions of second order partial differential equations. *Bulletin of the American Mathematical Society* 27(1), 1–67.

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.

Cui, X., L. Xu, and Y. Zeng (2015). Continuous time mean-variance portfolio optimization with piecewise state-dependent risk aversion. *Optimization Letters (Springer)* pp. 1–11.

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

Dang, D. and P. Forsyth (2016). Better than pre-commitment mean-variance portfolio allocation strategies: A semi-self-financing Hamilton–Jacobi–Bellman equation approach. *European Journal of Operational Research* (250), 827–841.

Dang, D., P. Forsyth, and Y.Li (2016). Convergence of the embedded mean-variance optimal points with discrete sampling. *Numerische Mathematik* (132), 271–302.

d'Halluin, Y., P. Forsyth, and K. Vetzal (2005). Robust numerical methods for contingent claims under jump diffusion processes. *IMA Journal of Numerical Analysis* (25), 87–112.

Forsyth, P. and K. Vetzal (2017). Dynamic mean variance asset allocation: Tests for robustness. *International Journal of Financial Engineering* 4:2. 1750021 (electronic).

Hu, Y., H. Jin, and X. Zhou (2012). Time-inconsistent stochastic linear-quadratic control. *SIAM Journal on Control and Optimization* 50(3), 1548–1572.

Kou, S. (2002). A jump-diffusion model for option pricing. *Management Science* 48(8), 1086–1101.

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

Li, D., X. Rong, and H. Zhao (2015a). Time-consistent reinsurance–investment strategy for a mean–variance insurer under stochastic interest rate model and inflation risk. *Insurance: Mathematics and Economics* 64, 28–44.

Li, D., X. Rong, and H. Zhao (2015b). Time-consistent reinsurance–investment strategy for an insurer and a reinsurer with mean–variance criterion under the cev model. *Journal of Computational and Applied Mathematics* 283, 142–162.

Li, Y. and Z. Li (2013). Optimal time-consistent investment and reinsurance strategies for mean–variance insurers with state dependent risk aversion. *Insurance: Mathematics and Economics* 53, 86–97.

Li, Y., H. Qiao, S. Wang, and L. Zhang (2015c). Time-consistent investment strategy under partial information. *Insurance: Mathematics and Economics* 65, 187–197.

Li, Z., Y. Zeng, and Y. Lai (2012). Optimal time-consistent investment and reinsurance strategies for insurers under Heston's SV model. *Insurance: Mathematics and Economics* 51, 191–203.

Liang, Z. and M. Song (2015). Time-consistent reinsurance and investment strategies for mean–variance insurer under partial information. *Insurance: Mathematics and Economics* 65, 66–76.

Lioui, A. (2013). Time consistent vs. time inconsistent dynamic asset allocation: Some utility cost calculations for mean variance preferences. *Journal of Economic Dynamics and Control* (37), 1066–1096.

Ma, K. and P. Forsyth (2016). Numerical solution of the Hamilton-Jacobi-Bellman formulation for continuous time mean variance asset allocation under stochastic volatility. *Journal of Computational Finance* 20:1, 1–37.

Mancini, C. (2009). Non-parametric threshold estimation models with stochastic diffusion coefficient and jumps. *Scandinavian Journal of Statistics* (36), 270–296.

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance* 7(1), 77–91.

Merton, R. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.

Oksendal, B. and A. Sulem (2005). *Applied Stochastic Control of Jump Diffusions*. Springer.

Pedersen, J. and G. Peskir (2017). Optimal mean-variance portfolio selection. *Mathematics and Financial Economics* (11), 137–160.

Ramezani, C. and Y. Zeng (2007). Maximum likelihood estimation of the double exponential jump-diffusion process. *Annals of Finance* 3(4), 487–507.

Sun, J., Z. Li, and Y. Zeng (2016). Precommitment and equilibrium investment strategies for defined contribution pension plans under a jump–diffusion model. *Insurance: Mathematics and Economics* (67), 158–172.

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

Vigna, E. (2016). On time consistency for mean-variance portfolio selection. *Working paper, Collegio Carlo Alberto* (476).

Vigna, E. (2017). Tail optimality and preferences consistency for intertemporal optimization problems. *Working paper, Collegio Carlo Alberto* (502).

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

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

Wang, J. and P. Forsyth (2011). Continuous time mean variance asset allocation: A time-consistent strategy. *European Journal of Operational Research* (209), 184–201.

Wei, J. and T. Wang (2017). Time-consistent mean-variance asset-liability management with random coefficients. *Insurance: Mathematics and Economics* (77), 84–96.

Wei, J., K. Wong, S. Yam, and S. Yung (2013). Markowitz's mean-variance asset-liability management with regime switching: A time-consistent approach. *Insurance: Mathematics and Economics* 53, 281–291.

Yu, P. (1971). Cone convexity, cone extreme points, and nondominated solutions in decision problem with multiobjectives. *Journal of Optimization Theory and Applications* (7), 11–28.

Zeng, Y. and Z. Li (2011). Optimal time-consistent investment and reinsurance policies for mean-variance insurers. *Insurance: Mathematics and Economics* 49(1), 145–154.

Zeng, Y., Z. Li, and Y. Lai (2013). Time-consistent investment and reinsurance strategies for mean–variance insurers with jumps. *Insurance: Mathematics and Economics* 52, 498–507.

Zhang, C. and Z. Liang (2017). Portfolio optimization for jump-diffusion risky assets with common shock dependence and state dependent risk aversion. *Optimal control applications and methods* (38), 229–246.

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

Zhou, Z., H. Xiao, J. Yin, X. Zeng, and L. Lin (2016). Pre-commitment vs. time-consistent strategies for the generalized multi-period portfolio optimization with stochastic cash flows. *Insurance: Mathematics and Economics* (68), 187–202.