

SEF Working paper: 02/2011  
January 2011

---

**On loss-avoiding lump-sum pension  
optimization with contingent targets**

Jeffrey Azzato, Jacek Krawczyk and  
Christopher Sissons

The Working Paper series is published by the School of Economics and Finance to provide staff and research students the opportunity to expose their research to a wider audience. The opinions and views expressed in these papers are not necessarily reflective of views held by the school. Comments and feedback from readers would be welcomed by the author(s).

Further copies may be obtained from:

The Administrator  
School of Economics and Finance  
Victoria University of Wellington  
P O Box 600  
Wellington 6140  
New Zealand

Phone: +64 4 463 5353

Email: [alice.fong@vuw.ac.nz](mailto:alice.fong@vuw.ac.nz)

**Working Paper 01/2011**

**ISSN 2230-259X (Print)**

**ISSN 2230-2603 (Online)**

# On loss-avoiding lump-sum pension optimisation with contingent targets

Jeffrey Azzato, Jacek B. Krawczyk\* and Christopher Sissons  
*Victoria University of Wellington, Wellington,  
New Zealand*

---

## Abstract

Consider a lump-sum pension fund problem, in which an agent deposits an amount with a fund manager up front and is later repaid a lump sum  $x(T)$  after time  $T$ . The fund manager may be both cautious in seeking a payoff  $x(T)$  meeting a certain target, but relaxed toward the possibility of exceeding this target. We use a computational method in stochastic optimal control (“SOCsol”) to find approximately-optimal decision rules for such “cautious-relaxed” fund managers. In particular, we examine fund optimisation problems in which the target is contingent upon market conditions such as inflation.

---

**Keywords:** Computational economics; pension funds; cautious-relaxed policies; approximating Markov decision chains; SOCsol.

JEL: C63, D92, G11

MSC: 93E20, 93E25, 90C39, 90C40, 90A09

---

## 1 Introduction

This paper<sup>1</sup> is about lump-sum pension fund problems, which arise when an agent pays an amount  $x_0$  to a fund manager, in exchange for an uncertain lump sum  $x(T)$  — the “pension” — at time  $T$ .

The paper’s aim is twofold. First, we discuss what such a fund manager’s objectives might be. This is because we believe that there is no criterion that managers would *always* want to optimise in a volatile market situation. In particular, we deem maximisation of one of the so-called “risk-averse” utility functions to be an unlikely objective,<sup>2</sup> despite its popularity in the literature. We show that

---

\*Corresponding author. Email: J.Krawczyk@vuw.ac.nz

<sup>1</sup>This paper draws from the conference paper presented at the 2009 *Quantitative Methods in Finance Conference (QMF)*, Sydney, NSW.

<sup>2</sup>These include the HARA (Hyperbolic Absolute Risk-Aversion) utility functions.

strategies maximising the expected present value of such a function can generate very “risky” right-skewed payoff distributions, in which “low” returns are relatively common. Accordingly, we contrast the shape of these distributions with those of the payoff spreadings obtainable by optimising one of three related loss-aversion (or avoidance) objectives. The payoff distribution of such a strategy can be left-skewed, delivering “high” returns more frequently than “low” ones.

Second, we use this loss-avoidance framework to model the pension fund problem with a target contingent on inflation or, more generally, on any observable exogenous stochastic process. We solve this problem numerically and comment on the applicability of the solution.

The problem of formulating an acceptable portfolio management strategy is well recognised in the *static* context. Markowitz (see [Mar59]) is credited with pioneering the classical mean-variance portfolio selection problem, whose solution balances a good average return against the likelihood of achieving it. Since his seminal work [Mar52], many authors have worked on extensions and alternatives that would allow for hedging against uncertainties and/or ensure an acceptable level of return (see, for example, [BLS04, dAJ04, Kra90, RU00]).

On the other hand, the main stream of research in *dynamic* portfolio management has followed the seminal works of Samuelson (see [Sam69]) and Merton (see [Mer69, Mer71]), concentrating on solutions to HARA problems. Such solutions typically provide an optimal strategy (either numerical or in closed form) that maximises expected present utility. The solutions are generically risk-sensitive (see, for example, [BSL97, FS00, MSS89]), but ignore the payoff and utility distribution consideration mentioned above.

Recently, problems on hedging and/or ensuring an acceptable return have also been investigated in the dynamic context. For example, [HRS96] discusses multi-period minimax hedging strategies, while [FR99] applies a risk-minimising hedging approach to dynamic strategy determination and [GRS04] uses mean-variance analysis for multistage portfolio management.<sup>3</sup> Most importantly for likelihood of payoff achievement, Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR) have been successfully used as optimisation constraints to help produce “certain payoff” management strategies in the dynamic context (see [BRU01, Yiu04]). In [BFLS04], optimal strategies maximising the probability of an agent’s wealth exceeding a target were computed, while [BF06] used a target-percentile risk measure to generate time-consistent optimal strategies.

Unfortunately, it is difficult to extend the Markowitz mean-variance approach to dynamic settings (whether multi-period or continuous time). This is because the payoff variance involves the inseparable (in the dynamic programming sense) term  $\mathbb{E}[(x(T))^2]$ , which is hard to analyse (see [ZL00, LZL02]).<sup>4</sup>

However, ensuring high mean returns can result in high return volatility (see [GRS04]), suggesting that less should be invested in risky assets in order to satisfy the VaR constraint (see [Yiu04]). Consequently, achieving returns that are both high *and* secure remains difficult. It seems (see [Kra08]) that maximising a

<sup>3</sup>The importance of dynamic one-side measures for return performance is discussed in [Maz04].

<sup>4</sup>The last two citations have overcome this difficulty, albeit for linear-quadratic models only.

“classical” risk-averse utility function subject to risk constraints reduces exposure to risk but still results in right-skewed payoff distributions.

This paper uses a “new” portfolio performance measure that is significantly asymmetric with respect to risk.<sup>5</sup> Similar measures are proposed in prospect theory (see [TK92, BKP04, JZ08] and the references therein). These measures each set a fund return target, penalise failure to achieve this target and reward over-achievement.

The differences between the new measures (explained in detail later) lie in the shapes of their penalising and rewarding components. Our measure strongly penalises target shortfall, while rewarding over-achievement mildly. This makes it *locally concave*, i.e. it is concave for wealth both below and above the target, reflecting overall *risk aversion*. The prospect-theoretic performance measures are S-shaped. I.e., they are convex for wealth below the target and concave for wealth above it, respectively reflecting *risk seeking* behaviour toward losses and *risk averse* behaviour toward gains. So while all three measures are *loss averse* (as they are each steeper below the target than above it), their attitudes to risk differ.

We show that optimisation of these target-based measures delivers strategies that can generate left-skewed payoff distributions having small VaR and CVaR. As these small values arise without imposing constraints (which may be necessary when using classical performance measures), we believe that these “new” performance measures may be good objective candidates for fund managers.

Given the complexity of portfolio problems with the non-concave utility measures, the only feasible solution method known to us was to use numerical optimisation. The solutions in this paper rely on a discretisation scheme inspired by Kushner (see [Kus90]),<sup>6</sup> which is implemented here in a suite of MATLAB<sup>®</sup> routines called SOCSol. See [KW97, AK06, AK08, Kra08] for explanations and some applications of this method.

The rest of this paper is organised as follows. In Section 2, we formulate a pension fund problem as a stochastic optimal control problem. Then, a simple-target problem (i.e., one not contingent on inflation) is used in Section 3 to illustrate the differences between the different performance measures. In Section 4, we solve a contingent-target problem. Section 5 provides some concluding remarks.

## 2 Pension fund problems as stochastic optimal control problems

A plausible situation in financial management is one in which an agent deposits an amount  $x_0$  with a fund manager at time 0, to be repaid at time  $T$  with a lump sum  $x(T)$ , which we call the *pension*. The pension  $x(T)$  depends upon both the

---

<sup>5</sup>This paper continues the line of research on lump-sum pension investment strategies initiated in [Kra01] and continued in [Kra03, Kra05, Kra08].

<sup>6</sup>See [KD01] for a complete treatment of Kushner’s Markov chain approximation method for stochastic optimal control. Also, see [PR02] for a controlled diffusion approximation method likewise motivated by [Kus90] but which differs from the approach used here.

investment policy  $\mu(x(t), t)$  ( $t \in [0, T]$ ) adopted by the fund manager and market conditions.<sup>7</sup> The latter are deterministically unpredictable, so are usually modelled via stochastic processes. Consequently, the pension is a random variable and the pension fund problem is inherently stochastic.

In determining an acceptable initial deposit  $x_0$ , the agent and manager need to address several practical considerations. Of these, the following seem to be the most pertinent.

- (i) What lump sum  $\bar{x}_T$  can the manager “promise” for a given initial deposit  $x_0$ ?
- (ii) Alternatively, what initial deposit  $x_0$  should be paid in exchange for a “promise” of the lump sum  $\bar{x}_T$ ?
- (iii) How should the “promise” be formulated: deterministically or in some probabilistic terms?

We attempt to address these considerations below in this paper.

The investment policy governing a fund depends on the fund manager’s objective function (or utility measure). Possible objectives include maximisation of the expected fund value, maximisation of the probability of achieving a target payoff, minimisation of shortfall, etc. Once an objective function is chosen, the manager’s policy can be computed as the solution to a stochastic optimal control problem determined by the objective. The solution routinely provides an optimal decision rule  $\mu(x(t), t)$ , which is crucial for the manager’s control of the portfolio. However, it should also provide more “practical” information about the resulting payoff distribution, as this allows the agent to decide what they can reasonably expect for their pension.

Knowledge of the distribution of  $x(T)$  is also helpful to the manager, as it helps describe the risks associated with obtaining a particular realisation of the objective. For example, the distribution may suggest that, for a given initial deposit  $x_0$  there is a “probable” lump sum  $\bar{x}_T$ , which the manager may choose to advertise as the pension target (subject to legislative constraints).

It is customary to model a pension fund problem like that described above using a form of Merton’s optimal portfolio selection model (see [Mer71]; further details are available on pp. 160–161 of [FR75]). As is commonly the case in the literature, we assume that the portfolio consists of two assets, one “risky” (e.g. shares) and the other “risk free” (e.g. cash). If the price  $p(t)$  per share of the risky asset changes according to the equation

$$dp(t) = \alpha p(t)dt + \sigma p(t)dw \quad (1)$$

where  $\alpha, \sigma > 0$  are constants and  $w$  is a one-dimensional standard Brownian motion, while the price  $q(t)$  per share of the risk-free asset changes according to the equation

$$dq(t) = rq(t)dt,$$

---

<sup>7</sup>The policy could be multidimensional, comprising a consumption rule, administration fee etc.

where  $r \in (0, \alpha)$  is a constant, then the fund value  $x(t)$  at time  $t \in [0, T]$  changes according to the stochastic differential equation

$$dx(t) = (1 - u(t))rx(t)dt + u(t)x(t)(\alpha dt + \sigma dw) - v(t)dt. \quad (2)$$

Here,  $u(t)$  and  $1 - u(t)$  respectively denote the fractions of the fund invested in the risky and risk-free assets at time  $t$ , and  $v(t)$  denotes the fund "consumption rate."

In [FR75], the manager wishes to find a two-dimensional strategy  $\mu(x, t) = (u(x, t), v(x, t))$  maximising the total expected discounted utility

$$J(x_0, \mu) = \mathbb{E} \left[ \int_0^T e^{-\rho t} g(v(t)) dt \mid x(0) = x_0 \right], \quad (3)$$

where  $\rho > 0$  is a force of discount and  $g(v(t))$  is the manager's instantaneous utility at time  $t$ , subject to the constraints

$$0 \leq u(t) \leq 1, \text{ and} \quad (4)$$

$$v(t) \geq 0 \quad (5)$$

holding for all  $t \in [0, T]$ .<sup>8</sup> In [FR75], no constraint is imposed on the wealth  $x(T)$  at time  $T$ , and the amount  $x_0$  of the initial deposit is taken as given.

If we augment the utility measure in (3) to include a final payoff function  $h(x(T))$ , giving

$$J_{aug}(x_0, \mu) = \mathbb{E} \left[ \int_0^T e^{-\rho t} g(v(t)) dt + e^{-\rho T} h(x(T)) \mid x(0) = x_0 \right], \quad (6)$$

then the problem of maximising (6) subject to (2), (4), (5) and other relevant constraints (e.g.  $x(t) \geq 0$ ) serves as a model of the pension fund problem. Indeed, setting  $v \equiv 0$  and maximising

$$J_*(x_0, u) := \mathbb{E} [h_*(x(T), \bar{x}_T) \mid x(0) = x_0] \quad (7)$$

in  $u$  defines a problem that captures the task of a fund manager.<sup>9</sup> Note that discounting in this context corresponds to multiplication of the utility measure (7) by a constant, and thus does not alter the manager's optimal strategy. Consequently, we omit the multiplier  $e^{-\rho T}$  from this and similar maximisation problems.

In the rest of this paper we also assume that a management fee having force  $cx(t)$  is charged, where  $c > 0$  is a constant. Consequently wealth evolves according to the equation

$$dx(t) = (1 - u(t))rx(t)dt + u(t)x(t)(\alpha dt + \sigma dw) - cx(t)dt. \quad (8)$$

<sup>8</sup>Constraint (4) means no short selling or borrowing. This restriction has been weakened in the literature; however, it may be reasonable in some situations.

<sup>9</sup>The expectation in (7) will be well defined under the implicit integrability assumption.



As mentioned in the Introduction, we distinguish between several managerial objectives, which we define below. Some will explicitly specify a target return  $\bar{x}_T$  to be achieved in time  $T$  from an initial outlay  $x_0$ . The target can be fixed (as in the problems of Section 3) or contingent on a stochastic process such as inflation (see Section 4).

We call the fund manager who seeks a strategy  $u_M$  maximising (7) with final payoff function

$$h_M(x(T), \bar{x}_T) := \frac{1}{\gamma}(x(T))^\gamma \quad (9)$$

subject to (4) and (8), where  $\gamma \in (0, 1)$  is a constant, the *Merton manager* **MM**. Note that the resulting utility measure  $J_M(x_0, u)$  does not explicitly depend on a fund target  $\bar{x}_T$ .

The so-called *cautious-relaxed manager* **CM** (see [Kra08, Kra05]) seeks a strategy  $u_C$  maximising (7) with final payoff function

$$h_C(x(T), \bar{x}_T) := \begin{cases} (x(T) - \bar{x}_T)^\kappa & \text{if } x(T) \geq \bar{x}_T, \text{ and} \\ -(\bar{x}_T - x(T))^a & \text{otherwise,} \end{cases} \quad (10)$$

subject to (4) and (8), where  $a > 1$  and  $\kappa \in (0, 1)$  are constants. The resulting utility measure  $J_C(x_0, u)$  only offers the fund manager a moderate incentive to exceed  $\bar{x}_T$  (as  $0 < \kappa < 1$ ), but punishes failure to reach this target substantially (as  $a > 1$ ).

The *prospect-theoretic manager* **PM** considered by [TK92] also seeks a strategy likely to result in funds meeting a target  $\bar{x}_T$ . However, this strategy  $u_P$  is obtained by maximising (7) with final payoff function

$$h_P(x(T), \bar{x}_T) := \begin{cases} B(x(T) - \bar{x}_T)^\kappa & \text{if } x(T) \geq \bar{x}_T, \text{ and} \\ -A(\bar{x}_T - x(T))^\zeta & \text{otherwise,} \end{cases} \quad (11)$$

where  $\kappa, \zeta \in (0, 1)$  and  $A > B > 0$  are constants. Here, the resulting utility measure  $J_C(x_0, u)$  also only moderately rewards the manager for exceeding  $\bar{x}_T$ , but does not severely penalise them for failing to reach the target (as  $0 < \zeta < 1$ ).

Finally, the manager considered by [BKP04], whom we call the *loss-averse manager* **LM**, also works toward a target. However, their strategy  $u_L$  is obtained by maximising (7) with final payoff function

$$h_L(x(T), \bar{x}_T) := \begin{cases} \frac{B(x(T))^\gamma}{\gamma} & \text{if } x(T) \geq \bar{x}_T, \text{ and} \\ \frac{A(x(T))^\gamma + (B - A)\bar{x}_T^\gamma}{\gamma} & \text{otherwise,} \end{cases} \quad (12)$$

where  $\gamma \in (0, 1)$  and  $A > B > 0$  are constants. The resulting utility measure  $J_L(x_0, u)$  has the same shape as  $J_P$ , but does not penalise small losses.

The final payoff functions for those managers aiming toward a target are shown in Figure 1 for the target  $\bar{x}_T = \$100\,000$  and a set of parameters explained and utilised in Section 3. We see that the **CM** manager behaves more conservatively than the other target-motivated managers, as the **CM** manager is



risk-averse in the “loss domain” below the target, as well as in the “profit domain” above it. The parameter values for **PM**, **LM** and **CM** utility measures are taken from the existing papers [BKP04] and [Kra05, Kra08], respectively. They are claimed (in [BKP04]) to reflect investors’ behaviour and proved (in [Kra05, Kra08]) to generate left skewed payoff distributions. The power of the **MM** utility function was calibrated to guarantee an interior solution for  $u(t)$  (see (4)).

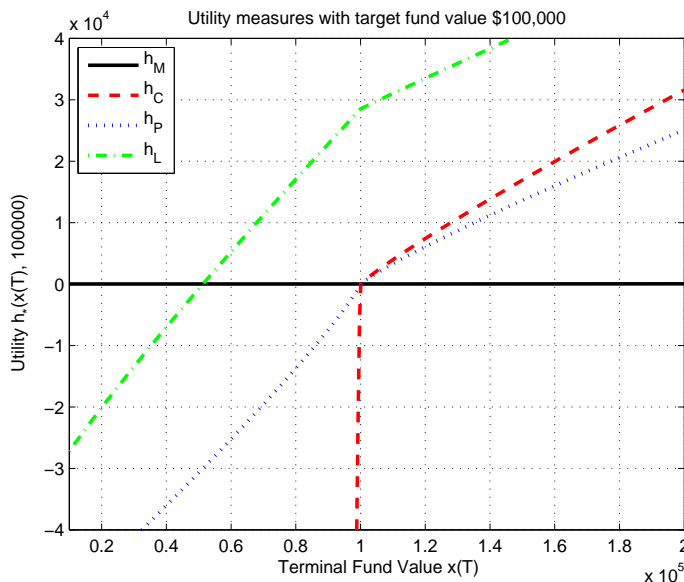


Figure 1: Risk-averse and target-seeking utility measures.

The key difference between the loss-aversion approach and the classical approach of Merton seems to lie in the setting of a distinguished point (the target) about which behaviour differs. For example, the **MM** utility measure  $h_M(x_T) = 20(x_T)^{0.05}$  assumed in Section 3 generates a smooth, almost flat line in Figure 1 at the given scale.

While each manager’s objective determines a stochastic optimal control problem, analytic solutions to these problems are frequently unavailable. Aside from the **MM** problem, no easily-interpretable feedback forms have been obtained for the optimal strategies. In particular, the expressions given in [BKP04] for the optimal wealth distributions and investment strategies for the **PM** and **LM** problems are based on a *state price density*, which may be a challenging choice of variable upon which to build a useable feedback solution. In contrast, current fund value and inflation (see Section 4) are *observable* state variables upon which it is sensible to base a feedback solution.

Moreover, we are unaware of a transparent procedure providing a closed form for the **CM** manager’s optimal strategy. As we wished to compare all four managers’ strategies and payoff distributions, and (perhaps more importantly) sought *feedback* strategies based on fund value, we solved the four problems using the same method. This is the numerical optimisation approach called SOCSol (see [KW97, AK06, AK08, Kra08]) that was mentioned in the Introduction.

### 3 A simple pension fund problem and its solutions

In this section we aim to explain the basic features of solutions obtained for the managerial objectives outlined in Section 2.

#### 3.1 Model parameters

Our interest lies in a link between the formulation of a manager's objective and the shape of resulting payoff distribution. As stated above, we generate the payoff distributions numerically for all four managers' problems to ensure consistency. The parameters used are given in Tables I–V.

$r$	$\alpha$	$\sigma$	$c$
5%	8.5%	20%	0.5%

Table I: Parameters for fund value dynamics (Equation (8)).

$\gamma$
0.05

Table II: Parameters for **MM** final payoff function (Equation (9)).

$a$	$\kappa$
1.5	0.9

Table III: Parameters for **CM** final payoff function (Equation (10)).

We observe that our managers have access to “bonds” that grow at the fixed rate  $r = 5\%$ , and that the risk premium is perceived to be 3.5% (so  $\alpha = r + 3.5\%$ ). The volatility of the uncertain asset is 20%. The managers charge a relatively small fee of 0.5%, as their policies do not entail “chasing the winners.” Each manager receives an initial outlay of  $x_0 = \$40,000$ , which they are to administer for  $T = 10$  years. The target-seeking managers assume that the target is static, being set at  $\bar{x}_{10} = \$100,000$  from the outset.

The **MM** manager is quite risk averse (as  $\gamma = 0.05$ ). This attitude is shared by the **CM** manager, who strongly wishes to avoid payoffs less than  $\bar{x}_T$  (as  $a = 1.5 > 1$ ). The **CM** manager also has no strong incentive to overshoot the target (since  $\kappa = 0.9 < 1$ ).

The other managers also feel the effect of a loss more strongly than that of an equivalent gain (as  $A > B$ ). This property of “losses looming larger than gains” has behavioural foundations; see [TK92]. Moreover, Tversky and Kahneman find in [TK92] that an agent's utility function is approximately a power function with an exponent of less than 1. The choices of  $\zeta$ ,  $\kappa$ ,  $\gamma$ ,  $A$  and  $B$  used here are based on their empirical estimates.

$A$	$B$	$\zeta$	$\kappa$
2.25	1	0.88	0.88

Table IV: Parameters for **PM** final payoff function (Equation (11)).

$A$	$B$	$\gamma$
2.25	1	0.88

Table V: Parameters for **LM** final payoff function (Equation (12)).

### 3.2 Management of the portfolios

Assume that the fund value process is described by Equation (8). Subject to these dynamics and the constraints given in Equation (4), the **MM** manager seeks a strategy  $u_M$  maximising  $J_M$  (see Equations (7) and (9)). Similarly, the other managers **CM**, **PM** and **LM** seek strategies  $u_C$ ,  $u_P$  and  $u_L$  maximising  $J_C$ ,  $J_P$  and  $J_L$  respectively (see Equations (10)–(12)).

We solved each manager's portfolio administration problem, determining optimal<sup>10</sup> strategies  $u_M$  and approximately optimal strategies  $u_C$ ,  $u_P$  and  $u_L$  in the feedback form i.e., as functions of observed fund value and time. The application of a strategy to  $x_0$  generates a final fund return for each given realisation of the Brownian motion  $w$ .

Figure 2 shows the managerial strategies  $u_M$ ,  $u_C$ ,  $u_P$  and  $u_B$ ; the corresponding fund value and strategy profiles are displayed in Figure 3. The final fund value realisations are those points in the upper subplots corresponding to time  $t = 10$ . Distributions of these final fund values are given in Figure 4.

Unfortunately, some of these strategies generate unacceptable final fund value distributions; see Figure 4. We comment on the features of the distributions in Section 3.3; here we observe that few fund managers would be happy to pay 40% of their clients below \$66,000 when investing solely in the secure asset would have yielded

$$x_S(t) = \$40,000 \exp(t(r - \text{"management fee"})) = |_{t=10} = \$62,732. \quad (14)$$

However, this is the case for the payoff distributions resulting from the **MM**, **PM** and **LM** managers' strategies (see Table VI).

All three loss-avoiding policies are dynamic, advocating that share holdings be inversely proportional to fund value for a variety of fund values. However, the **CM** policy differs markedly from the other two, as it distinguishes between two investment zones.

1. Fund values  $x(t) \in (0, x_S(t))$ , from which the target can only be reached

<sup>10</sup>The **MM** manager's problem has an easy analytical solution, namely

$$u_M(t) = \frac{\alpha - r}{\sigma^2(1 - \gamma)} = \frac{0.085 - 0.05}{0.2^2(1 - 0.05)} = \frac{35}{38} = 0.9211 \text{ (4 s.f.)} \quad (13)$$

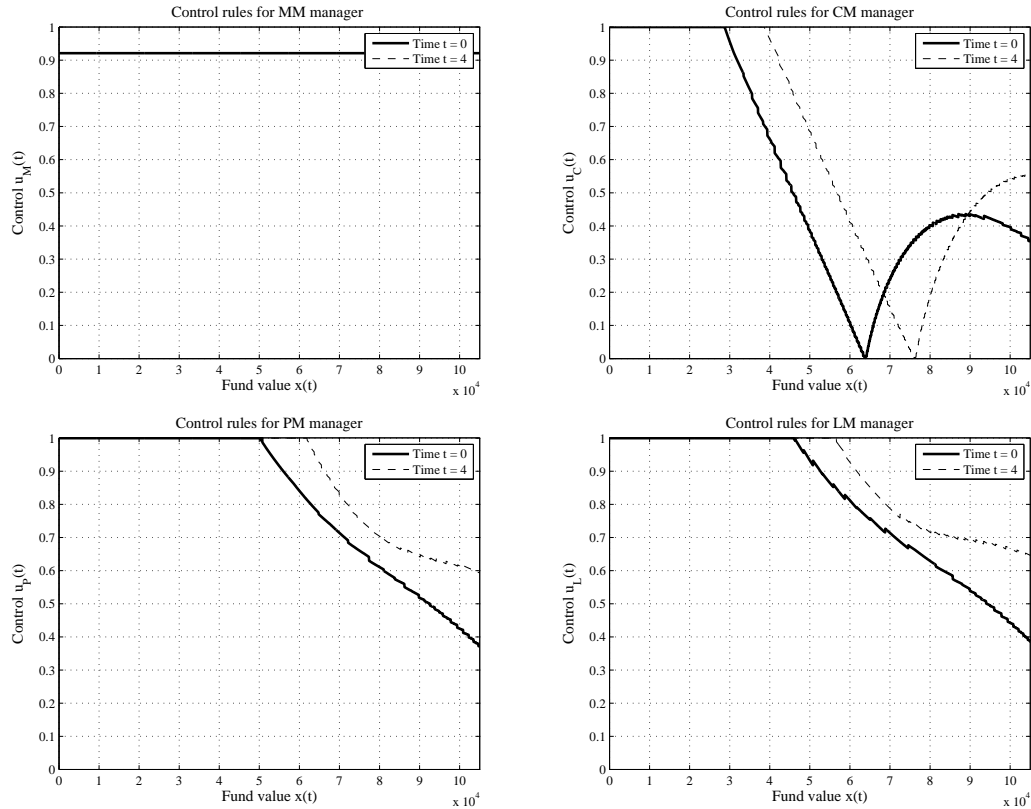


Figure 2: Portfolio management strategies for times 0 and 4.

by investing in the risky asset. Here, the manager must gamble to evade heavy penalties for falling short of the target.

2. Fund values  $x(t) > x_S(t)$ , from which the target can be reached by investing solely in the secure asset. In this zone, the manager maximises their reward for exceeding the target.

We call a **CM** manager whose fund value is in the first zone *cautious* (for they strive to avoid the consequences of underperforming), and one whose fund value lies in the second zone *relaxed* (for they enjoy the rewards of surpassing the target). We note that the second zone becomes unimportant in the *continuous* state-time setting, for the **CM** manager will set  $u_C(x_S(T-t)) = 0$ , thus withdrawing from the share market. This withdrawal from the share market results in the fund value never reaching the second zone as described above, but rather growing deterministically until it reaches the target at time  $T$ .

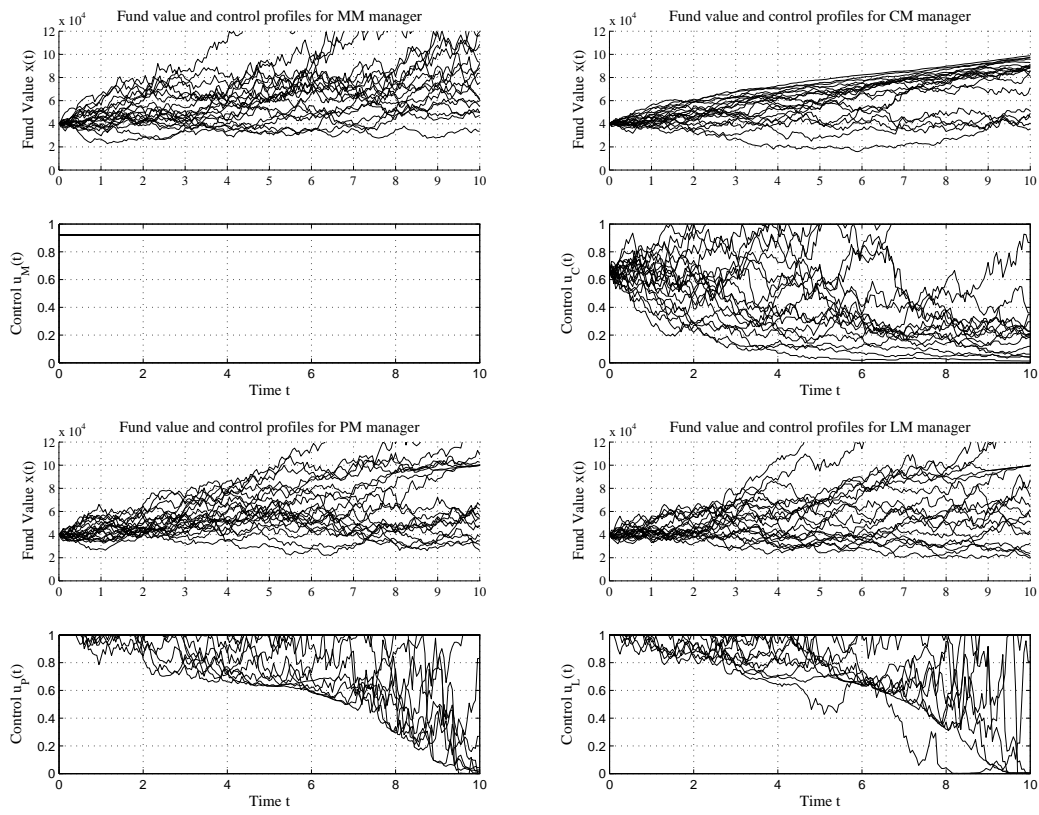


Figure 3: Fund value and strategy realisations.




Figure 4: Fund value distributions.

### 3.3 The effect of problem formulation on final return distribution

As is the case in computational economics, the results presented here (also see Section 4) are parameter-specific. The comments in this section should be interpreted in this light, with a view to further analysis of the final fund value distribution's sensitivity to the choice of utility measure. Nonetheless, the results appear to be *sufficient* to guide fund managers to choose a particular utility measure.

There is a vast difference between the concave (risk-averse) utility maximisation solutions of the **MM** manager and the target-dependent strategies of the other managers. The latter are "dynamic" in the sense that they propose actions depending on both the observed fund value and the *time-to-go* i.e., the remaining time before pension payment.

The **MM**, **PM** and **CM** managers' strategies each generate a right-skewed payoff distribution, while the **CM** manager's strategy yields a left-skewed payoff distribution. Table VI quantifies these facts. Note that the state bounds used for computing the **PM** and **LM** strategies mean that some of the tabulated values can strictly be taken only as lower bounds on the actual values; this is indicated in the table when applicable.

Manager	MM	CM	PM	LM
Mean of $x(10)$	\$86,546	\$75,018	> \$85,791	> \$86,024
Median of $x(10)$	\$73,514	\$83,272	\$82,290	\$80,161
40-th percentile of $x(10)$	\$63,610	\$77,648	\$65,592	\$64,479
Std. dev. of $x(10)$	\$53,994	\$21,649	> \$51,419	> \$53,980
Coeff. of skew. of $x(10)$	2.0426	-1.0014	> 1.7560	> 1.8558
$P(x(10) > \$100,000)$	0.2944	0	0.3205	0.1948
$P(x(10) > \$62,732)$	0.6081	0.7451	0.6216	0.6181
$P(x(10) < \$40,000)$	0.1451	0.1082	0.1640	0.1692
$P(x(10) < \$20,000)$	0.0135	0.0107	0.0207	0.0183

Table VI: Final fund return distribution statistics.

In addition the final four rows of the table give probabilities relative to the nominal target (\$100,000), earnings if investing solely in the secure asset (\$62,732), the initial deposit (\$40,000) and an arbitrary value below the initial deposit (\$20,000).

We note that with the exception of the **CM** policy, there is a probability of about 0.4 of earning less than the "secure" revenue  $x_S(T)$  that could be earned by investing solely in the risk-free asset (see (14)). Under the **CM** policy, this probability is about 0.25.

The final payoff falls short of the initial deposit  $x_0 = \$40,000$  with a probability of about 0.15 for the non-**CM** policies and about 0.1 for the **CM** policy. So, while a strategy maximising expected return (**MM**) is a very risky means of fund



management, the loss averse policies **PM** and **LM** are not much safer. We conclude that it is the cautious-relaxed policy **CM** that will best uphold the fund's yields in an acceptable range.

## 4 A contingent-target cautious-relaxed policy

### 4.1 A realistic target

We have seen that target-seeking portfolio management formulations are closer to what the managers want to achieve. Notwithstanding the good features of the prospect-theoretic policies (P and B) we believe that the cautious-relaxed strategy (C) delivers the final yield distributions, which should suit the *pension* fund managers best. We will now consider an extension to manager C's target seeking problem that makes it more realistic.

An alternative portfolio optimisation problem, with respect to those discussed in Section 3, could be one that allows for the target to be contingent on an exogenous stochastic process. We propose the following generalisation of problem  $J_C$ , (10):

$$J(x(0), p(0); u^*) = \sup_{u \in [0, 1]} \mathbb{E} \left( s(x_u(T), H(T)) \mid x(0) = x_0, p(0) = p_0 \right) \quad (15)$$

where

$$s(x_u(T), H(T)) = \begin{cases} (x_u(T) - H(T))^\kappa & \text{if } x_u(T) \geq H_T, \\ -(H(T) - x_u(T))^a & \text{otherwise} \end{cases} \quad (16)$$

$$0 < \kappa < 1, a > 1.$$

The above formulation is motivated by an observation that the amount  $\bar{x}_T$  aimed at in problem  $J_C$ , (10) does not have to be a fixed value. If the investor's objective is to hedge a certain liability  $H(T)$  at a given future time  $T$ , then  $\bar{x}_T$  may be identified as the final value  $H(T)$  of a "liability" process  $H(t)$ ; compare [FR99] for an optimisation model that computes risk-minimizing hedging strategies.

We can find in [BKP04] an attempt at dynamic updating of the fund's target  $H(T)$ . Their assumption is that the target may be driven by the same Brownian motion as the risky asset price (and wealth) is. In "real life", a contingent  $\bar{x}_T = H(T)$  may be measurable with respect to the  $\sigma$ -algebra generated by a stochastic process, largely unrelated to the wealth dynamics (8). For example, the overall price index (like CPI)  $\pi(t)$  fed in by inflation could be that process. Assuming that the price index process is a geometric Brownian motion<sup>11</sup> we get

$$d\pi(t) = \beta\pi(t)dt + \varphi\pi(t)dw_1, \quad t \in [0, T], \quad \pi(0) = 1 \quad (17)$$

where  $\beta$  is inflation (assume<sup>12</sup>  $\beta = 2.14\%$ ) and  $\varphi$  is the inflation "volatility" (assume  $\varphi = 0.7\%$ ). Allowing for (17) will enable us to "index" the final liability

<sup>11</sup>Inflation might have a seasonal component corresponding to the economic cycle. We ignore it in our GBM model that could though be extended to allow for the cycle.

<sup>12</sup>This and the next number come from the Reserve Bank of New Zealand data set for the period between 1996 and 2005.

by making it *contingent* on  $\pi(T)$ . Notice that  $w_1$  is a one-dimensional Brownian motion different from  $w$ .

Allowing for inflation in the process of striving to meet the target is a realistic feature of the model, which should help adjust the investment strategy  $u^*$  to different market situations. We propose the following model for the contingent target

$$H(T) = \tau\pi(T) \quad (18)$$

where  $\tau > 0$  is the nominal target (assume<sup>13</sup>  $\tau = 80735$ ). We know that price realisation of a geometric Brownian motion for a particular value of  $t \in [0, T]$  is a random variable distributed log-normally. Thus  $H(T)$  is a random variable distributed log-normally, *contingent* upon the inflation process.

In [Kra05] an analysis of a shortfall risk minimisation problem was conducted where the goal to not fall short of was contingent on the risky asset price process. The resulting portfolio management policies were obtained as time and state dependent, *looking* similar to those of Figure 2 albeit the time dependence was “inverted” (i.e., the policy advice was to invest more in earlier times than in Figure 2, for the same level of wealth<sup>14</sup>). Here, we solve several pension fund problems maximising the *cautious-relaxed* manager’s utility measure (15)–(16) for different assumptions on the degree of correlation between inflation and the risky asset price.

## 4.2 Contingent-target policies

Suppose  $w$  and  $w_1$  are two one-dimensional independent Brownian motions. Then, assuming that the risky asset price can be influenced by inflation, the share price process will follow

$$dp(t) = \alpha p(t)dt + \sqrt{1 - \rho^2} \sigma p(t)dw + \rho\sigma p(t)dw_1, \quad (19)$$

where  $\rho$  is correlation coefficient  $-1 \leq \rho \leq 1$ , rather than process (1). Consequently, the wealth equation will be modified to

$$dx(t) = (1 - u(t))rx(t)dt + u(t)x(t)(\alpha dt + \sigma\sqrt{1 - \rho^2} dw + \sigma\rho dw_1) - cx(t)dt. \quad (20)$$

Below, we will compute cautious-relaxed strategies  $u^*$  i.e., such that maximise (15)–(16) subject to (20), (17) and (4) for three situations:

- when there is no correlation between inflation and the risky asset price,  $\rho = 0$  (e.g.,  $u^*$  is investment in an overseas asset);
- when there is a negative correlation  $\rho < 0$  and

---

<sup>13</sup> $80735 \exp(10\beta) = 100000$  that is this target would match  $\bar{x}_T$  if there were no inflation volatility.

<sup>14</sup>However, the results cannot be directly compared because the volatility in [Kra05] is higher than in this paper.

- when there is a positive correlation  $\rho > 0$

**No correlation.** So, the portfolio manager maximises (15)–(16) subject to (20) (practically, (8)), (17), (4).

Figure 5 shows the approximately optimal policy  $u^*$  whose realisations are presented in Figure 7 (third panel) together with the two state variables' time profiles (top panels).

We notice that, in principle, the same kind of the manager's strategic approach to obtain  $H(T)$ , as was observed in the case of a non stochastic target shown in Figure 2 (included here as the *thinner* lines) is evident in Figure 5.

That is, the policy advice is that the further the current wealth  $x(t)$  is from the target, the more aggressive its pursuit should be. The lines drawn in the upper panel are for  $\pi = 1$ ; the lines drawn in the bottom panel concern wealth  $x = \$40\,000$ .




Figure 5: Market-dependent cautious-relaxed policies for  $\rho = 0$ .

The markable difference is in that the manager is never certain of achieving target  $H(T)$ , because it is driven by the stochastic (inflation) process. This causes the absence of  $u^* = 0$ , which was a feature of the policy advice for a

fixed target. Correspondingly, the wealth level for which the investor starts to be more “relaxed” about meeting the target and maximises the gains from exceeding it, is higher (about \$83 000 as opposed to \$64 000 for the fixed target). If there is less time to reach the target, the investment into the risky asset increases up to the “relaxation” point. We can also see that when inflation increases (for wealth=\$40 000), the manager buys more shares.

We have a full picture of the impact of inflation on the investment strategy in Figure 6 (for times 0 and 4). We can see that as the price level increases, the manager’s “aggressiveness” of pursuing the target diminishes. This is inferred from the lower gradient (in wealth) of the strategy for higher price levels. The same pattern prevails for a later time, see the right panel, albeit the commitment to buying shares increases.

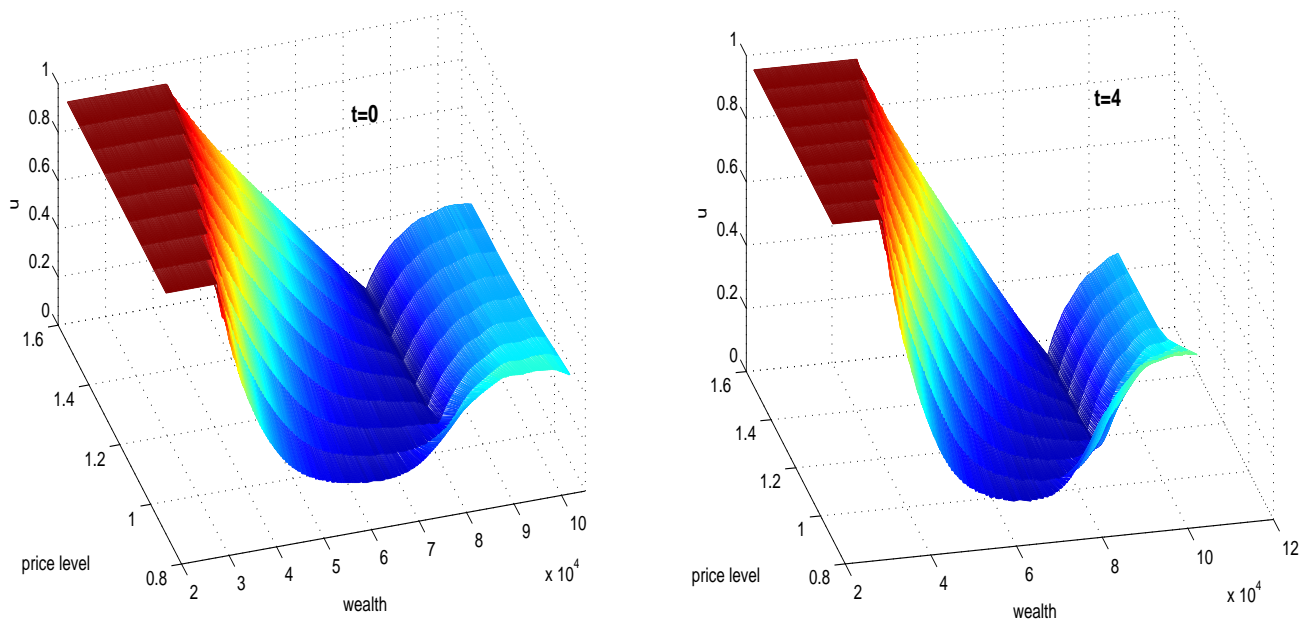


Figure 6: Market-dependent cautious-relaxed policies for different wealth and price levels ( $t = 0$ , and  $t = 4$ ;  $\rho = 0$ ).

Figure 7 presents a sample of 40 time profiles of the state variables (top two panels) and of the control (bottom). We can see that the wealth trajectories congregate at the top of the figure; however, the final value of wealth tends to be lower than in Figure 3. The goal however was contingent on inflation, so these lower absolute values cannot speak about goodness or badness of manager’s strategies.




Figure 7: Time profiles for the market-dependent cautious-relaxed policy ( $\rho = 0$ ).

We can see the histogram of the final payoff distribution in Figure 8 together with the distribution of the contingent target. We can appreciate that the manager's task can be interpreted as seeking an investment strategy that could generate a payoff that should match a target distributed log-normally.

Finally, figure 9 is the distribution of the ratio  $\frac{x(T)}{\tau\pi(T)}$  that tells us about the target satisfaction. It documents that meeting it is a difficult task, as only about 10% of the portfolio management cases end up with a final payoff exceeding the target. However, about 62% of cases exceed payoffs that are between 75% and 110% of the target. This is on par with the **CM** policy performance when the target is fixed, see Figure 4 (top right panel). A fixed-interest investment would generate  $x_S(10) = 62732.50$ , short of the target by any measure.

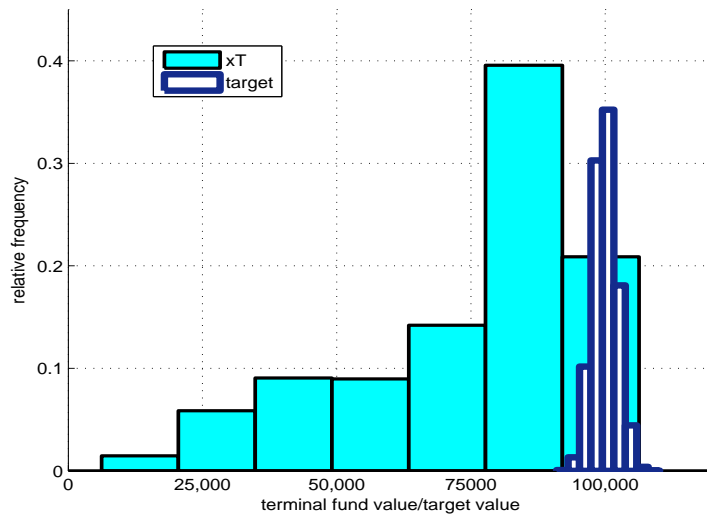


Figure 8: Final payoff distribution and the target distribution.




Figure 9: Distribution of target satisfaction.

We have also assessed the impact of the inflation *volatility*, measured by the parameter  $\phi$ , on the investment strategy. Figure 10 shows the approximately optimal strategies for  $\beta = 0$  i.e., as if the price level depended entirely on inflation volatility  $\phi$ . We see in the figure that if the price process followed just a random walk, the manager would have “gambled” less and have bought fewer shares (about 20% less relative to when the price process were a geometric Brownian motion.) The target satisfaction (not shown here) is not dissimilar to the previous cases.

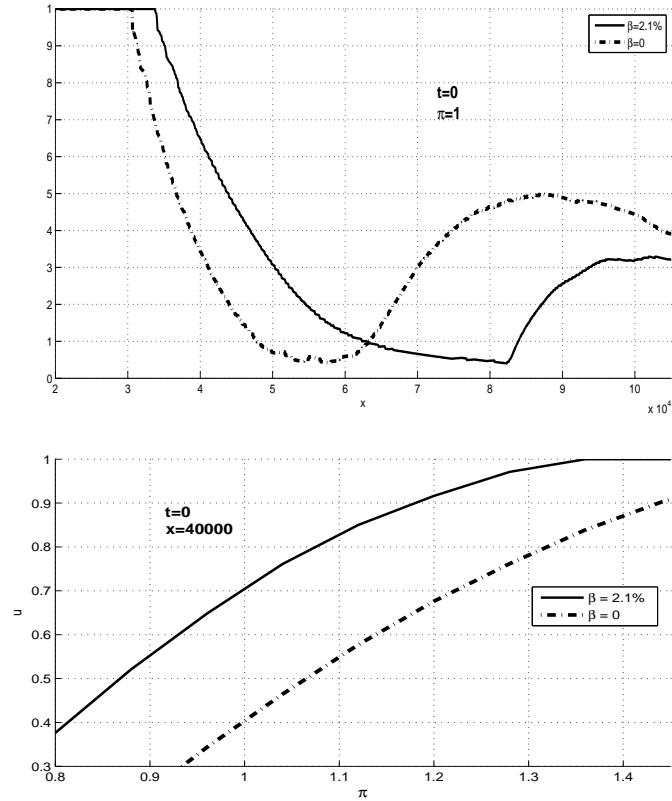


Figure 10: Market-dependent cautious-relaxed policies for  $\rho = 0$  and  $\beta = 0$ .



**A negative correlation.** Here, the portfolio manager maximises (15)–(16) subject to (20), (17), (4) and when  $\rho = -0.25$ .

Figure 11 shows the approximately optimal policies  $u^*$  for  $\rho = -0.25$  (dash-dotted lines); also for  $\rho = 0.25$  (dashed lines) and  $\rho = 0$  (solid thin lines, same as in Figure 5). The left panels are for  $t = 0$ ; the right ones for  $t = 4$ . The policy realisations for the case of negative correlation are presented in Figure 12 (bottom panel), along with the state variables' time profiles (top panels).

We notice that, in principle, the same kind of the manager's strategic approach to obtain  $H(T)$ , as was observed in the case of no correlation (see the *thin* lines), is evident in Figure 11. That is, the further the current wealth level  $x(t)$  is from the target, the more aggressive its pursuit.

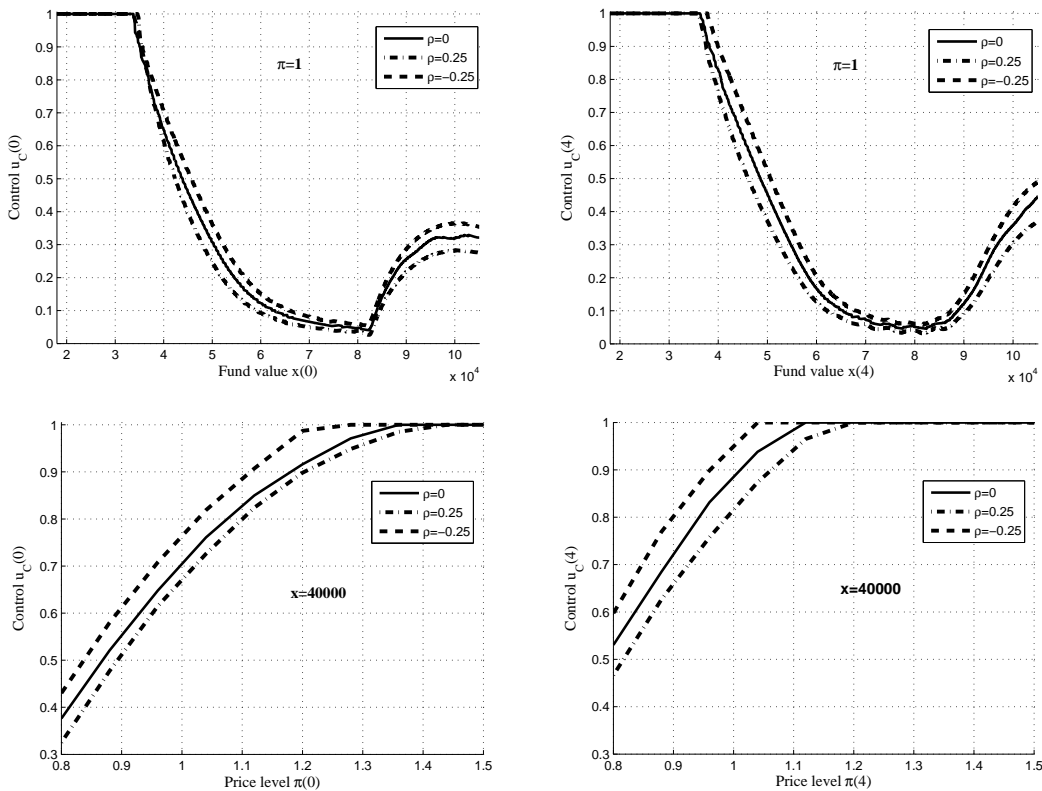


Figure 11: Market dependent cautious-relaxed policies for  $\rho = 0, \pm 0.25$ .

The difference between the policies obtained for the negative correlation between the shocks affecting inflation and the risky asset price and those, which allow for no correlation, is that now the manager buys more shares of the risky asset. This is not necessarily an intuitively obvious policy. Its explanation is that, in the zone where the target *cannot* be achieved through the sole investment in the secure asset, the negative correlation between the asset price and inflation (which drives the target level) pushes the manager to “gamble” more, than when there were no correlation.

With more funds and some certainty of achieving the target i.e., the other

zone, the manager simply seeks to maximise the payoff about the target. In plain language, the manager gambles in the first zone out of necessity and because they are greedy in the second.




Figure 12: Time profiles for the market-dependent cautious-relaxed policy ( $\rho = -0.25$ ).

The sample time-profiles shown in Figure 12 suggest that a correlation (negative in this case) between the risky asset price and inflation reduces the spread of the final yield, in comparison to the uncorrelated situation (see Figure 7).

We can see this in the histogram of the final payoff distribution in Figure 13, plotted together with the distribution of the contingent target. The modal bar reaches higher in this figure than in Figure 8 and the probability of receiving a payoff between \$63 000 and \$105 000 increases; the probability of receiving less than \$30 000 is negligible.

We can appreciate that the strategies that allow for this negative correlation help the manager's task to match the target, see Figure 14 where the ratio  $\frac{x(T)}{\tau\pi(T)}$  is displayed. We see that, under this correlation, the target satisfaction distribution is more left-skewed than was in Figure 9.

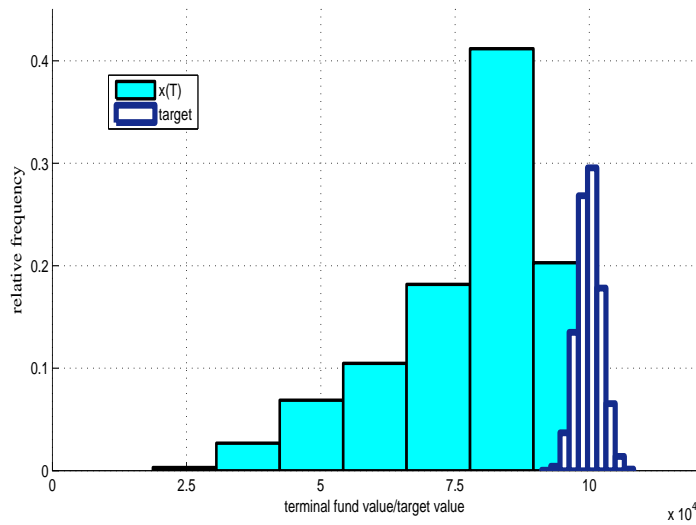


Figure 13: Final payoff distribution and the target distribution for  $\rho = -0.25$ .




Figure 14: Distribution of target satisfaction for  $\rho = -0.25$ .

**A positive correlation.** Finally, the portfolio manager maximises (15)–(16) subject to (20), (17), (4) and when  $\rho = 0.25$ .

The dash-dotted lines in Figure 11 show the approximately optimal policy  $u^*$  whose realisations are presented in Figure 15, (third panel) together with the state variables' time profiles (the top panels).

We notice (in the policy figure) that the strategies for both correlated cases (negatively and positively) are similar and analogous to those obtained for the uncorrelated case, see *thin* lines, copied from Figure 5 for  $\rho = 0$ . The general prescription: the further the current wealth  $x(t)$  is from the target, the more aggressive its pursuit is true for all cases. However, the difference between the "positively correlated" policy and the other ones is that now the manager buys less shares of the risky asset. This is a reflection of the manager's greater confidence that with this positive correlation between the risky asset price and inflation, which modulates the target, chances of reaching it are better than before. Hence, less "gambling" is required to achieve the objective.




Figure 15: Time profiles for the market-dependent cautious-relaxed policy ( $\rho = 0.25$ ).

However, the sample time-profiles shown in Figure 15 suggest that this positive correlation will increase the spread of the final yield, in comparison to the uncorrelated and negatively correlated cases. Consequently, this positive correlation case appears difficult to handle.

The final payoff distribution that results from the application of  $u^*$  policies from Figure 11 (dash-dotted lines) is presented in Figure 16 (together with the distribution of the contingent target, as before). We can appreciate that the probability of receiving a payoff between \$95 000 and \$105 000 has increased. However, so has the probability of a payoff below \$50 000 and, markedly, of one below \$30 000. The reported probability increase and decrease apply to both the uncorrelated and negatively correlated cases (qualitatively, the numerical probability values differ).

Finally, figure 17 is the distribution of the ratio  $\frac{x(T)}{\tau\pi(T)}$  that tells us about the target satisfaction. It confirms the above findings. About 25% of the final payoff realisations are within  $\pm 5\%$  of the target, which is a better result than for the other cases. However, the computed policy also delivers some low yields that were avoided in the other cases. About 75% of fund yields are above the fixed-interest investment payoff.




Figure 16: Final payoff distribution and the target distribution for  $\rho = -0.25$ .




Figure 17: Distribution of target satisfaction for  $\rho = -0.25$ .

## 5 Conclusion

A numerical optimisation method was used for the solutions to continuous-time stochastic optimal-control problems, reflecting several pension fund problems. A “common-sense” policy *bet aggressively if you are far from the target* has been quantified as an approximately optimal solution to problems of *cautious-relaxed* fund managers. The policy was shown to generate encouragingly high proportions of fund yields that meet a savings’ target in %80 or better. A model feature that helps to produce the above investment strategies with a client-friendly yield distribution appears to be an asymmetric utility function, which captures the investor’s strong loss avoidance attitude and a mild enjoyment from exceeding a target. We have also provided an analysis of other models showing advantages and disadvantages of their use.

As to the bullet point questions asked in Section 2 (p. 4), the following answers may be formulated.

- (i) and (ii) A relationship between  $\bar{x}_T$  and  $x_0$ , static or contingent on inflation, can be read from the fund yield histograms. The manager can advertise  $\bar{x}_T$  such that its realisation is 90% probable.
- (iii) Any “promised”  $\bar{x}_T$  is stochastic in nature. However, it could be formulated deterministically if the manager strongly believed that the “promised”  $\bar{x}_T$  is very close to the mode of the payoff distribution.

## Acknowledgment

Comments and suggestions by Toby Daglish, Graeme Guthrie, Martin Lally, Leigh Roberts, Wolfgang Runggaldier, Marek Rutkowski and Xun Yy Zhou are

gratefully acknowledged. All errors and omissions remain ours.

## References

- [AK06] Jeffrey D. Azzato and Jacek B. Krawczyk. SOCSol4L: An improved MATLAB<sup>®</sup> package for approximating the solution to a continuous-time stochastic optimal control problem. Technical report, School of Economics and Finance, Victoria University of Wellington, 2006. Available online at: <http://mpira.ub.uni-muenchen.de/10015/>. 3, 7
- [AK08] Jeffrey Azzato and Jacek B. Krawczyk. Applying a finite-horizon numerical optimization method to a periodic optimal control problem. *Automatica J. IFAC*, 44(6):1642–1651, 2008. DOI: doi:10.1016/j.automatica.2007.12.022. 3, 7
- [BF06] Kang Boda and Jerzy A. Filar. Time consistent dynamic risk measures. *Math. Methods Oper. Res.*, 63(1):169–186, Feb 2006. DOI: 10.1007/s00186-005-0045-1. 2
- [BFLS04] Kang Boda, J. A. Filar, Yuanlie Lin, and L. Spanjers. Stochastic target hitting time and the problem of early retirement. *IEEE Trans. Automat. Control*, 49(3):409–419, Mar 2004. DOI: 10.1109/TAC.2004.824469. 2
- [BKP04] Arjan B. Berkelaar, Roy Kouwenberg, and Thierry Post. Optimal portfolio choice under loss aversion. *Rev. Econ. Statist.*, 86(4):973–987, Nov 2004. DOI: doi:10.1162/0034653043125167. 3, 6, 7, 14
- [BLS04] Dimitris Bertsimas, Geoffry J. Lauprete, and Alexander Samarov. Shortfall as a risk measure: properties, optimization and applications. *J. Econ. Dynam. Control*, 28(7):1353–1381, Apr 2004. DOI: doi:10.1016/S0165-1889(03)00109-X. 2
- [BRU01] Erik Bogentoft, H. Edwin Romeijn, and Stanislav Uryasev. Asset/liability management for pension funds using CVaR constraints. *J. Risk Finance*, 3(1):57–71, 2001. DOI: 10.1108/eb043483. 2
- [BSL97] M. J. Brennan, E. S. Schwatz, and R. Lagnado. Strategic asset allocation. *J. Econ. Dynam. Control*, 21(8):1377–1403, Jun 1997. DOI: 10.1016/S0165-1889(97)00031-6. 2
- [dAJ04] Gustavo M. de Athayde and Renato G. Flôres Jr. Finding a maximum skewness portfolio — a general solution to three-moments portfolio choice. *J. Econ. Dynam. Control*, 28(7):1335–1352, Apr 2004. DOI: 10.1016/S0165-1889(02)00084-2. 2
- [FR75] Wendell H. Fleming and Raymond W. Rishel. *Deterministic and Stochastic Optimal Control*. Springer-Verlag, New York, 1975. 4, 5



- [FR99] R. F. Frey and W. J. Runggaldier. Risk-minimizing hedging strategies under restricted information: The case of stochastic volatility models observable only at discrete random times. *Math. Meth. Oper. Res.*, 50(2):339–350, Oct 1999. DOI: 10.1007/s001860050101. 2, 14
- [FS00] W. H. Fleming and S. J. Sheu. Risk-sensitive control and an optimal investment model. *Math. Finance*, 10(2):197–213, Apr 2000. DOI: 10.1111/1467-9965.00089. 2
- [GRS04] Nalan Gülpınar, Berç Rustem, and Reuben Settergren. Simulation and optimization approaches to scenario tree generation. *J. Econ. Dynam. Control*, 28(7):1291–1315, Apr 2004. DOI: 10.1016/S0165-1889(03)00113-1. 2
- [HRS96] M. A. Howe, B. Rustem, and M. J. P. Selby. Multi-period minimax hedging strategies. *European J. Oper. Res.*, 93(1):185–204, Aug 1996. DOI: 10.1016/0377-2217(95)00167-0. 2
- [JZ08] Hanqing Jin and Xun Yu Zhou. Behavioral portfolio selection in continuous time. *Mathematical Finance*, 18:385–426, 2008. 3
- [KD01] Harold J. Kushner and Paul Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Number 24 in Applications of Mathematics. Springer, New York, second edition, 2001. ISBN: 0-387-95139-3. 3
- [Kra90] Jacek B. Krawczyk. On variance constrained programming. *Asia-Pac. J. Oper. Res.*, 7(2):190–206, Nov 1990. 2
- [Kra01] Jacek B. Krawczyk. A Markovian approximated solution to a portfolio management problem. *ITEM.*, 1(1), 2001. Available at <http://www.item.woiz.polsl.pl/issue/journal1.htm>. 3
- [Kra03] J. B. Krawczyk. Numerical solutions to pension fund problems with realistic targets. Contributed paper at Quantitative Methods in Finance conference, Sydney, Australia, Dec 2003. 3
- [Kra05] J. B. Krawczyk. Numerical solutions to lump-sum pension problems that can yield left-skewed fund return distributions. In Christophe Deissenburg and Richard F. Hartl, editors, *Optimal Control and Dynamic Games*, number 7 in Advances in Computational Management Science, chapter 10, pages 155–176. Springer, New York, 2005. 3, 6, 7, 15
- [Kra08] Jacek B. Krawczyk. On loss-avoiding payoff distribution in a dynamic portfolio management problem. *J. Risk Finance*, 9(2):151–172, 2008. DOI: 10.1108/15265940810853913. 2, 3, 6, 7
- [Kus90] Harold J. Kushner. Numerical methods for stochastic control problems. *SIAM J. Control Optim.*, 28(5):999–1048, Sep 1990. DOI: 10.1137/0328056. 3

- [KW97] Jacek B. Krawczyk and Alistor Windsor. An approximated solution to continuous-time stochastic optimal control problems through Markov decision chains. Technical Report 9d-bis, School of Economics and Finance, Victoria University of Wellington, 1997. Available online at: <http://ideas.repec.org/p/wpa/wuwpc/9710001.html>. 3, 7
- [LZL02] Xun Li, Xun Yu Zhou, and Andrew E. B. Lim. Dynamic mean-variance portfolio selection with no-shorting constraints. *SIAM J. Control Optim.*, 40(5):1540–1555, 2002. DOI: 10.1137/S0363012900378504. 2
- [Mar52] Harry Markowitz. Portfolio selection. *J. Finance*, 7(1):77–91, Mar 1952. 2
- [Mar59] Harry M. Markowitz. *Portfolio Selection: Efficient Diversification of Investment*. Wiley, New York, 1959. 2
- [Maz04] Piera Mazzoleni. Risk measures and return performance: A critical approach. *European J. Oper. Res.*, 155(2):268–275, Jun 2004. DOI: 10.1016/S0377-2217(03)00085-7. 2
- [Mer69] Robert C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *Rev. Econ. Statist.*, 51(3):247–257, Aug 1969. 2
- [Mer71] Robert C. Merton. Optimum consumption and portfolio rules in a continuous-time model. *J. Econom. Theory*, 3(4):373–413, Dec 1971. 2, 4
- [MSS89] Randall Morck, Eduardo Schwartz, and David Stangeland. The valuation of forestry resources under stochastic prices and inventories. *J. Financ. Quant. Anal.*, 24(4):473–487, Dec 1989. 2
- [PR02] S. Pasquali and W. J. Runggaldier. Approximations of a controlled diffusion model for renewable resource exploitation. In Zhenting Hou, Jerzy A. Filar, and Anyue Chen, editors, *Markov Processes and Controlled Markov Chains*, pages 287–302. Kluwer Academic, Sep 2002. ISBN: 978-1402008030. 3
- [RU00] R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *J. Risk*, 2(3):21–41, 2000. 2
- [Sam69] Paul A. Samuelson. Lifetime portfolio selection by dynamic stochastic programming. *Rev. Econ. Statist.*, 51(3):239–246, Aug 1969. 2
- [TK92] Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertainty*, 5(4):297–323, Oct 1992. DOI: 10.1007/BF00122574. 3, 6, 8
- [Yiu04] K. F. C. Yiu. Optimal portfolios under a value-at-risk constraint. *J. Econ. Dynam. Control*, 28(7):1317–1334, Apr 2004. DOI: 10.1016/S0165-1889(03)00116-7. 2

- [ZL00] X. Y. Zhou and D. Li. Continuous-time mean-variance portfolio selection: A stochastic LQ framework. *Appl. Math. Optim.*, 42(1):19–33, 2000. DOI: [10.1007/s002450010003](https://doi.org/10.1007/s002450010003). [2](#)