

Accepted Manuscript

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

Pieter Van Staden, Duy-Minh Dang, Peter A. Forsyth

PII: S0167-6687(18)30125-2

DOI: <https://doi.org/10.1016/j.insmatheco.2018.08.003>

Reference: INSUMA 2487

To appear in: *Insurance: Mathematics and Economics*


Received date: March 2018

Revised date: August 2018

Accepted date: 27 August 2018

Please cite this article as: Van Staden P., Dang D.-M., Forsyth P.A., Time-consistent mean–variance portfolio optimization: A numerical impulse control approach. *Insurance: Mathematics and Economics* (2018), <https://doi.org/10.1016/j.insmatheco.2018.08.003>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



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 two main approaches to perform MV portfolio optimization, namely (i) the pre-commitment approach, and (ii) the time-consistent (or game theoretical) approach. It is

*Forsyth's research was supported by the Natural Sciences and Engineering Research Council of Canada (Grant number RGPIN-2017-03760).

[†]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

well-known that the pre-commitment approach typically yields *time-inconsistent* strategies (Basak and Chabakauri, 2010; Bjork and Murgoci, 2010; Dang and Forsyth, 2014; Li and Ng, 2000; Vigna, 2014; Wang and Forsyth, 2011; Zhou and Li, 2000). Specifically, for $0 \leq t < t' < u \leq T$, where $T > 0$ is the fixed horizon investment, the pre-commitment MV optimal strategy for time u , computed at time t , may not necessarily agree with the pre-commitment MV optimal strategy for the same time u , but computed at a later time t' . This time-inconsistency phenomenon is due to the fact that the variance term in the MV-objective is not separable in the sense of dynamic programming, and hence the corresponding MV portfolio optimization problem fails to admit the Bellman optimality principle.

The time-consistent approach addresses the problem of time-inconsistency of the MV optimal strategy by directly imposing a time-consistency constraint on the optimal control (Basak and Chabakauri, 2010; Bjork and Murgoci, 2010; Cong and Oosterlee, 2016; Wang and Forsyth, 2011). Specifically, the MV portfolio optimization problem is now constrained to ensure that, for any $0 \leq t < t' < u \leq T$, the optimal strategy for any time u , computed at time t' , must agree with the optimal strategy for the same time u , but computed at an earlier time t .¹ As a result, under this time-consistency constraint on the control, the corresponding MV portfolio optimization problem would admit the Bellman optimality principle, and hence, can be solved using dynamic programming. Without this time-consistency constraint, MV portfolio optimization would lead to a time-inconsistent optimal strategy, as in the case of the pre-commitment approach.² Throughout this paper, we refer to the time-consistency constrained optimization problem as the *time-consistent* MV problem.

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

¹We clearly distinguish this time-consistency constraint from investment constraints, such as leverage or solvency constraints, which do not affect the time-consistency of the optimal control.

²As an alternative to imposing a time-consistency constraint, the dynamical optimal approach proposed recently by Pedersen and Peskir (2017) deals with the time-inconsistency of the pre-commitment approach by recomputing the MV optimal strategy at each time instant t and controlled wealth value. This approach can therefore obtain time-consistent optimal controls by performing an infinite number of optimization problems. We refer the reader to Vigna (2017) for a more detailed discussion regarding the relationship of this approach to the standard pre-commitment and time-consistent approaches discussed here.

74 MV applications, such as asset-liability management, the explicit incorporation of insolvency consid-
75 erations is critical to ensure that the results are of any practical use. The analytical solutions in,
76 for example, Wei et al. (2013) and Wei and Wang (2017), while useful, necessarily assume trading
77 continues in the case of insolvency.

78 Moreover, in the time-consistent MV literature, it is typical for analytical techniques to allow for a
79 leverage ratio, i.e. the ratio of the investment in the risky asset to the total wealth, substantially larger
80 than a ratio that brokers would typically allow retail investors or financial regulators would likely allow
81 institutions to undertake in practice. More specifically, while a leverage ratio of around 1.5 times is
82 typically allowed in practice (for retail investors), some of the analytical techniques illustrated in the
83 available literature call for much larger leverage ratios, for example 2.4 times in Li et al. (2012), 3 times
84 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
85 14 times in Li et al. (2015a), none of which are practically feasible, and which only further increases
86 the probability of insolvency. In a number of publications, a leverage constraint is completely ignored,
87 such as Lioui (2013), and this potentially leads to misplaced economic conclusions. For example, it
88 is concluded in Lioui (2013) that the time-consistent strategy is preferred over the pre-commitment
89 strategy, since the latter requires “huge and unrealistic positions in risky assets; in some cases, the pre-
90 commitment strategy is more than 60 times the time consistent strategy”. However, such a conclusion
91 appears unconvincing, since the pre-commitment MV strategy’s positions in the risky asset would have
92 been significantly smaller, if a realistic leverage constraint had been incorporated into the problem
93 formulation.

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

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

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

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

117 The literature on numerical methods for time-consistent MV portfolio optimization is virtually
118 limited to the case of diffusion dynamics, i.e. Geometric Brownian Motion, for the risky asset, including
119 notable works of Cong and Oosterlee (2016); Wang and Forsyth (2011). However, it is well-documented
120 in the finance literature that jumps are often present in the price processes of risky assets (see, for
121 example, Cont and Tankov (2004); Ramezani and Zeng (2007)). Jump processes permit modelling

122 of non-normal asset returns and fat tails. We focus on jump-diffusions in this work, since previous
 123 studies indicate that mean-reverting stochastic volatility processes have a very small effect on the
 124 efficient frontier for long term (> 10 years) investors (Ma and Forsyth, 2016). Using a Monte Carlo
 125 approach, Cong and Oosterlee (2016) compare pre-commitment and time-consistent policies with
 126 leverage and bankruptcy constraints in the case of diffusion dynamics.³ In the present work, we
 127 go a step forward by considering both the continuous and discrete rebalancing versions of the time-
 128 consistent MV portfolio optimization problem with jump-diffusion dynamics for the risky asset and
 129 realistic investment constraints, such as transaction costs and different borrowing and lending interest
 130 rates. Moreover, we also provide a comprehensive comparison between the time-consistency and pre-
 131 commitment approaches, not only in terms of the resulting efficient frontiers, but also in terms of the
 132 optimal investment policies over time under the above-mentioned realistic context. Furthermore, our
 133 use of partial integro-differential equation (PIDE) methods for solution of the optimal control problem
 134 allows us to illustrate the strategies in terms of easy-to-interpret heat maps.

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

143 For time-consistent MV portfolio optimization with jump-diffusion dynamics, an impulse control
 144 approach can also be utilized to potentially achieve similar computational advantages. In the realistic
 145 context considered in this work, applying the popular method of Bjork et al. (2016); Bjork and Murgoci
 146 (2014), together with relevant results from Oksendal and Sulem (2005), the value function under an
 147 impulse control formulation can be shown to satisfy a strongly coupled, nonlinear system of equations,
 148 the so-called an extended Hamilton-Jacobi-Bellman (HJB) quasi-integro-variational inequality. This
 149 system of equations must be solved numerically, since a closed-form solution for it is not known to
 150 exist, except in special cases. However, it is not clear how such a very complex system of equations can
 151 be solved effectively numerically. As a result, in this case, the method of Bjork et al. (2016); Bjork and
 152 Murgoci (2014) does not appear to result in equations amenable for computational purposes. Hence,
 153 for numerical purposes, an alternative formulation of this problem is desirable.

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

162 The main contributions of this paper are as follows.

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

166 This approach results in only linear partial integro-differential equations (PIDEs) to solve be-
 167 tween intervention times, which is not only numerically simpler than the approach of Bjork et al.

³The bankruptcy constraint in (Cong and Oosterlee, 2016) is not quite the same as considered in this work.

(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 problem⁴ 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 have a strong strategic component - the investor/insurer may be more interested in overall global portfolio shifts from stocks to bonds and vice versa⁵, rather than the more secondary questions relating to risky asset basket compositions.

⁴As noted in the conclusion to this paper, we leave the investment-reinsurance problem for future work.

⁵It 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

207 Let $S(t)$ and $B(t)$ respectively denote the amounts (i.e. total dollars) invested in the risky and risk-
 208 free asset, at time $t \in [0, T]$, where $T > 0$ is the fixed horizon investment. Define $t^- = \lim_{\epsilon \downarrow 0} (t - \epsilon)$,
 209 $t^+ = \lim_{\epsilon \downarrow 0} (t + \epsilon)$, i.e. t^- (resp. t^+) as the instant of time before (resp. after) the (forward) time
 210 t . First, consider the risky asset. Let ξ be a random number representing a jump multiplier, with
 211 probability density function (pdf) $p(\xi)$. When a jump occurs, $S(t) = \xi S(t^-)$. As a specific example,
 212 we consider two jump distributions for ξ , namely the log-normal distribution (Merton, 1976) and the
 213 log-double-exponential distribution (Kou, 2002). Specifically, in the former case, $\log \xi$ is normally
 214 distributed, so that

$$215 \quad p(\xi) = \frac{1}{\xi \sqrt{2\pi\tilde{\gamma}^2}} \exp \left\{ -\frac{(\log \xi - \tilde{m})^2}{2\tilde{\gamma}^2} \right\}, \quad (2.1)$$

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

$$218 \quad p(\xi) = \nu \zeta_1 \xi^{-(\zeta_1+1)} \mathbb{I}_{[\xi \geq 1]} + (1 - \nu) \zeta_2 \xi^{\zeta_2-1} \mathbb{I}_{[0 \leq \xi < 1]}. \quad (2.2)$$

219 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
 220 a jump occurs, ν is the probability of an upward jump, and $(1 - \nu)$ is the probability of a downward
 221 jump. Furthermore, in this case, we have $E[\xi] = \frac{\nu \zeta_1}{\zeta_1 - 1} + \frac{(1 - \nu) \zeta_2}{\zeta_2 + 1}$.

222 In the context of pre-commitment MV analysis, the results in (Ma and Forsyth, 2016) indicate
 223 that the effects of mean-reverting stochastic volatility are unimportant for long-term (i.e. greater than
 224 10 years) investors. Hence we focus here on the effect of jump processes, as a major source of risk. In
 225 the absence of control, i.e. if we do not adjust the amount invested according to our control strategy,
 226 the amount S invested in the risky asset is assumed to follow the process

$$227 \quad \frac{dS(t)}{S(t^-)} = (\mu - \lambda \kappa) dt + \sigma dZ + d \left(\sum_{i=1}^{\pi(t)} (\xi_i - 1) \right). \quad (2.3)$$

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

232 It is assumed that the investor can earn a (continuously compounded) rate r_ℓ on cash deposits,
 233 and borrow at a rate of $r_b > 0$, with $r_\ell < r_b$. In the absence of control, the dynamics of the amount
 234 $B(t)$ invested in the risk-free asset are given by

$$235 \quad dB(t) = \mathcal{R}(B(t)) B(t) dt, \quad \text{where } \mathcal{R}(B(t)) = r_\ell + (r_b - r_\ell) \mathbb{I}_{[B(t) < 0]}. \quad (2.4)$$

236 We make the standard assumption that the real world drift rate μ of S is strictly greater than r_ℓ . Since
 237 there is only one risky asset, for a constant risk-aversion parameter, it is never MV-optimal to short
 238 stock. For the case of a risk aversion parameter inversely proportional to wealth, which we also will
 239 investigate in Section 5.5, we explicitly impose a short-selling restriction, as suggested in Bensoussan
 240 et al. (2014). Therefore, in all cases we allow only for $S(t) \geq 0$, $t \in [0, T]$. In contrast, we do allow
 241 short positions in the risk-free asset, i.e. it is possible that $B(t) < 0$, $t \in [0, T]$.

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

drives rotation from bonds into equities", by Kate Allen, which appeared in the Financial Times (FT) on January 16, 2018.

2.2 Dynamics of the controlled system

We denote by $X(t) = (S(t), B(t))$, $t \in [0, T]$, the multi-dimensional 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 C_t is defined as a double, possibly finite, sequence (Oksendal and Sulem, 2005)

$$C_t = \{t_1, t_2, \dots, t_n ; \eta_1, \eta_2, \dots, \eta_n, \dots\}_{n \leq n_{\max}} = \{\{t_n, \eta_n\}\}_{n \leq n_{\max}}, \quad n_{\max} \leq \infty. \quad (2.5)$$

Here, intervention times $t \leq t_1 < \dots < t_{n_{\max}} < T$ 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 $C_t^* = (\{t_n, \eta_n^*\})_{n \leq n_{\max}}$, $n_{\max} \leq \infty$, 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 intervention⁶ times (see for example equation (3.7) below). In what follows, we will consider both “continuous rebalancing” - see Section 5.2 (where, as $\max_n (t_n - t_{n-1}) \rightarrow 0$, 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

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

where the transaction costs have been taken into account.

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

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

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

⁶As 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”.

2.3 Admissible portfolios

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

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

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:

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

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

where

$$\Omega^\infty = [0, \infty) \times (-\infty, \infty). \quad (2.11)$$

The solvency condition can then be stated as

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

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

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

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

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

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 :

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

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

319 There are two important issues related to the pre-commitment MV problem (2.14). First, since
 320 variance does not satisfy the smoothing property of conditional expectation, dynamic programming
 321 cannot be applied directly to (2.14). To overcome this challenge, a technique is proposed in Li and
 322 Ng (2000); Zhou and Li (2000) to embed (2.14) in a new optimization problem, often referred to as
 323 the embedding problem, which can be solved using the dynamic programming principle. We refer the
 324 reader to Dang and Forsyth (2014); Dang et al. (2016); Wang and Forsyth (2010) for the numerical
 325 treatment of the problem as well as a discussion of technical issues.

326 It is well-known that, although dynamic programming can be used to solve the embedding problem,
 327 the obtained optimal controls remain time-inconsistent (see Bjork et al. (2016); Bjork and Murgoci
 328 (2014)). To explain the time-inconsistency issue further, with a slight abuse of notation, we denote
 329 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
 330 the pre-commitment approach, the “time-inconsistency” phenomenon means that, in general,

$$331 \quad \mathcal{C}_{t,u}^* \neq \mathcal{C}_{t',u}^*, \quad t' > t, \quad u \in [t', T]. \quad (2.16)$$

332 Simply put, (2.16) indicates that the optimal control for the same future time u , but computed at
 333 different prior times t and t' , are not necessarily the same. We conclude this subsection by referring
 334 the reader to Vigna (2014) an interesting alternative view of the notion of time-inconsistency.

335 2.4.2 Time-consistent approach

336 As discussed in Basak and Chabakauri (2010); Bjork et al. (2016); Bjork and Murgoci (2014); Hu
 337 et al. (2012), in the time-consistent approach, a “time-consistency” constraint is imposed on (2.14),
 338 giving the *time-consistent MV* (TCMV) problem as

$$339 \quad (TCMV_t(\rho)) : \quad V(s, b, t) = \sup_{\mathcal{C}_t \in \mathcal{A}} \left(E_{\mathcal{C}_t}^{x,t} [W(T)] - \rho \text{Var}_{\mathcal{C}_t}^{x,t} [W(T)] \right), \quad (2.17)$$

$$340 \quad \text{s. t.} \quad \mathcal{C}_{t,u}^* = \mathcal{C}_{t',u}^*, \quad \text{for all } t' \geq t \text{ and } u \geq t'. \quad (2.18)$$

341 Here, the time-consistency constraint (2.18) ensures that that the resulting optimal strategy for MV
 342 portfolio optimization is, in fact, time-consistent. As a result, the MV portfolio optimization (2.17)-
 343 (2.18) admits the Bellman optimality principle, and hence, dynamic programming can be applied
 344 directly to (2.17)-(2.18) to compute optimal controls and the TCMV efficient frontier. See, for example
 345 Wang and Forsyth (2011), for the pure-diffusion case.

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

353 *Remark 2.1.* (Game-theoretic perspective; notion of optimality). In Bjork and Murgoci (2014), the
 354 terminology “equilibrium” control is used as opposed to “optimal” control, since the time-consistent
 355 optimal control \mathcal{C}_t^* satisfies the conditions of a subgame perfect Nash equilibrium control. We will
 356 follow the example of Basak and Chabakauri (2010); Cong and Oosterlee (2016); Li and Li (2013);
 357 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 $U(s, b, t)$ and $Q(s, b, t)$ as follows

$$V(s, b, t) = U(s, b, t) - \rho Q(s, b, t) + \rho(U(s, b, t))^2, \quad (3.1)$$

where

$$U(s, b, t) = E_{C_t^*}^{x,t} [W(T)], \quad (3.2)$$

$$Q(s, b, t) = E_{C_t^*}^{x,t} [(W(T))^2], \quad (3.3)$$

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

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

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

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

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

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 $[0, T]$

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

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

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

$$402 \quad \mathcal{C} = \{ \eta_n \in \mathcal{Z} : n = 0, \dots, n_{\max} \}, \quad (3.8)$$

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

$$407 \quad \mathcal{C}_n = \{ \eta_n, \dots, \eta_{n_{\max}} \} \subseteq \mathcal{C} = \{ \eta_0, \dots, \eta_{n_{\max}} \}. \quad (3.9)$$

408 Subsequently, we use

$$409 \quad \mathcal{C}_n^* = \{ \eta_n^*, \dots, \eta_{n_{\max}}^* \} \quad (3.10)$$

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

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

$$414 \quad (TCMV_{t_n}(\rho)) : \quad V(s, b, t_n) = \sup_{\mathcal{C}_n \in \mathcal{A}} \left(E_{\mathcal{C}_n}^{x, t_n} [W(T)] - \rho \text{Var}_{\mathcal{C}_n}^{x, t_n} [W(T)] \right) \quad (3.11)$$

$$415 \quad \text{s.t. } \mathcal{C}_n = \{ \eta_n, \mathcal{C}_{n+1}^* \} := \{ \eta_n, \eta_{n+1}^*, \dots, \eta_{n_{\max}-1}^*, \eta_{n_{\max}}^* \} \quad (3.12)$$

416 where \mathcal{C}_{n+1}^* is optimal for problem $(TCMV_{t_{n+1}}(\rho))$.

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

$$423 \quad \mathcal{C}_n = \{ \eta, \eta_{n+1}^*, \dots, \eta_{n_{\max}}^* \} = \{ \eta, \mathcal{C}_{n+1}^* \} \quad (3.13)$$

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

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

$$429 \quad U(s, b, t_n) = E_{\eta_n^*}^{x, t_n} [U(S(t_{n+1}^-), B(t_{n+1}^-), t_{n+1})], \quad (3.14)$$

$$430 \quad Q(s, b, t_n) = E_{\eta_n^*}^{x, t_n} [Q(S(t_{n+1}^-), B(t_{n+1}^-), t_{n+1})], \quad (3.15)$$

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

$$433 \quad U(s, b, T) = U(s, b, t_{n_{\max}}) = W(s, b), \quad (3.16)$$

$$434 \quad Q(s, b, T) = Q(s, b, t_{n_{\max}}) = (W(s, b))^2. \quad (3.17)$$

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

$$436 \quad V(s, b, t_n) = \sup_{\eta \in \mathcal{Z}} \left\{ E_{\eta}^{x, t_n} [U(S(t_{n+1}^-), B(t_{n+1}^-), t_{n+1})] - \rho E_{\eta}^{x, t_n} [Q(S(t_{n+1}^-), B(t_{n+1}^-), t_{n+1})] \right. \\ 437 \quad \left. + \rho (E_{\eta}^{x, t_n} [U(S(t_{n+1}^-), B(t_{n+1}^-), t_{n+1})])^2 \right\}, \quad (3.18)$$

438 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
439 the discretized version of the intervention operator \mathcal{M} , defined in (3.6).

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

$$442 \quad \eta_n^* \in \arg \max_{\eta \in \mathcal{Z}} \left\{ E_{\eta}^{x, t_n} [U(S(t_{n+1}^-), B(t_{n+1}^-), t_{n+1})] - \rho E_{\eta}^{x, t_n} [Q(S(t_{n+1}^-), B(t_{n+1}^-), t_{n+1})] \right. \\ 443 \quad \left. + \rho (E_{\eta}^{x, t_n} [U(S(t_{n+1}^-), B(t_{n+1}^-), t_{n+1})])^2 \right\}. \quad (3.19)$$

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

446 3.2 Computation of expectations

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

$$448 \quad \bar{U}(s, b, \tau) = U(s, b, T - t), \quad \bar{Q}(s, b, \tau) = Q(s, b, T - t), \quad \bar{V}(s, b, \tau) = V(s, b, T - t), \quad (3.20)$$

449 and hence (3.1) becomes

$$450 \quad \bar{V}(s, b, \tau) = \bar{U}(s, b, \tau) - \rho \bar{Q}(s, b, \tau) + \rho (\bar{U}(s, b, \tau))^2 \quad (3.21)$$

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

$$452 \quad \{\tau_n = T - t_{n_{\max} - n} : n = 0, 1, \dots, n_{\max}\}. \quad (3.22)$$

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

$$455 \quad \hat{U}_{\eta}^n(s, b) = E_{\eta}^{x, \tau_n} [\bar{U}(S(\tau_{n-1}^+), B(\tau_{n-1}^+), \tau_{n-1}^+)], \quad (3.23)$$

$$456 \quad \hat{Q}_{\eta}^n(s, b) = E_{\eta}^{x, \tau_n} [\bar{Q}(S(\tau_{n-1}^+), B(\tau_{n-1}^+), \tau_{n-1}^+)]. \quad (3.24)$$

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

$$460 \quad \Psi_{\tau}(s, b, \tau) - \mathcal{L}\Psi(s, b, \tau) - \mathcal{J}\Psi(s, b, \tau) = 0 \quad (s, b, \tau) \in \mathcal{N} \times (\tau_{n-1}^+, \tau_n^-] \quad (3.25)$$

$$461 \quad \text{with initial condition } \Psi(s, b, \tau_{n-1}^+) = \bar{U}(s, b, \tau_{n-1}) \quad (3.26)$$

462 and

$$463 \quad \Phi_{\tau}(s, b, \tau) - \mathcal{L}\Phi(s, b, \tau) - \mathcal{J}\Phi(s, b, \tau) = 0 \quad (s, b, \tau) \in \mathcal{N} \times (\tau_{n-1}^+, \tau_n^-] \quad (3.27)$$

$$464 \quad \text{with initial condition } \Phi(s, b, \tau_{n-1}^+) = \bar{Q}(s, b, \tau_{n-1}) \quad (3.28)$$

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

$$466 \quad \bar{U}(s, b, 0) = W(s, b), \quad \bar{Q}(s, b, 0) = (W(s, b))^2. \quad (3.29)$$

467 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.
 468 Then, for a given arbitrary impulse $\eta \in \mathcal{Z}$, we obtain the ‘‘candidate’’ expectation values $\hat{U}_\eta^n(s, b)$ and
 469 $\hat{Q}_\eta^n(s, b)$ by

$$470 \quad \hat{U}_\eta^n(s, b) = \Psi(S(\tau_n^+), B(\tau_n^+), \tau_n^-), \quad (3.30)$$

$$471 \quad \hat{Q}_\eta^n(s, b) = \Phi(S(\tau_n^+), B(\tau_n^+), \tau_n^-), \quad (3.31)$$

472 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
 473 leverage constraint (2.13). Finally, using (3.30)-(3.31), we can find the optimal impulse value η_n^* via
 474 $\eta_n^* \in \arg \max_{\eta \in \mathcal{Z}} \left\{ \hat{U}_\eta^n(s, b) - \rho \hat{Q}_\eta^n(s, b) + \rho \left(\hat{U}_\eta^n(s, b) \right)^2 \right\}$.

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

$$477 \quad \bar{U}(s, b, \tau_n^-) = \bar{U}(0, W(s, b)e^{\mathcal{R}(s+b)\tau_n}, 0),$$

$$478 \quad \bar{Q}(s, b, \tau_n^-) = \bar{Q}(0, W(s, b)e^{\mathcal{R}(s+b)\tau_n}, 0), \quad (s, b) \in \mathcal{B}. \quad (3.32)$$

479 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** $n = 1, \dots, n_{\max}$ **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 $\eta_n^* \in \arg \max_{\eta \in \mathcal{Z}} \left\{ \hat{U}_\eta^n(s, b) - \rho \hat{Q}_\eta^n(s, b) + \rho \left(\hat{U}_\eta^n(s, b) \right)^2 \right\}$;
 - 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 $V(s, b, \tau_{n_{\max}}) = \bar{U}(s, b, \tau_{n_{\max}}) - \rho \bar{Q}(s, b, \tau_{n_{\max}}) + \rho (\bar{U}(s, b, \tau_{n_{\max}}))^2$;
-

480

481 *Remark 3.1.* (Convergence of numerical solution). Since the viscosity solution theory (Crandall et al.
 482 (1992)) does not apply in this case, we have no proof that Algorithm 3.1 converges to an appropriately
 483 defined (weak) solution of the corresponding extended HJB quasi-integrovariational inequality in the
 484 limit as $\Delta\tau \rightarrow 0$. However, we can show, as in Cong and Oosterlee (2016); Wang and Forsyth (2011),
 485 that our numerical solution converges to known analytical solutions available in special cases. Where
 486 no analytical solutions are available, the numerical PDE results are validated using Monte Carlo
 487 simulation.

488 4 Localization

489 4.1 Semi-Lagrangian timestepping scheme

490 Recall the definition of the operator \mathcal{L} , defined in (3.4). We observe that the PIDEs (3.25) and
 491 (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.
 492 involves partial derivatives with respect to both s and b . Direct implementation would be therefore
 493 computationally expensive.

494 With this in mind, we introduce the semi-Lagrangian timestepping scheme proposed in Dang and
 495 Forsyth (2014). The intuition behind the the semi-Lagrangian timestepping scheme is that, instead of
 496 obtaining the PIDEs by modelling the change (via Ito's lemma) in a test function $f(S(\tau), B(\tau), \tau)$
 497 with both S and B varying, we consider the Lagrangian derivative along the trajectory where B is
 498 held fixed over the length of the timestep. Specifically, we model the change in $f(S(\tau), B(\tau), \tau)$ with
 499 $(S(\tau), B(\tau) = b)$ for $\tau \in [\tau_{n-1}^+, \tau_n^-]$, with interest paid only at the end of the timestep, i.e. at time
 500 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
 501 settlement (payment or receipt) of interest due for the time interval $[\tau_{n-1}, \tau_n]$. Along this trajectory,
 502 the partial derivative of the test function $f(s, b, \tau)$ with respect to the b -variable is zero, resulting in
 503 a decoupling of the PIDE for every value of the b -variable.

504 We emphasize that the above argument is an intuitive explanation of the semi-Lagrangian scheme.
 505 In fact, we can prove rigorously that in the limit as $\Delta\tau \rightarrow 0$, this treatment converges to the case
 506 where interest is paid continuously.⁷ Moreover, this approach is also valid for discrete rebalancing,
 507 regardless of whether the interest is paid continuously or discretely.

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

$$509 \quad \Psi_b(s, b, \tau) = \Phi_b(s, b, \tau) = 0, \quad (s, b, \tau) \in \mathcal{N} \times (\tau_{n-1}^+, \tau_n^-],$$

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

$$511 \quad \mathcal{P}f(s, b, t) = (\mu - \lambda\kappa)sf_s + \frac{1}{2}\sigma^2s^2f_{ss} - \lambda f. \quad (4.1)$$

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

$$514 \quad \begin{aligned} \Psi_\tau(s, b, \tau) - \mathcal{P}\Psi(s, b, \tau) - \mathcal{J}\Psi(s, b, \tau) &= 0, & (s, b, \tau) \in \mathcal{N} \times (\tau_{n-1}^+, \tau_n^-] \\ \text{with initial condition } \Psi(s, b, \tau_{n-1}^+) &= \bar{U}(s, b, \tau_{n-1}), \end{aligned} \quad (4.2)$$

516 and

$$517 \quad \begin{aligned} \Phi_\tau(s, b, \tau) - \mathcal{P}\Phi(s, b, \tau) - \mathcal{J}\Phi(s, b, \tau) &= 0, & (s, b, \tau) \in \mathcal{N} \times (\tau_{n-1}^+, \tau_n^-] \\ \text{with initial condition } \Phi(s, b, \tau_{n-1}^+) &= \bar{Q}(s, b, \tau_{n-1}). \end{aligned} \quad (4.3)$$

519 The second consequence of semi-Lagrangian timestepping is that the calculation of the value of
 520 $S(\tau_n^-)$, 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
 521 the payment of interest at time τ_n :

$$522 \quad S(\tau_n^+) = \left(s + be^{\mathcal{R}(b)\Delta\tau}\right) - \eta - c_1 - c_2 \cdot |S(\tau_n^+) - s|. \quad (4.4)$$

⁷See Dang and Forsyth (2014) for the consistency proof in the context of the pre-commitment mean-variance problem.

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

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

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

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

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

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

where $r_{\max} = \max(r_b, r_\ell)$. Note that Ω_{s_0} is simply the boundary where $s = 0$, while $\Omega_{\mathcal{B}}$ is the localized insolvency region and Ω_{in} is the interior of the localized solvency region. The purpose of both Ω_{s^*} and $\Omega_{b_{\max}}$ is to act as buffer regions for the risky asset jumps and the risk-free asset interest payments, respectively, so that these events do not take us outside the computational grid (see Dang and Forsyth (2014) and d'Halluin et al. (2005)). 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).

Following the steps in Dang and Forsyth (2014), we have the following localized problem for Ψ :

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

Here,

$$\mathcal{J}_\ell f(s, b, \tau) = \lambda \int_0^{s_{\max}/s} f(\xi s, b, \tau) p(\xi) d\xi. \quad (4.11)$$

We briefly discuss each equation forming part of (4.10). The PIDE in Ω_{in} is essentially (4.2), with the localized jump operator \mathcal{J}_ℓ given in (4.11). The result in Ω_{s^*} is obtained as follows. Based on the initial condition (3.29), together with the definition of $W(s, b)$, we have the approximation $\Psi(s \rightarrow \infty, b, \tau = 0) \simeq (1 - c_2)s$, where c_2 is the proportional transaction cost. For an arbitrary $\tau \in [\tau_{n-1}^+, \tau_n^-]$, it is therefore reasonable to use the asymptotic form $\Psi(s \rightarrow \infty, b, \tau) \simeq A(\tau)s$. Provided that s^* in (4.5) is chosen sufficiently large so that this asymptotic form provides a reasonable approximation to Ψ in Ω_{s^*} , we substitute $\Psi(s, b, \tau) \simeq A(\tau)s$ into the PIDE (4.2) to obtain the corresponding equation for Ω_{s^*} in (4.10). Similar reasoning applies to the region $\Omega_{b_{\max}}$, except that the initial condition (3.29) now gives $\Psi(s, b \rightarrow \pm\infty, \tau = 0) \simeq b$, which leads to the asymptotic form $\Psi(s, |b| > |b_{\max}|, \tau) \simeq C(s, \tau)b$ to be used in $\Omega_{b_{\max}}$. Setting $b = b_{\max}$ and $b = -b_{\max}$ (which is inside Ω rather than $\Omega_{b_{\max}}$), the computed solution in Ω can be used to obtain the approximation for Ψ in $\Omega_{b_{\max}}$ shown above. Finally, at $s = 0$, the PIDE (4.2) degenerates into the result shown for Ω_{s_0} , while for $\tau = \tau_{n-1}$, we have the initial condition from (4.2) applicable to all $(s, b) \in \Omega$. More details on this approach be found in Dang and Forsyth (2014).

Using similar arguments, the localized problem for Φ can be obtained as follows:

$$\begin{aligned}
\Phi_\tau(s, b, \tau) - \mathcal{P}\Phi(s, b, \tau) - \mathcal{J}_\ell\Phi(s, b, \tau) &= 0, & (s, b, \tau) \in \Omega_{in} \times [\tau_{n-1}^+, \tau_n^-] \\
\Phi_\tau(s, b, \tau) - [2\mu + \sigma^2 + \lambda\kappa_2]\Phi(s, b, \tau) &= 0, & (s, b, \tau) \in \Omega_{s^*} \times [\tau_{n-1}^+, \tau_n^-], \\
\Phi(s, b, \tau) - \bar{Q}(0, b, \tau_{n-1}) &= 0, & (s, b, \tau) \in \Omega_{s_0} \times [\tau_{n-1}^+, \tau_n^-], \\
\Phi(s, |b| > |b_{\max}|, \tau) - \left(\frac{b}{b_{\max}}\right)^2 \Phi(s, \text{sgn}(b)b_{\max}, \tau) &= 0, & (s, b, \tau) \in \Omega_{b_{\max}} \times [\tau_{n-1}^+, \tau_n^-] \\
\text{with } \Phi(s, b, \tau = \tau_{n-1}) - \bar{Q}(s, b, \tau_{n-1}) &= 0, & (s, b) \in \Omega.
\end{aligned} \tag{4.12}$$

We solve the localized problems (4.10)-(4.12) 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 $\{s_i : i = 1, \dots, i_{\max}\}$ and b -direction $\{b_j : j = 1, \dots, j_{\max}\}$, 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 $\mathcal{Z}_h = \{b_j : j = 1, \dots, j_{\max}\} \cap \mathcal{Z}$ to be the discretization of the admissible impulse space. We use linear interpolation onto the computational grid if the spatial point $(s_i, b_j e^{\mathcal{R}(b_j)\Delta\tau})$, 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}_ℓ in (4.11) 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 :

$$\mathcal{Y}_h = \bigcup_{\rho \geq 0} \left(\sqrt{\left(Var_{C_0^*}^{t=0} [W(T)]\right)_h}, \left(E_{C_0^*}^{t=0} [W(T)]\right)_h \right)_\rho, \tag{4.13}$$

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 $\{\rho_{\min}, \dots, \rho_{\max}\} \subset [0, \infty)$, executing Algorithm 3.1 gives us the following quantities:

$$U_0(W_{init}) \simeq \left(E_{C_0^*}^{(s=0, b=W_{init}), t=0} [W(T)]\right)_h, \quad Q_0(W_{init}) \simeq \left(E_{C_0^*}^{(s=0, b=W_{init}), t=0} [(W(T))^2]\right)_h,$$

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

$$\left(Var_{C_0^*}^{t=0} [W(T)]\right)_h = Q_0(W_{init}) - (U_0(W_{init}))^2, \quad \left(E_{C_0^*}^{t=0} [W(T)]\right)_h = U_0(W_{init}). \tag{4.14}$$

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 (4.12). 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 $\mathcal{O}(1/h^5)$.

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 $0 = \tilde{t}_0 \leq \tilde{t}_1 < \dots < \tilde{t}_{m_{\max}} < T$, where t_0 is the inception of the investment. With the change of variable $\tau = T - t$, the set of intervention times become

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

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 $\tilde{\tau}_m^-$, $m = 1, \dots, m_{\max}$, possibly using multiple timesteps for the solution of the corresponding PIDE, to obtain $\Psi(s, b, \tilde{\tau}_m^-)$ and $\Phi(s, b, \tilde{\tau}_m^-)$. 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.⁸ 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*,⁹ augmented by National Bureau of Economic Research (NBER) short-term government bond yields for *1926:1 - 1933:12*¹⁰ 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 .
- 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.¹¹

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

⁸More 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.

⁹See <http://research.stlouisfed.org/fred2/series/TB3MS>.

¹⁰See <http://www.nber.org/databases/macroeconomy/contents/chapter13.html>.

¹¹CPI 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>.

635 Mancini (2011); Mancini (2009) has been used, as applied in Dang and Forsyth (2016); Forsyth and
 636 Vetzal (2017), for the calibration. Specifically, if $\Delta\hat{X}_i$ denotes the i th inflation-adjusted, detrended
 637 log return in the historical risky asset index time series, we identify a jump in period i if

$$638 \quad \left| \Delta\hat{X}_i \right| > \alpha\hat{\sigma}\sqrt{\Delta t}, \quad (5.1)$$

639 where $\hat{\sigma}$ is the estimate of the diffusive volatility, Δt is the time period over which the log return has
 640 been calculated, and α is the “threshold parameter” for identifying a jump. Distinguishing between
 641 “up” and “down” jumps for the Kou model is achieved using upward and downward jump indicators -
 642 see Forsyth and Vetzal (2017) for further details, including the simultaneous estimation of the diffusive
 643 volatility. We will use $\alpha = 3$ in what follows - in other words, we would only detect a jump in the
 644 historical time series if the (absolute, inflation-adjusted, and detrended) log return in that period
 645 exceeds 3 standard deviations of the “geometric Brownian motion change”, which is a very unlikely
 646 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).

Parameters	Models		
	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
\tilde{m} (log jump multiplier mean)	n/a	-0.0700	n/a
$\tilde{\gamma}$ (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

648 5.2 Convergence analysis

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

652 5.2.1 Analytical solutions

653 Analytical solutions for the time-consistent problem are available if the risky asset follows GBM (see
 654 Basak and Chabakauri (2010)) or any of the commonly-encountered jump models, including the Mer-
 655 ton and Kou models (see Bjork and Murgoci (2010) and Zeng et al. (2013)), under the following
 656 assumptions: (i) continuous rebalancing of the portfolio, (ii) trading continues in the event of in-
 657 solvency, (iii) no investment constraints or transaction costs, and (iv) same lending and borrowing
 658 rate ($= r$). Under these assumptions, the efficient frontier solution is given by

$$659 \quad E_{C_0^*}^{t=0} [W(T)] = W(0) e^{rT} + \frac{1}{2\rho} \left[\frac{(\mu - r)^2}{\sigma^2 + \lambda\kappa_2} \right] T,$$

$$660 \quad Stdev_{C_0^*}^{t=0} [W(T)] = \frac{1}{2\rho} \left(\frac{\mu - r}{\sqrt{\sigma^2 + \lambda\kappa_2}} \right) \sqrt{T}, \quad (5.2)$$

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

662 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

663
664 Table 5.3 illustrates the numerical convergence analysis for an initial wealth of $W(0) = 100$,
665 maturity $T = 10$ years, and scalarization parameter $\rho = 0.005$. For illustrative purposes, we assume
666 the risky asset follows the Merton model - qualitatively similar results are obtained if the Kou or GBM
667 models are assumed. The “Error” column shows the difference between the analytical solution and
668 the PDE solution, while the “Ratio” column shows the ratio of successive errors for each increase in
669 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

670

671 5.2.2 Monte Carlo validation

672 Consider now the following case where analytical solutions are *not* available: we assume discrete
673 periodic rebalancing of the portfolio at the end of each year, with liquidation in the event of insolvency,
674 and a maximum allowable leverage ratio of $q_{\max} = 1.5$. Additionally, we assume the risky asset follows
675 the Kou model, with initial wealth of $W(0) = 100$, maturity $T = 20$ years, and scalarization parameter
676 $\rho = 0.0014$. For the numerical PDE solution, using 7,280 equal timesteps, and 1,121 and 2,209 s -grid
677 and b -grid nodes, respectively, we obtain the following approximations to the expectation and standard
678 deviation:

$$679 \left(E_{C_0^*}^{t=0} [W(T)], Stdev_{C_0^*}^{t=0} [W(T)] \right) = (544.58, 400.20). \quad (5.3)$$

680 At each timestep of our numerical PDE procedure, we output and store the computed optimal strategy
681 for each discrete state value. We then carry out Monte Carlo simulations for the portfolio (using the
682 specified parameters) from $t = 0$ to $t = T$, rebalancing the portfolio in accordance with the stored
683 PDE-computed optimal strategy at each discrete rebalancing time. If necessary, we use interpolation to
684 determine the optimal strategy for a given state value. We then compare the Monte Carlo computed
685 means and standard deviations of the terminal wealth with the corresponding values computed by
686 the numerical PDE method, given in (5.3). The results are shown in Table 5.4. Note that, for the

687 MC method, due to the possibility of insolvency, it is not possible to take finite timesteps between
 688 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%

688 We observe that, as the number of Monte Carlo simulations and timesteps increase, the Monte
 689 Carlo computed means and standard deviations converge to the corresponding values computed by
 690 the numerical PDE method, given in (5.3).
 691

692 5.3 Time-consistent MV efficient frontiers

693 In this subsection, we study time-consistent MV efficient frontiers. In particular, we consider the
 694 impact of investment constraints and other assumptions, including transaction costs, we construct five
 experiments as outlined in Table 5.5.

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

695 We highlight the following:
 696

- 697 • The interest rates for Experiments 4 and 5 were obtained by assuming that the approximate
 698 relationship between current interest rates paid on margin accounts in relation to current 3-
 699 month US T-bill rates¹², also holds in relation to the historically observed 3-month US T-bill
 700 rates used to obtain the constant rate of 0.00623 (see Table 5.1).

¹²The 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.

- 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).

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, for illustrative purposes, we use very fine temporal and spatial timestep sizes, namely 7,280 equal timesteps, and 561 and 1,105 s -grid and b -grid nodes, respectively. With these very fine stepsizes, the calculation of the mean and the standard deviation of a point on the efficient frontier, i.e. corresponds to one ρ value, takes about two hours to obtain.¹³ Since different points on the efficient frontier, can be computed in parallel, it takes about the same amount time to trace out an entire efficient frontier. However, for practical purposes, much coarser stepsizes can be used, and hence significantly less computation time can be achieved. For example, we can obtain a mean and standard deviation with a relative error of less than 10% of the respective results reported below in only about 10 minutes, if we use half the number of partition points in both the s -grid and b -grid, and assume weekly, instead of daily, interest payments. The algorithm, therefore, allows for the computation of the solution within a very reasonable time.

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.




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

¹³The algorithm was coded in C++ and run on a server with 12 physical cores (+12 hyper-threaded cores), namely 2 x Intel E5-2667 6-core 2.90 GHz with 256GB RAM.

729 in (5.2) to guide our intuition. Note that (5.2) can be re-arranged to give the expected value in terms
730 of the standard deviation,

$$731 \quad E_{C_0^*}^{t=0} [W(T)] = W(0) e^{rT} + \left(\frac{\mu - r}{\sqrt{\sigma^2 + \lambda\kappa_2}} \right) \sqrt{T} \cdot \left(Stdev_{C_0^*}^{t=0} [W(T)] \right). \quad (5.4)$$

732 Fixing a standard deviation value on the efficient frontier, we observe that the effect of model
733 choice on the associated expected value on the efficient frontier is entirely due to the multiplier
734 $(\mu - r) / \sqrt{\sigma^2 + \lambda\kappa_2}$ in (5.4). With calibrated process parameters as given in Table 5.1, we have
735 combinations of parameters as given in Table 5.6. In particular, we conclude that the multiplier
736 $(\mu - r) / \sqrt{\sigma^2 + \lambda\kappa_2}$ is lower for the Kou model, due to the higher variance of the log-double exponential
737 distribution of the jump multipliers (resulting in a higher value of $\kappa_2 = \mathbb{E}[(\xi - 1)^2] = Var(\xi) + \kappa^2$)
738 compared to the that of the lognormal distribution in the case of the Merton model. We also note
739 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
$\kappa_2 = \mathbb{E}[(\xi - 1)^2]$	0.0000	0.0365	0.0844
$(\mu - r) / \sqrt{\sigma^2 + \lambda\kappa_2}$	0.4046	0.4103	0.3612

740
741 Returning to the results shown in Figure 5.1 where no analytical solutions are available, we conclude
742 the following. With the exception of parameters affecting the jump distribution, the other model
743 parameters (drift, diffusive volatility, jump intensity) of the Kou and Merton models in Table 5.1 are
744 very similar. Since the jump multipliers have a higher variance in the Kou model compared to the
745 Merton model (both calibrated to the same data), then for a given level of expected terminal wealth,
746 the Kou model results in a larger standard deviation of the terminal wealth. Consequently, the efficient
747 frontier is lower for the Kou model than for the Merton model. Furthermore, similar multiplier values
748 for the GBM and Merton models (observed above) imply that the relatively higher diffusive volatility
749 of the GBM model has a similar effect as the incorporation of jumps using the Merton model over this
750 long investment time horizon, resulting in similar efficient frontiers for the GBM and Merton models.

751 5.3.2 Investment constraints

752 The effect of investment constraints on the time-consistent MV efficient frontiers are shown in Figure
753 5.2 for the Kou model only, since the results for other models are qualitatively similar.

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

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

765 such a strategy. By comparison, the effect of additionally introducing transaction costs (Experiment
766 5) is relatively negligible.

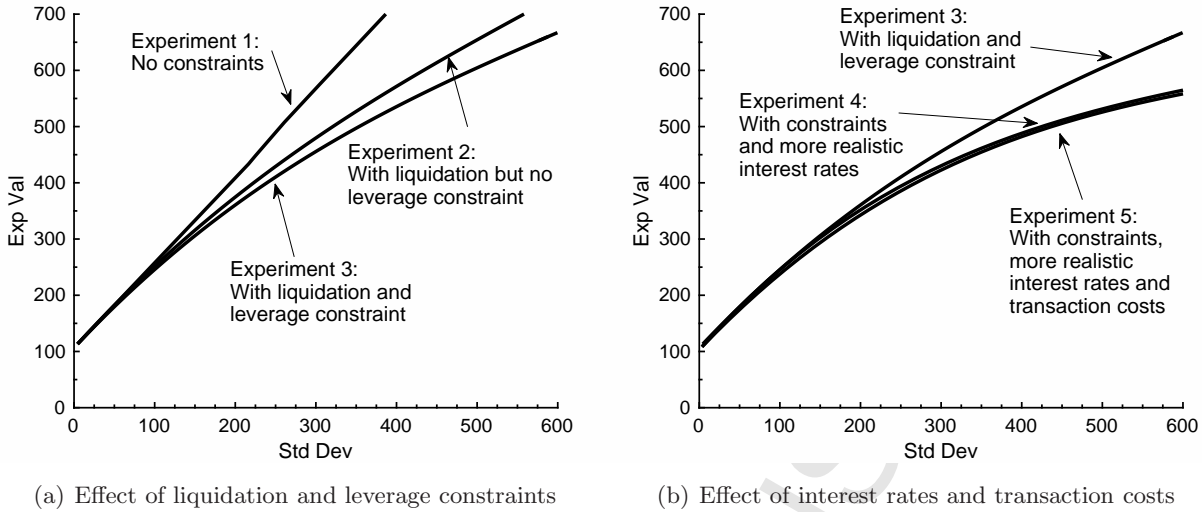


Figure 5.2: Time-consistent MV efficient frontiers - Kou model: Effect of investment constraints

767 5.4 Time-consistent MV vs. Pre-commitment MV strategies

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

777 5.4.1 Combined investment constraints

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

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

788 5.4.2 Leverage constraint

789 Next, we focus on the impact of the leverage constraint. Figure 5.4 illustrates the effect of dif-
790 ferent maximum leverage constraint q_{\max} assumptions on the efficient frontiers associated with the
791 pre-commitment and time-consistent MV problems. (In these tests, the solvency constraint is also
792 imposed.) Since leverage may not be allowed for pension fund investments, we also consider the effect

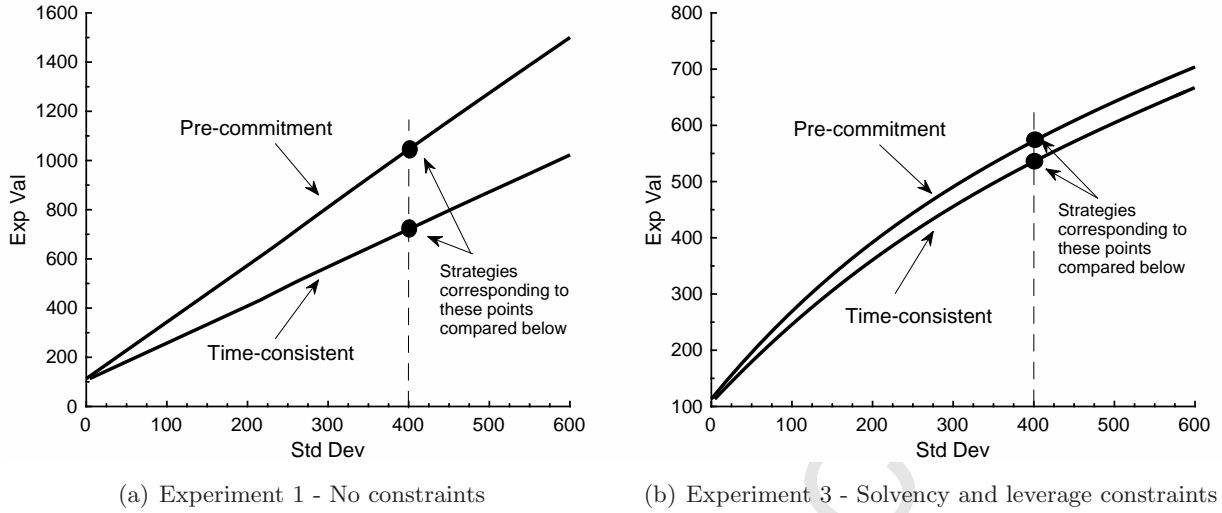


Figure 5.3: Pre-commitment MV vs. Time-consistent MV efficient frontiers - Kou model

793 of setting $q_{\max} = 1$ (so that the fraction of total wealth invested in the risky asset may not exceed
 794 one) in Experiment 3.

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

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

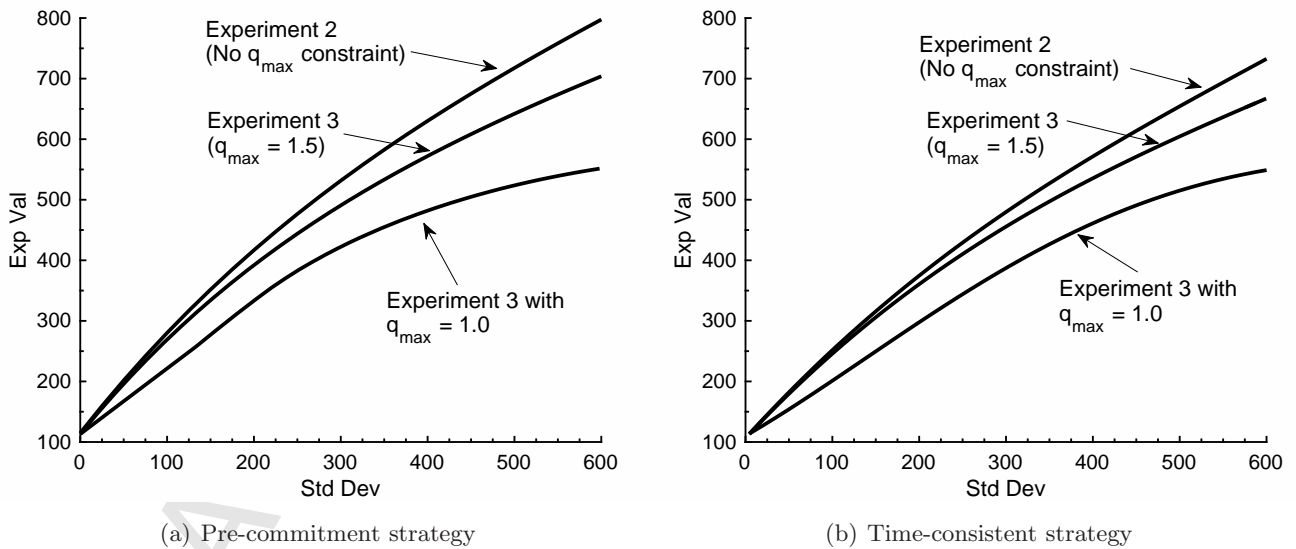


Figure 5.4: Pre-commitment MV vs. Time-consistent MV - Kou model: Effect of maximum leverage constraint q_{\max} .

804

805 5.4.3 Comparison of optimal controls

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

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

817 The results of this study are summarized in Figure 5.5 and Figure 5.6, where we show the median
 818 (50th percentile), as well as the 25th and 75th percentiles, of the distribution of the MV-optimal
 819 fraction of wealth invested in the risky asset over time.

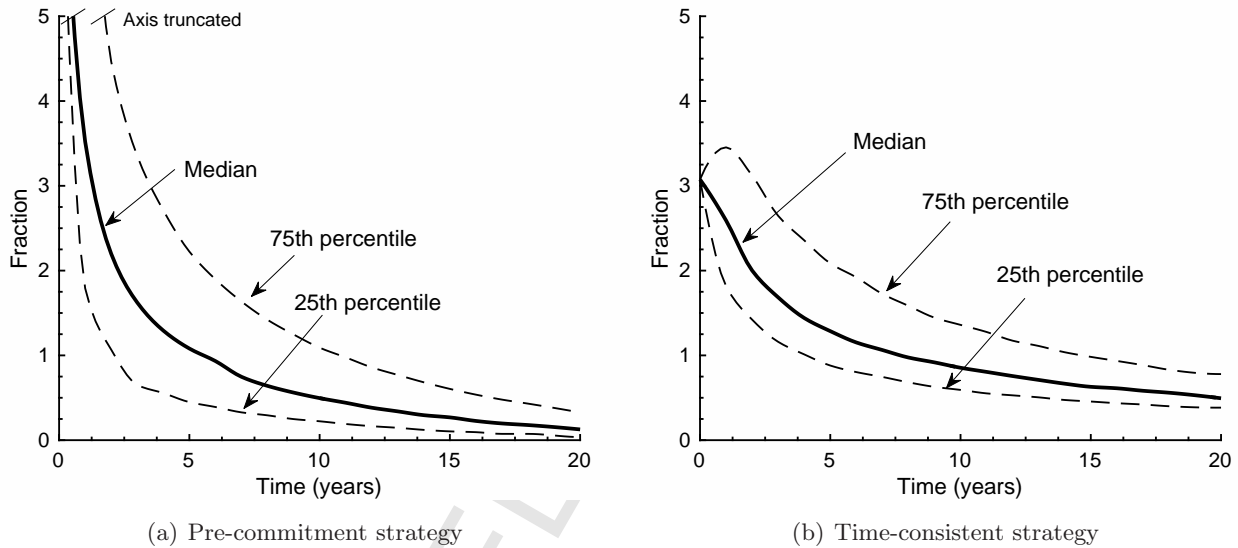


Figure 5.5: MV-optimal fraction of wealth in the risky asset: Kou model, Experiment 1, standard deviation of terminal wealth equal to 400.

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

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

835 If we impose liquidation in the event of insolvency, as well as a maximum leverage ratio of $q_{\max} =$

836 1.5, i.e. Experiment 3, Figure 5.6 shows that the resulting MV-optimal fraction of wealth invested
 837 in the risky asset changes substantially compared to Figure 5.5. In particular, we observe that the
 838 fraction invested in the risky asset for the pre-commitment strategy (Figure 5.6(a)) is more strongly
 839 affected by the maximum leverage constraint than the fraction for the time-consistent strategy (Figure
 840 5.6(b)). While this only considers only one point on the efficient frontier, where the standard deviation
 841 of terminal wealth is equal to 400, we have observed the higher sensitivity of the pre-commitment
 842 strategy to the maximum leverage constraint across the efficient frontier in Figure 5.4. This is due to
 843 the very large pre-commitment MV-optimal investment in the risky asset required during the early
 844 stages of the investment time period in order to achieve the implicit wealth target. On the other hand,
 845 it is interesting to observe that the pre-commitment strategy at the 25th percentile shows a very rapid
 846 de-risking compared to the time-consistent strategy.

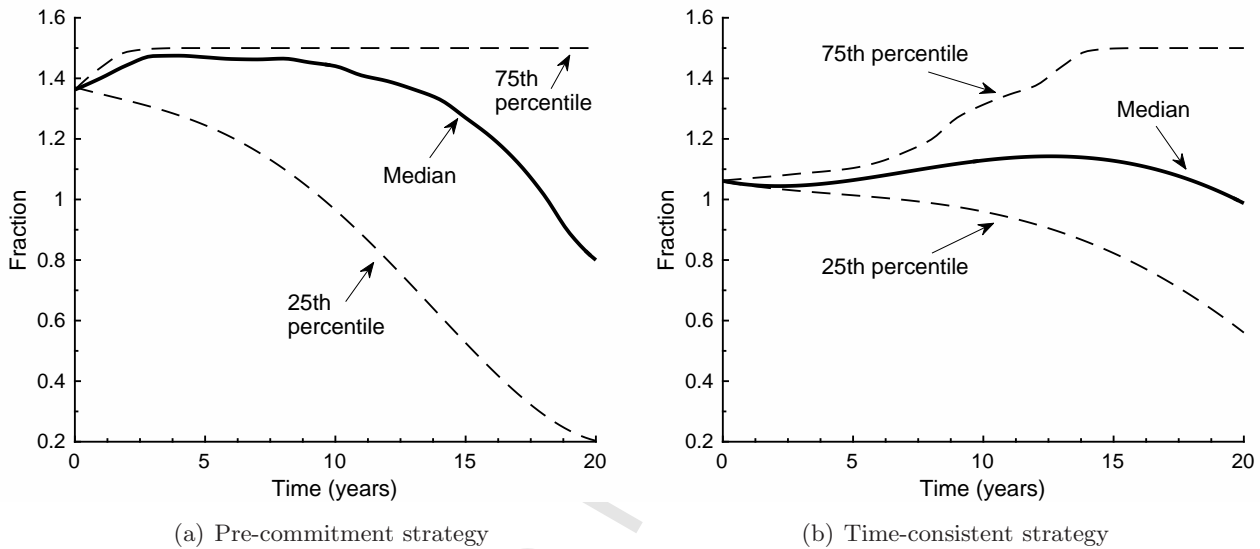


Figure 5.6: MV-optimal fraction of wealth in the risky asset: Kou model, Experiment 3, standard deviation of terminal wealth equal to 400.

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

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

862 In contrast, in the case of the time-consistent optimal control (Figure 5.7(b)), there are a number of
 863 qualitative similarities to the pre-commitment optimal control (Figure 5.7(a)), but also key differences.
 864 Both of the strategies are contrarian, in the sense that all else being equal, investment in the risky
 865 asset is increased if its returns in the past have been unfavourable. However, compared to the pre-

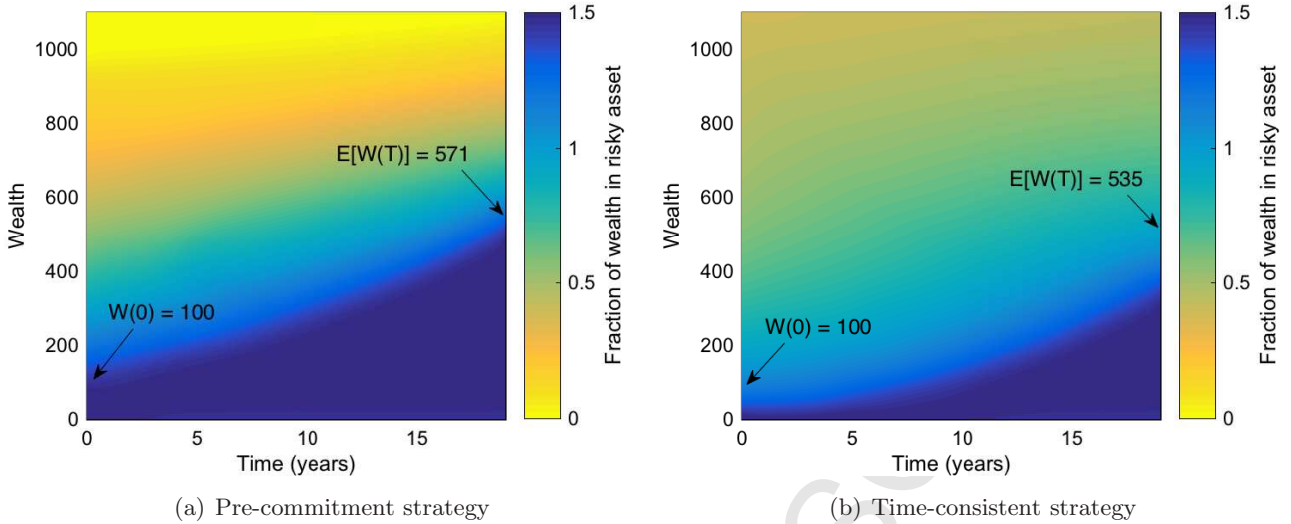


Figure 5.7: Optimal control as a fraction of wealth in risky asset: Kou model, Experiment 3, standard deviation of terminal wealth equal to 400.

866 commitment optimal control, the time-consistent optimal control requires generally higher investment
 867 in the risky asset if past returns have been favourable (resulting higher wealth), and lower investment
 868 in the risky asset if past returns have been unfavourable (resulting in lower wealth). Even if the risky
 869 asset performs extremely well, the time-consistent strategy never calls for zero exposure to the risky
 870 asset. Figure 5.7 also shows why the pre-commitment strategy would be more heavily impacted if the
 871 maximum leverage ratio is reduced; the time-consistent strategy calls for generally lower leverage, and
 872 would therefore be less sensitive to the maximum leverage constraint.

873 5.5 Effect of a wealth-dependent scalarization parameter

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

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

$$887 \sup_{C_t \in \mathcal{A}} \left(E_{C_t}^{x,t} [W(T)] - \rho(W(t)) \text{Var}_{C_t}^{x,t} [W(T)] \right) \quad (5.5)$$

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

$$889 \rho(W(t)) = \frac{\theta}{W(t)}, \quad \theta > 0 \quad (5.6)$$

890 where for each θ , we obtain a point on the resulting efficient frontier. The use of a wealth-dependent
 891 scalarization parameter has been popular in time-consistent MV literature within the non-constraint

892 setting, especially in insurance-related applications (see for example Zeng and Li (2011), Wei et al.
893 (2013), Li and Li (2013), as well as Liang and Song (2015)).

894 As discussed in Bensoussan et al. (2014), in a discrete-time setting, the choice (5.6) implies that the
895 shorting of stock might be MV-optimal. As such, the optimal wealth process may take on negative
896 values, potentially giving rise to a negative risk-aversion parameter. This would in turn cause the
897 MV objective (5.5) to become unbounded and the optimal control to exhibit economically irrational
898 decision making. For these reasons, following Bensoussan et al. (2014), we also impose a no short-
899 selling constraint on the risky asset in this section.

900 While some modifications to (5.6) are also considered in literature (for example, allowing θ to be
901 time-dependent), we explore the effect of using the definition (5.6) in our setting, specifically because
902 this simple case reveals how a seemingly reasonable definition of a wealth-dependent scalarization
903 parameter, when used in combination with investment constraints and liquidation in the event of
904 bankruptcy, can result in conclusions that are not economically reasonable.

905 Given Algorithm 3.1, implementing a wealth-dependent scalarization parameter such as (5.6) is
906 straightforward, since we simply replace ρ in the algorithm with $\rho(W(s, b)) = \theta/W(s, b)$, where
907 $W(s, b)$ is given by equation (2.8), without any further changes required. Varying $\theta > 0$ in this case
908 traces out the efficient frontier.

909 We consider Experiment 3 in Table 5.5 (in other words we impose both liquidation in bankruptcy
910 and a leverage constraint), since - as pointed out in Wang and Forsyth (2011) - allowing for negative
911 wealth in equation (5.6) would lead to inappropriate risk aversion coefficients. In Figure 5.8, the
912 efficient frontier obtained with a constant scalarization parameter ρ is compared with the efficient
913 frontier obtained with wealth-dependent scalarization parameter of the form (5.6). We observe a
914 similar result as in Wang and Forsyth (2011), where the case of continuous controls and no jumps was
915 investigated: the resulting time-consistent MV efficient frontier with a wealth-dependent scalarization
916 parameter is significantly lower than that obtained using a constant scalarization parameter. In
917 other words, given an acceptable level of risk as measured by variance, a strategy based on the wealth-
918 dependent scalarization parameter given by (5.6) would result in much lower expected terminal wealth,
and is therefore less efficient from a MV-optimization perspective.

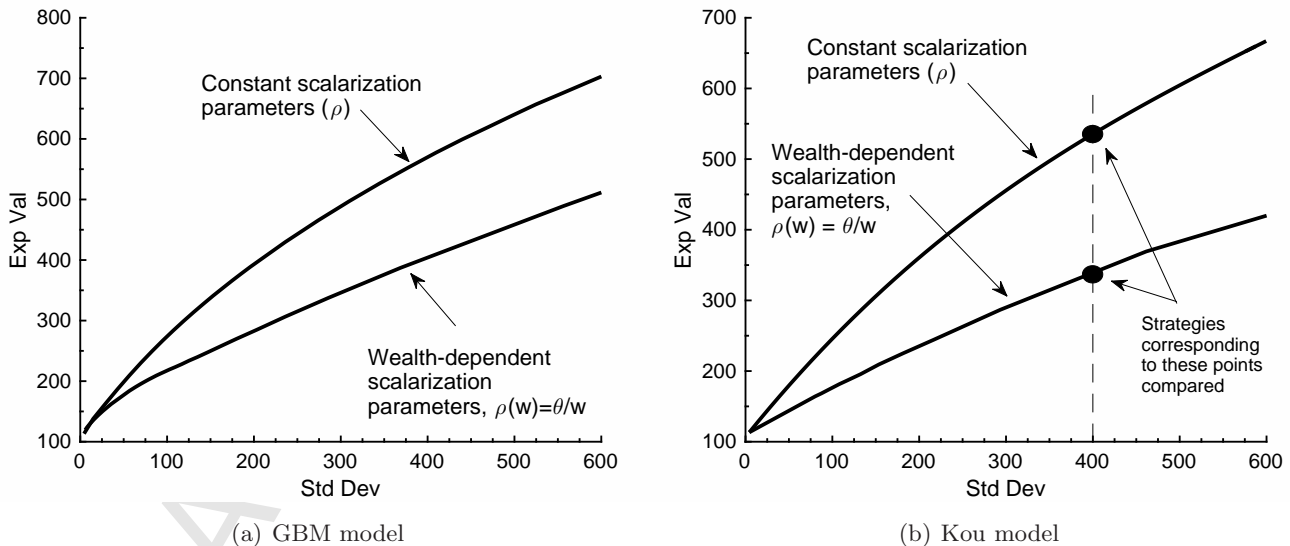



Figure 5.8: Time-consistent MV efficient frontiers - Experiment 3 (solvency and leverage constraints): Effect of using a constant scalarization parameter vs. using a wealth-dependent scalarization parameter of the form $\rho(w) = \theta/w$.

919

920 We now further compare the optimal trading strategies for the Kou model in both scenarios,

921 namely a constant scalarization parameter and a wealth-dependent scalarization parameter of the
 922 form (5.6). In this case, we pick two points on the efficient frontiers corresponding to a standard
 923 deviation of terminal wealth equal to 400, as highlighted in Figure 5.8(b). In Figure 5.9, we now
 924 compare the resulting MV-optimal strategies corresponding to these points. Specifically, proceeding
 925 as in Subsection 5.4.3, using Monte Carlo simulations and rebalancing the portfolio in accordance
 926 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.



(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.

927

928

929 Figure 5.9 (a) compares the median of the time-consistent MV-optimal fraction of wealth in the
 930 risky asset in both scenarios.¹⁴ Figure 5.9 (b) illustrates the heatmap of the time-consistent MV-
 931 optimal control (as the fraction of wealth invested in the risky asset) as a function of time and wealth
 932 in the case of a wealth-dependent scalarization parameter of the form (5.6). The heatmap for the
 933 time-consistent MV-optimal control in the case of a constant scalarization parameter (also for the
 934 Kou model, Experiment 3, and a standard deviation of terminal wealth equal to 400) is provided in
 Figure 5.7(b).

935

936

937 We make the following interesting observations. While the increase in exposure to the risky asset
 938 over time has been observed in the case of the wealth-dependent risk aversion parameter in the setting
 939 of no jumps, constraints or bankruptcy (see, for example, Bjork et al. (2014)), in the case of realistic
 940 investment constraints this is even more dramatic. Such observed dramatic impact can be explained
 941 as follows. The form of the wealth-dependent risk aversion in (5.6) implies that the risk aversion is
 942 inversely related to wealth. As such, it is possible (and indeed observed in Figure 5.9 (a)) that the
 943 investment in the risky asset can be zero until wealth has grown sufficiently to make an investment
 944 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.

945

946

We note the surprisingly undesirable discontinuity in the optimal control closer to maturity (e.g.
 $t_n \geq 15$ (years)) in Figure 5.9 (b). Specifically, the investment in the risky asset transitions very

¹⁴For the constant scalarization scenario, this corresponds to the median line in Figure 5.6(b).

947 quickly from zero to the maximum investment possible, despite the continuity of risk aversion in
 948 wealth implied by (5.6). This contrasts with the case of a constant scalarization parameter $\rho(w) = \rho$,
 949 where a similar discontinuity is not observed (see Figure 5.7 (b)). In the appendix, we explain this
 950 undesirable behavior of the optimal control by showing that, as the intervention time $t_n \rightarrow T$, there is
 951 a very fast transition in the fraction of wealth invested in the risky asset from zero, when $w = 0$, to a
 952 nonzero value when $w > 0$. In addition, it is also shown in the appendix that, with the set of realistic
 953 parameters used in this experiment, this fast transition is very dramatic, namely a jump from zero to
 954 $q_{\max} = 1.5$, as observed in Figure 5.9 (b). Finally, we note that for $w = 0$, there should always be
 955 a “yellow strip”, i.e. zero investment in the risky asset, for all t_n , which, as noted above, should become
 956 infinitesimal as $t_n \rightarrow T$. Since any numerical scheme can only approximate this infinitesimal strip (as
 957 $t_n \rightarrow T$) by some finite size (as in Figure 5.9 (b)), it is expected the approximated strip shrinks as the
 958 mesh is refined. Although not reported herein, we note that this shrinkage was indeed observed.

959 While the economic merits of such a strategy depends on the particular application, it is unlikely
 960 to be economically reasonable in institution-related applications of MV optimization (such as in the
 961 case of pension funds or insurance). Specifically, relatively low investments in the risky asset during
 962 early years (due to high risk aversion resulting from relatively lower wealth levels) might result in
 963 lower terminal wealth - indeed, the expectation of terminal wealth is substantially lower with wealth-
 964 dependent scalarization parameter of the form (5.6) - which in turn might make it harder to fund
 965 liabilities, while the increase in risky asset exposure over time does not actually reduce the variance
 966 of terminal wealth (compared to the case of a constant ρ).

967 Therefore, in contrast to, for example Li and Li (2013), we conclude that a wealth-dependent scalar-
 968 ization parameter defined by (5.6) does not appear well-suited for obtaining realistic time-consistent
 969 MV optimal strategies in the presence of investment constraints, since the resulting terminal wealth
 970 is less MV-efficient (as compared with the results obtained using a constant scalarization parameter),
 971 while the steady increase in risk exposure over time might be undesirable in many applications of
 972 time-consistent MV optimization.

973 6 Conclusions

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

985 The economic implications of investment constraints on the efficient frontiers and on the resulting
 986 optimal controls have been explored in detail. The numerical results illustrate that these realistic
 987 considerations can have a substantial impact on the efficient frontiers and associated optimal controls,
 988 resulting in economically plausible conclusions. In addition, the results from the time-consistent
 989 problem are compared to those of the pre-commitment problem, leading to the conclusion that the
 990 time-consistent problem is less sensitive to the maximum leverage constraint than the pre-commitment
 991 problem. In addition, we explored the consequences of implementing a popular form of a wealth-
 992 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.

A Appendix

In this appendix we investigate the behavior of the control as the intervention time $t_n \rightarrow T$, for both the choices $\rho(w) = \rho$ (a constant) and $\rho(w) = \theta/w$ with $w \geq 0$. For the purposes of this discussion, we fix a small $\Delta t_n > 0$, let $t_n = T - \Delta t_n$. We set transaction costs equal to zero, and both lending and borrowing rates equal to the risk-free rate r . At time t_n^- , the system is assumed to be in state $x = (s, b)$, implying that $W(t_n^-) = s + b = w$; at rebalancing time t , the investor chooses an admissible impulse η_n that solves

$$\sup_{\eta_n \in \mathcal{Z}} (E_{\eta_n}^{x, t_n} [W(T)] - \rho(w) \cdot \text{Var}_{\eta_n}^{x, t_n} [W(T)]). \quad (\text{A.1})$$

Also recall from (2.6) that, applying the impulse η_n at time t_n gives $B(t_n) = \eta_n$ and $S(t_n) = w - \eta_n$.

We briefly consider admissible values of η_n . Note that $w = 0$ corresponds to insolvency at time t_n^- (see definition (2.10)), in which case any existing investments in the risky asset has to be liquidated, resulting in zero wealth being invested in the risky asset at time t_n , so that the optimal control is $\eta_n^* \equiv w$, or equivalently, the fraction of wealth invested in the risky asset is zero.

For the rest of this appendix, we therefore restrict our attention to the case of $w > 0$. In this setting, the leverage constraint with $q_{\max} = 1.5$ and the short-selling prohibition constraint on the risky asset give rise to the following range for the admissible impulse η_n

$$\left\{ \begin{array}{l} S(t_n)/w = (w - \eta_n)/w \leq q_{\max} = 1.5 \\ S(t_n) = w - \eta_n \geq 0 \end{array} \right\} \Rightarrow -\frac{1}{2}w \leq \eta_n \leq w, \quad \text{with } w > 0. \quad (\text{A.2})$$

For a chosen admissible impulse η_n at time t_n , i.e. $B(t_n) = \eta_n$ and $S(t_n) = w - \eta_n$, the portfolio is not rebalanced again during the time interval $[t_n, T]$. Assuming Δt_n is sufficiently small, we approximate $W(T)$ by $W(t_n) + \Delta W$, where the increment ΔW is given by

$$\Delta W := [(\mu - \lambda\kappa)S(t_n) + rB(t_n)]\Delta t_n + \sigma S(t_n)\sqrt{\Delta t_n}\hat{Z} + S(t_n)\sum_{i=1}^{\pi[t_n, T]}(\xi_i - 1) \quad (\text{A.3})$$

with $\hat{Z} \sim \text{Normal}(0, 1)$, and $\pi[t_n, T]$ denoting the number of jumps in the interval $[t_n, T]$. Substituting $B(t_n) = \eta_n$ and $S(t_n) = w - \eta_n$ into (A.3) gives the following approximations

$$\begin{aligned} E_{\eta_n}^{x, t_n} [W(T)] &\simeq E_{\eta_n}^{x, t_n} [w + \Delta W] = (1 + \mu\Delta t_n)w - (\mu - r)\eta_n\Delta t_n, \\ \text{Var}_{\eta_n}^{x, t_n} [W(T)] &\simeq \text{Var}_{\eta_n}^{x, t_n} [w + \Delta W] = (\eta_n - w)^2(\sigma^2 + \lambda\kappa_2)\Delta t_n. \end{aligned} \quad (\text{A.4})$$

Case 1: $\rho(w) = \rho$

For $\rho(w) = \rho > 0$ constant in (A.1), we see from (A.4) that the variance term $-\rho \text{Var}_{\eta_n}^{x, t_n} [W(T)]$ is quadratic in w , while the expected value term $E_{\eta_n}^{x, t_n} [W(T)]$ is linear in w . Therefore, as $w \downarrow 0$, the $E_{\eta_n}^{x, t_n} [W(T)]$ term dominates, so that the objective (A.1) can be approximated as $\sup_{\eta_n \in \mathcal{Z}} (E_{\eta_n}^{x, t_n} [W(T)])$,

1022 leading an investor to invest all wealth in the risky asset for very low levels of $w > 0$. Conversely,
 1023 as $w \rightarrow \infty$, the variance term $-\rho \text{Var}_{\eta_n}^{x,t_n} [W(T)]$ dominates, so that the investor's objective (A.1)
 1024 effectively becomes $\sup_{\eta_n \in \mathcal{Z}} (-\rho \cdot \text{Var}_{\eta_n}^{x,t_n} [W(T)])$, resulting in all wealth being invested in the risk-free
 1025 asset for very large $w > 0$. This is illustrated in the heatmap of optimal controls in the case of
 1026 a constant scalarization parameter (see Figure 5.7 (b)) - observe the decreasing fraction of wealth
 1027 invested in the risky asset as wealth increases.

1028 **Case 2: $\rho(w) = \theta/w, \theta > 0$**

1029 In this case the variance term in (A.4) becomes

$$1030 \quad -\frac{\theta}{w} \cdot \text{Var}_{\eta_n}^{x,t_n} [W(T)] \simeq -\frac{\theta}{w} \cdot (\eta_n - w)^2 (\sigma^2 + \lambda\kappa_2) \Delta t_n, \quad (\text{A.5})$$

1031 which is no longer quadratic in w . The intuition and argument explaining the results for a constant
 1032 ρ therefore cannot be applied to this case in a straightforward way. Instead, using (A.4), we obtain

$$1033 \quad \frac{d}{d\eta_n} \left[E_{\eta_n}^{x,t_n} [W(T)] - \frac{\theta}{w} \cdot \text{Var}_{\eta_n}^{x,t_n} [W(T)] \right]$$

$$1034 \quad \simeq -(\mu - r) \Delta t_n + 2\theta (\sigma^2 + \lambda\kappa_2) \Delta t_n - 2 \left(\frac{\theta}{w} \right) (\sigma^2 + \lambda\kappa_2) \Delta t_n \cdot \eta_n \quad (\text{A.6})$$

$$1035 \quad \leq \left[-(\mu - r) + 3\theta (\sigma^2 + \lambda\kappa_2) \right] \Delta t_n, \quad \text{for } -\frac{1}{2}w \leq \eta_n \leq w, \quad w > 0, \quad (\text{A.7})$$

1036 where the upper bound (A.7) on the derivative follows from the bound on η_n in (A.2). Re-arranging
 1037 (A.7), we see that if $\theta < \theta_{crit}$, where

$$1038 \quad \theta_{crit} := \frac{(\mu - r)}{3(\sigma^2 + \lambda\kappa_2)}, \quad (\text{A.8})$$


1039 then the upper bound (A.7) is strictly negative for admissible impulse η_n which satisfies (A.2). Hence,
 1040 the objective function is strictly decreasing in admissible impulse η_n as $t_n \rightarrow T$. As such, the optimal
 1041 impulse is always $\eta_n^* = -\frac{1}{2}w$. That is, it is always optimal to invest the minimum amount η_n^* in the
 1042 risk-free asset, or equivalently, to invest the maximum amount $q_{max}w$ in the risky asset. In summary,
 1043 for $\rho(w) = \theta/w$ and $\theta < \theta_{crit}$,

$$1044 \quad \theta < \theta_{crit} \implies \frac{w - \eta_n^*}{w} = q_{max}, \quad \text{for } w > 0, \quad \text{as } t_n \rightarrow T. \quad (\text{A.9})$$

1045 For $w = 0$, the fraction of wealth invested in the risky asset is zero, as discussed previously.

1046 Now consider the particular case of the parameters used to obtain the MV-optimal control for the
 1047 case of $\rho(w) = \theta/w$, illustrated in Figure 5.9 (b). The figure is based on the θ -value of $\theta = 0.082$
 1048 (chosen because the required standard deviation of terminal wealth is achieved), and assumes the Kou
 1049 model for the risky asset dynamics, so we use the relevant parameters in Table 5.1 and Table 5.6 to
 1050 calculate $\theta_{crit} = 0.5359$. Therefore, since $\theta < \theta_{crit}$ in this particular case, the discontinuity in the ratio
 1051 (A.9) explains the very fast transition of the fraction of wealth invested in the risky asset from zero,
 1052 when $w = 0$, to q_{max} , when $w > 0$, as $t_n \rightarrow T$, observed in Figure 5.9 (b).

1053 The role of θ in (A.6) and the subsequent conclusion (A.9) should be highlighted. If $\theta \geq \theta_{crit}$, the
 1054 result (A.9) may not necessarily hold, since larger θ in $\rho(w) = \theta/w$ has the effect of increasing the
 1055 overall level of risk aversion associated with any value of $w > 0$. As $t_n \rightarrow T$, we still expect to see a
 1056 very fast transition from zero investment in the risky asset for $w = 0$ to some nonzero investment in
 1057 the risky asset for $w > 0$, but we do not expect the fraction of wealth invested in the the risky asset
 1058 to be necessarily equal to the maximum possible (q_{max}). This is illustrated in Figure A.1 below.



(a) $\theta = 0.1222$ ($\theta < \theta_{crit}$, but not as small as θ in Figure 5.9 (b))

(b) $\theta = 1.004$ ($\theta \geq \theta_{crit}$)

Figure A.1: Effect of using a using different θ values in the definition of a wealth-dependent scalarization parameter of the form $\rho(w) = \theta/w$. The results are based on the same parameters used in Section 5.5 - Kou model, Experiment 3 (solvency and leverage constraints) - and can be compared with Figure 5.9 (b).

References

- 1059
- 1060 Alia, I., F. Chighoub, and A. Sohail (2016). A characterization of equilibrium strategies in continuous-
- 1061 time mean-variance problems for insurers. *Insurance: Mathematics and Economics* (68), 212–223.
- 1062 Basak, S. and G. Chabakauri (2010). Dynamic mean-variance asset allocation. *Review of Financial*
- 1063 *Studies* 23, 2970–3016.
- 1064 Bensoussan, A., K. C. Wong, S. C. P. Yam, and S. P. Yung (2014). Time-consistent portfolio selection
- 1065 under short-selling prohibition: From discrete to continuous setting. *SIAM Journal on Financial*
- 1066 *Mathematics* 5, 153–190.
- 1067 Bjork, T., M. Khapko, and A. Murgoci (2016). A theory of Markovian time-inconsistent stochastic
- 1068 control in continuous time. *Working paper* .
- 1069 Bjork, T. and A. Murgoci (2010). A general theory of Markovian time inconsistent stochastic control
- 1070 problems. *Working paper* Available at <http://ssrn.com/abstract=1694759>.
- 1071 Bjork, T. and A. Murgoci (2014). A theory of Markovian time-inconsistent stochastic control in
- 1072 discrete time. *Finance and Stochastics* (18), 545–592.
- 1073 Bjork, T., A. Murgoci, and X. Zhou (2014). Mean-variance portfolio optimization with state-dependent
- 1074 risk aversion. *Mathematical Finance* (1), 1–24.
- 1075 Cong, F. and C. Oosterlee (2016). On pre-commitment aspects of a time-consistent strategy for a
- 1076 mean-variance investor. *Journal of Economic Dynamics and Control* 70, 178–193.
- 1077 Cont, R. and C. Mancini (2011). Nonparametric tests for pathwise properties of semi-martingales.
- 1078 *Bernoulli* (17), 781–813.

- 1079 Cont, R. and P. Tankov (2004). *Financial modelling with jump processes*. Chapman and Hall / CRC
1080 Press.
- 1081 Crandall, M., H. Ishii, and P. Lions (1992). User's guide to viscosity solutions of second order partial
1082 differential equations. *Bulletin of the American Mathematical Society* 27(1), 1–67.
- 1083 Cui, X., D. Li, S. Wang, and S. Zhu (2012). Better than dynamic mean-variance: Time inconsistency
1084 and free cash flow stream. *Mathematical Finance* 22(2), 346–378.
- 1085 Cui, X., L. Xu, and Y. Zeng (2015). Continuous time mean-variance portfolio optimization with
1086 piecewise state-dependent risk aversion. *Optimization Letters (Springer)* pp. 1–11.
- 1087 Dang, D. and P. Forsyth (2014). Continuous time mean-variance optimal portfolio allocation under
1088 jump diffusion: A numerical impulse control approach. *Numerical Methods for Partial Differential*
1089 *Equations* 30, 664–698.
- 1090 Dang, D. and P. Forsyth (2016). Better than pre-commitment mean-variance portfolio allocation
1091 strategies: A semi-self-financing Hamilton–Jacobi–Bellman equation approach. *European Journal*
1092 *of Operational Research* (250), 827–841.
- 1093 Dang, D., P. Forsyth, and Y. Li (2016). Convergence of the embedded mean-variance optimal points
1094 with discrete sampling. *Numerische Mathematik* (132), 271–302.
- 1095 d'Halluin, Y., P. Forsyth, and K. Vetzal (2005). Robust numerical methods for contingent claims
1096 under jump diffusion processes. *IMA Journal of Numerical Analysis* (25), 87–112.
- 1097 Forsyth, P. and K. Vetzal (2017). Dynamic mean variance asset allocation: Tests for robustness.
1098 *International Journal of Financial Engineering* 4:2. 1750021 (electronic).
- 1099 Hu, Y., H. Jin, and X. Zhou (2012). Time-inconsistent stochastic linear-quadratic control. *SIAM*
1100 *Journal on Control and Optimization* 50(3), 1548–1572.
- 1101 Kou, S. (2002). A jump-diffusion model for option pricing. *Management Science* 48(8), 1086–1101.
- 1102 Li, D. and W.-L. Ng (2000). Optimal dynamic portfolio selection: multi period mean variance formu-
1103 lation. *Mathematical Finance* 10, 387–406.
- 1104 Li, D., X. Rong, and H. Zhao (2015a). Time-consistent reinsurance–investment strategy for a mean-
1105 variance insurer under stochastic interest rate model and inflation risk. *Insurance: Mathematics*
1106 *and Economics* 64, 28–44.
- 1107 Li, D., X. Rong, and H. Zhao (2015b). Time-consistent reinsurance–investment strategy for an insurer
1108 and a reinsurer with mean–variance criterion under the cev model. *Journal of Computational and*
1109 *Applied Mathematics* 283, 142–162.
- 1110 Li, Y. and Z. Li (2013). Optimal time-consistent investment and reinsurance strategies for mean-
1111 variance insurers with state dependent risk aversion. *Insurance: Mathematics and Economics* 53,
1112 86–97.
- 1113 Li, Y., H. Qiao, S. Wang, and L. Zhang (2015c). Time-consistent investment strategy under partial
1114 information. *Insurance: Mathematics and Economics* 65, 187–197.
- 1115 Li, Z., Y. Zeng, and Y. Lai (2012). Optimal time-consistent investment and reinsurance strategies for
1116 insurers under Heston's SV model. *Insurance: Mathematics and Economics* 51, 191–203.

- 1117 Liang, Z. and M. Song (2015). Time-consistent reinsurance and investment strategies for mean-
1118 variance insurer under partial information. *Insurance: Mathematics and Economics* 65, 66–76.
- 1119 Lioui, A. (2013). Time consistent vs. time inconsistent dynamic asset allocation: Some utility cost
1120 calculations for mean variance preferences. *Journal of Economic Dynamics and Control* (37), 1066–
1121 1096.
- 1122 Ma, K. and P. Forsyth (2016). Numerical solution of the Hamilton-Jacobi-Bellman formulation for
1123 continuous time mean variance asset allocation under stochastic volatility. *Journal of Computational*
1124 *Finance* 20:1, 1–37.
- 1125 Mancini, C. (2009). Non-parametric threshold estimation models with stochastic diffusion coefficient
1126 and jumps. *Scandinavian Journal of Statistics* (36), 270–296.
- 1127 Markowitz, H. (1952). Portfolio selection. *The Journal of Finance* 7(1), 77–91.
- 1128 Merton, R. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of*
1129 *Financial Economics* 3, 125–144.
- 1130 Oksendal, B. and A. Sulem (2005). *Applied Stochastic Control of Jump Diffusions*. Springer.
- 1131 Pedersen, J. and G. Peskir (2017). Optimal mean-variance portfolio selection. *Mathematics and*
1132 *Financial Economics* (11), 137–160.
- 1133 Ramezani, C. and Y. Zeng (2007). Maximum likelihood estimation of the double exponential jump-
1134 diffusion process. *Annals of Finance* 3(4), 487–507.
- 1135 Sun, J., Z. Li, and Y. Zeng (2016). Precommitment and equilibrium investment strategies for defined
1136 contribution pension plans under a jump–diffusion model. *Insurance: Mathematics and Economics*
1137 (67), 158–172.
- 1138 Vigna, E. (2014). On efficiency of mean-variance based portfolio selection in defined contribution
1139 pension schemes. *Quantitative Finance* 14(2), 237–258.
- 1140 Vigna, E. (2016). On time consistency for mean-variance portfolio selection. *Working paper, Collegio*
1141 *Carlo Alberto* (476).
- 1142 Vigna, E. (2017). Tail optimality and preferences consistency for intertemporal optimization problems.
1143 *Working paper, Collegio Carlo Alberto* (502).
- 1144 Wang, J. and P. Forsyth (2008). Maximal use of central differencing for Hamilton–Jacobi–Bellman
1145 PDEs in finance. *SIAM Journal on Numerical Analysis* (46), 1580–1601.
- 1146 Wang, J. and P. Forsyth (2010). Numerical solution of the Hamilton-Jacobi-Bellman formulation for
1147 continuous time mean variance asset allocation. *Journal of Economic Dynamics and Control* 34,
1148 207–230.
- 1149 Wang, J. and P. Forsyth (2011). Continuous time mean variance asset allocation: A time-consistent
1150 strategy. *European Journal of Operational Research* (209), 184–201.
- 1151 Wei, J. and T. Wang (2017). Time-consistent mean-variance asset-liability management with random
1152 coefficients. *Insurance: Mathematics and Economics* (77), 84–96.
- 1153 Wei, J., K. Wong, S. Yam, and S. Yung (2013). Markowitz’s mean-variance asset-liability management
1154 with regime switching: A time-consistent approach. *Insurance: Mathematics and Economics* 53,
1155 281–291.

- 1156 Yu, P. (1971). Cone convexity, cone extreme points, and nondominated solutions in decision problem
1157 with multiobjectives. *Journal of Optimization Theory and Applications* (7), 11–28.
- 1158 Zeng, Y. and Z. Li (2011). Optimal time-consistent investment and reinsurance policies for mean-
1159 variance insurers. *Insurance: Mathematics and Economics* 49(1), 145–154.
- 1160 Zeng, Y., Z. Li, and Y. Lai (2013). Time-consistent investment and reinsurance strategies for mean-
1161 variance insurers with jumps. *Insurance: Mathematics and Economics* 52, 498–507.
- 1162 Zhang, C. and Z. Liang (2017). Portfolio optimization for jump-diffusion risky assets with common
1163 shock dependence and state dependent risk aversion. *Optimal control applications and methods* (38),
1164 229–246.
- 1165 Zhou, X. and D. Li (2000). Continuous time mean variance portfolio selection: a stochastic LQ
1166 framework. *Applied Mathematics and Optimization* 42, 19–33.
- 1167 Zhou, Z., H. Xiao, J. Yin, X. Zeng, and L. Lin (2016). Pre-commitment vs. time-consistent strate-
1168 gies for the generalized multi-period portfolio optimization with stochastic cash flows. *Insurance:
1169 Mathematics and Economics* (68), 187–202.