SIAM J. FINANCIAL MATH.
Vol. 11, No. 2, pp. 358–384

Multiperiod Mean Conditional Value at Risk Asset Allocation: Is It Advantageous to Be Time Consistent?*

Peter A. Forsyth†

Abstract. We formulate the multiperiod, time consistent mean-CVAR (conditional value at risk) asset allocation problem in a form amenable to numerical computation. Our numerical algorithm can impose realistic constraints such as no shorting, no leverage, and discrete rebalancing. We focus on long term (i.e., 30 year) strategies, which would be typical of an investor in a defined contribution pension plan. A comparison with precommitment mean-CVAR strategies shows that adding the time consistent constraint compares unfavorably with the pure precommitment strategy. Since the precommitment strategy computed at time zero is identical to a time consistent strategy based on an alternative objective function, the time zero precommitment mean-CVAR strategy is implementable in this case. Hence it would seem that there is little to be gained from enforcing time consistency.

Key words. multi-period mean CVAR, time consistent, precommitment, asset allocation

AMS subject classifications. 91G, 65N06, 65N12, 35Q93

DOI. 10.1137/19M124650X

1. Introduction.

Long term investors, such as those saving for retirement in a defined contribution (DC) pension are motivated by asset allocation strategies which are optimal under multi-period criteria. Initial work on multi-period strategies under mean variance criteria was carried out in (Li and Ng, 2000; Zhou and Li, 2000). More recent issues are addressed in (Basak and Chabakauri, 2010; Bjork and Murgoci, 2010; Wang and Forsyth, 2011; Bjork and Murgoci, 2014; Bjork, Murgoci, and Zhou, 2014). It is important to distinguish between two categories of multi-period optimal control strategies.

Precommitment strategies (Li and Ng, 2000; Zhou and Li, 2000) are globally optimal when viewed from the initial time, but these strategies are not *time consistent*. Consider an asset allocation problem with a fixed stopping time T. We compute the optimal strategy, as a function of the state variables, at time zero. Now, suppose we recompute the strategy at some later time t, $0 < t < T$. For a given state of the system, the strategy we compute at this later time may not agree with the strategy computed at time zero. This has led many investigators to label precommitment mean-variance strategies as *nonimplementable*, since the investor has an incentive to deviate from the precommitment strategy at $t > 0$. However, in the mean-variance case, this objection is perhaps not well thought out. For any precommitment mean-variance optimal strategy, there exists a parameter W^*, such that this

*Received by the editors February 25, 2019; accepted for publication (in revised form) November 7, 2019; published electronically April 13, 2020.
https://doi.org/10.1137/19M124650X
Funding: The author acknowledges support from the Natural Sciences and Engineering Research Council of Canada (NSERC), grant RGPIN-2017-03760.

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

strategy is also optimal for the strategy which minimizes quadratic shortfall (Cui et al., 2012; Dang and Forsyth, 2016)

$$E\left[(\min(W_T - W^*,0))^2\right], \tag{1.1}$$

where W_T is the accumulated wealth at $t = T$, and $E[\cdot]$ is the expectation operator. For a fixed W^*, we can determine the optimal policy for objective function (1.1) using dynamic programming. Hence, if we fix the target W^*, the precommitment policy is the time consistent strategy for objective function (1.1). In fact, as noted in Vigna (2017),

> *"The equivalence between [precommitment] mean-variance criterion and the target-based approach is one of the characteristics that make the mean-variance preferences appealing with respect to other types of preferences.... For the average pension fund member it is easier to select a wealth target rather than an abstract index."*

This makes the objective function (1.1) with a fixed target W^* very useful in practice. The objective function (1.1), with a fixed W^*, has been termed the *mean-variance induced utility maximization problem* in Strub, Li, and Cui (2019), to emphasize that this objective function does in fact generate a time consistent policy.

In an effort to force time consistency, while retaining the mean-variance objective function when viewed at times $t > 0$, several authors have developed techniques to ensure this property (Basak and Chabakauri, 2010; Bjork and Murgoci, 2010; Wang and Forsyth, 2011; Bjork and Murgoci, 2014; Bjork, Murgoci, and Zhou, 2014; Van Staden, Dang, and Forsyth, 2018; Landriault et al., 2018). Since we can view time consistent mean-variance strategies as precommitment strategies with an additional constraint, it is immediately obvious that time consistent strategies are not globally optimal as seen at time zero. In fact, as noted by Cong and Oosterlee (2016a,b), the precommitment mean-variance strategy is consistent with a fixed target, but not with a risk aversion attitude. Conversely, time consistent mean-variance strategies are consistent with a fixed risk aversion, but not with a fixed target. Further discussion concerning the equivalence of precommitment mean variance and the time consistent policy which minimizes the target based objective function (1.1) can be found in Vigna (2014); Menoncin and Vigna (2017). The merits and demerits of time consistent and precommitment policies are also discussed in Vigna (2017). Some other interesting problems with time consistency are noted in Bensoussan, Wong, and Yam (2019), in the case of a wealth dependent risk aversion parameter. Suffice to say, we cannot dismiss precommitment policies out of hand.

An interesting alternative for using variance to measure risk is conditional value at risk (CVAR). CVAR at level α is simply the average of the worst α fraction of outcomes, hence, is a measure of tail risk. Precommitment mean-CVAR strategies were developed in Miller and Yang (2017); Gao et al. (2017). As we shall see, although precommitment mean-CVAR strategies are formally time inconsistent, these strategies are identical at time zero to a linear target shortfall policy with a fixed target. This alternative objective function generates a time consistent policy. Hence these strategies are in fact implementable. Under this alternative objective function, the investor has no incentive to deviate from the strategy computed at time zero.

The objective of this article is to formulate the time-consistent mean-CVAR problem into a form which is amenable to computation. We do this by expanding the state space to

include the local value at risk (VAR) as an independent variable. We develop a numerical technique for solving this problem. Since we use numerical methods, we can consider realistic constraints, such as no leverage, no shorting, and discrete rebalancing. We then compare the time consistent and precommitment mean-CVAR solutions for a long term investment problem. The stochastic process parameters are determined by fitting to 90 years of market data.

We should mention that our approach differs from the method in Cui et al. (2019), where the authors reduce the time consistent mean-CVAR problem to the solution of a convex program based on a finite Monte Carlo sampling of return paths. In particular, Cui et al. (2019) assumes a lump sum investment, where the control (in terms of amount in each asset) is shown to be a piecewise linear function of wealth. The lump sum assumption does not hold in general for DC plan investors, who typically make periodic contributions to an investment portfolio.

Our main conclusion is that imposing a time consistent constraint on the solution of a mean-CVAR problem appears to result in an investment strategy with undesirable properties. On the other hand, based on the cumulative distribution function of the final wealth, the time consistent linear target shortfall strategy (which coincides with precommitment mean-CVAR at time zero) is superior in terms of tail risk reduction compared to the time consistent mean-CVAR policy.

We note that there are many possible objective functions which can be used to determine strategies for asset allocation. A nonexhaustive list includes time consistent multi-period mean variance, minimizing quadratic shortfall, deterministic strategies, and various traditional utility functions. A study of several strategies is carried out in Forsyth and Vetzal (2019). Our objective in this paper is deliberately narrow. We focus on mean-CVAR policies, and argue that forcing time consistency in this case is undesirable. This narrows down (somewhat) the choice of possible objective functions which should be considered.

2. Formulation.

For simplicity we assume that there are only two assets available in the financial market, namely, a risky asset and a risk-free asset. In practice, the risky asset would be a broad market index fund.

The investment horizon is T. S_t and B_t, respectively, denote the *amounts* invested in the risky and risk-free assets at time t, $t \in [0, T]$. In general, these amounts will depend on the investor's strategy over time, as well as changes in the unit prices of the assets. In the absence of an investor determined control (i.e., cash injections or rebalancing), all changes in S_t and B_t result from changes in asset prices. In this case (absence of control), we assume that S_t follows a jump diffusion process. Let $S_{t^-} = S(t - \epsilon), \epsilon \to 0^+$, i.e., t^- is the instant of time before t, and let ξ be a random number representing a jump multiplier. When a jump occurs, $S_t = \xi S_{t^-}$. Allowing discontinuous jumps allows us to explore the effects of severe market crashes on the risky asset holding. We assume that $\log(\xi)$ follows a double exponential distribution (Kou, 2002; Kou and Wang, 2004). If a jump occurs, p_{up} is the probability of an upward jump, while $1 - p_{up}$ is the chance of a downward jump. The density function for $y = \log(\xi)$ is

$$f(y) = p_{up}\eta_1 e^{-\eta_1 y}\mathbf{1}_{y\geq 0} + (1 - p_{up})\eta_2 e^{\eta_2 y}\mathbf{1}_{y<0} \ . \tag{2.1}$$

For future reference, note that

$$E[y = \log \xi] = \frac{p_{up}}{\eta_1} - \frac{(1-p_{up})}{\eta_2}\ ,\ E[\xi] = \frac{p_{up}\eta_1}{\eta_1 - 1} + \frac{(1-p_{up})\eta_2}{\eta_2 + 1},$$
$$E[(\xi-1)^2] = \frac{p_{up}\eta_1}{\eta_1 - 2} + \frac{(1-p_{up})\eta_2}{\eta_2 + 2} - 2\left(\frac{p_{up}\eta_1}{\eta_1 - 1} + \frac{(1-p_{up})\eta_2}{\eta_2 + 1}\right) + 1\ . \tag{2.2}$$

In the absence of control, S_t evolves according to

$$\frac{dS_t}{S_{t^-}} = (\mu - \lambda_\xi \kappa_\xi)\ dt + \sigma\, dZ + d\left(\sum_{i=1}^{\pi_t}(\xi_i - 1)\right),$$
$$\kappa_\xi = E[\xi - 1]\ , \tag{2.3}$$

where μ is the (uncompensated) drift rate, σ is the volatility, dZ is the increment of a Wiener process, π_t is a Poisson process with positive intensity parameter λ_ξ, and ξ_i are independent and identically distributed positive random variables having distribution (2.1). Moreover, ξ_i, π_t, and Z are assumed to all be mutually independent.

We focus on jump diffusion models for long term equity dynamics since sudden drops in the equity index just before retirement can have a devastating impact on retirement portfolios. Since we consider discrete rebalancing, the jump process models the cumulative effects of large market drops between rebalancing times. Previous studies show that stochastic volatility effects are small for the long term investor (Ma and Forsyth, 2016). This can be traced to the fact that stochastic volatility models are mean reverting, with typical mean reversion times of less than one year.

In the absence of control, we assume that the dynamics of the amount B_t invested in the risk-free asset are

$$dB_t = rB_t\, dt, \tag{2.4}$$

where r is the (constant) risk-free rate.

Remark 2.1 (parsimonious model). Equations (2.3) and (2.4) are very simplified models of real stock and bond processes. However, tests of the controls determined using these parsimonious model processes on bootstrapped historical market data gives good results for a variety of objective functions (Forsyth, Vetzal, and Westmacott, 2019). This suggests that the parsimonious model (2.3)–(2.4) seems sufficient for the purposes of generating an asset allocation strategy for the long term investor.

We define the investor's total wealth at time t as

$$\text{Total wealth } \equiv W_t = S_t + B_t. \tag{2.5}$$

We impose the constraints that shorting stock and using leverage (i.e., borrowing) are not permitted, which would be typical of a DC plan retirement savings account.

Properties 2.1 (constant coefficients). *In the following, we will assume that the stochastic process parameters $r, \mu, \sigma, \lambda_\xi, p_u, \eta_1, \eta_2$ are constants, independent of (S,B,t). This also implies that κ_ξ is a constant as well.*

3. Notational conventions. To avoid subscript clutter, in the following, we will occasionally use the notation $S_t \equiv S(t), B_t \equiv B(t)$, and $W_t \equiv W(t)$. Let the inception time of the investment be $t_0 = 0$. We consider a set $\mathcal{T}$ of pre-determined *rebalancing times*,

$$\mathcal{T} \equiv \{t_0 = 0 < t_1 < \cdots < t_M = T\}. \tag{3.1}$$

For simplicity, we specify $\mathcal{T}$ to be equidistant with $t_i - t_{i-1} = \Delta t = T/M$, $i = 1, \ldots, M$. At each rebalancing time t_i, $i = 0, 1, \ldots, M-1$, the investor (i) injects an amount of cash q_i into the portfolio, and then (ii) rebalances the portfolio. At $t_M = T$, the portfolio is liquidated. In the following, given a time dependent function $f(t)$, we will use the shorthand notation

$$f(t_i^+) \equiv \lim_{\epsilon \to 0^+} f(t_i + \epsilon), \qquad f(t_i^-) \equiv \lim_{\epsilon \to 0^+} f(t_i - \epsilon)\ . \tag{3.2}$$

We assume that there are no taxes or other transaction costs, so that the condition

$$W(t_i^+) = W(t_i^-) + q_i \tag{3.3}$$

holds. Typically, DC plan savings are held in a tax advantaged account, with no taxes triggered by rebalancing. With infrequent (e.g., yearly) rebalancing, we also expect transaction costs to be small.

We denote by $X(t) = (S(t), B(t))$, $t \in [0,T]$, the multidimensional controlled underlying process, and by $x = (s, b)$ the state of the system. Let the rebalancing control $p_i(\cdot)$ be the fraction invested in the risky asset at the rebalancing date t_i, i.e.,

$$p_i\left(X(t_i^-)\right) = p\left(X(t_i^-), t_i\right) = \frac{S(t_i^+)}{S(t_i^+) + B(t_i^+)}\ . \tag{3.4}$$

Note that the controls depend on the state of the investment portfolio, before the rebalancing occurs, i.e., $p_i(\cdot) = p\left(X(t_i^-), t_i\right)\big) = p\left(X_i^-, t_i\right)$, $t_i \in \mathcal{T}$, where $\mathcal{T}$ is the set of rebalancing times. More specifically, in our case, we find the optimal strategies amongst all strategies with constant wealth (after injection of cash), so that

$$\begin{aligned} p_i(\cdot) &= p(W(t_i^+), t_i), \\ & W(t_i^+) = S(t_i^-) + B(t_i^-) + q_i, \\ & S(t_i^+) = S_i^+ = p_i(W_i^+)\ W_i^+ \ \ ; \ \ B(t_i^+) = B_i^+ = (1 - p_i(W_i^+))\ W_i^+\ . \end{aligned} \tag{3.5}$$

Let $\mathcal{Z}$ represent the set of admissible values of the control $p_i(\cdot)$. An admissible control $\mathcal{P} \in \mathcal{A}$, where $\mathcal{A}$ is the admissible control set, can be written as

$$\mathcal{P} = \{p_i(\cdot) \in \mathcal{Z}\ :\ i = 0, \ldots, M-1\}. \tag{3.6}$$

We also define $\mathcal{P}_n \equiv \mathcal{P}_{t_n} \subset \mathcal{P}$ as the tail of the set of controls in $[t_n, t_{n+1}, \ldots, t_{M-1}]$, i.e.,

$$\mathcal{P}_n = \{p_n(\cdot), \ldots, p_{M-1}(\cdot)\}\ . \tag{3.7}$$

4. Definition of CVAR. Let $g(W_T)$ be the probability density function of wealth W_T at $t = T$. Let

$$\int_{-\infty}^{W_\alpha^*} g(W_T)\, dW_T = \alpha, \tag{4.1}$$

i.e., $Pr[W_T > W_\alpha^*] = 1 - \alpha$. We can interpret W_α^* as the VAR at level α. The CVAR at level α is then

$$\text{CVAR}_\alpha = \frac{\int_{-\infty}^{W_\alpha^*} W_T\ g(W_T)\ dW_T}{\alpha}, \tag{4.2}$$

which is the mean of the worst α fraction of outcomes. Typically $\alpha \in \{.01, .05\}$. Note that the definition of CVAR in (4.2) uses the probability density of the final wealth distribution, not the density of *loss*. Hence, in our case, a larger value of CVAR (i.e., a larger value of average worst case terminal wealth) is desired.

Given an expectation under control $\mathcal{P}$, $E_{\mathcal{P}}[\cdot]$, as noted by Rockafellar and Uryasev (2000), CVAR_α can be alternatively written as

$$\begin{aligned} \text{CVAR}_\alpha &= \sup_{W^*} E_{\mathcal{P}}\left[W^* + \frac{1}{\alpha}(W_T - W^*)^-\right], \\ &\qquad (W_T - W^*)^- \equiv \min(W_T - W^*, 0)\ . \end{aligned} \tag{4.3}$$

The admissible set for W^* in (4.3) is over the set of possible values for W_T. Using this equivalent definition of CVAR_α, as noted by Miller and Yang (2017) and Strub et al. (2017), the mean-CVAR problem can be expressed as (for a given scalarization parameter $\kappa > 0$)

$$\sup_{\mathcal{P}}\left\{\sup_{W^*} E_{\mathcal{P}}\left[W^* + \frac{1}{\alpha}(W_T - W^*)^- + \kappa W_T\right]\right\}\ . \tag{4.4}$$

In some cases, it is useful to interchange the $\sup\sup(\cdot)$ in (4.4), as suggested in Gao et al. (2017); Miller and Yang (2017). This allows us to rewrite the objective function (4.4) as

$$\sup_{W^*}\left\{\sup_{\mathcal{P}} E_{\mathcal{P}}\left[W^* + \frac{1}{\alpha}(W_T - W^*)^- + \kappa W_T\right]\right\}, \tag{4.5}$$

and to solve the inner optimization problem using an HJB equation (Dang and Forsyth, 2014; Forsyth and Labahn, 2019). Standard numerical methods can then be used to solve the outer optimization problem.

5. Time consistent mean CVAR. With these notational conventions, for a given scalarization parameter $\kappa > 0$ and an intervention time t_n, we define the scalarized time consistent

mean-CVAR problem $(TCMC_{t_n}(\kappa))$ in terms of the value function $J(s,b,t)$ as follows:

$$(TCMC_{t_n}(\kappa)): \quad J\left(s,b,t_n^-\right) = \max_{\mathcal{P}_n \in \mathcal{A}} \sup_{W^*} \left\{ E_{\mathcal{P}_n}^{X_n^+,t_n^+} \left[W^* + \frac{1}{\alpha}(W_T - W^*)^- + \kappa W_T \right. \right.$$

$$\left. \left. \Big| X(t_n^-) = (s,b) \right] \right\} \tag{5.1}$$

$$\text{s.t. } \mathcal{P}_n = \left\{p_n(\cdot), \mathcal{P}_{n+1}^*\right\} = \left\{p_n(\cdot), p_{n+1}^*(\cdot), \ldots, p_{M-1}^*(\cdot)\right\}, \tag{5.2}$$

$$\text{where } \mathcal{P}_{n+1}^* \text{ is optimal for problem } \left(TCMC_{t_{n+1}}(\kappa)\right)$$

$$\text{subject to} \quad \begin{cases} (S_t, B_t) \text{ follow processes (2.3)–(2.4)}, \quad t \notin \mathcal{T}, \\ W_n^+ = s + b + q_n\,, \;\; X_n^+ = (S_n^+, B_n^+), \\ S_n^+ = p_n(\cdot) W_n^+\,, \;\; B_n^+ = (1 - p_n(\cdot)) W_n^+\,, \\ p_n(\cdot) \in \mathcal{Z} = [0,1]. \end{cases} \tag{5.3}$$

Remark 5.1 (replacement of sup by max). Note that we have replaced $\sup_{\mathcal{P}}$ by $\max_{\mathcal{P}}$ in (5.1). Since the admissible value set $\mathcal{Z}$ is compact, this amounts to assuming continuity of the value function with respect to the controls. We can avoid this assumption by taking the max over the upper semicontinuous envelope of the value function, but this would add unpleasant heavy notation.

Remark 5.2 (time consistent constraint). Time consistency is enforced via the constraint (5.2). This approach, which explicitly enforces the time consistent constraint, is similar in spirit to the methods used in Wang and Forsyth (2011); Van Staden, Dang, and Forsyth (2018); Landriault et al. (2018) for the mean-variance case, and in the mean-CVAR case in Cui and Shi (2015).

Remark 5.3 (constraints). Note that we enforce no leverage, no shorting by requiring that $p_n \in \mathcal{Z} = [0,1]$ in (5.3). Due to the leverage constraint imposed in (5.3), this optimization problem is well-posed without adding an additional constraint on the terminal wealth (Gao et al., 2017).

Remark 5.4 (admissible set for W^*). For all the examples in this work, we have that the initial wealth is nonnegative, $q_i \geq 0$, and the constraints (5.3) are imposed. This implies that $W_T, W^* \in [0, \infty)$.

5.1. Expanded state space formulation. In order to develop a numerical algorithm for problem $(TCMC_{t_n}(\kappa))$, we follow the usual strategy of embedding the original problem in a higher dimensional space. We lift the state space to $\hat{X} = (s,b,W^*)$, and define an auxiliary function $V(s,b,W^*,t)$, which is given by

$$V\left(s,b,W^*,t_n^-\right) = E_{\mathcal{P}_n}^{\hat{X}_n^+,t_n^+} \left[W^* + \frac{1}{\alpha}(W_T - W^*)^- + \kappa W_T \Big| \hat{X}(t_n^-) = (s,b,W^*) \right] \tag{5.4}$$

$$\text{subject to} \quad \begin{cases} (S_t, B_t) \text{ follow processes (2.3)–(2.4)}, \quad t \notin \mathcal{T}, \\ W_n^+ = s + b + q_n\,, \;\; \hat{X}_n^+ = (S_n^+, B_n^+, W^*), \\ S_n^+ = p_n(\cdot) W_n^+\,, \;\; B_n^+ = (1 - p_n(\cdot)) W_n^+\,. \end{cases} \tag{5.5}$$

The optimal control $p_n(w)$ at time t_n is then determined from

$$p_n(w) = \arg\max_{p' \in \mathcal{Z}} \left\{ \sup_{W^*} V(w\ p', w\ (1-p'), W^*, t_n^+) \right\} . \tag{5.6}$$

The solution is advanced (backwards) across time t_n by

$$\begin{aligned} V(s,b,W^*,t_n^-) &= V\big(\ w^+\ p_n(w^+), w^+\ (1-p_n(w^+)\ ), W^*, t_n^+\ \big), \\ w^+ &= s+b+q_n\ . \end{aligned} \tag{5.7}$$

At $t=T$, we have

$$V(s,b,W^*,T^+) = W^* + \frac{1}{\alpha}\big((s+b)-W^*\big)^- + \kappa(s+b)\ . \tag{5.8}$$

For $t \in (t_{n-1}, t_n)$, there are no external cash flows, discounting, or controls applied, hence, the tower property gives for $h < (t_n - t_{n-1})$

$$V(s,b,W^*,t) = E\Big[V(S(t+h), B(t+h), W^*, t+h)\Big| S(t)=s, B(t)=b\Big],\ \ t \in (t_{n-1}, t_n - h)\ . \tag{5.9}$$

Applying Ito's lemma for jump processes (2.3)–(2.4), and letting $h \to 0$ gives the partial integro-differential equation (PIDE) for $V(s,b,W^*,t)$ for $t \in (t_{n-1}, t_n)$ (Tankov and Cont, 2009):

$$V_t + \frac{\sigma^2 s^2}{2} V_{ss} + (\mu - \lambda_\xi \kappa_\xi) s V_s - \lambda_\xi V + rbV_b + \lambda_\xi \int_{-\infty}^{+\infty} V(e^y s, b, t) f(y)\ dy\ = 0\ . \tag{5.10}$$

Remark 5.5 (form of PIDE). Between rebalancing dates, W^* can be regarded as a constant parameter, i.e., as in (5.9). Interdependence of the solution on W^* comes about indirectly due to the control equation (5.6).

Proposition 5.1 (equivalence of formulation (5.4)–(5.10)). *Define*

$$J\left(s,b,t_n^-\right) = \sup_{W^*} V(s,b,W^*,t_n^-)\ , \tag{5.11}$$

then formulation (5.5)–(5.10) *is equivalent to problem* ($TCMC_{t_n}(\kappa)$).

Proof. Replace $V(s,b,W^*,t_n^-)$ in (5.11) by the expressions in (5.4)–(5.10), and apply these expressions recursively backwards in time, using condition (5.8) at $t=T$. We then obtain (5.1)–(5.3). ∎

5.2. Computation of $E[W_T]$. Given an initial wealth $(s,b) = (0,W_0)$, and the optimal controls $\mathcal{P}^*$, then the above method can be used to determine

$$\begin{aligned} J_0 = J\left(0,W_0,t_0^-\right) = \sup_{W^*} \Big\{ & E_{\mathcal{P}^*}^{X_0^+, t_0^+} \left[W^* + \frac{1}{\alpha}(W_T - W^*)^- + \kappa W_T \right] \\ & \Big| X(t_0^-) = (0,W_0) \Big\}\ . \end{aligned} \tag{5.12}$$

We can determine CVAR_α from

$$\text{CVAR}_\alpha = J_0 - \kappa E^{X_0^+,t_0^+}_{\mathcal{P}^*}[W_T] \,, \tag{5.13}$$

which means we need to compute $E^{X_0^+,t_0^+}_{\mathcal{P}^*}[W_T]$ separately. This is easily done. Define the function

$$U(s,b,T^+) = (s+b), \tag{5.14}$$

where at $t_n \in \mathcal{T}$

$$\begin{aligned} U(s,b,t_n^-) &= U\big(\ w^+\ p_n(w^+), w^+\ (1-p_n(w^+)\ ), t_n^+\ \big), \\ w^+ &= s+b+q_n \end{aligned} \tag{5.15}$$

with $p_n(w)$ being the optimal control from (5.6). For $t \in (t_{n-1}, t_n)$, U satisfies the PIDE

$$U_t + \frac{\sigma^2 s^2}{2} U_{ss} + (\mu - \lambda_\xi \kappa_\xi) s U_s - \lambda_\xi U + rbU_b + \lambda_\xi \int_{-\infty}^{+\infty} U(e^y s, b, q, t) f(y)\, dy \ = 0 \,, \tag{5.16}$$

so that

$$E^{X_0^+,t_0^+}_{\mathcal{P}^*}[W_T] = U(0, W_0, t_0^-) \,. \tag{5.17}$$

6. Precommitment mean-CVAR. For a given scalarization parameter κ and intervention times t_n, the precommitment mean-CVAR problem $(PCMC_{t_0}(\lambda))$ is given in terms of the value function $\hat{J}(s,b,t_0^-)$:

$$\begin{aligned} (PCMC_{t_0}(\kappa)): \qquad \hat{J}\left(s,b,t_0^-\right) = \max_{\mathcal{P}_0 \in \mathcal{A}} \sup_{W^*} \Bigg\{ E^{X_0^+,t_0^+}_{\mathcal{P}_0} \Bigg[W^* + \frac{1}{\alpha}(W_T - W^*)^- + \kappa W_T \\ \Bigg| X(t_0^-) = (s,b) \Bigg] \Bigg\} \end{aligned} \tag{6.1}$$

$$\text{subject to} \quad \begin{cases} (S_t, B_t) \text{ follow processes (2.3)–(2.4)}, \quad t \notin \mathcal{T}, \\ W_\ell^+ = s+b+q_\ell \,, \ X_\ell^+ = (S_\ell^+, B_\ell^+) \,, \\ S_\ell^+ = p_\ell(\cdot) W_\ell^+ \,, \ B_\ell^+ = (1-p_\ell(\cdot)) W_\ell^+ \,, \\ p_\ell(\cdot) \in \mathcal{Z} = [0,1] \,, \ \ell = 0, \ldots, M-1 \,. \end{cases} \tag{6.2}$$

Remark 6.1 (relation to $(TCMC(\kappa))$). Note that compared to (5.1)–(5.3) we have dropped the time consistent constraint in (6.1)–(6.2), and specified that the optimal allocation is determined at time t_0.

Remark 6.2 (interchange $\max_{\mathcal{P}_n \in \mathcal{A}} \sup_{W^*}\{\cdot\}$). Observe that we can interchange the max sup in (6.1), as in (4.5), to obtain

$$\hat{J}\left(s,b,t_0^-\right) = \max_{\mathcal{P}_0 \in \mathcal{A}} \sup_{W^*} \left\{ E^{X_0^+,t_0^+}_{\mathcal{P}_0} \left[W^* + \frac{1}{\alpha}(W_T - W^*)^- + \kappa W_T \middle| X(t_0^-) = (s,b) \right] \right\} \tag{6.3}$$

$$= \sup_{W^*} \max_{\mathcal{P}_0 \in \mathcal{A}} \left\{ E^{X_0^+,t_0^+}_{\mathcal{P}_0} \left[W^* + \frac{1}{\alpha}(W_T - W^*)^- + \kappa W_T \middle| X(t_0^-) = (s,b) \right] \right\} . \tag{6.4}$$

6.1. Expanded state space formulation. In order to develop a numerical algorithm for problem $(PCMC_{t_n}(\kappa))$, we again embed the original problem in a higher dimensional space. This method is essentially the method in (Miller and Yang, 2017). We will also use the form of the value function as defined in (6.4). We lift the state space to $\hat{X} = (s,b,W^*)$, and define an auxiliary function $\hat{V}(s,b,W^*,t)$, which is given by

$$\hat{V}\left(s,b,W^*,t_n^-\right) = E_{\mathcal{P}_n}^{\hat{X}_n^+,t_n^+}\left[W^* + \frac{1}{\alpha}(W_T - W^*)^- + \kappa W_T \middle| \hat{X}(t_n^-) = (s,b,W^*)\right] \tag{6.5}$$

$$\text{subject to} \quad \begin{cases} (S_t, B_t) \text{ follow processes (2.3)–(2.4)}, & t \notin \mathcal{T}, \\ W_n^+ = s + b + q_n\,, \;\; \hat{X}_n^+ = (S_n^+, B_n^+, W^*), \\ S_n^+ = p_n(\cdot)W_n^+\,, \;\; B_n^+ = (1 - p_n(\cdot))W_n^+\,. \end{cases} \tag{6.6}$$

The optimal control $p_n(w,W^*)$ at time t_n is then determined from

$$p_n(w,W^*) = \arg\max_{p' \in \mathcal{Z}} \left\{ \hat{V}(w\ p', w\ (1-p'), W^*, t_i^+) \right\}\,. \tag{6.7}$$

The solution is advanced (backwards) across time t_n by

$$\begin{aligned} \hat{V}(s,b,W^*,t_n^-) &= \hat{V}\big(\ w^+\ p_n(w^+,W^*), w^+\ (1 - p_n(w^+,W^*)\ ), W^*, t_n^+\ \big), \\ w^+ &= s + b + q_n\ . \end{aligned} \tag{6.8}$$

Remark 6.3 (dependence on W^*). Since we do not impose a time consistent constraint in (6.1)–(6.2), this has the effect of decoupling the solution as a function of W^*. This can be seen by examining (6.7) in contrast to (5.6).

At $t = T$, we have

$$\hat{V}(s,b,W^*,T^+) = W^* + \frac{1}{\alpha}\big((s+b) - W^*\big)^- + \kappa(s+b)\ . \tag{6.9}$$

The usual argument gives the PIDE for $\hat{V}(s,b,W^*,t)$ for $t \in (t_{n-1}, t_n)$ as

$$\hat{V}_t + \frac{\sigma^2 s^2}{2}\hat{V}_{ss} + (\mu - \lambda_\xi \kappa_\xi)s\hat{V}_s - \lambda_\xi \hat{V} + rb\hat{V}_b + \lambda_\xi \int_{-\infty}^{+\infty} \hat{V}(e^y s, b, q, t) f(y)\ dy\ = 0\ . \tag{6.10}$$

Proposition 6.1 (equivalence of formulation (6.5–6.10)). *Define*

$$\hat{J}\left(s,b,t_0^-\right) = \sup_{W'} \hat{V}(s,b,W',t_0^-)\ , \tag{6.11}$$

then formulation (6.5)–(6.10)) is equivalent to problem $(PCMC_{t_0}(\kappa))$.

Proof. Replace $\hat{V}(s,b,W',t_n^-)$ in (6.11) by the expressions in (6.5)–(6.10), use (6.9), and recursively work backwards in time; then we obtain (6.1)–(6.2), by interchanging the sup max in the final step. ∎

Remark 6.4 (time inconsistency). Define

$$W^*(s,b,t_n) = \arg\max_{W'} \left(\hat{V}(s,b,W',t_n^-) \right)^{u^*} , \tag{6.12}$$

where $(\cdot)^{u^*}$ refers to the upper semicontinuous envelope of the argument, as a function of W' with fixed (s,b). Define

$$\hat{J}\left(s,b,t_n^-\right) = \sup_{W'} \hat{V}(s,b,W',t_n^-) \ , \tag{6.13}$$

which is the precommitment problem started at time t_n. Then (6.12) shows that $W^*(s,b,t_n)$ depends on t_n and the current state (s,b). This is the source of the time inconsistency.

Note that in Strub et al. (2017), the authors show that the precommitment policy is time consistent if we allow κ and α in (6.1) to depend on time in a specified manner, when reoptimizing at $t_n > 0$.

In the following, we will use the notation $W^*(t_0) \equiv W(0, W_0, t_0)$, where it is understood that $W^*(t_0)$ is evaluated at the initial wealth level W_0.

Proposition 6.2 (precommitment strategy equivalence to time consistent policy for an alternative objective function). *The precommitment mean-CVAR strategy $\mathcal{P}^*$ determined by solving $\hat{J}(0, W_0, t_0^-)$ (with $W^*(t_0)$ from (6.12)) is the time consistent strategy for the equivalent problem TCEQ (with fixed $W^*(t_0)$), with value function $\tilde{J}(s,b,t)$ defined by*

$$(TCEQ_{t_n}(\kappa\alpha)): \quad \tilde{J}\left(s,b,t_n^-\right) = \sup_{\mathcal{P}_n \in \mathcal{A}} \left\{ E_{\mathcal{P}_n}^{X_n^+, t_n^+} \left[(W_T - W^*(t_0))^- + (\kappa\alpha) W_T \middle| X(t_n^-) = (s,b) \right] \right\} . \tag{6.14}$$

Proof. This follows since we can regard $W^*(t_0)$ as a constant in objective function (6.5), and $\alpha > 0$, which is then equivalent to (6.14). With a fixed value of $W^*(t_0)$, the objective function (6.14) is a simple expectation, hence, we can determine $\mathcal{P}^*$ by dynamic programming, which is clearly time consistent. ∎

To be precise, we define an *implementable strategy* in terms of the controls relevant to this paper:

Definition 6.3 (implementable strategy). *A strategy is implementable if there is no incentive to deviate from the strategy computed at the initial time. More precisely, let $p^{t_\ell}(w,t_m)$ be the optimal control at time t_m, computed at time t_ℓ, $t_m \geq t_\ell$, $t_m, t_\ell \in \mathcal{T}$. An implementable strategy is such that*

$$p^{t_q}(w,t_n) = p^{t_\ell}(w,t_n) \quad \forall t_n, \ \forall w \ , \quad t_n \geq t_q \geq t_\ell \ . \tag{6.15}$$

Corollary 6.4 (alternative objective function TCEQ: An implementable strategy). *The following linear target shortfall strategy (TCEQ), based on precommitment mean CVAR, is implementable:*

- *At $t = t_0$, solve for the precommitment control from (6.1)–(6.2). As a by-product, we obtain $W^*(t_0)$.*
- *Using this fixed value of $W^*(t_0)$, we solve problem (6.14) for all $t > t_0$.*

Remark 6.5 (time consistent problem TCEQ). We can alternatively regard problem $TCEQ_{t_n}(\kappa\alpha)$ in (6.14) as the fundamental objective function. In order to determine an appropriate value for $W^*(t_0)$,we solve the precommitment problem at time zero. This generates $W^*(t_0)$ from (6.12). Since the precommitment solution is only used at time zero to determine $W^*(t_0)$, the time inconsistency of $PCMC$ is irrelevant for problem $TCEQ_{t_n}(\kappa\alpha)$.

Remark 6.6 (intuitive appeal of TCEQ). Problem ($TCEQ_{t_n}(\kappa\alpha)$) is a target based objective function, which should have great intuitive appeal to DC plan investors (Vigna, 2017). $W^*(t_0)$ can be interpreted as a disaster level of terminal real wealth.

We learn from Bjork and Murgoci (2010) that for any time inconsistent problem having a particular form, where we force time consistency, there is a different unconstrained utility function which produces a time consistent problem having the same controls. In other words, enforcing time consistency means that the investor is not preferences consistent to his original preferences (Vigna, 2017). It is not clear that the general results in Bjork and Murgoci (2010) can be directly applied to the mean-CVAR case. However, in our case, it is easy to see that the time inconsistent precommitment mean-CVAR problem has the same controls as the time consistent problem $TCEQ_{t_n}(\kappa\alpha)$.

To summarize, problem $TCEQ_{t_n}(\kappa\alpha)$ with fixed $W^*(t_0)$ is an implementable strategy, whose controls coincide with the precommitment mean-CVAR strategy at time zero. However, this strategy is not mean-CVAR efficient for later periods. We emphasize that requiring time consistency effectively changes the objective function, either from forcing the time consistent constraint, or from fixing $W^*(t_0)$. However, as we shall see from our numerical examples, problem $TCEQ_{t_n}(\kappa\alpha)$ generates a strategy which has some intuitively reasonable properties in terms of tail risk, and so is worthy of serious consideration.

6.2. Computation of $E[W_T]$. As for problem (TCMC), we need to determine $E[W_T]$ in order to recover CVAR from the value function. Define the function

$$\hat{U}(s,b,W^*,T^+) = (s+b), \tag{6.16}$$

where, at $t_n \in \mathcal{T}$,

$$\begin{aligned} \hat{U}(s,b,W^*,t_n^-) &= \hat{U}\big(\ w^+\ p_n(w^+,W^*), w^+\ (1-p_n(w^+,W^*)\ ), t_n^+\ \big), \\ w^+ &= s+b+q_n \end{aligned} \tag{6.17}$$

with $p_n(w,W^*)$ being the optimal control from (6.7). For $t \in (t_{n-1}^+, t_n^-)$, U satisfies the PIDE

$$\hat{U}_t + \frac{\sigma^2 s^2}{2}\hat{U}_{ss} + (\mu - \lambda_\xi \kappa_\xi)s\hat{U}_s - \lambda_\xi \hat{U} + rb\hat{U}_b + \lambda_\xi \int_{-\infty}^{+\infty} \hat{U}(e^y s, b, q, t) f(y)\, dy \ = 0\ , \tag{6.18}$$

so that

$$E_{\mathcal{P}^*}^{X_0^+, t_0^+}[W_T] = \hat{U}(0, W_0, W^*(t_0), t_0^-)\ . \tag{6.19}$$

7. Limit as rebalancing interval $\Delta t \to 0$. Recall that $\Delta t = t_{i+1} - t_i$, i.e., the rebalancing time interval. It is, of course, possible in our numerical algorithm to let $q_i = \hat{q}\Delta t$, where $\hat{q}$ is a rate of cash injection, and then take the limit as $\Delta t \to 0$, i.e., the time between rebalancing tends to zero. Intuitively, we expect that this would converge to the continuously rebalanced portfolio.

In the precommitment case, we can regard W^* as fixed in (6.5)–(6.10), hence, this PIDE falls under the scope of viscosity solution theory (Crandall, Ishii, and Lions, 1992). We can use the results in Barles and Souganidis (1991) to prove that our numerical algorithm converges to the viscosity solution of problem (6.5)–(6.10). The final precommitment solution is obtained by maximizing over W^* in (6.11) at t_0^-. We conjecture that by discretizing W^* with parameter h, using an exhaustive search to find the global maximum, taking the limit as $h \to 0$, and using the stability properties of viscosity solutions, we can prove that the entire numerical algorithm converges to the solution of problem (6.1) in the limit as the rebalancing interval $\Delta t \to 0$. However, this is a lengthy procedure, which is beyond the scope of this work.

Assume for the moment that we restrict attention to a situation where there are no jumps in process (2.3). Then, in the time consistent case, if we take the limit of continuous rebalancing, as noted in He and Jiang (2019), it is not obvious that in this limit we obtain the system of PDEs as described in Bjork and Murgoci (2010), which are based on an equilibrium game theoretic concept. In addition, this system of PDEs falls outside the scope of viscosity solution theory, hence, we cannot use standard techniques to prove convergence of the numerical algorithm as $\Delta t \to 0$.

Nevertheless, it is arguably more realistic to consider discrete rebalancing, i.e., Δt is finite and fixed. In this case, the PIDEs between rebalancing dates are linear, and standard methods can be used to prove convergence of the numerical PIDE solution technique. The entire problem is then a sequence of linear PIDE solutions. At each rebalancing time, new initial conditions are determined by solving an optimization problem. Hence, in the discrete rebalancing case, it is straightforward (although lengthy) to show convergence of the numerical methods for both precommitment and time consistent strategies.

Note as well that in the discrete rebalancing case, there are no difficulties in defining the appropriate equilibrium strategy (i.e., forcing the time consistent constraint) (He and Jiang, 2019). As a result, we focus on the discrete rebalancing case in this work.

8. Scaling property of the time consistent mean CVAR control. We consider the degenerate case where a lump sum investment is made, and no cash is injected at rebalancing times. In other words, at $t = 0$, $(s,b) = (0,W_0)$ with W_0 being the initial lump sum, and $q_i = 0\ \forall i$. For ease of analysis, we will use the formulation of problem $(TCMC_{t_n}(\kappa))$ as given in section 5.1. Before stating our main result, the following lemmas will be useful.

Lemma 8.1 (properties of solution of (5.10)). *If Property 2.1 holds, then, given a scalar $\lambda > 0$, if*

$$V(\lambda s, \lambda b, \lambda W^*, t_i^-) = \lambda V(s, b, W^*, t_i^-)\ , \tag{8.1}$$

then

$$V(\lambda s, \lambda b, \lambda W^*, t_{i-1}^+) = \lambda V(s, b, W^*, t_{i-1}^+)\ . \tag{8.2}$$

Proof. Changing variables in PIDE (5.10) to $x = \log s$, and noting that the transformed PIDE now has constant coefficients (from Property 2.1), then we can write the solution of PIDE (5.10) as

$$V(s, b, W^*, t_{i-1}^+) = \int_{-\infty}^{\infty} \mathcal{G}\left(\log s - x', \Delta t\right) V(e^{x'}, be^{r\Delta t}, W^*, t_i^-)\, dx', \tag{8.3}$$
$$\Delta t = t_i - t_{i-1}\ ,$$

where $\mathcal{G}(x - x', \Delta t)$ is the Green's function of the PIDE (5.10), excluding the rbV_b term (Garroni and Menaldi, 1992; Forsyth and Labahn, 2019), after transforming to $x = \log s$ coordinates. Let $x' = \log s'$ in (8.3) to give

$$V(s, b, W^*, t_{i-1}^+) = \int_0^{\infty} \mathcal{G}\left(\log(s/s'), \Delta t\right) V(s', be^{r\Delta t}, W^*, t_i^-)\, \frac{ds'}{s'}\ . \tag{8.4}$$

For $\lambda > 0$ we have

$$\begin{aligned} V(\lambda s, \lambda b, \lambda W^*, t_{i-1}^+) &= \int_0^{\infty} \mathcal{G}\left(\log(\lambda s/s'), \Delta t\right) V(s', \lambda be^{r\Delta t}, \lambda W^*, t_i^-)\, \frac{ds'}{s'} \\ &= \int_0^{\infty} \mathcal{G}\left(\log(s/s''), \Delta t\right) V(\lambda s'', \lambda be^{r\Delta t}, \lambda W^*, t_i^-) \frac{ds''}{s''} \qquad (s' = \lambda s'') \\ &= \lambda \int_0^{\infty} \mathcal{G}\left(\log(s/s'), \Delta t\right) V(s', be^{r\Delta t}, W^*, t_i^-)\, \frac{ds'}{s'} \qquad \text{(from (8.1))} \\ &= \lambda V(s, b, W^*, t_{i-1}^+)\ . \end{aligned} \tag{8.5}$$

■

Across rebalancing dates $t_n^+ \to t_n^-$, we have the following result.

Lemma 8.2. *Suppose the time consistent solution satisfies (5.7), where $p_n(w)$ is given by (5.6). If $q_n = 0$, and, for any $\lambda > 0$*

$$V(\lambda s, \lambda b, \lambda W^*, t_n^+) = \lambda V(s, b, W^*, t_n^+)\ , \tag{8.6}$$

then

$$p_n(w) = p_n(\lambda w)\ , \tag{8.7}$$
$$V(\lambda s, \lambda b, \lambda W^*, t_n^-) = \lambda V(s, b, W^*, t_n^-)\ . \tag{8.8}$$

Proof. From (5.6)

$$\begin{aligned} p_n(w) &= \arg\max_{p' \in \mathcal{Z}} \left\{ \sup_{W^*} V(w\ p', w\ (1 - p'), W^*, t_i^+) \right\} \\ &= \arg\max_{p' \in \mathcal{Z}} \left\{ \sup_{W^*} \lambda V(w\ p', w\ (1 - p'), W^*, t_i^+) \right\} \qquad (\lambda > 0) \\ &= \arg\max_{p' \in \mathcal{Z}} \left\{ \sup_{W^*} V(\lambda w\ p', \lambda w\ (1 - p'), \lambda W^*, t_i^+) \right\} \qquad \text{(from 8.6)} \\ &= \arg\max_{p' \in \mathcal{Z}} \left\{ \sup_{W'} V(\lambda w\ p', \lambda w\ (1 - p'), W', t_i^+) \right\} \qquad (W' = \lambda W^*\ ;\ W^* \in [0, \infty)) \\ &= p_n(\lambda w)\ . \end{aligned} \tag{8.9}$$

Next, we have, from (5.7), noting that $w = s + b$,

$$\begin{aligned} V(\lambda s, \lambda b, \lambda W^*, t_n^-) &= V\big(\ \lambda w\ p_n(\lambda w), \lambda w\ (1 - p_n(\lambda w)\ ), \lambda W^*, t_n^+\ \big) \\ &= V\big(\ \lambda w\ p_n(w), \lambda w\ (1 - p_n(w)\ ), \lambda W^*, t_n^+\ \big) \qquad \text{(from 8.9)} \\ &= \lambda V\big(\ w\ p_n(w), w\ (1 - p_n(w)\ ), W^*, t_n^+\ \big) \qquad \text{(from 8.6)} \\ &= \lambda V(s, b, W^*, t_n^-)\ . \end{aligned} \tag{8.10}$$

We now have our final result.

Theorem 8.3. *If a lump sum investment is made (i.e., $q_n = 0\ \forall n$), Property 2.1 holds, and the time consistent mean-CVAR strategy is determined by (5.5)–(5.11), then the optimal control for problem $(TCMC_{t_n}(\kappa))$ at each time t_n is independent of wealth w, that is,*

$$p_n(w) = p(t_n)\ ,\ n = 0, \ldots, M-1, \tag{8.11}$$

which also implies that $V(\lambda s, \lambda b, \lambda W^*, t) = \lambda V(s,b,W^*,t)$.

Proof. From (5.8) we have, for constant $\lambda > 0$,

$$V(\lambda\ s, \lambda\ b, \lambda\ W^*, T^+) = \lambda V(s, b, W^*, T^+)\ . \tag{8.12}$$

Apply Lemmas 8.2 and 8.1 recursively. Then we have that (8.9) holds for all n. Hence, for any $\lambda > 0$

$$p_n(\lambda w) = p_n(w) = p(t_n)\ , \tag{8.13}$$

and $V(\lambda s, \lambda b, \lambda W^*, t) = \lambda V(s,b,W^*,t)$. ∎

Remark 8.1 (significance of Theorem 8.3). In our numerical examples, we will consider only the practical case where the initial investment is zero, and the investor adds a fixed amount (real) to the portfolio at each rebalancing date, which is at odds with the assumptions of Theorem 8.3. However, suppose at time $t_n \in \mathcal{T}$,

$$w = (s + b) \gg \sum_{i=n}^{i=M-1} e^{-r(T-t_i)} q_i\ . \tag{8.14}$$

In other words, we examine points in the state space where the future discounted value of the cash injections is small compared to the current wealth. In this case, we can expect that $p_n(w)$ is only weakly dependent on wealth.

If we are, in fact, interested in a pure lump sum investment, then Theorem 8.3 can be used to reduce the dimensionality of the problem, with resulting computational efficiency.

Remark 8.2 (result in Cui et al. (2019)). For the lump sum investment case, Cui et al. (2019) show that the amount of wealth invested in the risky asset is a piecewise linear function of the current wealth. The problem in Cui et al. (2019) is posed in terms of specifying a minimum value of $E[W_T]$ which introduces a nonscaled variable into the problem. In contrast, we use the scalarization method in (5.1), which does not introduce any nonscaled variables. In any case, from the analysis above, we can see that a simple form for the control is unlikely to exist for the case of periodic contributions to the portfolio, which we verify in our numerical computations.

Remark 8.3 (generality of Theorem 8.3). Suppose we have a portfolio of M risky assets with s_i being the amount in asset i. Let $\mathbf{x}_i = \log s_i$. A multiasset analogue of Theorem 8.3 will hold if the multidimensional Green's function for the expected value PIDE has the property that

$$\mathcal{G}(\mathbf{x}, \mathbf{x}', \Delta t) = \mathcal{G}(\mathbf{x} - \mathbf{x}', \Delta t) \ . \tag{8.15}$$

A relevant example of where this would hold is a scenario where all assets follow correlated jump diffusion processes with constant coefficients. Note that it is commonplace for practitioners to model the returns of a constant maturity bond index directly, i.e., without postulating an interest rate process. In fact, it is usual to model the return of a constant maturity bond index by a geometric Brownian motion (Lin, MacMinn, and Tian, 2015). Hence the standard practitioner modeling approach, where the assets are stock and bond indices, falls under this case.

In this multiasset scenario, we would then conclude that in the case of a lump sum investment, the optimal policy for the time consistent mean-CVAR objective is deterministic. Due to this property, we conjecture that the behavior of the mean-CVAR policy under these assumptions will be similar to that reported in our numerical results section.

9. Numerical methods.

9.1. Time consistent mean CVAR. For problem ($TCMC_{t_n}(\kappa)$), our starting point is the formulation in section 5.1. We discretize the state space using n_x equally spaced nodes in the $x = \log s$ direction on a finite localized computational domain $[x_{\min}, x_{\max}]$. We discretize the domain $[0, b_{\max}]$ using an unequally spaced grid with n_b nodes and, similarly, we discretize the domain $[0, W^*_{\max}]$ using n_w nodes (unequally spaced). We use the Fourier methods discussed in Forsyth and Labahn (2019) to solve PIDE (5.10) between rebalancing times. To minimize localization effects and wraparound errors, we extend the computational domain for $x < x_{\min}$ and $x > x_{\max}$ and assume a constant value for the solution in the extended domain as described in Forsyth and Labahn (2019). This effectively adds artificial boundary conditions on the localized domain boundary. This localization error can be made small by selecting $|x_{\min}|, x_{\max}$ sufficiently large. In the b and W^* directions, we localize the problem (i.e., add artificial boundary conditions) by capping the solution values at the $(b_{\max}, W^*_{\max})$ values. The error in regions of interest can be made small by selecting sufficiently large values of $(b_{\max}, W^*_{\max})$. In contrast to the precommitment case, we can determine the optimal strategy by solving only a single, three dimensional PIDE. Of course this solution is not exact; errors can be made small by refining the grid.

At rebalancing times $t_n \in \mathcal{T}$, we discretize $p \in [0,1]$ using n_b equally spaced nodes, and then evaluate the right-hand side of (5.6) using linear interpolation. We then solve the optimization problem (5.6) using an exhaustive search over the discretized p values and the discretized W^* grid.

Once the optimal control is determined, we then use this control to determine the solution for $E[W_T]$ in section 5.2. Similarly, we use a Fourier method to advance the solution between rebalancing times.

9.2. Precommitment mean CVAR. For problem ($PCMC_{t_n}(\kappa)$), we start with the formulation in section 6.1. We will use the approach described in Miller and Yang (2017). This

method is based on (6.4). We solve the outer optimization problem, maximization with respect to W^*, by solving a sequence of inner optimization problems, which require optimizing with respect to the rebalancing controls $\mathcal{P}$.

For the inner optimization problems (i.e., we regard W^* as fixed) we proceed as follows. We discretize in the $x = \log s$ direction using n_x equally spaced nodes on the domain $[x_{\min}, x_{\max}]$, and n_b nodes in the b direction on the domain $[0, b_{\max}]$. We use the same Fourier methods as described in section 9.1 to advance the solution between rebalancing times. We use the same localization methods (i.e., add artificial boundary conditions) as used in the time-consistent case (see Forsyth and Labahn (2019) for details).

At rebalancing times $t_n \in \mathcal{T}$, we discretize $p \in [0,1]$ using n_b equally spaced nodes, and then evaluate the right-hand side of (6.17) using linear interpolation. We then solve the optimization problem (6.17) using an exhaustive search over the discretized p values (recall that W^* is fixed in this case).

The outer optimization problem in (6.4) can be written in terms of $\hat{V}(s,b,W^*,t)$ as

$$\hat{J}(0, W_0, t_0^-) = \sup_{W'} \hat{V}(0, W_0, W', t_0^-), \tag{9.1}$$

where each evaluation of $\hat{V}(\cdot)$ requires solution of problem (6.5).

We carry out the maximization in (9.1) by using a sequence of grids $n_x \times n_b$ to solve problem (6.5). On the coarsest grid, we discretize W^* and solve problem (6.5) for each discrete value of W^*. We then determine the $\max_{W^*}$ by exhaustive search. We use this value of W^* as a starting point for a one dimensional optimization algorithm on a sequence of finer grids. The solution on the coarse grid is inexpensive, and the fine grid optimization solutions do not require many iterations since we have a good starting estimate.

Note that there is no guarantee that we have found the global maximum since the problem is not guaranteed to be convex.[1] However, we have made a few tests by carrying a grid search on the finest grid, which suggests that we do indeed have the globally optimal solution.

10. Numerical example.

10.1. Market parameters.

The data and the method used to fit the parameters for process (2.3) are described in (Dang and Forsyth, 2016; Forsyth and Vetzal, 2017; Forsyth, Vetzal, and Westmacott, 2019). Briefly, we fit the parameters for process (2.3) from monthly market data for the sample period of 1926:1 to 2017:12. We use the Center for Research in Security Prices (CRSP) deciles (1–10) index. This is a total return value-weighted index of US stocks. We also use one month Treasury bill (T-bill) returns, over the period 1926:1 to 2017:12, for the risk-free asset. We adjust the returns for inflation by using the US CPI index, so all returns are real.

Table 1 provides the resulting annualized parameter estimates for the double exponential jump diffusion given in (2.1). The average real one-month T-bill rate for the period 1926:1-2017:12 was $r = .00464$.

[1] In Miller and Yang (2017), the precommitment mean-CVAR problem is posed in terms of $\log W_T$, which is then shown to result in a convex outer optimization over W^*. However, we have posed the problem in terms of W_T, which seems more natural, since *"you can only spend wealth dollars, not returns."*

Table 1

Estimated annualized parameters for the double exponential jump diffusion model given in (2.1) applied to the value-weighted CRSP deciles (1–10) index, deflated by the CPI. Sample period 1926:1 to 2017:12.

μ	σ	λ	p_{up}	η_1	η_2
.0884	.1451	.3370	.2581	4.681	5.600

Table 2

Investment scenario and model parameters.

Investment parameters	
Expiry time T	30 years
Initial wealth	0
Rebalancing frequency	yearly
Cash injection $\{q_i\}_{i=0,\ldots,29}$	20,000
Model parameters	
Real interest rate r	.00464
Equity process parameters	Table 1
Numerical parameters	
$x_{\max}$	$\log(10^5)+8$
$x_{\min}$	$\log(10^5)-8$
$b_{\max} = W^*_{\max}$	5×10^8

10.2. Investment scenario. Studies have shown that earnings for a typical employee increase rapidly until the age of 35, then increase slowly thereafter, until a few years before retirement, and then decreasing as fewer hours are worked in the transition to retirement (Cocco, Goems, and Maenhout, 2005; Ruppert and Zanella, 2015).

As a motivating example, we consider a 35 year old investor saving for retirement in a DC pension plan. We assume that the investor has a constant (real) salary of \$100,000 per year, and the total employee-employer contribution to a tax advantaged DC plan account is 20% of (real) salary per year. The investor plans to retire at age 65.

To be more precise, in our modeling context, we assume that the investor has zero initial wealth, and injects \$20,000 per year (real) into the portfolio at times $t = 0, 1, \ldots, 29$ years. The investment horizon is $T = 30$ years with annual rebalancing. Further details are given in Table 2.

11. Tests of convergence. In Table 3 we show the results for the scenario in Table 2, but using a default strategy of rebalancing to a constant weight of $p = 0.4$ in equities at each rebalancing date. In the following, we will use the results based on $N_{sim} = 2.56 \times 10^6$. From Table 3, the numerical results indicate that at least three digits are correct.

A typical glide path strategy for DC plan investors might have an allocation to equities of $p = 0.8$ at the initial time, declining to $p = 0.0$ at retirement, for a time averaged allocation to equities of about $p = 0.4$. The logic behind this is that an investor can take on more risk

Table 3

Convergence test, rebalance to constant weight $p = .4$ in stocks. Parameters in Table 2. N_{sim} is the number of Monte Carlo simulations. The numbers in brackets are the standard errors at the 99% confidence level. Units: thousands of dollars (real).

N_{sim}	$E[W_T]$	CVAR (5%)	$Median[W_T]$
1.6×10^5	1160 (2.7)	577	1083
6.4×10^5	1162 (1.3)	598	1084
2.56×10^6	1162 (.7)	598	1084

Table 4

Convergence test, precommitment mean CVAR. Parameters in Table 2. The Monte Carlo method used 2.56×10^6 simulations. The numbers in brackets are the standard errors at the 99% confidence level. $\kappa = 0.1, \alpha = .05$. Grid refers to the grid used to solve the HJB PDE: $n_x \times n_b$, where n_x is the number of nodes in the $\log S$ direction, and n_b is the number of nodes in the B direction. Units: thousands of dollars (real).

HJB equation				Monte Carlo		
Grid	$E[W_T]$	CVAR (5%)	W^*	$E[W_T]$	CVAR (5%)	$Median[W_T]$
512×333	2503	674.6	785.5	2485 (6.2)	679.2	1143
1024×665	2452	680.8	801.4	2447 (6.2)	682.0	1080
2048×1329	2434	682.3	806.8	2433 (6.2)	682.6	1067

when she is younger, and less risk when older. However, in Forsyth and Vetzal (2019), it is shown that the final wealth distribution of a glide path strategy is virtually indistinguishable from a constant weight strategy having the same time average equity allocation. Hence we use a default strategy of a constant weight $p = 0.4$ as representative of commonly suggested allocations.

Table 4 shows a convergence test for solution of the precommitment mean-CVAR problem $(PCMC_{t_n}(\kappa))$, (6.1)–(6.2). Examining the difference between the results on the two finest grids, suggests that the solution is accurate to within about one percent. The optimal controls are computed and stored, and then used as input to Monte Carlo simulations. The Monte Carlo results are also shown in Table 4. We choose the parameter κ in (6.1) so that $Median[W_T]$ is approximately the same as for the constant weight $p = 0.4$ strategy ($\kappa = 0.1$). In addition, we set $\alpha = .05$. The outer minimization iteration requires 15 HJB equation solves on the coarse grid, 12 HJB solves on the medium grid, and 11 HJB solves on the finest grid. The coarse grid estimates for W^* were used as starting values for the finer grid iterations. The total cumulative CPU time for all the coarse, medium, and fine grid solutions was 548 sec on a desktop.

Table 5 shows a convergence test for solution of the time consistent mean-CVAR problem $(TCMC_{t_n}(\kappa))$, (5.2)–(5.3). The optimal controls are computed and stored, and then used as input to Monte Carlo simulations. The Monte Carlo results are also shown in Table 5. We choose the parameter κ in (5.2) so that $Median[W_T]$ is approximately the same as for the

Table 5

Convergence test, time consistent mean CVAR. Parameters in Table 2. The Monte Carlo method used 2.56×10^6 *simulations.* $\kappa = 2.5, \alpha = .05$. *Grid refers to the grid used to solve the HJB PDE:* $n_x \times n_b \times n_w$, *where* n_x *is the number of nodes in the* $\log S$ *direction,* n_b *is the number of nodes in the* B *direction, and* n_w *is the number of nodes in the* W^* *direction. The numbers in brackets are the standard errors at the* 99% *confidence level. Units: thousands of dollars (real).*

HJB equation			Monte Carlo		
Grid	$E[W_T]$	CVAR (5%)	$E[W_T]$	CVAR (5%)	$Median[W_T]$
$256 \times 309 \times 309$	1746	495.7	1728 (1.9)	502.6	1381
$512 \times 617 \times 617$	1455	519.6	1450 (1.9)	520.4	1168
$1024 \times 1233 \times 1233$	1343	529.8	1341 (1.9)	530.0	1079

constant weight $p = 0.4$ strategy ($\kappa = 2.5$). Again, we set $\alpha = .05$. Note that in contrast to the precommitment results, convergence is considerably slower here, as the mesh is refined. In this case, since W^* is an independent variable in the grid, only one HJB equation solve is required at each grid refinement. The CPU time on the finest grid was 7152 sec on a desktop. Note that the finest grid has $\simeq 10^9$ nodes, so storage becomes an issue for finer grids.

We verified that the artificial boundary conditions did not affect the computed solution in any meaningful way by carrying out the following tests. For both precommitment and time consistent cases, we carried out tests replacing $[x_{\min}, x_{\max}]$ by $[x_{\min} - 2, x_{\max} + 2]$, b_{max} by $b_{\max} \times 10$, and for the time consistent case, we replaced $W^*_{\max}$ by $W^*_{\max} \times 10$. In all cases, this resulted in changes to the solution in at most the fifth digit.

Remark 11.1 (slower convergence: time consistent case). In the precommitment case, we can rewrite (6.7)–(6.8) as

$$
\begin{aligned}
V(s,b,W^*,t_n^-) &= \max_{p' \in \mathcal{Z}} V\left(w^+ \; p', w^+ \; (1-p') \, , W^*, t_n^+ \right) \, , \\
w^+ &= s + b + q_n \, ,
\end{aligned}
\tag{11.1}
$$

which means that the value function $V(s,b,W^*,t_n^-)$ is maximized at each point in (s,b,W^*) space. Intuitively, this means that even if the control is comparatively inaccurate, the value function solution is still reasonably accurate, since the control is an extreme point of the right-hand side of (11.1). In contrast, for the time consistent case, we can see from (5.6)–(5.7) that, in general, the value function $V(s,b,W^*,t_n^-)$ does not maximize the right-hand side of (5.7) at all points in (s,b,W^*) space. Hence the extreme point property is lost, and we can expect slower convergence as the mesh is refined.

Remark 11.2 (effect of errors in the time consistent case). Table 5 indicates that the difference between the two finest grids suggests errors on the order of 9 per cent. However, storage limitations precluded using a finer grid. Assuming first order convergence, then the extrapolated exact solution from Table 5 would give $(CVAR,\ Median) = (540, 990)$. This suggests that the exact solution for, say, Median = 1067 will have $CVAR,\ Median \leq 540$, which

is significantly worse than $(CVAR,\ Median) = (683, 1067)$ for the precommitment case (i.e., the Median-CVAR frontier is shifted significantly to the left for the time consistent strategy compared to the precommitment policy). Note as well that $(CVAR,\ Median) = (598, 1084)$ for the constant weight $p = 0.4$ strategy. In other words, the numerical solution errors in the summary statistics for the time consistent case are significantly less than the differences between the summary statistics for the time consistent policy and the other two strategies.

12. Numerical results.

12.1. Basis of comparison. As a reminder to the reader, we reiterate the following facts. Time consistent mean-CVAR strategies are constrained mean-CVAR efficient at all times, but not efficient in terms of objective function $TCEQ$ (6.14) at any time. Nor is the time consistent mean-CVAR policy efficient in the unconstrained mean-CVAR sense at the initial time. Conversely, the precommitment mean-CVAR strategy is efficient in terms of objective (6.14) at all times, unconstrained mean-CVAR efficient only at the initial time, and is never constrained mean-CVAR efficient.

It is then difficult to identify a comparison criterion based on efficiency. Both the precommitted policy and the time consistent policy are efficient in terms of some criteria and not in others. Both strategies are time consistent in terms of their specific objective functions.

It is common in the actuarial literature to regard the application of different objective functions simply as a means of generating different investment strategies. This leads to the idea of comparing different investment objective functions in terms of the statistical evolution of portfolio wealth (Blake, Wright, and Zhang, 2013; Vigna, 2014; Donnelly et al., 2015; Graf, 2017). This is quite natural from an actuarial viewpoint, since a major concern is the probabilistic time evolution of assets and liabilities. Our numerical examples are motivated by long term strategies for holders of DC pension plans. It is therefore relevant to consider an actuarial viewpoint when comparing strategies.

To this end, we will focus on the statistical properties of the evolution of the investor's wealth, under both strategies. We are making the assumption that the investor is essentially agnostic as to the philosophy behind the objective function, and is only concerned with the probabilistic evolution of the total portfolio wealth. Of course, alternative views are possible, but we consider this approach to have intuitive appeal.

12.2. Results. In the following, for the time consistent and precommitment policies, we compute the controls using the finest grids in Tables 4 and 5. All Monte Carlo computations used 2.56×10^6 simulations.

Figures 1 and 2 compare the precommitment and time consistent mean-CVAR strategies in terms of (a) percentiles of accumulated wealth, (b) percentiles fraction in equities, and (c) control heat maps. The heat maps show the optimal control as a function of realized wealth and time. The optimal fraction in the risky asset can be determined by comparing the color code at point (W_t, t) on the map with the legend on the right-hand side.

Contrast Figure 1(c) with Figure 2(c). We can see that the time consistent heat map is somewhat poorly defined (contours not sharply delineated), which is expected from Remark 11.1. In particular, from Remark 8.1, we expect that for large values of W, then the time consistent heat map contours should become straight vertical lines (i.e., control independent

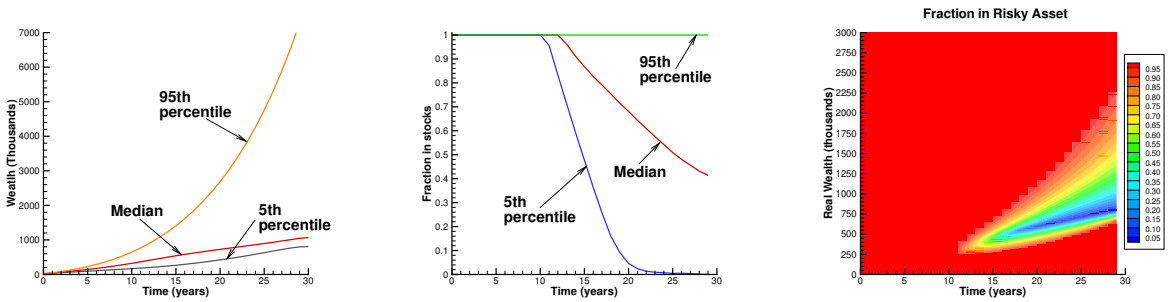

(a) Percentiles of accumulated wealth.

(b) Percentiles of optimal fraction in equities.

(c) Optimal control heat map.

Figure 1. *Precommitment mean-CVAR, parameters in Table 2. Optimal control computed solving problem (6.1)–(6.2). Percentiles are of real wealth and the optimal fraction invested in equities. Statistics based on* 2.56×10^6 *Monte Carlo simulation runs.*

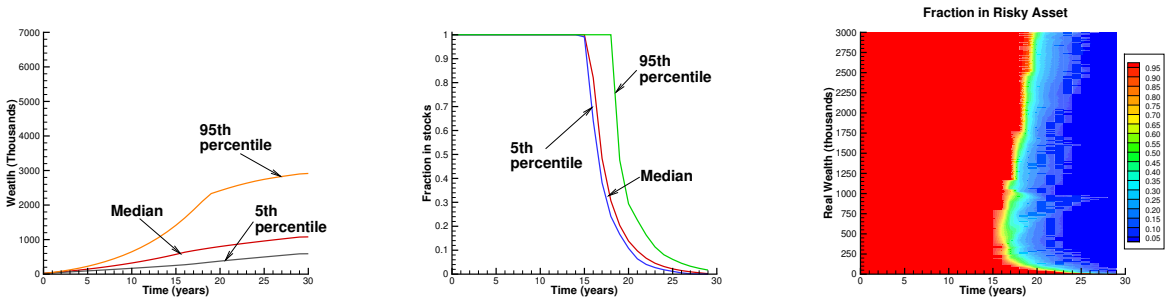

(a) Percentiles of accumulated wealth.

(b) Percentiles of optimal fraction in equities.

(c) Optimal control heat map.

Figure 2. *Precommitment mean-CVAR, parameters in Table 2. Optimal control computed solving problem (5.1)–(5.3). Percentiles are of real wealth and the optimal fraction invested in equities. Statistics based on* 2.56×10^6 *Monte Carlo simulation runs.*

of W). Due to storage limitations, it was not possible to compute the time consistent solution with a finer grid (the finest grid in Table 5 had $\simeq 10^9$ nodes).

As a point of comparison, we show the percentiles of accumulated wealth for the constant weight $p = 0.4$ case in Figure 3.

A more precise comparison of all three strategies is shown in Figure 4, which shows the cumulative distribution functions for the cumulative wealth W_T. By design, all three strategies have approximately the same $Median[W_T]$, which can be verified by noting that all curves intersect at $Prob(W_T < W) = 0.5$. The investor has contributed a total of \$600,000 (real) over the thirty years. Therefore, any values of $W < 600{,}000$ should be regarded as a poor result.

The time consistent strategy is dominated by the constant weight strategy for W below the median, which is a poor result given that (TCMC) is attempting to minimize tail risk. The precommitment strategy has a sharp decrease in the cumulative distribution function for

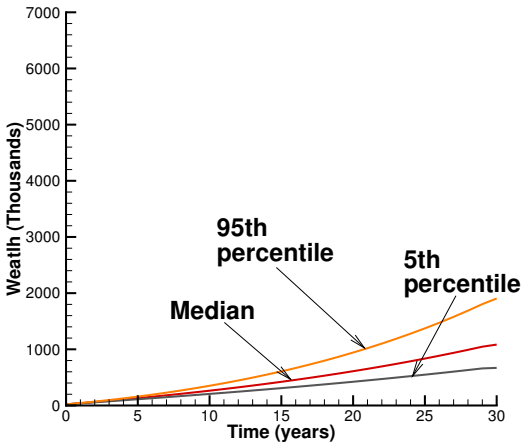

Figure 3. *Constant weight $p = 0.4$ in equities, parameters in Table 2. Percentiles are of wealth.*

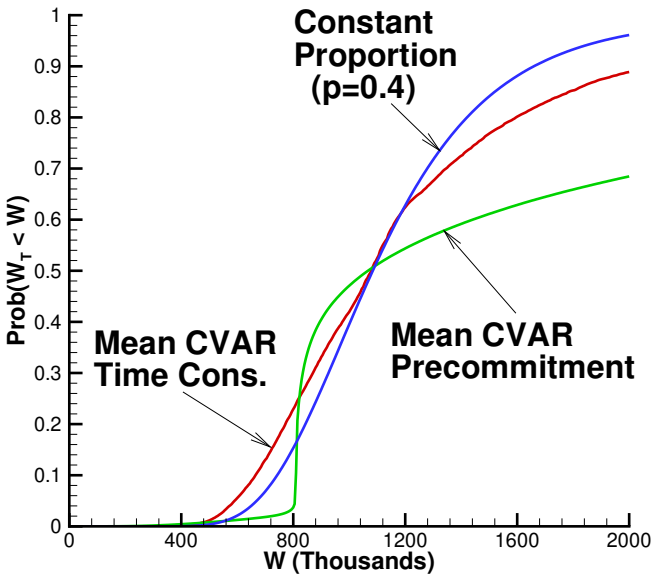

Figure 4. *Comparison of cumulative distribution functions for W_T.*

$W < 800$, which is reflected in the fact that the CVAR for the precommitment strategy is the largest of all three strategies (recall that with definition (4.2) for CVAR, a larger CVAR value has less risk). In terms of reduction of tail risk, as observed at $t = 0$, the precommitment mean-CVAR strategy is clearly superior to the other strategies. Rather surprisingly, the time consistent mean-CVAR strategy appears to be the least effective of all the three strategies (smallest CVAR of all strategies). The precommitment strategy dominates the other strategies for wealth levels $W < 800{,}000$ and $W > 1{,}100{,}000$.

12.3. Summary of numerical results. Perhaps the key result for explaining the poor performance of time consistent mean-CVAR strategies is Theorem 8.3 and Remark 8.1. From Remark 8.1, and Figures 2(b)–2(c), we can see that the time consistent control is only weakly dependent on the current wealth (over a wide range of wealth values). This is simply a consequence of forcing time consistency. Essentially, all strategies computed at $t > 0$, compute the tail risk always with respect to current wealth, which may be very large or very small compared with wealth at previous times.

In contrast, the induced time consistent strategy $TCEQ_{t_n}(\kappa\alpha)$ is always based on $W^*(t_0)$. In any realistic situation, $W^*(t_0)$ is based on the initial wealth targets, hence, this strategy responds to how well the current wealth is on track to meet the target.

We would expect that the time consistent strategy behaves similarly to a deterministic strategy (i.e., nonresponsive to realized wealth). This is consistent with the results in Figure 4. The poor performance of deterministic strategies was also observed in Forsyth and Vetzal (2019), where, under mean-variance criteria, it is shown that deterministic strategies offer virtually no improvement over constant weight strategies (for fixed $E[W_T]$).

From a practical point of view, it would appear that applying a time consistent constraint to multi-period mean-CVAR policies results in a strategy which is counterintuitive. At time zero, we have some idea of what we desire as a minimum final wealth. Fixing this shortfall target for all $t > 0$ makes intuitive sense. If we have a billion dollars, we are probably only concerned with the probability of a final wealth being less than, say, 50 million dollars. On the other hand, if we have one million dollars, then we are probably concerned about ending up with less than 500,000. In other words, our intuitive shortfall target is *not* a constant proportion of current wealth. This intuition contrasts with time consistent strategies, which adjust the effective shortfall target in response to current wealth. This does not appear to generate a reasonable strategy. On the other hand, precommitment mean-CVAR strategies at time zero are equivalent to a time consistent strategy which fixes the shortfall target at the initial time. This equivalent objective function (linear target shortfall) appears to produce an intuitively reasonable strategy.

12.4. Extension to multiple risky assets. As usual, numerical solution of an HJB PIDE is restricted to problems with three or fewer dimensions, due to computational complexity considerations. This means that it is practical to determine the optimal strategy for two risky assets and one risk-free asset in the precommitment case. On the other hand, we can solve the time consistent mean-CVAR problem for only a single risky asset and one risk-free asset. However, it is possible to use a machine learning approach to determine optimal strategies for portfolios with large numbers of risky assets (Li and Forsyth, 2019).

This assumes that we rebalance at discrete times. However, if we take the limit and assume continuous rebalancing, the PIDE problem for the precommitment case has dimension one, hence, is computationally feasible. We conjecture that a similar reduction to a system of one dimensional PIDEs is also possible for the time consistent case as well. Continuous rebalancing, of course, is not realistic in practice.

13. Conclusions. Precommitment strategies have been widely criticized for being "non-implementable." However, in the mean-CVAR case, we know from Corollary 6.4, that since precommitment mean-CVAR is equivalent to a linear target shortfall strategy with a fixed

shortfall target W^*, then this strategy is time consistent in terms of this alternative objective function, hence, is implementable.

Since all strategies in this study can be considered to be time consistent (when viewed in terms of the appropriate objective function), we can then rank the various strategies in terms of the cumulative distribution functions of the terminal wealth. Forcing a time consistent constraint on a mean-CVAR strategy produces a final distribution function which has a larger left tail risk than a simple constant weight strategy (with the same median value of final wealth).

Note that forcing a time consistent constraint for the mean-variance problem has been shown to have undesirable consequences in some cases (Bensoussan, Wong, and Yam, 2019). Hence, it is perhaps not surprising to see a similar effect for the mean-CVAR objective function.

Time consistent mean-CVAR policies behave in a manner very similar to deterministic (i.e., only a function of time) strategies. This offers little (if any) improvement compared to the standard constant weight strategy.

In contrast, the precommitment mean-CVAR strategy (which coincides with the time consistent linear target shortfall strategy at time zero) does minimize the left tail risk, compared to the other strategies.

Finally, in agreement with (Vigna, 2017; Bensoussan, Wong, and Yam, 2019), our results indicate that simply forcing a time consistent constraint onto a precommitment policy, without considering the economic ramifications, may lead to strategies with undesirable characteristics.

Declarations of interest. The author has no conflicts of interest to declare.

REFERENCES

G. Barles and P. Souganidis (1991), *Convergence of approximation schemes for fully nonlinear equations*, Asymptot. Anal., 4, pp. 271–283.

S. Basak and G. Chabakauri (2010), *Dynamic mean-variance asset allocation*, Rev. Financ. Stud., 23, pp. 2970–3016.

A. Bensoussan, K. C. Wong, and S. C. P. Yam (2019), *A paradox in time-consistency in the mean-variance problem?*, Finance Stoch., 23, pp. 173–207.

T. Bjork and A. Murgoci (2010), *A General Theory of Markovian Time Inconsistent Stochastic Control Problems*, preprint, SSRN, 1694759.

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

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

D. Blake, D. Wright, and Y. Zhang (2013), *Target-driven investing: Optimal investment strategies in defined contribution plans under loss aversion*, J. Econom. Dynam. Control, 37, pp. 195–209.

J. F. Cocco, F. J. Goems, and P. J. Maenhout (2005), *Consumption and portfolio choice over the life cycle*, Rev. Financ. Stud., 18, pp. 491–533.

F. Cong and C. Oosterlee (2016a), *Multi-period mean-variance porfolio optimization based on Monte Carlo simulation*, J. Econom. Dynam. Control, 64, pp. 23–38.

F. Cong and C. Oosterlee (2016b), *On the pre-commitment aspects of a time-consistent strategy for a mean-variance investor*, J. Econom. Dynam. Control, 70, pp. 178–193.

M. Crandall, H. Ishii, and P. Lions (1992), *User's guide to viscosity solutions of second order partial differential equations*, Bull. Amer. Math. Soc. (N.S.), 27, pp. 1–67.

X. Cui, J. Gao, Y. Shi, and S. Zhu (2019), *Time-consistent and self-coordination strategies for multi-period mean-conditional value-at-risk portfolio selection*, European J. Oper. Res., 276, pp. 781–789.

X. Cui, D. Li, S. Wang, and S. Zhu (2012), *Better than dynamic mean-variance: Time inconsistency and free cash flow stream*, Math. Finance, 22, pp. 346–378.
X. Cui and Y. Shi (2015), *Multiperiod mean-CVaR portfolio selection*, in Modelling, Computation and Optimization in Information Systems and Management Sciences, H. A. Le Thi, T. Pham Dinh, and N. T. Nguyen, eds., Springer, Cham, Switzerland, pp. 293–304.
D.-M. Dang and P. A. Forsyth (2014), *Continuous time mean-variance optimal portfolio allocation under jump diffusion: A numerical impulse control approach*, Numer. Methods Partial Differential Equations, 30, pp. 664–698.
D.-M. Dang and P. A. Forsyth (2016), *Better than pre-commitment mean-variance portfolio allocation strategies: A semi-self-financing Hamilton-Jacobi-Bellman equation approach*, European J. Oper. Res., 250, pp. 827–841.
C. Donnelly, R. Gerrard, M. Guillen, and J. P. Nielsen (2015), *Less is more: Increasing retirement gains by using an upside terminal wealth constraint*, Insurance Math. Econom., 64, pp. 259–267.
P. Forsyth and G. Labahn (2019), *$\epsilon-$monotone Fourier methods for optimal stochastic control in finance*, J. Comput. Finance, 22, pp. 25–71.
P. Forsyth and K. Vetzal (2019), *Optimal asset allocation for retirement savings: Deterministic vs. time consistent adaptive strategies*, Appl. Math. Finance, 26, pp. 1–37.
P. Forsyth, K. Vetzal, and G. Westmacott (2019), *Management of portfolio depletion risk through optimal life cycle asset allocation*, N. Amer. Actuar. J., 23, pp. 447–468.
P. A. Forsyth and K. R. Vetzal (2017), *Dynamic mean variance asset allocation: Tests for robustness*, Int. J. Financ. Eng., 4, 1750021, https://doi.org/10.1142/S2424786317500219.
J. Gao, K. Zhou, D. Li, and X. Cao (2017), *Dynamic Mean-LPM and Mean-CVaR portfolio optimization in continuous-time*, SIAM J. Control Optim., 55, pp. 1377–1397.
M. G. Garroni and J. L. Menaldi (1992), *Green Functions for Second Order Parabolic Integro-Differential Problems*, Longman Scientific, New York.
S. Graf (2017), *Life-cycle funds: Much ado about nothing?*, Eur. J. Finance, 23, pp. 974–998.
X. D. He and Z. Jiang (2019), *On the Equilibrium Strategies for Time Inconsistent Problems in Continuous Time*, preprint, SSRN, 3308274.
S. G. Kou (2002), *A jump-diffusion model for option pricing*, Manag. Sci., 48, pp. 1086–1101.
S. G. Kou and H. Wang (2004), *Option pricing under a double exponential jump diffusion model*, Manag. Sci., 50, pp. 1178–1192.
D. Landriault, B. Li, D. Li, and V. R. Young (2018), *Equilibrium strategies for the mean-variance investment problem over a random horizon*, SIAM J. Financial Math., 9, pp. 1046–1073.
D. Li and W.-L. Ng (2000), *Optimal dynamic portfolio selection: Multiperiod mean-variance formulation*, Math. Finance, 10, pp. 387–406.
Y. Li and P. A. Forsyth (2019), *A data driven neural network approach to optimal asset allocation for target based defined contribution pension plans*, Insurance: Mathematics and Economics, 86, pp. 189–204.
Y. Lin, R. MacMinn, and R. Tian (2015), *De-risking defined benefit plans*, Insurance Math. Econom., 63, pp. 52–65.
K. Ma and P. A. Forsyth (2016), *Numerical solution of the Hamilton-Jacobi-Bellman formulation for continuous time mean variance asset allocation under stochastic volatility*, J. Comput. Finance, 20, pp. 1–37.
F. Menoncin and E. Vigna (2017), *Mean-variance target based optimisation for defined contribution pension schemes in a stochastic framework*, Insurance Math. Econom., 76, pp. 172–184.
C. W. Miller and I. Yang (2017), *Optimal control of conditional value-at-risk in continuous time*, SIAM J. Control Optim., 55, pp. 856–884.
R. T. Rockafellar and S. Uryasev (2000), *Optimization of conditional value-at-risk*, J. Risk, 2, pp. 21–42.
P. Ruppert and G. Zanella (2015), *Revisiting wage, earnings, and hours profiles*, J. Monetary Econom., 72, pp. 114–130.
M. Strub, D. Li, and X. Cui (2019), *An Enhanced Mean-Variance Framework for Robo-Advising Applications*, preprint, SSRN, 3302111.
M. Strub, D. Li, X. Cui, and J. Gao (2017), *Discrete-Time Mean-CVaR Portfolio Selection and Time-Consistency Induced Term Structure of the CVaR*, preprint, SSRN, 3040517.

P. Tankov and R. Cont (2009), *Financial Modelling with Jump Processes*, Chapman and Hall/CRC, New York.

P. M. Van Staden, D.-M. Dang, and P. A. Forsyth (2018), *Time-consistent mean-variance portfolio optimization: A numerical impulse control approach*, Insurance Math. Econom., 83, pp. 9–28.

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

E. Vigna (2017), *Tail Optimality and Preferences Consistency for Intertemporal Optimization Problems*, Working paper 502, Collegio Carlo Alberto, Università Degli Studi di Torino, Turin, Italy.

J. Wang and P. A. Forsyth (2011), *Continuous time mean variance asset allocation: A time-consistent strategy*, European J. Oper. Res. 209, pp. 184–201.

X. Y. Zhou and D. Li (2000), *Continuous-time mean-variance portfolio selection: A stochastic LQ framework*, Appl. Math. Optim. 42, pp. 19–33.