International Journal of Theoretical and Applied Finance
Vol. 20, No. 7 (2017) 1750049 (26 pages)

DOI: 10.1142/S0219024917500492

ON ROBUST MULTI-PERIOD PRE-COMMITMENT AND TIME-CONSISTENT MEAN-VARIANCE PORTFOLIO OPTIMIZATION

F. CONG[*,‡] and C. W. OOSTERLEE[*,†,§]

[*]*Delft Institute of Applied Mathematics*
Delft University of Technology, Mekelweg 4
Delft, 2628 CD, The Netherlands

[†]*Centrum Wiskunde & Informatica, Science Park 123*
Amsterdam, 1098 XG, The Netherlands
[‡]*fcong1989@gmail.com*
[§]*c.w.oosterlee@cwi.nl*

Received 19 February 2017
Revised 25 June 2017
Accepted 12 September 2017
Published 7 November 2017

We consider robust pre-commitment and time-consistent mean-variance optimal asset allocation strategies, that are required to perform well also in a worst-case scenario regarding the development of the asset price. We show that worst-case scenarios for both strategies can be found by solving a specific equation each time step. In the unconstrained asset allocation case, the robust pre-commitment as well as the time-consistent strategy are identical to the corresponding robust myopic strategies, by which investors perform robust portfolio control only for one time step and conduct a risk-free strategy afterwards. In the experiments, the robustness of pre-commitment and time-consistent strategies is studied in detail. Our analysis and numerical results indicate that the time-consistent allocation strategy is more stable when possible incorrect assumptions regarding the future asset development are modeled and taken into account. In some situations, the time-consistent strategy can even generate higher efficient frontiers than the pre-commitment strategy (which is counter-intuitive), because the time-consistency restriction appears to protect an investor in such a situation.

Keywords: Robust optimization; mean-variance optimal asset allocation; target-based strategy; time-consistent strategy; model prediction error.

1. Introduction

After the introduction of the concept of static mean-variance portfolio optimization in Markowitz (1952), portfolio optimization theory has received a lot of attention from academics as well as from the industry. This optimization problem which is based on two criteria is popular with practitioners because it has a clear and very informative target function, which explicitly contains a profit term, a risk term and the trade-off between these. From the perspective of academic research, this mean-variance framework forms the basis for many interesting research directions. One potential way to generalize Markowitz's mean-variance strategy is to take dynamic optimal asset allocation into consideration.

Solving a dynamic mean-variance optimization problem is not a trivial task, however. Due to the nonlinearity of the variance operator, the dynamic mean-variance optimization problem cannot be solved directly via the Bellman dynamic programming principle (Bellman 1957). To tackle this issue, two possible directions are recommended in the literature. One is based on placing the dynamic mean-variance problem into a dynamic linear-quadratic (LQ) optimization context, using an embedding technique (Zhou & Li 2000, Li & Ng 2000), and the other direction is to impose a time-consistency restriction, which also serves as a condition, at all intermediate time steps (Basak & Chabakauri 2010).

In Basak & Chabakauri (2010), the strategy corresponding to the LQ problem was termed the *pre-commitment* strategy and the strategy obtained by introducing time-consistency conditions was termed a *time-consistent* strategy. Related research on the pre-commitment portfolio optimization has been performed in Li *et al.* (2002), Zhu *et al.* (2004), Bielecki *et al.* (2005), Fu *et al.* (2010) and Cui *et al.* (2014a), where constraints on control variables were introduced in the dynamic optimization process. Extensions of the time-consistent strategy have been made by modeling the trade-off parameter between risk and reward as a state-dependent variable, see Hu *et al.* (2012), Cui *et al.* (2014b), Björk *et al.* (2014) and Cui *et al.* (2015). Pre-commitment as well as time-consistent strategies constitute important parts of pension management problems. Discussions about the pre-commitment strategy in a defined contribution pension scheme can be found for example in Haberman & Vigna (2002), Gerrard *et al.* (2004), Vigna (2014) and Sun *et al.* (2016). For the time-consistent investment and re-insurance problems, we refer to Zeng & Li (2011) and Liang & Song (2015). Numerical algorithms for solving the pre-commitment and the time-consistency problems have been introduced in Wang & Forsyth (2010, 2011), where the solution is obtained by solving the corresponding Hamilton–Jacobi–Bellman partial differential equation, and in Cong & Oosterlee (2016a, 2016b), where a method based on using simulation and least-squares regression is introduced. With these numerical algorithms, dynamic mean-variance problems including various kinds of constraints on the allocations can be solved efficiently.

The work in the above mentioned papers is based on the assumption that the market evolves exactly as the model prescribes. This may be questionable in reality, since we can only estimate model parameters from historical data. As mentioned in Best & Grauer (1991), Black & Litterman (1992) and Britten-Jones (1999), designing an investment strategy based on historical data may lead to significant losses. One possible way to tackle this problem is to take *model uncertainty* into account and to consider *robust variants* of the optimal asset allocation problem. Garlappi *et al.* (2006) extend the classical mean-variance portfolio optimization problem to allow for the possibility of multiple priors and to incorporate aversion to uncertainty regarding the estimated expected returns. They find that the max–min problem faced by an investor who cares about parameter uncertainty can be reduced to a maximization-only problem, where the estimated expected returns are adjusted to reflect parameter uncertainty. A robust utility optimization problem where the risky asset dynamics follow a diffusion process with misspecified trend and volatility coefficients is considered by Tevzadze *et al.* (2013). They propose an explicit characterization of the solution which was given by the Hamilton–Jacobi–Bellman–Isaacs equation. In Perret-Gentil & Victoria-Feser (2005), it is suggested to solve the mean-variance portfolio selection problem using statistically robust estimates. Noticing that accurately estimating returns may be a difficult task, Kuhn *et al.* (2009) recommend to replace the original uncertain return process by a tractable one. Maccheroni *et al.* (2006) model a robust dynamic utility optimization problem by introducing a penalty function on all possible probabilistic models.

Another common way to introduce model uncertainty is to consider an approach, in which the corresponding optimal strategy is required to perform well even in a so-called 'worst-case scenario'. In Tütüncü & Koenig (2004), a worst-case static mean-variance problem is transformed into a saddle-point problem and is solved using an interior-point algorithm. In Gülpınar & Rustem (2007), the authors implement a scenario tree to represent stochastic aspects and introduce uncertainty into a multi-period mean-variance portfolio problem. For a general discussion on robust optimization, we refer to Bertsimas *et al.* (2011), where a static robust mean-variance optimization is discussed as a special case. For more pointers to aspects of robust portfolio problems, we refer to a review paper Kim *et al.* (2014) and the references therein.

To our knowledge, robust pre-commitment and time-consistent mean-variance optimization problems have not yet been extensively discussed. In this paper, we will address this issue.

We start our work by analyzing the robust versions of the pre-commitment and time-consistent asset allocation problems. Following El Ghaoui & Nilim (2005) and Iyengar (2005), we consider an independent structure for parameter uncertainty, which makes the Bellman backward programming principle feasible within the robust dynamic optimization context. Without any constraints on the asset allocations, analytic solutions can be derived. We show that the worst-case scenarios are generated by solving a specific equation at each time step for both the pre-commitment and the time-consistent strategies. The optimal robust pre-commitment and time-consistent strategies are identical to the corresponding robust myopic strategies, where an investor derives the optimal allocation for one upcoming time period without taking future optimal allocations into account. Robustness can be introduced into the pre-commitment and the time-consistent strategies without drastically increasing the computational complexity.

The robustness of the pre-commitment and the time-consistent strategies is examined in particular when model prediction errors occur, meaning that the assumptions on the behavior of the stochastic asset process do not reflect accurately the actually observed asset path (in the future). We find that the time-consistent strategy appears more stable than the pre-commitment strategy.

In the numerical section, we test the robustness of both strategies using the algorithms proposed in Cong & Oosterlee (2016a, 2016b), that are feasible for both unconstrained and constrained optimization problems. We show that in the case of an unexpectedly poor market the time-consistent strategy can be superior to the pre-commitment strategy. In such a situation, the constrained pre-commitment strategy even yields a higher frontier than the unconstrained pre-commitment strategy, since the constraints on the allocations act as a form of "protection" when the model prediction is inaccurate. In the case of two risky assets and one risk-free asset, we check how the portfolio mean-variance frontier will be impacted given an inaccurate correlation prediction for the risky assets. The unconstrained pre-commitment strategy appears vulnerable to such prediction inaccuracies, whereas the time-consistent strategy appears robust.

The paper is structured as follows. In Sec. 2, we describe the pre-commitment and the time-consistent strategies and their robust counterparts. Analysis for both strategies is performed in Sec. 3, where the robustness of the time-consistent strategy is also studied. Numerical results are presented in Sec. 4. We conclude in the last section.

2. Problem Formulation

2.1. *Multi-period mean-variance portfolio*

In this section, we describe the dynamic portfolio optimization problem with mean-variance criteria. We assume that the financial market is defined on a probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{0\leq t\leq T}, \mathbb{P})$ with finite time horizon $[0,T]$. The state space Ω is the set of all realizations of the financial market within the time horizon $[0,T]$. $\mathcal{F}$ is the sigma algebra of events till time T, i.e. $\mathcal{F} = \mathcal{F}_T$. We assume that the filtration $\{\mathcal{F}_t\}_{0\leq t\leq T}$ is generated by the price processes of the financial market and augmented with the null sets of $\mathcal{F}$.

We consider a portfolio consisting of $n+1$ assets, one risk-free and n risky. We assume that the portfolio can be traded at discrete opportunities,[a] $t \in \{0, \Delta t, \ldots, T-\Delta t\}$, before terminal time T. At the initial time $t_0 = 0$, an investor decides a trading strategy to maximize the expectation of the terminal wealth and to minimize the investment risk. Formally, the investor's problem is given by

$$\hat{V}_0(W_0) = \max_{\{\hat{\mathbf{x}}_t\}_{t=0}^{T-\Delta t}} \{\mathbb{E}[W_T \,|\, W_0] - \lambda \cdot \mathrm{Var}[W_T \,|\, W_0]\} \tag{2.1}$$

with $\hat{V}$ the value function, subject to the wealth restriction:

$$W_{t+\Delta t} = W_t \cdot (\hat{\mathbf{x}}_t' \mathbf{R}_t^e + R_f), \quad t = 0, \Delta t, \ldots, T - \Delta t.$$

[a]The re-balancing times are equidistantly distributed and the number of re-balancing opportunities before terminal time T equals M. The time step Δt between two re-balancing days is $\frac{T}{M}$.

Without loss of generality, we assume that $\{W_t\}_{t=0}^{T-\Delta t}$ will not be zero in the multi-period optimization process. We use notations with hats in Eq. (2.1), because we will reserve the plain notations for the formulation of the robust optimization problem, which forms the main part of this paper. $\hat{\mathbf{x}}_t = [\hat{x}_t(1), \hat{x}_t(2), \ldots, \hat{x}_t(n)]'$ denotes the asset allocations of the investor's wealth in the risky assets in the time period $[t, t+\Delta t)$. The prime sign denotes the vector transpose. The admissible investment strategy $\hat{\mathbf{x}}_t$ is assumed to be adapted to $\mathcal{F}_t$. The risk aversion attitude of the investor is denoted by λ, which is a trade-off factor between maximizing the profit and minimizing the risk. R_f is the return of the risk-free asset in one time step, which is assumed to be constant for simplicity, and $\mathbf{R}_t^e = [R_t^e(1), R_t^e(2), \ldots, R_t^e(n)]'$ denotes the vector of excess returns of the risky assets during $[t, t+\Delta t)$. We assume that the excess returns $\{\mathbf{R}_t^e\}_{t=0}^{T-\Delta t}$ are sequentially independent. At each time point t, $\mathbf{R}_t^e$ is supposed to follow a distribution with determined parameters. Extending the problem to a situation where the distribution parameters are uncertain constitutes the robust counterpart of a dynamic mean-variance problem, which we will elaborate on in Sec. 2.2.

The difficulty of solving the dynamic mean-variance problem is caused by the nonlinearity of conditional variances, i.e. $\text{Var}[\text{Var}[W_T \mid \mathcal{F}_t] \mid \mathcal{F}_s] \neq \text{Var}[W_T \mid \mathcal{F}_s]$, $s \leq t$, which makes the well-known dynamic programming valuation approach, Bellman (1957), not directly applicable. To tackle this problem, there are basically two viable approaches: one is to use an embedding technique and replace the dynamic mean-variance problem by a dynamic quadratic optimization problem (Zhou & Li 2000, Li & Ng 2000), and the other is to introduce a time-consistency restriction as an additional condition into the backward programming approach (Basak & Chabakauri 2010, Wang & Forsyth 2011).

Following the first path, we can formulate the dynamic quadratic problem as in Eq. (2.2):

$$\hat{J}_0(W_0) = \min_{\{\hat{\mathbf{x}}_t\}_{t=0}^{T-\Delta t}} \{\mathbb{E}[(W_T - \gamma)^2 \mid W_0]\}, \tag{2.2}$$

where we use $\hat{J}$ to denote the value function. By assigning different values to the parameter γ and solving the corresponding problems, we can trace out points on an efficient frontier. In Li & Ng (2000), it is proved that this efficient frontier is the same as the one obtained by solving Eq. (2.1) with the trade-off parameter λ taking different values. An advantage of considering dynamic quadratic problem (2.2) is that the Bellman dynamic programming principle can be applied and the problem can therefore be solved in a backward recursive fashion. Since parameter γ in (2.2) acts as an investment target in the dynamic quadratic problem, this kind of optimization problem is also termed *target-based optimization* by Haberman & Vigna (2002) and Gerrard *et al.* (2004).

It is mentioned in Basak & Chabakauri (2010) that prescribing a determined target for an investor will cause a time-inconsistency. To solve this problem, they suggest to take a time-consistency restriction into account, which forms the other path for solving the dynamic mean-variance problem. Since a time-consistency restriction can also be treated as a condition in the dynamic programming framework, the dynamic mean-variance problem with time-consistency conditions can also be solved in a backward recursive manner. Basak & Chabakauri (2010) call the strategy obtained by following the first path the *pre-commitment strategy* and the one achieved by following the second path the *time-consistent strategy*. For a dynamic mean-variance optimization problem, we formally define these two strategies as follows.

Definition 2.1 (Pre-commitment strategy). The pre-commitment strategy $\{\hat{\mathbf{x}}_t^{*\mathrm{pc}}\}_{t=0}^{T-\Delta t}$ is defined by the optimal control for Eq. (2.2).

Definition 2.2 (Time-consistent strategy). The time-consistent strategy $\{\hat{\mathbf{x}}_t^{*\mathrm{tc}}\}_{t=0}^{T-\Delta t}$ is defined by the optimal control for Eq. (2.1) with an additional time-consistency condition requiring that $\{\hat{\mathbf{x}}_t^{*\mathrm{tc}}\}_{t=\tau}^{T-\Delta t}$ also constitutes an optimal control for:

$$\hat{V}_\tau^{\mathrm{tc}}(W_\tau) = \max_{\{\hat{\mathbf{x}}_t\}_{t=\tau}^{T-\Delta t}} \{\mathbb{E}[W_T \mid W_\tau] - \lambda \cdot \mathrm{Var}[W_T \mid W_\tau]\}, \tag{2.3}$$

for $\tau = \Delta t, 2 \cdot \Delta t, \ldots, T - \Delta t$.

2.2. *The robust counterpart*

In the discussion above, we assumed that the excess returns $\mathbf{R}_t^e$ of the risky assets follow a distribution of determined parameters. However, this assumption may not be realistic. Since we can only assess the returns of risky assets by means of historical data, the estimated distribution parameters may be biased and may not necessarily reflect the dynamics of the risky assets in the future. Deriving an investment strategy based on these estimated parameters can lead to significant losses as pointed out by Best & Grauer (1991) and Black & Litterman (1992). To make the investment strategy more reliable, an estimation error in the parameters can be taken into account. To this end, we consider here the *robust counterpart* to a dynamic mean-variance optimization problem. In that case, an investor has a rival, i.e. *nature*, that gives rise to difficulties in the optimization process. For the mean-variance problem, we assume that the rival specifies a mean vector and a covariance matrix of the excess returns $\mathbf{R}_t^e$ of the risky assets at time $t \in \{0, \Delta t, \ldots, T - \Delta t\}$, that are respectively denoted by $\mathbf{u}_t$ and $\sum_t$.

We only assume the uncertainty sets of $\mathbf{u}_t$ and $\sum_t$ to be bounded and nonempty. This implies that the uncertainty set can be an ellipsoidal set as discussed in Goldfarb & Iyengar (2003) or it can be a separable set as considered in Halldórsson & Tütüncü (2003). Besides, we assume that $(\mathbf{u}_t, \sum_t)$, $t = 0, \Delta t, \ldots, T - \Delta t$, is not stationary. If we consider the robust optimization problem as a game of two players, the investor and nature, this latter assumption implies that nature does not necessarily choose the same adverse strategy at each time step. This leads to a time-varying uncertainty model as termed by El Ghaoui & Nilim (2005).

In order to make the Bellman programming principle feasible for the robust dynamic optimization problem, we prescribe the *rectangularity assumption* as proposed in Iyengar (2005). In this paper, the rectangularity assumption[b] is defined by:

Assumption 2.1 (Rectangularity). The choice of $\mathbf{u}_t$ and $\sum_t$ at time t, does not restrict the choice of $\mathbf{u}_s$ and $\sum_s$ at time $s \in \{0, \Delta t, \ldots, T - \Delta t\} \backslash \{t\}$.

As mentioned in Iyengar (2005), since the sources of uncertainty in different time periods are typically independent of each other, the rectangularity assumption, which is also an independence assumption, is appropriate for finite horizon Markovian problems in most cases. However, if we would consider a stochastic volatility asset model or a time series asset model, the rectangularity assumption does not hold any more.

For the pre-commitment and the time-consistent strategies as formed in Sec. 2.1, we establish their robust counterparts as follows.

Definition 2.3 (Robust pre-commitment strategy). The robust pre-commitment strategy $\{\mathbf{x}_t^{*\text{pc}}\}_{t=0}^{T-\Delta t}$ is defined by the optimal control for the following dynamic programming problem:

$$J_t(W_t) = \min_{\mathbf{x}_t} \max_{\mathbf{u}_t, \sum_t} \{\mathbb{E}[J_{t+\Delta t}(W_t \cdot (\mathbf{x}_t' \mathbf{R}_t^e + R_f)) \,|\, W_t]\}, \quad t = 0, \Delta t, \ldots, T - \Delta t \tag{2.4}$$

with the terminal condition $J_T(W_T) = (W_T - \gamma)^2$. Here the maximization operator indicates that the worst-case scenario in reality is taken into account.

Definition 2.4 (Robust time-consistent strategy). The robust time-consistent strategy $\{\mathbf{x}_t^{*\text{tc}}\}_{t=0}^{T-\Delta t}$ is defined by the optimal control for

$$V_0(W_0) = \max_{\{\mathbf{x}_t\}_{t=0}^{T-\Delta t}} \min_{\{\mathbf{u}_t, \sum_t\}_{t=0}^{T-\Delta t}} \{\mathbb{E}[W_T \,|\, W_0] - \lambda \cdot \text{Var}[W_T \,|\, W_0]\} \tag{2.5}$$

with an additional time-consistency condition requiring that $\{\mathbf{x}_t^{*\text{tc}}\}_{t=\tau}^{T-\Delta t}$ also constitutes an optimal solution for:

$$V_\tau^{\text{tc}}(W_\tau) = \max_{\{\mathbf{x}_t\}_{t=\tau}^{T-\Delta t}} \min_{\{\mathbf{u}_t, \sum_t\}_{t=\tau}^{T-\Delta t}} \{\mathbb{E}[W_T \,|\, W_\tau] - \lambda \cdot \text{Var}[W_T \,|\, W_\tau]\}, \tag{2.6}$$

for $\tau = \Delta t, 2 \cdot \Delta t, \ldots, T - \Delta t$. Here the minimization operator indicates that the worst-case scenario in reality is taken into account.

In order to meet the duality condition as proposed in Halldórsson & Tütüncü (2003) for a min–max mean-variance optimization problem, we require the asset allocations to be *loosely bounded*, i.e. the allocation $\mathbf{x}_t$ at each time t satisfies $-M \leq \mathbf{x}_t \leq M$ for a large positive number M, where the inequality sign is in element-wise

[b]For a more general definition, we refer the readers to Iyengar (2005).

sense. We use the term “loosely bounded” to emphasize that this restriction does not have an impact on the choice of the optimal asset allocations since the positive M can be chosen sufficiently large. Therefore, if we assume the asset allocations to be loosely bounded, the optimal control performed by an investor can still be obtained by solving the first-order conditions.

Remark 2.1. It should be emphasized that we form the robust pre-commitment strategy by directly imposing parameter uncertainty on a plain pre-commitment strategy as defined in Sec. 2.1. Without the robustness requirement, the equivalence between the pre-commitment strategy and the optimal dynamic mean-variance strategy has been established in Li & Ng (2000). However, it is not yet clear whether the robust pre-commitment strategy is equivalent to the robust dynamic mean-variance strategy. In this paper we will consider the robust pre-commitment strategy as described in Definition 2.3.

3. Analysis in the Unconstrained Case

Within the framework presented in the last section, we can derive an analytic solution for the optimal robust pre-commitment and the optimal robust time-consistent strategy by the Bellman dynamic programming principle. In the former case, we consider the value function iteration in our proof, while in the latter case we make use of an essential property of a time-consistent control. Meanwhile, we also generate the adverse choices taken by nature at each time step. Similar to our findings in Cong & Oosterlee (2016a, 2016b), we observe that, for either a pre-commitment or a time-consistent investor, the robust dynamic mean-variance strategy is the same as a corresponding robust myopic strategy.

In our derivation, we assume that the expectation of the excess return of any risky asset is larger than the return of the risk-free asset. This means that we do not consider the trivial case where investment is never in the risky assets. Also we make the assumption that the covariance of the excess returns of the risky assets is positive definite.

3.1. *Robust pre-commitment strategy*

We first consider the robust pre-commitment strategy, which has been formed in a recursive setting as in Definition 2.3. The optimal control for the robust pre-commitment strategy can be described by the following proposition.

Proposition 3.1. *For the robust pre-commitment optimization problem as in Definition* 2.3, *an investor at time t with wealth W_t has the following optimal control*:

$$\mathbf{x}_t^{*\mathrm{pc}}(W_t) = \frac{\gamma - W_t R_f^{(T-t)/\Delta t}}{W_t R_f^{(T-t)/\Delta t - 1}} \cdot \left(\sum_t^{*} + \mathbf{u}_t^* \cdot \mathbf{u}_t^{*\prime} \right)^{-1} \cdot \mathbf{u}_t^*, \quad t = 0, \Delta t, \ldots, T - \Delta t \tag{3.1}$$

with the parameters $\{\mathbf{u}_t^*, \sum_t^*\}$ *for the worst-case scenario solving a minimization problem*:

$$\left\{\mathbf{u}_t^*, \sum_t^*\right\} = \arg\min_{\mathbf{u}_t, \sum_t} \left\{\mathbf{u}_t' \sum_t^{-1} \mathbf{u}_t\right\}. \tag{3.2}$$

Proof. At time step T, the value function is known as:

$$J_T(W_T) = (W_T - \gamma)^2.$$

At time step $T - \Delta t$, the value function can be calculated by:

$$\begin{aligned} &J_{T-\Delta t}(W_{T-\Delta t}) \\ &\quad = \min_{\mathbf{x}_{T-\Delta t}} \max_{\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}} \{\mathbb{E}[J_T(W_{T-\Delta t} \cdot (\mathbf{x}'_{T-\Delta t}\mathbf{R}^e_{T-\Delta t} + R_f)) \,|\, W_{T-\Delta t}]\} \\ &\quad = \min_{\mathbf{x}_{T-\Delta t}} \max_{\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}} \{\mathbb{E}[(W_{T-\Delta t} \cdot (\mathbf{x}'_{T-\Delta t}\mathbf{R}^e_{T-\Delta t} + R_f) - \gamma)^2 \,|\, W_{T-\Delta t}]\}. \end{aligned}$$

Based on Lemma 2.3 in Halldórsson & Tütüncü (2003), perfect duality holds, i.e. changing the order of minimization and maximization does not influence the value of the value function. Therefore, we have

$$\begin{aligned} J_{T-\Delta t}(W_{T-\Delta t}) = \max_{\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}} \min_{\mathbf{x}_{T-\Delta t}} \{\mathbb{E}[(W_{T-\Delta t} \\ \cdot (\mathbf{x}'_{T-\Delta t}\mathbf{R}^e_{T-\Delta t} + R_f) - \gamma)^2 \,|\, W_{T-\Delta t}]\}. \end{aligned}$$

We define a new function $F_{T-\Delta t}(\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}, W_{T-\Delta t})$ by

$$\begin{aligned} &F_{T-\Delta t}\left(\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}, W_{T-\Delta t}\right) \\ &\quad := \min_{\mathbf{x}_{T-\Delta t}} \{\mathbb{E}[(W_{T-\Delta t} \cdot (\mathbf{x}'_{T-\Delta t}\mathbf{R}^e_{T-\Delta t} + R_f) - \gamma)^2 \,|\, W_{T-\Delta t}]\}, \end{aligned} \tag{3.3}$$

where $\mathbf{u}_{T-\Delta t}$ and $\sum_{T-\Delta t}$ influence the distribution of $\mathbf{R}^e_{T-\Delta t}$. The value function $J_{T-\Delta t}(W_{T-\Delta t})$ can then be written as:

$$J_{T-\Delta t}(W_{T-\Delta t}) = \max_{\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}} \left\{F_{T-\Delta t}(\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}, W_{T-\Delta t})\right\}.$$

For a given set of parameters $\{\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}, W_{T-\Delta t}\}$, the optimization problem with respect to $\mathbf{x}_{T-\Delta t}$ as shown in Eq. (3.3) constitutes a smooth and convex optimization problem. Therefore, by solving the first-order-conditions for the optimality, we get:

$$\begin{aligned} \mathbf{x}^*_{T-\Delta t}\left(\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}, W_{T-\Delta t}\right) = \frac{\gamma - W_{T-\Delta t}R_f}{W_{T-\Delta t}} \cdot \left(\sum_{T-\Delta t} + \mathbf{u}_{T-\Delta t}\mathbf{u}'_{T-\Delta t}\right)^{-1} \\ \cdot \mathbf{u}_{T-\Delta t}. \end{aligned} \tag{3.4}$$

Inserting the optimal control in Eq. (3.4) into Eq. (3.3) gives us:

$$F_{T-\Delta t}\left(\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}, W_{T-\Delta t}\right)$$
$$= (\gamma - W_{T-\Delta t}R_f)^2 \cdot \left(1 - \mathbf{u}'_{T-\Delta t} \cdot \left(\sum_{T-\Delta t} + \mathbf{u}_{T-\Delta t}\mathbf{u}'_{T-\Delta t}\right)^{-1} \cdot \mathbf{u}_{T-\Delta t}\right)$$

and

$$J_{T-\Delta t}(W_{T-\Delta t}) = \max_{\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}} \left\{(\gamma - W_{T-\Delta t}R_f)^2 \cdot \left(1 - \mathbf{u}'_{T-\Delta t} \cdot \left(\sum_{T-\Delta t} + \mathbf{u}_{T-\Delta t}\mathbf{u}'_{T-\Delta t}\right)^{-1} \cdot \mathbf{u}_{T-\Delta t}\right)\right\}.$$

Therefore, the optimal adverse policy taken by nature should solve the minimization problem:

$$\left\{\mathbf{u}^*_{T-\Delta t}, \sum_{T-\Delta t}^{*}\right\} = \arg \min_{\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}} \left\{\mathbf{u}'_{T-\Delta t} \cdot \left(\sum_{T-\Delta t} + \mathbf{u}_{T-\Delta t}\mathbf{u}'_{T-\Delta t}\right)^{-1} \cdot \mathbf{u}_{T-\Delta t}\right\}.$$

By the Sherman–Morrison formula, Sherman & Morrison (1950), we can simplify the optimization target and the optimal adverse policy should satisfy:

$$\left\{\mathbf{u}^*_{T-\Delta t}, \sum_{T-\Delta t}^{*}\right\} = \arg \min_{\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}} \left\{\mathbf{u}'_{T-\Delta t} \sum_{T-\Delta t}^{-1} \mathbf{u}_{T-\Delta t}\right\}.$$

So, the proposition is verified at time $T - \Delta t$ and the value function at time $T - \Delta t$ reads:

$$J_{T-\Delta t}(W_{T-\Delta t}) = (\gamma - W_{T-\Delta t}R_f)^2 \cdot \left(1 - \mathbf{u}^{*\prime}_{T-\Delta t} \cdot \left(\sum_{T-\Delta t}^{*} + \mathbf{u}^*_{T-\Delta t}\mathbf{u}^{*\prime}_{T-\Delta t}\right)^{-1} \cdot \mathbf{u}^*_{T-\Delta t}\right).$$

Since the policy of nature satisfies the rectangularity assumption, the factor $1 - \mathbf{u}^{*\prime}_{T-\Delta t} \cdot (\sum^{*}_{T-\Delta t} + \mathbf{u}^*_{T-\Delta t}\mathbf{u}^{*\prime}_{T-\Delta t})^{-1} \cdot \mathbf{u}^*_{T-\Delta t}$ will not influence the optimality at time $T - 2\Delta t$. Therefore, we can finalize the proof by mathematical induction. □

In our proof, in the inner optimization, i.e. the optimization with respect to the asset allocations, we solve the first-order-conditions to obtain the optimality. This is feasible since the asset allocations are assumed to be loosely bounded. In the outer optimization, i.e. choosing the optimal adverse policy, we do not explicitly solve the problem. However, since the rectangularity assumption holds, this will not affect the backward programming process.

3.2. *Robust time-consistent strategy*

Here, we consider the optimal control for the time-consistent strategy. Different from the derivation for the robust pre-commitment strategy, we do not utilize the value function iteration, which will be complicated in the time-consistent case. Instead, by benefiting from the special structure of a time-consistent optimal control, we focus on generating the optimal controls directly. A similar approach is considered in Cong & Oosterlee (2016b). Our findings in the robust time-consistent case can be described by the following proposition.

Proposition 3.2. *For the robust time-consistent optimization problem in Definition* 2.4, *an investor at time t with wealth W_t has the following optimal control*:

$$\mathbf{x}_t^{*\text{tc}}(W_t) = \frac{\sum_t^{*-1} \mathbf{u}_t^*}{2\lambda W_t R_f^{(T-t)/\Delta t-1}}, \quad t = 0, \Delta t, \ldots, T-\Delta t, \tag{3.5}$$

with the parameters $\{\mathbf{u}_t^*, \sum_t^*\}$ *for the worst-case scenario solving the minimization problem* (3.2).

Proof. It is not difficult to prove that the proposition is correct at time $T-\Delta t$:

$$\mathbf{x}_{T-\Delta t}^{*\text{tc}}(W_{T-\Delta t}) = \frac{\sum_{T-\Delta t}^{*-1} \mathbf{u}_{T-\Delta t}^*}{2\lambda W_{T-\Delta t}}, \tag{3.6}$$

and the optimal adverse policy reads:

$$\left\{\mathbf{u}_{T-\Delta t}^*, \sum_{T-\Delta t}^*\right\} = \arg \min_{\mathbf{u}_{T-\Delta t}, \sum_{T-\Delta t}} \left\{\mathbf{u}_{T-\Delta t}' \sum_{T-\Delta t}^{-1} \mathbf{u}_{T-\Delta t}\right\}.$$

Assume that at time $T-2\Delta t$ an investor has wealth $W_{T-2\Delta t}$, then, after performing some control at time $T-2\Delta t$, the corresponding terminal wealth is given by

$$W_T = W_{T-2\Delta t} \cdot (\mathbf{x}_{T-2\Delta t}' \mathbf{R}_{T-2\Delta t}^e + R_f) \cdot ((\mathbf{x}_{T-\Delta t}^{*\text{tc}})' \cdot \mathbf{R}_{T-\Delta t}^{e*} + R_f),$$

where $\mathbf{x}_{T-\Delta t}^{*\text{tc}}$ and $\mathbf{R}_{T-\Delta t}^{e*}$ indicate that the investor will take future optimality into account while designing the optimal control at time $T-2\Delta t$. Further we have:

$$\begin{aligned} W_T &= W_{T-2\Delta t} \cdot (\mathbf{x}_{T-2\Delta t}' \mathbf{R}_{T-2\Delta t}^e + R_f) \cdot ((\mathbf{x}_{T-\Delta t}^{*\text{tc}})' \cdot \mathbf{R}_{T-\Delta t}^{e*} + R_f) \\ &= \frac{1}{2\lambda} \mathbf{u}_{T-\Delta t}^{*\prime} \sum_{T-\Delta t}^{*-1} \mathbf{R}_{T-\Delta t}^{e*} + W_{T-2\Delta t} \cdot (\mathbf{x}_{T-2\Delta t}' \mathbf{R}_{T-2\Delta t}^e + R_f) \cdot R_f \\ &= K_{T-\Delta t} + W_{T-2\Delta t} \cdot (\mathbf{x}_{T-2\Delta t}' \mathbf{R}_{T-2\Delta t}^e + R_f) \cdot R_f, \end{aligned}$$

the second equality is valid since the optimal control is as in Eq. (3.6), which indicates that multiplying the wealth at time $T-\Delta t$ with the optimal time-consistent control should yield a determined number. In the last line, we define the factor $K_{T-\Delta t}$, which only contains elements from time step $T-\Delta t$. Since we have the rectangularity assumption and the excess returns are assumed to be sequentially independent, the optimal control at time $T-2\Delta t$ can be obtained by ignoring factor $K_{T-\Delta t}$ and by solving:

$$\max_{\mathbf{x}_{T-2\Delta t}} \min_{\mathbf{u}_{T-2\Delta t}, \sum_{T-2\Delta t}} \{\mathbb{E}[W_{T-2\Delta t} \cdot (\mathbf{x}'_{T-2\Delta t}\mathbf{R}^e_{T-2\Delta t} + R_f) \cdot R_f \mid W_{T-2\Delta t}]$$

$$- \lambda \cdot \mathrm{Var}[W_{T-2\Delta t} \cdot (\mathbf{x}'_{T-2\Delta t}\mathbf{R}^e_{T-2\Delta t} + R_f) \cdot R_f \mid W_{T-2\Delta t}]\}.$$

By solving this static max-min optimization problem, we can verify the proposition at time $T-2\Delta t$. At the remaining time steps, we can prove the proposition by mathematical induction. The key point in the proof is that, after taking the structure of an optimal control into account, the robust time-consistent problem shares the same optimal control as a robust myopic problem, which can be solved elegantly. □

For the time-consistent case, we do not give details about how the two-layer optimization problem can be solved, because the basic machinery is the same as in the pre-commitment case. We find that the robust pre-commitment and the robust time-consistent strategies share some common features. We will elaborate on them in Sec. 3.3.

3.2.1. *Robust mean-variance efficiency*

In the real world, it is impossible to determine the mean $\mathbf{u}_t$ and the covariance matrix $\sum_t$ for the future excess return at time step t, since only one realization can be observed. Usually, the *basic strategy* is that we specify the excess returns at different time steps to follow a stationary distribution with pre-determined mean $\tilde{\mathbf{u}}$ and covariance matrix $\tilde{\sum}$, i.e. we assume $\mathbf{u}_t = \tilde{\mathbf{u}}$ and $\sum_t = \tilde{\sum}$, for $t = 0, \Delta t, \ldots, T-\Delta t$. Assuming the excess returns to be stationary may not be very restrictive, but prescribing the mean and the covariance may be questionable. For example, if in reality the risky assets follow a stationary distribution with a mean value $\hat{\mathbf{u}}$ and covariance matrix $\hat{\sum}$, that are significantly different from $\tilde{\mathbf{u}}$ and $\tilde{\sum}$, then the corresponding optimal control on the portfolio will certainly be different from the control resulting from the basic strategy. Reflecting on this, the question is: how much different is a portfolio managed under the basic strategy from a portfolio managed under an *optimal strategy* (where we mean by optimal strategy, the case where accurate asset information is used in the model)? As implied by Proposition 3.3, in some situations the basic strategy can generate the same mean-variance efficient frontier as an optimal strategy.

Proposition 3.3. *For two investors, Investor A assuming that the excess returns of risky assets have mean* $\tilde{\mathbf{u}}$ *and covariance* $\tilde{\sum}$ *and Investor B assuming the mean to be* $\hat{\mathbf{u}}$ *and the covariance* $\hat{\sum}$, *their time-consistent strategies generate portfolios with* *the same Sharpe ratio (Sharpe 1966) given the risk-free portfolio as the benchmark, if*

- (*Condition* 1) *their strategies are performed in the same market,*
- (*Condition* 2) $\tilde{\sum}^{-1}\tilde{\mathbf{u}} = K \cdot \hat{\sum}^{-1}\hat{\mathbf{u}}$, *where* $K \in \mathbb{R}^+$.

Proof. The proof is straightforward. We first generate the optimal time-consistent controls of these two investors and then cast the controls on the portfolios starting with the same amount of wealth.

Following a similar derivation as in the proof of Proposition 3.2, we obtain the optimal asset allocations for Investors A and B, as:

$$\mathbf{x}_t^A(W_t) = \frac{\tilde{\sum}^{-1}\tilde{\mathbf{u}}}{2\lambda W_t R_f^{(T-t)/\Delta t-1}}, \quad \mathbf{x}_t^B(W_t) = \frac{\hat{\sum}^{-1}\hat{\mathbf{u}}}{2\lambda W_t R_f^{(T-t)/\Delta t-1}},$$

at time points $t = 0, \Delta t, \ldots, T - \Delta t$.

For Investor A, imposing the optimal control on a portfolio starting with initial wealth W_0 will generate terminal wealth:

$$W_T^A = W_0 R_f^{T/\Delta t} + \frac{\tilde{\sum}^{-1}\tilde{\mathbf{u}}}{2\lambda} \cdot (\mathbf{R}_0^e + \mathbf{R}_{\Delta t}^e + \cdots + \mathbf{R}_{T-\Delta t}^e). \tag{3.7}$$

For Investor B, the terminal wealth can be written in a similar fashion as:

$$W_T^B = W_0 R_f^{T/\Delta t} + \frac{\hat{\sum}^{-1}\hat{\mathbf{u}}}{2\lambda} \cdot (\mathbf{R}_0^e + \mathbf{R}_{\Delta t}^e + \cdots + \mathbf{R}_{T-\Delta t}^e). \tag{3.8}$$

Note that the excess returns $\{\mathbf{R}_t^e\}_{t=0}^{T-\Delta t}$ are random numbers, that indicate the movement of the risky assets in the market. Since we assume that two investors perform their strategies in the same market, the same notations for the excess returns are used in Eqs. (3.7) and (3.8).

The Sharpe ratio of the portfolio managed by Investor A is given by

$$\begin{aligned} S^A &= \frac{\mathbb{E}[W_T^A] - W_0 R_f^{T/\Delta t}}{\sqrt{\mathrm{Var}[W_T^A]}} \\ &= \frac{\mathbb{E}\left[\left(\tilde{\sum}^{-1}\tilde{\mathbf{u}}\right) \cdot (\mathbf{R}_0^e + \mathbf{R}_{\Delta t}^e + \cdots + \mathbf{R}_{T-\Delta t}^e)\right]}{\sqrt{\mathrm{Var}\left[\left(\tilde{\sum}^{-1}\tilde{\mathbf{u}}\right) \cdot (\mathbf{R}_0^e + \mathbf{R}_{\Delta t}^e + \cdots + \mathbf{R}_{T-\Delta t}^e)\right]}}. \end{aligned}$$

In a similar way, we can calculate the Sharpe ratio of the portfolio managed by Investor B. Since we have the assumption that $\tilde{\sum}^{-1}\tilde{\mathbf{u}} = K \cdot \hat{\sum}^{-1}\hat{\mathbf{u}}$, we obtain $S^A = S^B$, i.e. the portfolios managed by Investor A and Investor B share the same Sharpe ratio. □

As shown in Eqs. (3.7) and (3.8), since two investors have different views on the market, the amounts of their terminal wealth are different. However, when their insights on the market satisfy a special condition, Proposition 3.3 implies that their portfolios will share the same return-to-risk ratio. It also means that the mean-variance efficient frontier generated by Investor A with a basic guess of the market will be identical to that generated by Investor B with expert knowledge.

Remark 3.1. Based on Proposition 3.3, we can conclude that for a time-consistent investor managing a portfolio with one risky asset and one risk-free asset, no matter what is the estimation of the return of the risky asset, the strategy is always mean-variance efficient.

Remark 3.2. Björk *et al.* (2014) proposed a time-consistent strategy with a state-dependent risk aversion parameter. If the market forecast is incorrect and the state-dependent risk aversion is adopted instead of the constant risk aversion parameter, the corresponding time-consistent strategy may become less stable. However, as shown in Wang & Forsyth (2011) and Cong & Oosterlee (2016b), the time-consistent strategy with state-dependent risk aversion typically generates a lower mean-variance efficient frontier than the time-consistent strategy with constant risk aversion.

3.3. *Some reflections*

Based on the derivations in Secs. 3.1 and 3.2, we give some insights respectively on the adverse policy of nature and on the optimal policy of an investor in a dynamic mean-variance optimization framework.

On the optimal adverse policy. At a given time step, the optimal policy of nature solves the same optimization problem for either the pre-commitment or the time-consistent problem. In the case of one risky asset, the optimization problem is intuitive: the expectation of the excess return is assumed to be as low as possible and the variance of the excess return as high as possible.

Although the policy of nature may be strictly constrained, we note that constraints on the policy of nature do not influence the smoothness of the value function in the pre-commitment case and we assume that also the smoothness of the value function in the time-consistent case is not affected. At a given time step, the policy of nature is independent of the amount of wealth held by the investor.

Moreover, if we assume that at each time step the feasible sets for nature are identical, we see that the optimal adverse policy is to take a stationary strategy although it is not required to perform in this manner.

On the optimal policy of an investor. At each time step, for either a robust pre-commitment or a robust time-consistent investor, the optimal asset allocations can be obtained by solving the corresponding robust myopic problem, where a myopic investor is assumed to perform an optimal control only for a next time period and to adopt the risk-free strategy afterwards. For a nonrobust pre-commitment or a time-consistent strategy, similar results have been found respectively in Cong & Oosterlee (2016a, 2016b). One application of this is that, in the unconstrained case, the optimal control policy of an investor can be generated efficiently in a forward fashion.

A recently published paper by Pınar (2016) also deals with the robust dynamic mean-variance problem. In that model setting, the author also finds that the robust dynamic strategy is identical to a myopic strategy. Our discussion differs from this in two aspects. First, the investment risk in Pınar (2016) is required to be bounded at each time step and is therefore neither related to the pre-commitment nor to the time-consistent case discussed in this paper. Secondly, as a consequence of choosing different model settings, our optimal asset allocations are not same. The optimal choice of nature is included in deriving our optimal strategy, while the optimal strategy in Pınar (2016) does not take this into account.

As discussed in Wang & Forsyth (2011) and Cong & Oosterlee (2016b), when the nominal mean and variance of the excess returns are correctly predicted, the pre-commitment strategy usually generates a higher mean-variance efficient frontier than the time-consistent strategy, since the time-consistent strategy is restricted by the time-consistency constraint. However, our derivations in Sec. 3.2.1 suggest that, when the market does not perform as expected, *a time-consistent strategy appears to be more stable* and has the potential to yield higher efficient frontiers than a pre-commitment strategy. In the special case where a portfolio consists of one risky asset and one risk-free asset and the risky asset's return has a stationary distribution, any time-consistent mean-variance strategy is guaranteed to generate the same mean-variance efficient frontier as a robust time-consistent strategy.

Of course, all our derivations in Secs. 3.1 and 3.2.1 are based on the assumption that the asset allocations are loosely bounded. If we impose strict constraints on the policy of an investor, the value functions will be nonsmooth and our assumption will be violated. In that case, the optimal strategy of an investor is not the same as an optimal myopic strategy any more, but we can implement numerical methods similar to those in Cong & Oosterlee (2016a, 2016b) to generate constrained solutions (as is done in the numerical section to follow).

Remark 3.3. As in Cong & Oosterlee (2016b), when there is periodic money withdrawal from (or injection in) the portfolio, a pre-commitment strategy will be very sensitive to the amount of withdrawal. When the amount is not as expected, a time-consistent strategy may generate a higher efficient frontier than a pre-commitment strategy.

4. Numerical Experiments

In the previous sections, we discussed the optimal allocations and the worst-case scenarios for the robust mean-variance optimization problem. A robust mean-variance strategy suggests that an investor should take the worst-case scenario into account and adopt a conservative strategy. In this paper, we assume that the uncertainty sets of parameters at different time steps are identical. Our proof then implies that the worst-case market, which solves Eq. (3.2) for either an unconstrained pre-commitment or a time-consistent investor, can be generated by the same model parameters at each time step. In the constrained case, we make the conjecture that this choice of model parameters still yields the worst-case scenario. Since the purpose of constructing the worst-cast scenario is to challenge an investor to achieve good performance of the management strategy, we believe that, even if this conjecture is not correct, this choice of model parameters has a significant influence on the performance of a portfolio. With these criteria for choosing model parameters, we can compare the unconstrained and the constrained cases in the same framework.

Different from Gülpınar & Rustem (2007) and Kuhn *et al.* (2009), we do not consider an uncertainty parameter set in our tests, since a fixed point in the uncertainty set can already give us the worst-case scenario (see Propositions 3.1 and 3.2). We perform numerical tests to examine the robustness of the proposed mean-variance strategies. We assume that a model prediction error is present in our test cases. Two scenarios are considered, one is the model scenario and the other the real-world observed scenario. We derive the asset allocations from the model scenario, adopt those allocations and analyze their impact for the observed real-world scenario. When the model scenario appears conservative, the corresponding strategy can be seen as a robust strategy.

Numerical Algorithm. We utilize the numerical method in Cong & Oosterlee (2016a, 2016b) to solve the pre-commitment and the time-consistent problems, respectively. This numerical method consists of two phases, a sub-optimal solution is first generated in the forward phase and subsequently updating is performed in the backward phase to improve the solution. The backward phase is only necessary for the constrained case. After iterating the forward-backward process for several times, we obtain highly satisfactory results. In our tests, we always choose the myopic strategy as the initial guess and report the result obtained after three backward iterations.

Since we wish to check how a strategy performs if the observed real-world market does not appear to be as expected by the model assumptions, our simulation-based optimization algorithms are slightly adjusted. In the forward phase, we generate paths by using the dynamics of the observed real market; in the backward phase, we update the path-wise controls only based on the model information.

4.1. 1*D problem*

Test Setup. We first perform our numerical tests in the one-dimensional case, where the portfolio contains one risky asset and one risk-free asset. We choose geometric Brownian motion as the dynamics of the risky asset and assume that the log-returns of the risky asset are governed by volatility σ and mean $r_f + \xi \cdot \sigma$. Here, r_f is the log-return of the risk-free asset and ξ is the market price of risk. When we consider constrained optimization scenarios, we impose a bounded

Table 1. Parameter setting.

Risk-free rate r_f	0.03
Investment duration T (year)	1
Re-balancing opportunities	12
Initial wealth W_0	1
Leverage constraint $[x_{\min}, x_{\max}]$	$[0, 1]$

leverage constraint $x \in [x_{\min}, x_{\max}]$ on the portfolio allocations. The values of the parameters are presented in Table 1.

In the following tests, σ and ξ represent the parameters used in the model whereas σ_{real} and ξ_{real} are the parameters in the observed real-world market; the difference between these values we call the model prediction error.

We design numerical experiments with the following questions in mind:

- Is a time-consistent strategy sensitive to model prediction errors?
- When an unexpectedly poor market is encountered, how are efficient frontiers, generated by a pre-commitment or a time-consistent strategy, affected?

In order to answer these questions, various choices for (σ, ξ) and $(\sigma_{\text{real}}, \xi_{\text{real}})$ are made in the following tests.

Robust Mean-Variance Efficiency of Time-Consistent Policy. We first check whether the time-consistent strategy is sensitive to model prediction errors (i.e. real-world parameters $\neq$ model parameters). We assume that the volatility of the real-world market is the same as that indicated by the model, $\sigma_{\text{real}} = \sigma = 0.15$. We consider two model settings, $\xi = 0.1$ and $\xi = 1$, and two market settings, $\xi_{\text{real}} = 0.1$ and $\xi_{\text{real}} = 2$. When we choose $\xi = 0.1$, it yields a robust strategy since the management strategy is generated in the worst-case scenario. The choices of ξ_{real} can be explained as follows: $\xi_{\text{real}} = 0.1$ means that the worst-case scenario indeed appears and $\xi_{\text{real}} = 2$ indicates a good market where the risky asset yields a high return.

As shown in Fig. 1, in the unconstrained as well as the constrained situation, the efficient frontiers generated by the time-consistent strategy are not sensitive to the model prediction errors. The locations of the efficient frontiers only depend on the real-world market parameters. When the real-world market is booming, the time-consistent efficient frontiers are high. When the market is poor, the time-consistent frontiers are low.

Unexpectedly Poor Market. In this numerical test, we check the performance of the pre-commitment and the time-consistent strategies when the market is not as expected. We consider two scenarios, one with fixed mean $(r_f + \sigma \cdot \xi)$ and unexpected (high) volatility[c] of the risky asset return and the other with fixed volatility

[c] In case of unexpected volatility, we choose the volatility of the real market to be 50%, which is similar to the scenario happening after the 2008 financial crisis in the American market Manda (2010).

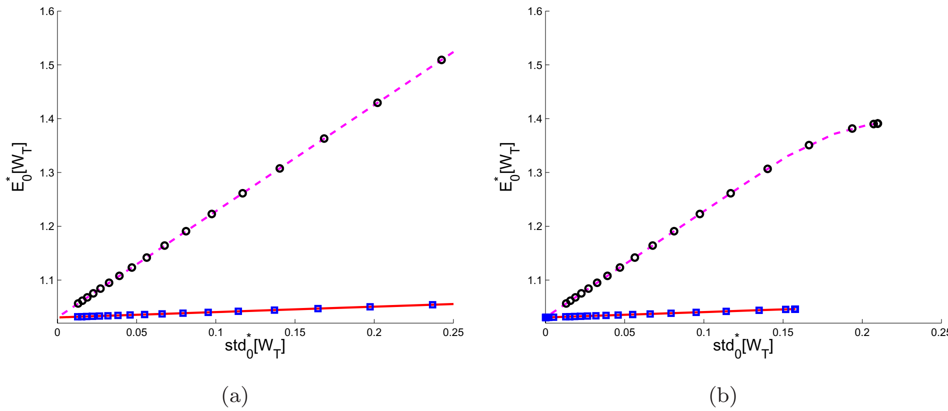

Fig. 1. By choosing different pairs of (ξ, ξ_{real}), we obtain four efficient frontiers generated by the time-consistent strategy. They are respectively represented by: the red line $(\xi, \xi_{\text{real}}) = (0.1, 0.1)$, the magenta dashed line $(\xi, \xi_{\text{real}}) = (0.1, 2)$, the squares $(\xi, \xi_{\text{real}}) = (1, 0.1)$ and the circles $(\xi, \xi_{\text{real}}) = (1, 2)$. (a) Unconstrained case and (b) with leverage constraint.

Table 2. Parameters for modeling an unexpectedly poor market.

Set I: $\sigma = 0.1$, $\xi = 0.33$, $\sigma_{\text{real}} = 0.5$, $\xi_{\text{real}} = 0.066$.
Set II: $\sigma = 0.15$, $\xi = 1$, $\sigma_{\text{real}} = 0.15$, $\xi_{\text{real}} = 0.1$.

and unexpected (low) mean of the risky asset return. In these two scenarios, the parameters are chosen as in Table 2.

As seen in our first experiment, the time-consistent strategy is not sensitive to the model prediction and yields efficient frontiers that can also be achieved by a time-consistent investor with correct market information. When a bounded leverage constraint is introduced, the time-consistent frontier tends to be somewhat lower than in the unconstrained case.

A surprising finding is that, when the model prediction is inconsistent with the real market, the pre-commitment strategy may generate lower efficient frontiers than a time-consistent strategy. According to Wang & Forsyth (2011) and Cong & Oosterlee (2016b), if the market moves according to the model prediction, a pre-commitment strategy generates a higher frontier than a time-consistent strategy. This is due to the fact that time-consistency can be regarded as a constraint on a pre-commitment strategy. However, in case of an unexpectedly poor market, the time-consistency constraint may *protect an investor*, while a pre-commitment investor may suffer from the poor market. As presented in Fig. 2, in both situations, the unconstrained pre-commitment strategy generates the lowest efficient frontier. When the constraint is introduced into the pre-commitment strategy, the efficient frontier gets higher than in the unconstrained case. This is not difficult to understand. When a model yields the correct prediction, introducing a constraint forms a restriction on a portfolio; when a model generates an incorrect prediction,

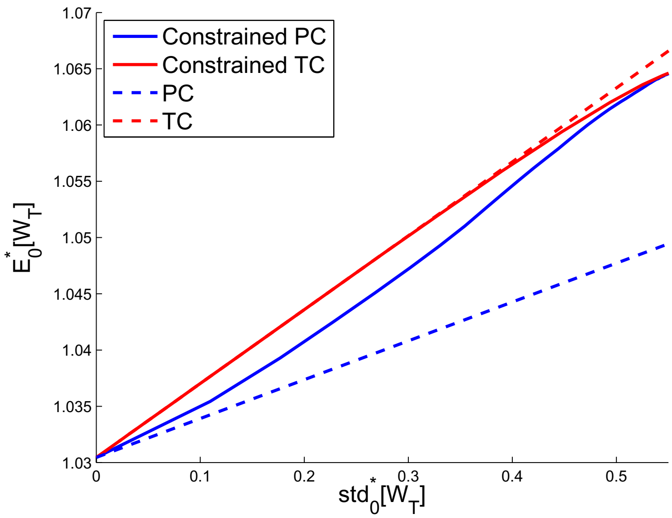

(a)

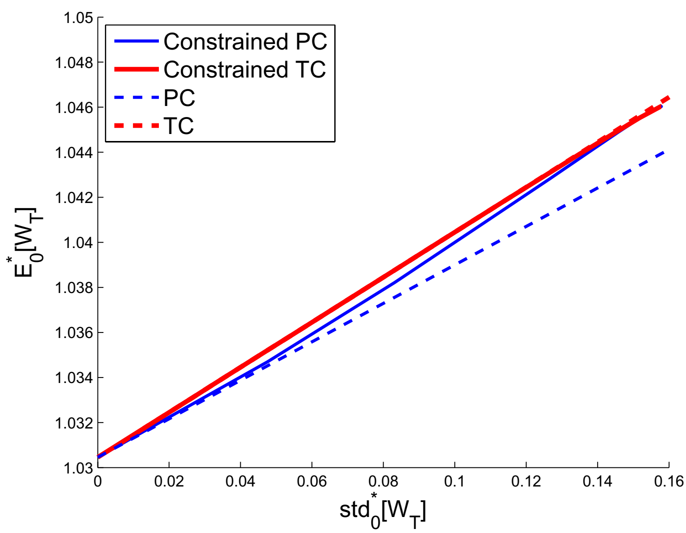

(b)

Fig. 2. Comparing the performance of a pre-commitment and a time-consistent strategy when the market is unexpectedly poor. In the unexpected volatility case, we choose parameters from Set I in Table 2. In the unexpected mean case, parameters from Set II are used. (a) Unexpected volatility and (b) unexpected mean.

the constraint acts as a "protection". Therefore, when market movement is not as anticipated, the constrained pre-commitment strategy may perform better than its unconstrained counterpart. The "protection" by means of the portfolio constraints is also reported in Jagannathan & Ma (2003).

We also see that the constrained pre-commitment and time-consistent frontiers coincide at their right ends. When an investor pursues a very high return, the maximal possible allocation will be set at almost all time steps. Therefore, in this situation, a constrained pre-commitment strategy will be similar to a constrained time-consistent strategy.

Remark 4.1. When we consider an unexpectedly booming market (for example $\sigma_{\text{real}} = \sigma$ and $\xi_{\text{real}} > \xi$), the pre-commitment strategy generates a higher efficient frontier than the time-consistent strategy. In this scenario, introducing constraints makes both the pre-commitment and the time-consistent efficient frontiers lower, as expected.

Although the time-consistent strategy can generate a robust efficient frontier, it does not mean that the time-consistent investor, who designs a strategy with expected return $\mathbb{E}_0^*[W_T] = d$ in mind, will achieve this amount of return in an unexpected market. When the market is worse than expected, the pre-commitment investor will attempt to reach the predetermined target by taking more risk. Meanwhile, the time-consistent investor may just be satisfied with a lower mean return associated with less risk. In this case, it is interesting to examine both strategies by checking the probability of getting less than, say, 90% of the predetermined return d. This shortfall probability also reflects the robustness of an investment strategy.

As shown in Fig. 3, when the market is as expected in the model, both the time-consistent strategy and the pre-commitment strategy lead to low shortfall probabilities. When an unexpectedly poor market occurs, the shortfall probabilities increase. It is not easy to say which strategy is more robust with respect to the shortfall probability criterion. According to Fig. 3(b), the time-consistent strategy has a higher probability of generating the terminal wealth lower than 90% of the desired level. However, if we consider the probability of getting the terminal wealth lower than 60% of the desired target, the time-consistent strategy appears to be less risky as shown in Fig. 3(c).

4.2. 2*D problem with unexpected correlation*

An advantage of using the simulation-based numerical algorithms from Cong & Oosterlee (2016a, 2016b) is that they can be generalized to higher-dimensional scenarios. In this part, we consider a portfolio with two risky assets and one risk-free asset. In terms of model uncertainty, we consider a scenario where the correlation between two risky assets is not as predicted. The parameters for risky assets A and B are respectively shown in Table 3, where ρ denotes the correlation between the two risky assets in the model and ρ_{real} denotes the observed real-world correlation in the market. In the constrained case, we consider bounded leverage constraints $[x_{\min}, x_{\max}] = [0, 0.5]$ on both risky assets. For the other parameters, we set them as in Table 1.

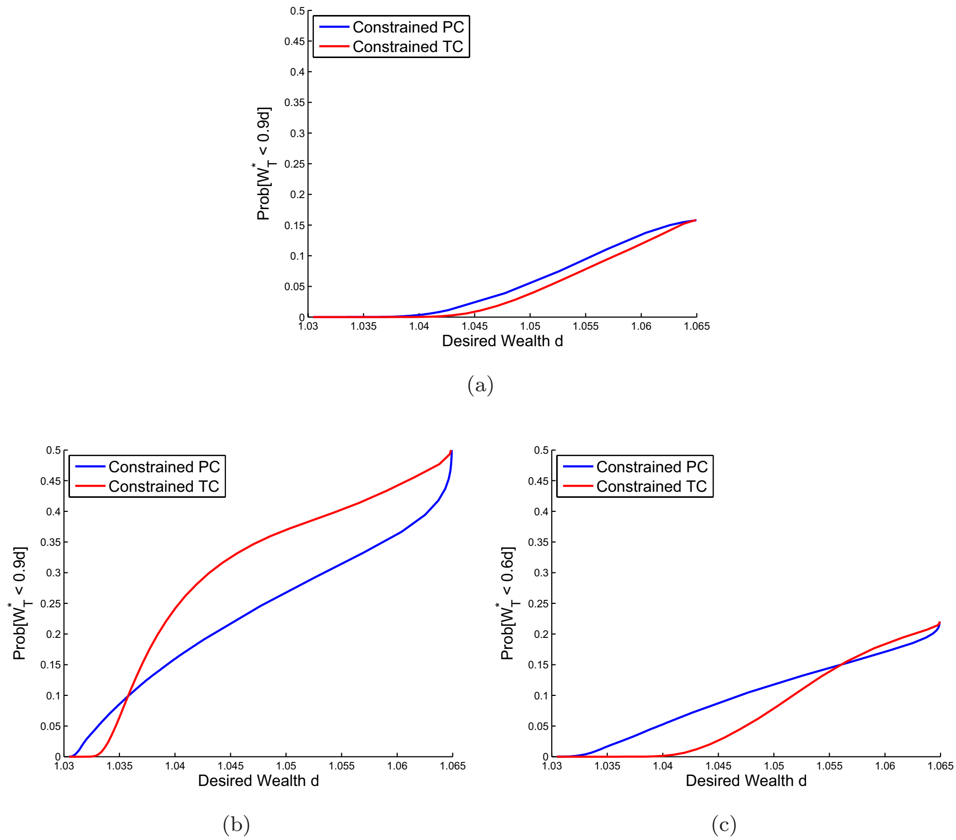

Fig. 3. The shortfall probability of the pre-commitment and the time-consistent strategies in the constrained case. The x-axis displays the desired target wealth and the y-axis the shortfall probability. For "an unexpectedly poor market", we use the parameters from Set I in Table 2. For "an expected market", the same model parameters are used and the real-world market parameters are assumed to be identical to the model parameters. (a) In an expected market, (b) in an unexpectedly poor market and (c) in an unexpectedly poor market.

Table 3. Parameters for two risky assets A and B.

$\sigma_A = 0.2,\ \xi_A = 0.5,\ \sigma_B = 0.4,\ \xi_B = 0.5$
$\rho = -0.9,\ \rho_{\text{real}} = 0.9.$

In the unconstrained case, we first check what kind of results a correlation prediction error can bring. We compare the efficient frontiers generated by the pre-commitment and time-consistent strategies, when the correlation is inaccurately predicted, as shown in Table 3, as well as the frontiers generated by both strategies when the observed real-world correlation is predicted accurately ($\rho = \rho_{\text{real}}$). As shown in Fig. 4, when accurate information is available, the pre-commitment frontier is slightly higher than the time-consistent frontier. However, when the model

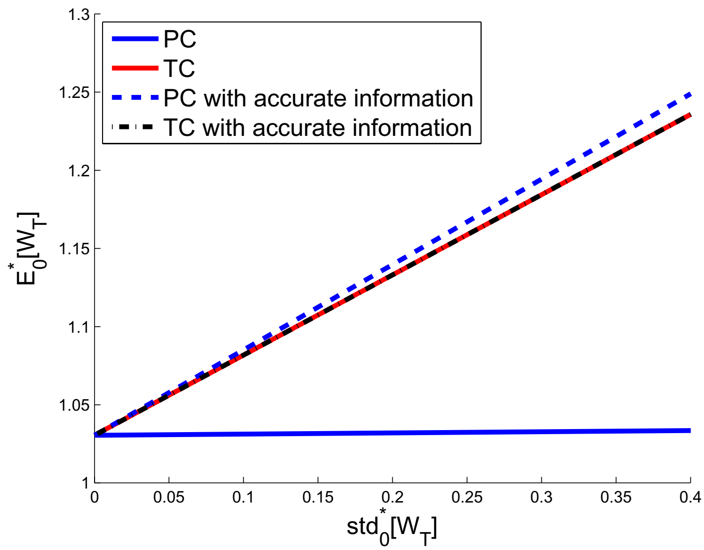

Fig. 4. Comparison of the efficient frontiers generated by the pre-commitment and the time-consistent strategies in cases with or without accurate information of the correlation between risky assets.

correlation is not accurate, the pre-commitment strategy degrades, while the time-consistent strategy does not change significantly in terms of mean-variance efficiency. The time-consistent strategy with inaccurate information generates almost the same frontier as the one with accurate information. When we consider a test case with a longer investment horizon or fewer re-balancing opportunities, the difference between these two time-consistent frontiers will be more pronounced. However, in general, the time-consistent strategy is more robust than the pre-commitment strategy in terms of an inaccurate prediction of the assets correlation.

In Fig. 5, we show the frontiers generated by the pre-commitment and the time-consistent strategies in the unconstrained and the constrained cases, when the asset correlation is not correctly predicted by the model. When the constraints are introduced on the allocations, the pre-commitment frontier increases and the time-consistent frontier decreases. In our test setting, when the mean return is not large, the time-consistent frontier is higher than its pre-commitment counterpart.

However, please note that when the misprediction of the correlation is not very significant, for example $[\rho, \rho_{\text{real}}] = [-0.2, 0.2]$, the constrained pre-commitment strategy still generates higher frontiers. This is as expected, since the pre-commitment strategy should generate higher frontiers when the market information is known exactly.

In Fig. 6, we show the asset allocations of both strategies over time. We choose a scenario where both strategies generate mean returns that are equal to 1.13. As shown in Fig. 5, this is approximately the point where the time-consistent and the

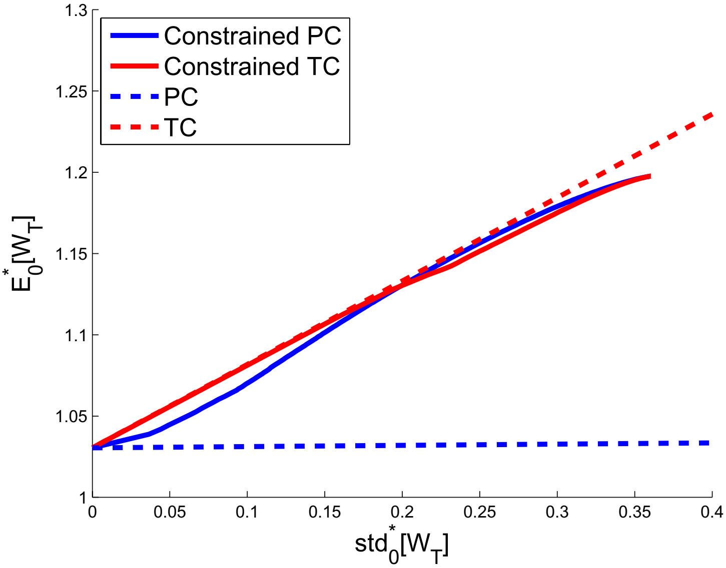

Fig. 5. Comparison of the efficient frontiers when the correlations between risky assets are inaccurately predicted.

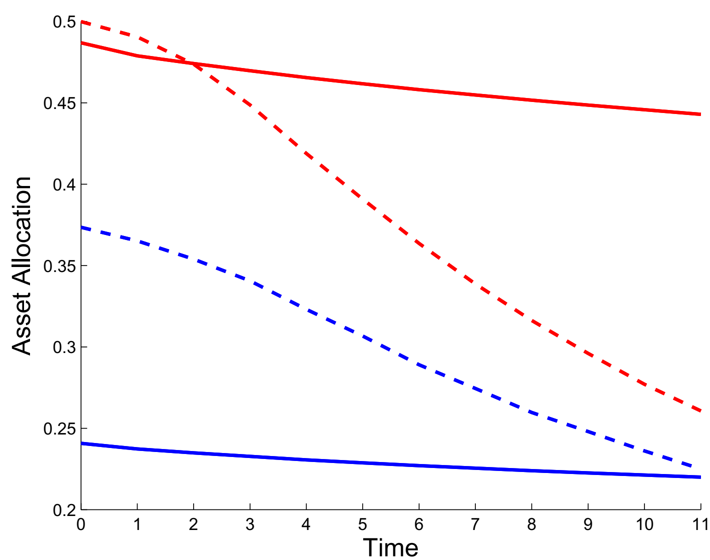

Fig. 6. Comparison of asset allocations in the two-dimensional case. The red lines represent the allocations in asset A and the blue lines the allocations in asset B, which exhibits higher volatility than asset A. The straight lines denote the allocations for the time-consistent strategy and the dashed lines for the pre-commitment strategy.

pre-commitment frontiers cross. The presented allocations are the average values of the allocations on all simulated paths. Although both strategies generate returns with the same mean and also similar variances, their allocations are significantly different. By adopting the pre-commitment strategy, an investor assigns more money to the risky assets initially and shifts to risk-free asset allocations at the end of the investment period. This is due to the fact that this investor has a target in mind and close to the target she may not take risk to achieve higher wealth levels. For a time-consistent investor, the optimal asset allocations are quite different. Since the time-consistent investor is not satisfied with a target, the strategy does not reduce to a risk-free strategy. Initially, a time-consistent investor is more risk-averse than a pre-commitment investor; however, at the end, the time-consistent investor appears to be more risk-seeking.

5. Conclusion

In this paper, we considered the robust pre-commitment and the robust time-consistent mean-variance optimization problems. In the *unconstrained case*, a specific equation for determining the worst-case scenario was derived for the robust pre-commitment and the robust time-consistent strategies. At a given time step, the optimal allocations generated by both strategies are then identical to their myopic counterparts, where an investor derives the optimal allocation for one upcoming time period, assuming that a risk-free strategy will be taken in the future.

The robustness of the pre-commitment and the time-consistent strategies is checked. Our analysis and the corresponding numerical experiments suggest that a time-consistent strategy appears to be more robust in terms of model prediction errors. When an unexpectedly poor market is encountered, the time-consistent strategy may generate higher efficient frontiers than the pre-commitment strategy. Introducing constraints into the robust pre-commitment strategy can even increase the frontiers, since the constraints may serve as a "protection". In the two-dimensional case, the influence of inaccurately predicting the correlation between risky asset returns was examined. Again we found that the pre-commitment strategy may be vulnerable to such prediction errors and constraints on asset allocations can increase a pre-commitment frontier. Meanwhile, the time-consistent strategy still performed in a robust way. We checked the asset allocations of both strategies when they generate similar mean-variance pairs, and found that a pre-commitment investor prefers to bear more risk at the beginning of an investment period while a time-consistent investor appears to be more risk-seeking than a pre-commitment investor at the end of the period.

References

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

R. Bellman (1957) *Dynamic Programming.* US: Princeton University Press.

D. Bertsimas, D. B. Brown & C. Caramanis (2011) Theory and applications of robust optimization, *SIAM Review* **53** (3), 464–501.

M. J. Best & R. R. Grauer (1991) Sensitivity analysis for mean-variance portfolio problems, *Management Science* **37** (8), 980–989.

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

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

F. Black & R. Litterman (1992) Global portfolio optimization, *Financial Analysts Journal* **48** (5), 28–43.

M. Britten-Jones (1999) The sampling error in estimates of mean-variance efficient portfolio weights, *The Journal of Finance* **54** (2), 655–671.

F. Cong & C. W. Oosterlee (2016a) Multi-period mean-variance portfolio optimization based on Monte-Carlo simulation, *Journal of Economic Dynamics and Control* **64** (1), 23–38.

F. Cong & C. W. Oosterlee (2016b) On pre-commitment aspects of a time-consistent strategy for a mean-variance investor, *Journal of Economic Dynamics and Control* **70** (1), 178–193.

X. Cui, J. Gao, X. Li & D. Li (2014a) Optimal multi-period mean-variance policy under no-shorting constraint, *European Journal of Operational Research* **234** (2), 459–468.

X. Cui, X. Li, D. Li & Y. Shi (2014b) Time consistent behavior portfolio policy for dynamic mean-variance formulation, Available at SSRN 2480299.

X. Cui, L. Xu & Y. Zeng (2015) Continuous time mean-variance portfolio optimization with piecewise state-dependent risk aversion, *Optimization Letters* **10**(8), 1681–1691.

L. El Ghaoui & A. Nilim (2005) Robust solutions to Markov decision problems with uncertain transition matrices, *Operations Research* **53** (5), 780–798.

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

L. Garlappi, R. Uppal & T. Wang (2006) Portfolio selection with parameter and model uncertainty: A multi-prior approach, *The Review of Financial Studies* **20** (1), 41–81.

R. Gerrard, S. Haberman & E. Vigna (2004) Optimal investment choices post-retirement in a defined contribution pension scheme, *Insurance*: *Mathematics and Economics* **35** (2), 321–342.

D. Goldfarb & G. Iyengar (2003) Robust portfolio selection problems, *Mathematics of Operations Research* **28** (1), 1–38.

N. Gülpınar & B. Rustem (2007) Worst-case robust decisions for multi-period mean–variance portfolio optimization, *European Journal of Operational Research* **183** (3), 981–1000.

S. Haberman & E. Vigna (2002) Optimal investment strategies and risk measures in defined contribution pension schemes, *Insurance*: *Mathematics and Economics* **31** (1), 35–69.

B. V. Halldórsson & R. H. Tütüncü (2003) An interior-point method for a class of saddle-point problems, *Journal of Optimization Theory and Applications* **116** (3), 559–590.

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

G. N. Iyengar (2005) Robust dynamic programming, *Mathematics of Operations Research* **30** (2), 257–280.

R. Jagannathan & T. Ma (2003) Risk reduction in large portfolios: Why imposing the wrong constraints helps, *The Journal of Finance* **58** (4), 1651–1684.

J. H. Kim, W. C. Kim & F. J. Fabozzi (2014) Recent developments in robust portfolios with a worst-case approach, *Journal of Optimization Theory and Applications* **161** (1), 103–121.

D. Kuhn, P. Parpas, B. Rustem & R. Fonseca (2009) Dynamic mean-variance portfolio analysis under model risk, *Journal of Computational Finance* **12** (91115), 7.

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

X. Li, X. Zhou & A. E. Lim (2002) Dynamic mean-variance portfolio selection with no-shorting constraints, *SIAM Journal on Control and Optimization* **40** (5), 1540–1555.

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

F. Maccheroni, M. Marinacci & A. Rustichini (2006) Ambiguity aversion, robustness, and the variational representation of preferences, *Econometrica* **74** (6), 1447–1498.

K. Manda (2010) Stock market volatility during the 2008 financial crisis, PhD Thesis, Stern School of Business, New York University.

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

C. Perret-Gentil & M.-P. Victoria-Feser (2005) Robust mean-variance portfolio selection, Available at SSRN 721509.

M. Ç. Pınar (2016) On robust mean-variance portfolios, *Optimization*, doi:10.1080/02331934.2015.1132216.

W. F. Sharpe (1966) Mutual fund performance, *The Journal of Business* **39** (1), 119–138.

J. Sherman & W. J. Morrison (1950) Adjustment of an inverse matrix corresponding to a change in one element of a given matrix, *The Annals of Mathematical Statistics* **21** (1), 124–127.

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

R. Tevzadze, T. Toronjadze & T. Uzunashvili (2013) Robust utility maximization for a diffusion market model with misspecified coefficients, *Finance and Stochastics* **17** (3), 535–563.

R. H. Tütüncü & M. Koenig (2004) Robust asset allocation, *Annals of Operations Research* **132** (1–4), 157–187.

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

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

J. Wang & P. A. Forsyth (2011) Continuous time mean variance asset allocation: A time-consistent strategy, *European Journal of Operational Research* **209** (2), 184–201.

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

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

S. Zhu, D. Li & S. Wang (2004) Risk control over bankruptcy in dynamic portfolio selection: A generalized mean-variance formulation, *IEEE Transactions on Automatic Control* **49** (3), 447–457.